User Manual

XCubeSAN Software Manual

For SANOS 4.0

Part Number: QSUM1602A
Published: May 2016
Edition: 1.0

Edition 1.0 (May 2016)

This edition applies to QSAN XCubeSAN and firmware version 1.0.0 (SANOS 4.0). Note that this document was produced based on beta code and some screens may change when it becomes generally available.

Firmware Versions

XCubeSAN Series FW 1.0.0 (SANOS 4.0)
XCubeDAS Series FW 1.0.0

Login Information

Management IP Address: DHCP
Username: admin
Password: 1234

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

Notices

This document could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the products. All features, functionality, and product specifications are subject to change without prior notice or obligation. Document contained herein is subject to change without notice.

The references in this document to non-QSAN websites are provided for convenience only. They do not in any manner serve as an endorsement of those websites. The documents at those websites are not part of the materials for QSAN products. Using those websites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products.

All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

QSAN, the QSAN logo, and qsan.com are trademarks or registered trademarks of QSAN Technology, Inc. All products and trade names used in this document are trademarks or registered trademarks of their respective companies.

Contents

Figures

Tables

Preface

About This Manual

This manual is the introduction of QSAN storage system and it aims to help users know the operations of the disk array system easily. Information contained in this manual has been reviewed for accuracy, but not for product warranty. Information and specification will be changed without further notice. For any update information, please visit www.qsan.com and your contact windows.

Before reading this manual, it assumes that you are familiar with computer skills such as hardware, storage concepts and network technology. It also assumes you have basic knowledge of Redundant Array of Independent Disks (RAID), Storage Area Network (SAN), Fibre Channel (FC), Internet SCSI (iSCSI), Serial-attached SCSI (SAS), and Serial ATA (SATA) technologies.

Technical Support

Thank you for using QSAN Technology, Inc. products; if you have any questions, please contact QSAN Support. We will reply to you as soon as possible.

- Website: http://www.qsan.com/en/contact_support.php
- Email: support@qsan.com (09:00 GMT+8 ~ 18:00 GMT+8, 09:00 GMT ~ 18:00 GMT)
- Skype ID: qsan.support (09:00 GMT+8 ~ 18:00 GMT+8, 09:00 GMT ~ 18:00 GMT)

Notes and Cautions

This manual uses the following symbols to draw attention to important safety and operational information.

NOTE:

NOTE provides helpful information, guidelines, or suggestions for performing tasks more effectively.

CAUTION:

CAUTION indicates that failure to take a specified action could result in damage to the software or hardware.

Conventions

The following table describes the typographic conventions used in this manual.

Conventions	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click **OK** button.
<Italic>	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: copy *<source-file> <target-file>*.
[ ] square brackets	Indicates optional values. Example: [a | b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a | b } indicates that you must choose either a or b.
| vertical bar	Indicates that you have a choice between two or more options or arguments.
/ Slash	Indicates all options or arguments.
underline	Indicates the default value. Example: [a | b]

1. SANOS Overview

Thank you for purchasing QSAN Technology, Inc. products. QSAN XCubeSAN Storage is a high-performance storage solutions combining outstanding performance with high reliability, availability, flexibility, and manageability. Moreover, it is designed for a variety of purposes, allowing you to perform the enterprise storage features with the web-based SANOS. This chapter provides an overview of the SANOS and includes a brief explanation of storage terminology.

1.1. Introduction to SANOS 4.0

SANOS 4.0 is the heart and soul of XCubeSAN product line. With 64 bit architecture and many patented technologies, SANOS 4.0 stays ahead of wave in competition to deliver the performance, availability, reliability, data protection, security and scalability required for enterprise SAN storage. Version 4.0 has optimized for hybrid storage to support auto tiering and SSD caching. Integrated support for VMware VAAI and Hyper-V ODX allow XCubeSAN series to perform more efficiently and run more VMs (Virtual Machine) in virtualized environment.

1.1.1. Software Architecture

In figure 1-1, it shows the SANOS 4.0 software architecture in one controller.

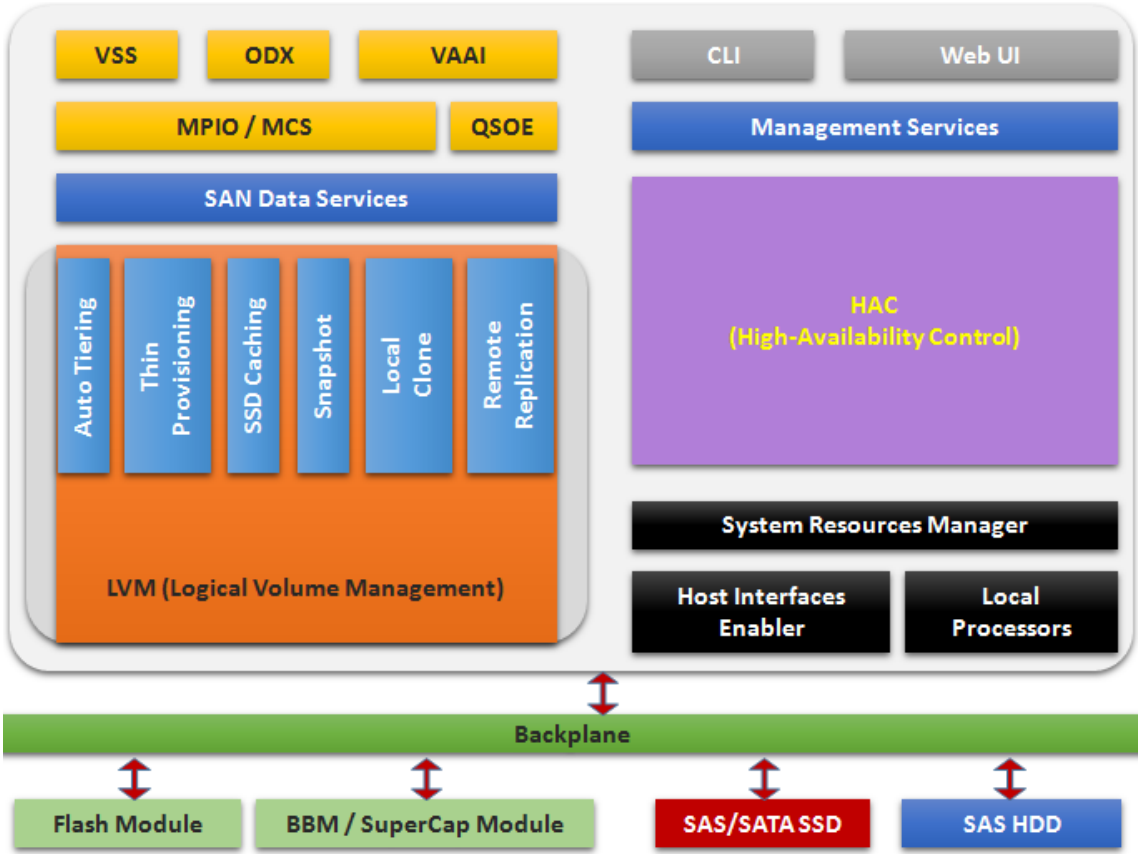

Figure 1-1 SANOS 4.0 Software Architecture

The data path is back and forth between host interfaces and disk drives via backplane. The LVM (Logical Volume Management) is the core to handle these data between them. It provides a method of allocating space on mass-storage devices that is more flexible than conventional partitioning schemes. In particular, a volume manager can concatenate, stripe together or otherwise combine partitions into larger virtual ones that administrators can re-size or move, potentially without interrupting system use. Of course, it may rely on system resources manager to arrange the processor time slot and scheduled tasks.

Base on LVM technology, it develops the features of auto tiering, thin provisioning, SSD caching. Also data backup functions of snapshot, local clone, and remote replication. Upon LVM, QSOE (QSAN Storage Offload Engine) which we are proud to develop increases the performance of IOPS and throughputs. Besides, SANOS also supports some data services such as MPIO (Multi-Path I/O), MC/S (Multi Connections per Session), Microsoft® VSS (Volume Shadow Copy Service), Microsoft ODX (Offloaded Data Transfer), VMware® VAAI (VMware vSphere® Storage APIs –. Array Integration). In the periphery, there are management service daemon to provide user interface operations, C2F (Cache to Flash) mechanism includes BBM (Battery Backup Module) / SCM (SuperCap Module) and flash module to protect cache data from power shortage. Enterprise storage features are all in a box.

1.1.2. High Availability Architecture

In dual-controller architecture, HAC (High Availability Control) is an important module to communicate with both controllers. In figure 1-2, it shows the whole picture of high availability software architecture. HAC has heartbeat mechanism to detect whether the other controller is alive. Cache mirror synchronizes the memory data of both controllers including the status and storage configuration.

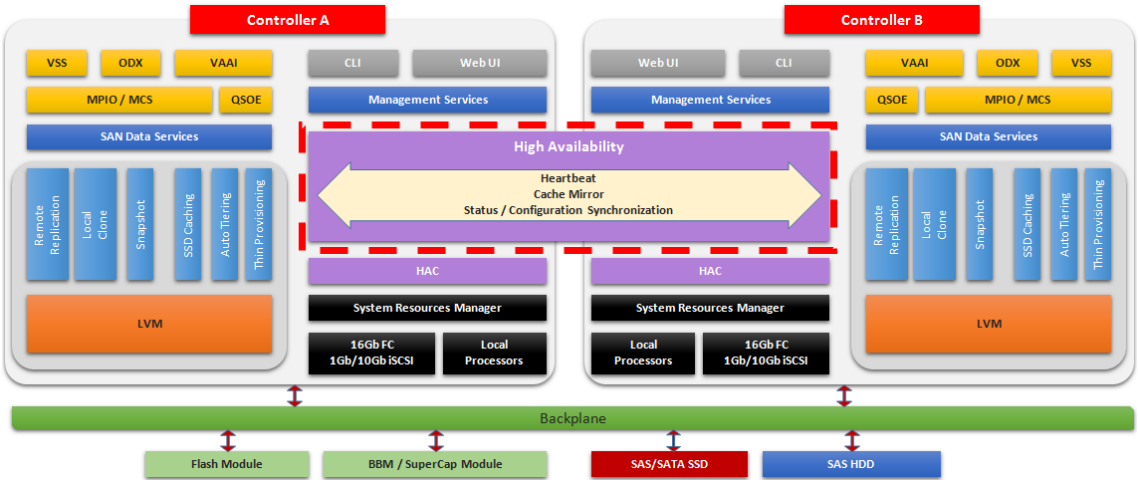

Figure 1-2 SANOS 4.0 for High Availability

Although cache mirror always synchronizes the memory data of both controllers, it will take over very soon when single one controller fails. In addition, zero system downtime for the dual-controller system takes place during firmware upgrading period.

1.1.3. Software Features

SANOS 4.0 features the following benefits:

- Brings enterprise storage technology for the administrators to do the data plans for different circumstances
- Enhanced QSOE 2.0 to improve the system performance higher
- Auto tiering feature is lower the total cost of ownership, increase performance by intelligently managing data placement, and get the best Cost/Performance combination according to application
- SSD caching feature is good to increase the applications access performances
- Thin provisioning uses virtualization technology for over-scribed provisioning, dynamic size and allocate on demand, just-in-time and just-enough
- Includes built-in data backup services for protecting the data from most unpredictable accidents
- Easy setup and management
- Simple integration into the server environment and application
- Virtualization improvements

In summary, the XCubeSAN system uses latest SANOS 4.0 to ship the overwhelming performance and advanced enterprise features.

1.2. Software Terminology

In this section, we introduce the terms that are used for the XCubeSAN storage system throughout this manual.

1.2.1. RAID

RAID is the abbreviation of Redundant Array of Independent Disks. The basic idea of RAID is to combine multiple drives together to form one large logical drive. This RAID drive obtains performance, capacity and reliability than a single drive. The operating system detects the RAID drive as a single storage device.

The disk drives in storage are referred to as members of the array. Each array has a RAID level. RAID levels provide different degrees of redundancy and performance. They also have different restrictions regarding the number of members in the array.

The following RAID levels are available:

- RAID 0 (Striping, no redundancy)

  RAID 0 consists of striping, without mirroring or parity. It has more than two disks in parallel, as a large-capacity disk. The capacity of a RAID 0 volume is the sum of the capacities of the disks. There is no added redundancy for handling disk failures. Thus, failure of one disk causes the loss of the entire RAID 0 volume. Striping distributes the contents roughly equally among all member disks, which makes concurrent read or write operations on the

multiple disks and results in performance improvements. The concurrent operations make the throughput of most read and write operations equal to the throughput of one disk multiplied by the number of disks. Increased throughput is the big benefit of RAID 0, at the cost of increased vulnerability to drive failures.

- RAID 1 (Mirroring between two disks)

  RAID 1 consists of data mirroring, without parity or striping. Data is written identically to two drives, thereby producing a mirrored set of drives. Thus, any read request can be serviced by any member drive. Write throughput is always slower because every drive must be updated, and the slowest drive limits the write performance. The array continues to operate as long as at least one drive is functioning.

- N-way Mirror (Mirroring between N disks)

  It's an extension of RAID 1 level. Data is written identically to N drives, thereby producing a N-way mirrored set of drives

- RAID 3 (Striping, can survive one disk drive fault, with parity on a dedicated disk drive)

  RAID 3 consists of byte-level striping with dedicated parity. All disk spindle rotation is synchronized and data is striped such that each sequential byte is on a different drive. Parity is calculated across corresponding bytes and stored on a dedicated parity drive. The data disperses on a different hard drive, even if you want to read short information, it may need all the hard drive to work. So this is more suitable for large amounts of reading data.

- RAID 5 (Striping, can survive one disk drive fault, with interspersed parity over the member disk drives)

  RAID 5 consists of block-level striping with distributed parity. It requires at least three disk drives. Upon failure of a single drive, subsequent reads can be calculated from the distributed parity such that no data is lost. RAID 5 is seriously affected by the general trends regarding array rebuild time and the chance of disk drive failure during rebuild. Rebuilding an array requires reading all data from all disks, opening a chance for a second disk drive failure and the loss of the entire array.

- RAID 6 (Striping, can survive two disk drive faults, with interspersed parity over the member disk drives)

  RAID 6 consists of block-level striping with double distributed parity. It requires a minimum of four disks. Double parity provides fault tolerance up to two failed disk drives. This makes larger RAID groups more practical, especially for high-availability systems, as large-capacity drives take longer to restore. As with RAID 5, a single drive failure results in reduced performance of the entire array until the failed drive has been replaced. With a RAID 6 array, using drives from multiple sources and manufacturers, it is possible to mitigate most of the problems associated with RAID 5. The larger the drive capacities and the larger the array size, the more important it becomes to choose RAID 6 instead of RAID 5.

- RAID 0+1 (RAID 1 on top of RAID 0)

  RAID 0+1 creates a second striped set to mirror a primary striped set. The array continues to operate with one or more drives failed in the same mirror set, but if drives fail on both sides of the mirror, the data on the RAID system is lost.

- RAID 10 (RAID 0 on top of RAID 1)

  RAID 10 creates a striped set from a series of mirrored disk drives. The array can sustain multiple drive losses so long as no mirror loses all its drives.

- RAID 30 (RAID 3 on top of RAID 0)

  RAID 30 is the combination of RAID 3 and RAID 0, do RAID 3 first, further RAID 0. It is composed of multiple sets of RAID 3 stripe access to each other. Because RAID 30 is based on RAID 3 which requires at least three disk drives, therefore RAID 30 is constituted a plurality RAID 3, at least six disk drives. RAID 30 can still operate when appearing a damaged disk drive in disk group of RAID 3. But if any one group of RAID 3 appears two or two or more disk drives damaged, the entire RAID 30 will fail.

- RAID 50 (RAID 5 on top of RAID 0)

  RAID 50 is the combination of RAID 5 and RAID 0, do RAID 5 first, further RAID 0. The concept is the same as RAID 30. RAID 50 is requires at least six disk drives. Since RAID 50 constitutes stripe of multiple disk group of RAID 5, it has higher performance than RAID 5, but capacity utilization is lower than RAID5.

- RAID 60 (RAID 6 on top of RAID 0)

  RAID 60 is the combination of RAID 6 and RAID 0, do RAID 6 first, further RAID 0. In other words, it accesses stripes for more than two groups of RAID 6. RAID 6 needs to have at least four disk drives, so the minimum requirement of RAID 60 is eight disk drives.

  RAID 60 can tolerate maximum two damage disk drives in any disk group of RAID 6, while it still maintain the operating; but as long as it damages three drives in any group of RAID 6, the entire RAID 60 will fail. Of course, the probability of this case is quite low.

  Compared to a simple RAID 6, RAID 60 binds stripes through multiple of RAID 6. Therefore it is has higher performance. However, high threshold usage and the capacity utilization rate are major problems.

Summary of the RAID levels are on the following.

Table 1-1 RAID Levels

Type	Description	Min. No. of Drives
RAID 0	Disk striping.	1
RAID 1	Disk mirroring over two disks.	2
N-way mirror	Extension to RAID 1 level. It has N copies of the disk.	N
RAID 3	Striping with parity on the dedicated disk.	3
RAID 5	Striping with interspersed parity over the member disks.	3
RAID 6	2-dimensional parity protection over the member disks.	4
RAID 0+1	Mirroring of the member RAID 0 volumes.	4
RAID 10	Striping over the member RAID 1 volumes.	4
RAID 30	Striping over the member RAID 3 volumes.	6

RAID 50	Striping over the member RAID 5 volumes.	6
RAID 60	Striping over the member RAID 6 volumes.	8

SAN OS supports hot spare drives. When a member disk drive in array fails, the system automatically replaces the failed member with a hot spare drive and rebuilds the array to restore its redundancy. Candidate and spare drives can be manually exchanged with array members. For more information, see Rebuild section in Advanced Pool Administration chapter.

1.2.2. Pools

A storage pool is a collection of disk drives SAN OS defines three pool types. They are listed on the following.

- Thick or fat provisioning pool, we use the term of thick provisioning pool on the following.
- Thin provisioning pool
- Auto tiering pool

We describe thick provisioning pool here, and introduce thin provisioning and auto tiering pool in next Software Features section. A storage pool is grouped to provide capacity for volumes. Volumes are then allocated out of the storage pool and are mapped to LUN which can be accessed by a host system. The storage architecture of a thick provisioning pool is on the following.

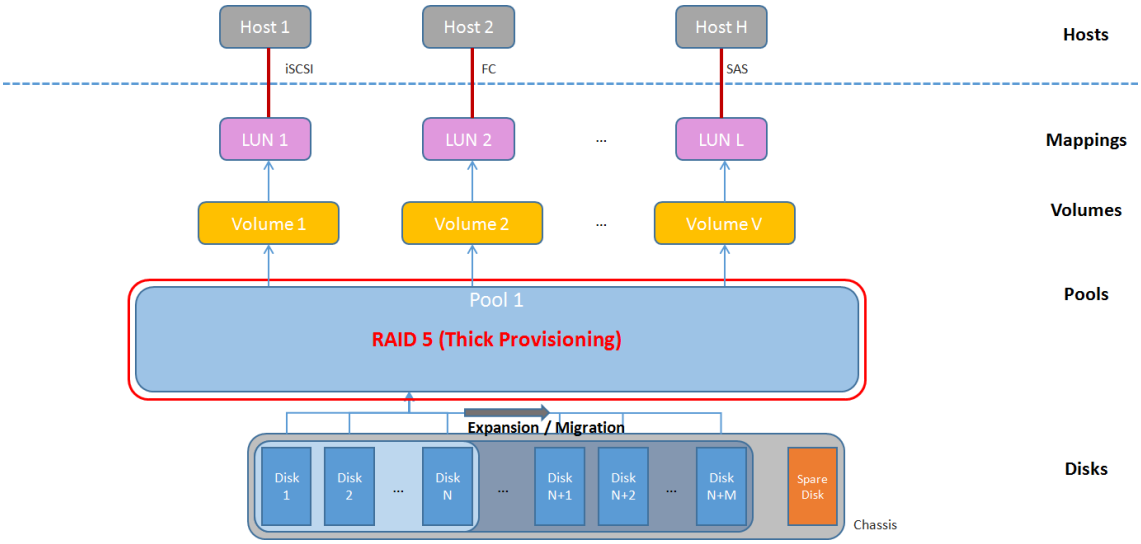

Figure 1-3 Storage Architecture of Thick Provisioning Pool

A thick provisioning pool contains up to 64 disk drives. For more information about pool operation, please see Configuring Storage Pools section in Storage Management chapter.

Table 1-2 Thick Provisioning Pool Parameters

Item	Value

Maximum Disk Drive Quantity in a Pool (Include Dedicated Spares)	64
Maximum Pool Quantity per System	64
Maximum Dedicated Spare Quantity in a Pool	8

Disk drives can be added to a thick provisioning pool at any time to increase the capacity of the pool. Or disk drives can be added to upgrade the RAID level in a thick provisioning pool. These operations are called migrate. For example, it can migrate RAID 5 to RAID6 by adding disk drives because RAID 6 needs one more disk space for parity.

1.2.3. Volumes

A volume is a logical disk that is presented to a host system. The capacity is provided by pool. Each pool can be divided into several volumes. The volumes in one pool share the same RAID level, but may have different volume capacity.

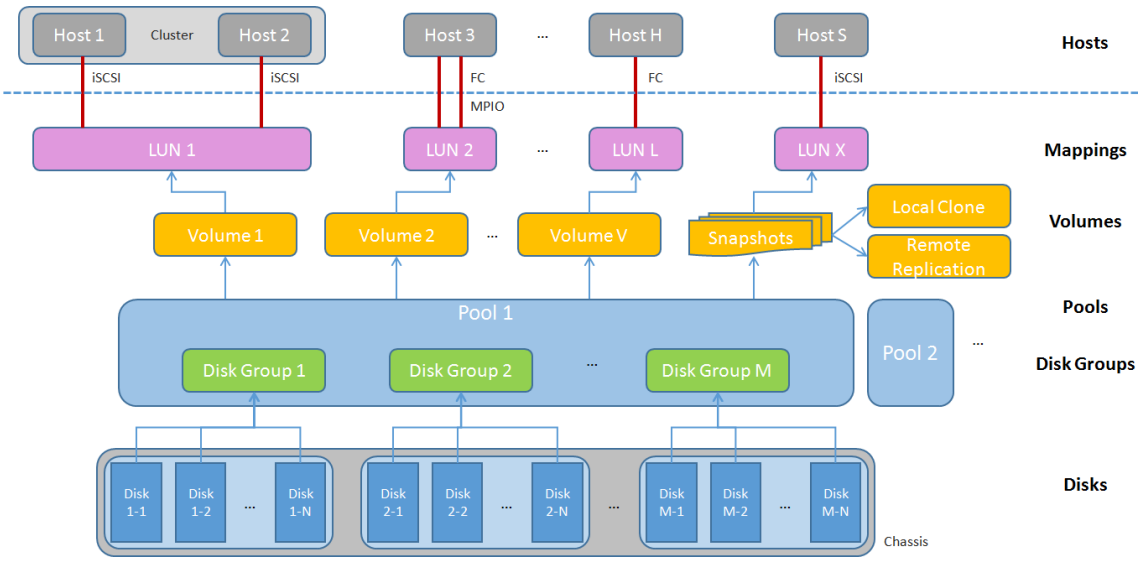

Figure 1-4 Volume and LUN Relationship in Storage Architecture

A pool contains up to 96 volumes and a system can contain up to 4,096 volumes including snapshot volumes. For more information about volume operation, please see Configuring Volume section in Storage Management chapter.

Table 1-3 Volumes Parameters

Item	Value
Maximum Volume Quantity in a Pool	96
Maximum Volume Quantity Per System (include Snapshot Volumes)	4,096
Maximum Host Number Per Volume	16
Provisioning Granularity	1 GB

Maximum Thin Volume Capacity	128 TB

Volume is a basic unit for data backup. Base on the volume, SAN OS provides snapshot, local clone, and remote replication functions. These terms will be described in next Software Features section.

1.2.4. LUN

LUN (Logical Unit Number) is a number used to identify a logical unit, which is a device addressed by the SCSI protocol or Storage Area Network protocols which encapsulate SCSI, such as Fibre Channel or iSCSI.

The LUN is not the only way to identify a logical unit. There is also the SCSI Device ID, which identifies a logical unit uniquely in the world. Labels or serial numbers stored in a logical unit's storage volume often serve to identify the logical unit. However, the LUN is the only way for an initiator to address a command to a particular logical unit, so initiators often create, via a discovery process, a mapping table of LUN to other identifiers.

There is one LUN which is required to exist in every target: LUN 0. The logical unit with LUN 0 is special in that it must implement a few specific commands, most notably Report LUNs, which is how an initiator can find out all the other LUNs in the target. But LUN 0 need not provide any other services, such as a storage volume. Many SCSI targets contain only one logical unit (so its LUN is necessarily 0). Others have a small number of logical units that correspond to separate physical devices and have fixed LUNs. A large storage system may have up to thousands of logical units, defined logically, by administrative command, and the administrator may choose the LUN or the system may choose it.

Table 1-4 LUN Parameters

Item	Value
Maximum Quantity of LUN	4,096

In Figure 1-2, two hosts with cluster that access the same disk volume may use the same LUN. So does the host with MPIO. For more information about LUN operation, please see LUN Mappings section in Storage Management chapter.

1.2.5. iSCSI

iSCSI (Internet SCSI) is a protocol which encapsulates SCSI (Small Computer System Interface) commands and data in TCP/IP packets for linking storage devices with servers over common IP infrastructures. iSCSI provides high performance SANs over standard IP networks like LAN, WAN or the Internet.

IP SANs are true SANs (Storage Area Networks) which allow several servers to attach to an infinite number of storage volumes by using iSCSI over TCP/IP networks. IP SANs can scale the storage capacity with any type and brand of storage system. In addition, it can be used by any type of network (Ethernet, Fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet) and combination of operating systems (Microsoft Windows, Linux, Solaris, Mac, etc.) within the SAN

network. IP-SANs also include mechanisms for security, data replication, multi-path and high availability.

Storage protocol, such as iSCSI, has “two ends” in the connection. These ends are initiator and target. In iSCSI, we call them iSCSI initiator and iSCSI target. The iSCSI initiator requests or initiates any iSCSI communication. It requests all SCSI operations like read or write. An initiator is usually located on the host side (either an iSCSI HBA or iSCSI SW initiator).

The target is the storage device itself or an appliance which controls and serves volumes. The target is the device which performs SCSI command or bridge to an attached storage device.

Each iSCSI node, that is, an initiator or target, has a unique IQN, which can have a size of up to 255 bytes. The IQN is formed according to the rules that were adopted for Internet nodes. The IQNs can be abbreviated by using a descriptive name, which is known as an alias. An alias can be assigned to an initiator or a target. For more information, please see Configure for iSCSI Connectivity section in Host Configuration chapter.

1.2.6. Fibre Channel

Fibre channel started use primarily in the supercomputer field, but has become the standard connection type for storage area networks (SAN) in enterprise storage.

The target is the storage device itself or an appliance which controls and serves volumes or virtual volumes. The target is the device which performs SCSI commands or bridges to an attached storage device.

Fibre Channel is the traditional method used for data center storage connectivity. The XCubeSAN supports FC connectivity at speeds of 4, 8, and 16 Gbps. Fibre Channel Protocol is used to encapsulate SCSI commands over the FC network. Each device on the network has a unique 64 bit World Wide Port Name. For more information, please see Configure for Fibre Channel Connectivity section in Host Configuration chapter.

1.2.7. SAS

Serial-attached SCSI offers advantages over older parallel technologies. The cables are thinner, and the connectors are less bulky. Serial data transfer allows the use of longer cables than parallel data transfer.

The target is the storage device itself or an appliance which controls and serves volumes or virtual volumes. The target is the device which performs SCSI command or bridge to an attached storage device.

XCubeSAN supports direct SAS host attachment providing easy-to-use, affordable storage needs. Each SAS port device has a worldwide unique 64-bit SAS address and operates at 12 Gbps.

1.3. Software Features

In this section, we describe the software features of SAN OS.

1.3.1. Thin Provisioning

Thin provisioning, in a shared-storage environment, provides a method for optimizing utilization of available storage. It relies on on-demand allocation of blocks of data versus the traditional method of allocating all the blocks up front. This methodology eliminates almost all whitespace which helps avoid the poor utilization rates, that occur in the traditional storage allocation method where large pools of storage capacity are allocated to individual servers but remain unused (not written to). This traditional model is often called "fat" or "thick" provisioning.

With thin provisioning, storage capacity utilization efficiency can be automatically driven up towards 100% with very little administrative overhead. Organizations can purchase less storage capacity up front, defer storage capacity upgrades in line with actual business usage, and save the operating costs (electricity and floor space) associated with keeping unused disk capacity spinning.

A thin provisioning pool is a collection of disk groups which contain disk drives. The same, a storage pool is grouped to provide capacity for volumes. Volumes are then allocated out of the storage pool and are mapped to LUN which can be accessed by a host system. The storage architecture of a thin provisioning pool is on the following.

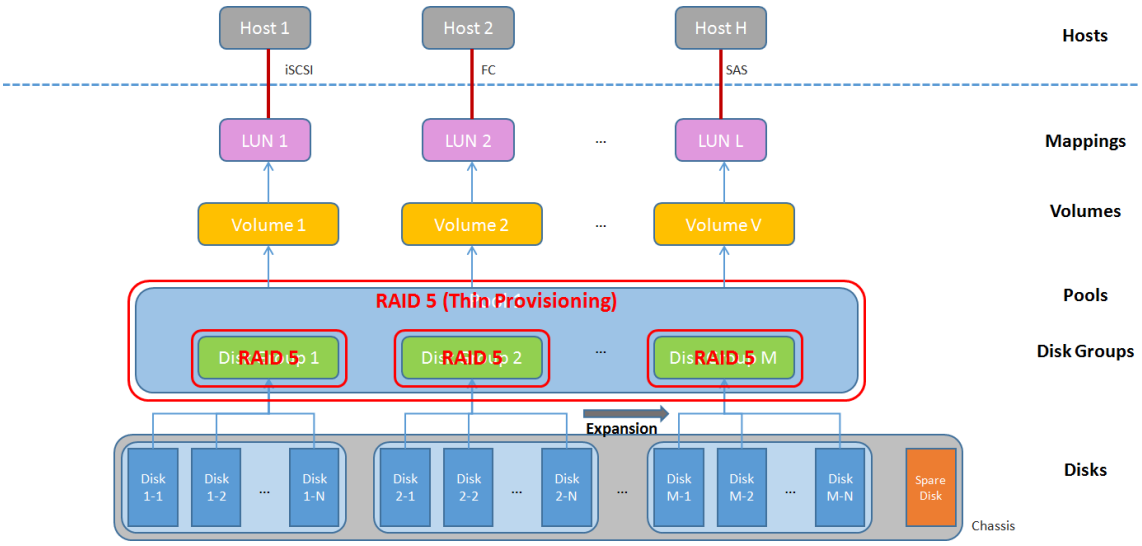

Figure 1-5 Storage Architecture of Thin Provisioning Pool

Disk groups which contain disk drives can be added to a thin provisioning pool at any time to increase the capacity of the pool. For simplifying usage and better performance, every disk group must have the same quantity of disk drives. A thin provisioning pool can have up to 32 disk groups with each disk group containing up to 8 disk drives. The maximum capacity of each disk group is 64TB. So the maximum capacity in a system is 256TB. For more information, please see Thin Provisioning chapter.

Table 1-5 Thin Provisioning Pool Parameters

Item	Value
Maximum Disk Group Quantity in a Pool	32

Maximum Disk Drive Quantity in a Disk Group	8
Maximum Disk Drive Quantity in a Pool	256 (= 32 x 8)
Maximum Capacity of a Disk Group	64 TB
Maximum Thin Provisioning Pool Capacity per System	256 TB

NOTE:

Thick Provisioning vs. Thin Provisioning: The efficiency of thin or thick provisioning is a function of the use case, not of the technology. Thick provisioning is typically more efficient when the amount of resource used very closely approximates to the amount of resource allocated. Thin provisioning offers more efficiency where the amount of resource used is much smaller than allocated, so that the benefit of providing only the resource needed exceeds the cost of the virtualization technology used.

1.3.2. Auto Tiering

Auto Tiering is the automated progression or demotion of data across different tiers (types) of storage devices and media. The movement of data takes place in an automated way with the help of software and is assigned to the related media according to performance and capacity requirements. It also includes the ability to define rules and policies that dictate if and when data can be moved between the tiers, and in many cases provides the ability to pin data to tiers permanently or for specific periods of time.

New created auto tiering pool is based on thin provisioning technology. Each tier is a disk group. Every tier must have the same disk quantity when creating a tiering pool as a base unit of disk group. Maximum quantity of disk in a disk group is 8. For example, you can select 4 SSDs, 4 SAS HDDs, and 4 NL-SAS HDDs. The storage architecture of an auto tiering pool is on the following.

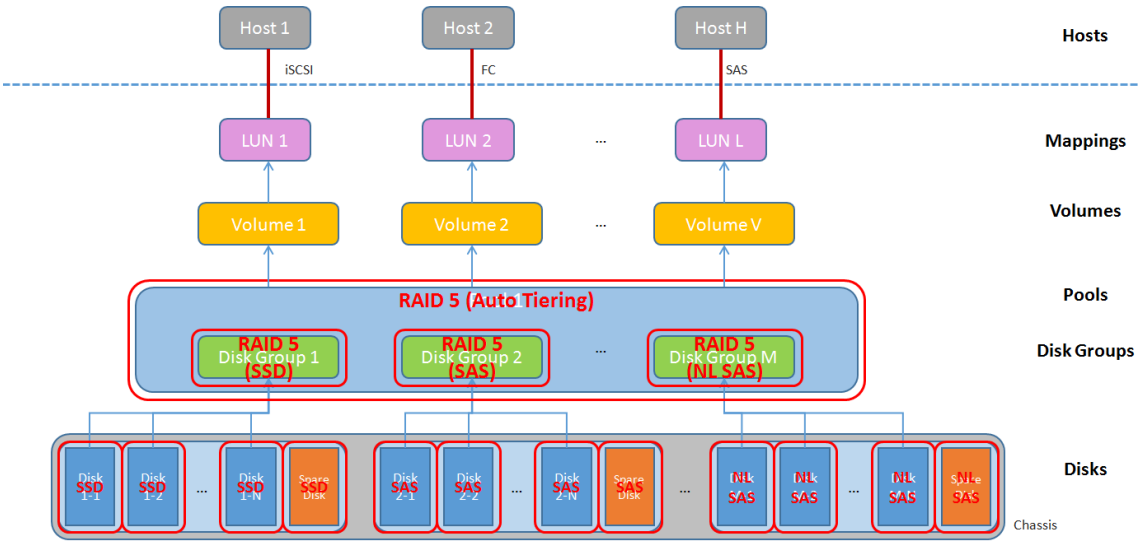

Figure 1-6 Storage Architecture of Auto Tiering Pool

Any tier (disk group) which contains disk drives can be added to an auto tiering at any time to increase the capacity of the pool. The same rules as thin provisioning and better performance consideration, every new added tier (disk group) must have the same quantity of disk drives. A thin provisioning pool can have up to 32 disk groups with each disk group containing up to 8 disk drives. For tiering, the quantity of lowest tier (disk group) must be larger or equal to the higher ones. The maximum capacity of each disk group is 64TB. So the maximum capacity in a system is 256TB.

Table 1-6 Auto Tiering Pool Parameters

Item	Value
Maximum Tiers	4
Maximum Tier (Disk Group) Quantity in a Pool	32
Maximum Disk Drive Quantity in a Tier (Disk Group)	8
Maximum Disk Drive Quantity in a Pool	256 (= 32 x 8)
Maximum Capacity of a Tier (Disk Group)	64 TB
Maximum Thin Provisioning Pool Capacity per System	256 TB

In our advanced implementations, auto tiering includes the ability to select policies that dictate if and when data can be moved between the tiers, and in many cases provides the ability to pin data to tiers permanently or for specific periods of time. For more information, please see Auto Tiering chapter.

1.3.3. SSD Caching

Smart Response Technology (also known as SSD caching) allows a solid-state drive to function as cache for a hard disk drive (HDD). It is secondary cache that improves performance by keeping frequently accessed data on SSDs where they are read far more quickly than from HDDs. One or more SSDs can be assigned to a single volume to provide the SSD cache. Note that the capacity allocated to the cache is not available for regular data storage. Currently, the maximum SSD cache capacity in a system is 3.6TB.

Table 1-7 SSD Caching Parameters

Item	Value
Maximum SSD Capacity Support per System	3.6 TB

For more information, please see SSD Caching chapter.

NOTE:

Auto Tiering vs. SSD Caching: A key difference between tiering and caching is that tiering moves data to SSD instead of copying it, both from slower storage to faster storage and vice versa. Caching is essentially a one-way transaction. When the cache is done with the data it was accelerating it

simply nullifies it instead of copying it back to HDD. The important difference between moves and copies is that a cache does not need to have the redundancy that tiering does. Tiering stores the only copy of data for potentially a considerable period of time so it needs to have full data redundancy like RAID or mirroring.

Total storage capacity in auto tiering is a sum of all individual tier capacities whereas in caching, the cache capacity does not add to the overall slower storage capacity. This is one of the key differences. In addition, SSD caching affects rapider than auto tiering because auto tiering will be affected by relocation the data in a period of time. So SSD caching warm-up timeframe is usually minutes/hours whereas tiering warm-up is usually in days.

1.3.4. Snapshot

A volume snapshot is the state of a system at a particular point in time. It uses the copy-on-write on the volumes to copy the changed blocks, just before they are going to be overwritten. These snapshots have backup data at a point in time. When disaster takes place, (e.g. virus attack, data corruption, human errors and so on), the important data can be restored through the state of any time which was previously captured in case.

Besides snapshot and rollback, SAN OS provides more than these basic functions. The data of taken snapshots can be reviewed by user via mapping a LUN in advance, the benefit of this could be see whether the snapshot is taken successfully, and which snapshot will be rollback if disaster happens. In addition, we provide writable snapshot which extends the copy-on-write approach by disassociating any blocks modified within the snapshot from their "parent" blocks in the original volume for development and testing. And also schedule snapshots to back up data automatically.

The maximum snapshot quantity per volume is 64, and the maximum volume quantity for snapshot is also 64. So a system can have 4,096 snapshots.

Table 1-8 Snapshot Parameters

Item	Value
Maximum Snapshot Quantity per Volume	64
Maximum Volume Quantity for Snapshot	64
Maximum Snapshot Quantity per System	4,096 (= 64 x 64)
Maximum Snapshot Space Capacity of a Thin Provisioning Volume	128 TB

For more information, please see Managing Snapshots section in Data Backup chapter.

1.3.5. Local Clone

Local clone function has another physical data copy as the original volume. It's an asynchronous local clone base on snapshot technology. At the beginning, copy all data from the source volume to target. It is also called full copy. Afterwards, take a snapshot on source volume and then copy

delta data to perform the incremental copy. In addition, we also provide schedule local clone function which will execute a clone task automatically.

For more information, please see Managing Local Clones section in Data Backup chapter.

1.3.6. Remote Replication

Remote replication function prevents primary site failure by replicating data to the remote sites. It's convenient to implement over Ethernet. It's an asynchronous replication base on snapshot technology. Same as local clone, remote replication must have full copy and then incremental copy. Besides schedule remote replication, we also provide traffic-shaping to manage network bandwidth used by replication on rush hours, to leave the network bandwidth for critical transactions to use.

Although full copy through limited network bandwidth spends very long time, we provide local volume cloning and then covert to remote replication. This can reduce lots of time required at the initialization of the remote replication

For more information, please see Managing Remote Replication section in Data Backup chapter.

1.3.7. Fast Rebuild

When executing rebuild, the Fast Rebuild feature skips any partition of the volume where no write changes have occurred, it will focus only on the parts that have changed. This mechanism may reduce the amount of time needed for the rebuild task. It also reduces the risk of RAID failure cause of reducing the time required for the RAID status from degraded mode to healthy. At the same time, it frees up CPU resources more quickly to be available for other I/O and demands. For more information, please refer to Fast Rebuild White Paper. The document is available at this website:

https://www.qsan.com/en/download.php?cid=&keywords=Fast+Rebuild&act=query

2. Initial Configuration

This chapter provides a description of the initial configuration steps for the system.

2.1. Configuration Planning

In this section, we introduce the configuration planning; we recommend that they are prepared before using.

2.1.1. Configuration Planning Check List

First, we list the check list on the following table and take examples. You can take this list to prepare to configure in advance. In next sections, we will describe more information of these items.

System Settings

Table 2-1 System Settings Check List

Item	Example	Planning
System Name	XCubeSAN	
Admin Password	1234	
NTP	pool.ntp.org	
Management IP Address at Controller 1	IP: 192.168.1.234 SM: 255.255.255.0 GW: 192.168.1.254	IP: SM: GW:
Management IP Address at Controller 2 (Option)	IP: 192.168.1.235 SM: 255.255.255.0 GW: 192.168.1.254	IP: SM: GW:
DNS	8.8.8.8	
Mail-from Address	admin@company.com	
Mail-to Address 1	user1@company.com	
Entity Name	Iqn.2016-05.com.company	
SMTP Server	smtp@company.com	
Syslog Server (Option)	syslog.company.com	
SNMP Trap Address 1 (Option)	snmp.company.com	

iSCSI Ports Configuration

Table 2-2 iSCSI Ports Check List

Item	Example	Planning

iSCSI Port Onboard LAN1 IP Address at Controller 1	IP: 10.10.1.1 SM: 255.255.255.0 GW: 10.10.1.254 Default Gateway	IP: SM: GW:
iSCSI Port Onboard LAN2 IP Address at Controller 1	IP: 10.10.2.1 SM: 255.255.255.0 GW: 10.10.2.254	IP: SM: GW:
iSCSI Port Onboard LAN1 IP Address at Controller 2	IP: 10.20.1.1 SM: 255.255.255.0 GW: 10.20.1.254 Default Gateway	IP: SM: GW:
iSCSI Port Onboard LAN2 IP Address at Controller 2	IP: 10.20.2.1 SM: 255.255.255.0 GW: 10.20.2.254	IP: SM: GW:
Entity Name	Iqn.2004-08.com.qsan	
iSNS IP Address (Option)	10.1.1.1	
CHAP Accounts (Option)	Username: chap1 Secret: 123456789012	Username: Secret:

Fibre Channel Ports Configuration

Table 2-3 Fibre Channel Ports Check List

Item	Example	Planning
Fibre Channel Slot1 at Controller 1	Link Speed: Automatic Topology: Point-to-Point	Link Speed: Topology:
Fibre Channel Slot1 at Controller 2	Link Speed: Automatic Topology: Point-to-Point	Link Speed: Topology:

Pools Configuration

Table 2-4 Pools Check List

Item	Example	Planning
Pool Name	P1	
Pool Type	Auto Tiering	
Disks	SSD: 4x 100GB SAS: 4x 600GB NL-SAS: 4x 4,000GB	SSD: SAS: NL-SAS:
RAID Level	RAID 5	
Estimate Capacity	14,100GB (= 100GB x 3 + 600GB x 3 + 4,000GB x 3)	

Volumes Configuration

Table 2-5 Volumes Check List

Item	Example	Planning
Volume Name	V1-AP1	
Capacity	8,000GB	
Snapshot Space	1,600GB	
Volume Type	For RAID Usage	
SSD Caching	SSD: 1x 400GB	SSD:

Item	Example	Planning
Volume Name	V2-AP2	
Capacity	2,000GB	
Volume Type	For Backup Usage	
Snapshot Space	400GB	
SSD Caching	SSD: None	SSD:

LUN Mappings Configuration

Table 2-6 LUN Mappings Check List

Item	Example	Planning
Volume Name	V1-AP1	
Protocol	iSCSI	
Target	0	
LUN	LUN 0	

Item	Example	Planning
Volume Name	V2-AP2	
Protocol	FCP	
Target	2000001378123456	
LUN	LUN 0	

2.1.2. Hardware Configuration Planning

After verifying that you have the storage and all of the peripheral equipment that you purchased, it is important to carry out proper planning before the actual physical installation. Install the hardware as described in XCubeSAN Hardware Manual. The document is available at the website:

https://www.qsan.com/en/download.php?cid=&keywords=XCubeSAN+Hardware+Manual&act=query

Please make sure all the cables (including power, Ethernet, fibre channel, and SAS cables) are connected properly.

2.1.3. Management IP Configuration Planning

Prepare a storage system configuration plan by the network administrator. The plan should include network information for the management port. If using static IP addresses, please prepare a list of the static IP addresses, the subnet mask, and the default gateway. If using dual-controller, please prepare double settings. The management IP address is associated with one of the controller. Should this controller go offline (planned or unplanned), the management IP address fails over to the other controller.

2.1.4. iSCSI Configuration Planning

Although SAN is a private network environment, IP addresses of iSCSI data ports should be planned, too. Optionally, CHAP security information, including CHAP username and secret should be prepared.

For an iSCSI dual-controller system, recommend to install an iSNS server on the same storage area network. For better usage of failover, the host must logon the target twice (both controller 1 and controller 2), and then the MPIO should setup automatically. More advanced, all the connections among hosts (with clustering), switches, and the dual-controller are recommended as redundant as below.

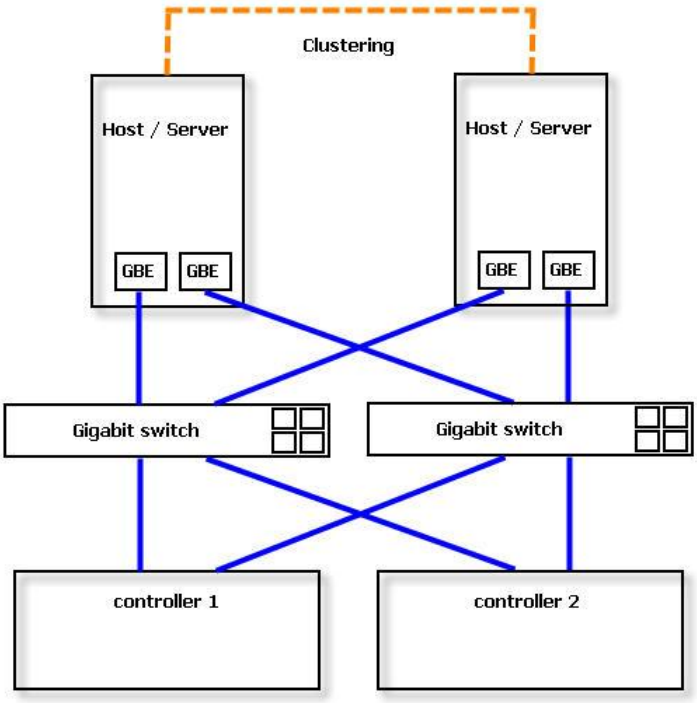

Figure 2-1 Dual-controller Topology

2.1.5. Fibre Channel Configuration Planning

SAN can be used with a fibre channel host configuration. User should consider the fibre channel typology, either direct-attach to HBA or connection with fibre channel switch. User should know the WWNN / WWPN of HBA or switch for fibre channel access control.

2.1.6. Storage Configuration Planning

This is the most important part to use the SAN storage. In chapter 1 Overview, we describe thick provisioning, thin provisioning, SSD Caching, auto tiering, and their comparison. You should know their benefits. Depend on your application, estimated volume capacity, and disk failure risk; you should have a well storage plan. Complete all storage configuration planning tasks before you configure the system.

Planning the Disks

To plan the disk drives, determine:

- The cost/performance ratio of the disk type.
- Estimate the total capacity of the system.
- Whether the disk drives are SSDs.
- Quantity of spare disks.

Planning the Pools

Necessary activities:

- The pool name.
- According to the pool and disk characteristics, such as total capacity, budget, performance, and reliability, determine a pool type (thick, thin provisioning, or auto tiering) and a RAID level of storage systems to use. For example, create a pool which is grouped by RAID 10 for performance and RAID 5 in another storage pool for achieve.
- Plan to extent the capacity of the storage pool in the future.

Planning the Volumes

A volume is a member of the pool. Before you create a volume, determine:

- The volume name.
- The volume capacity. If you want to use snapshot, reserve some snapshot space from pool. We suggest reserving 20% of the volume capacity at least.
- Whether the cache mode for the volume is either write-back or write-through. The default is Enable Cache Mode (Write-back cache).
- Plan the volume capacity for local clone or remote replication.

2.1.7. Miscellaneous Configuration Planning

During the initial setup of the system, the initial configuration wizard asks for various information to be prepared during the installation process. The information in the following checklist is helpful to have before the initial configuration is performed.

- System name allows maximum 32 characters.
- Administrator password allows maximum 16 characters.
- The date and time can be manually entered, but to keep the clock synchronized, recommend using a NTP (network time protocol) service.
- DNS (Domain Name Service) server address provides a means to translate FQDN (Fully Qualified Domain Name) to IP address. Some notification services need DNS setting.

After the initial configuration, you might want to configure the notification.

- SMTP server and mail address to direct alerts the notifications to user.
- Syslog server address to log the events.
- SNMP (Simple Network Management Protocol) trap to send system event logs to SNMP trap agent.

2.2. System Management

The web user interface is used to configure, manage, and troubleshoot the storage system. It is used to configure RAID arrays and volumes, map volumes to a LUN, replace and rebuild failed disk drives, and expand the volumes.

It allows for troubleshooting and management tasks, such as checking the status of the storage server components, updating the firmware, and monitor storage server.

The web user interface also offers advanced functions, such as thin provisioning, auto tiering, SSD caching, snapshots, local clones, and remote replication.

This section describes system management using the web user interface, serial console, and SSH (secure shell) remote access.

2.2.1. Web User Interface

A web browser is used for graphic user interface access. It supports most common web browsers.

Supported Browsers

The approved client browsers include:

- Google Chrome 45 and newer (Windows, Mac)
- Firefox 39 and newer (Windows)
- Internet Explorer 10 and 11 (Windows)
- Safari 5.1 and newer (Mac)

Be sure to connect the LAN cable to the management port of the system.

Accessing the Management Web User Interface

The default management IP address is DHCP (Dynamic Host Configuration Protocol), we will describe how to find the IP address later in First-time Setup section. If you have already known the management IP address, please enter the IP address into your browser to display the authentication screen.

- http(s)://<*Host name* or *IP Address*>

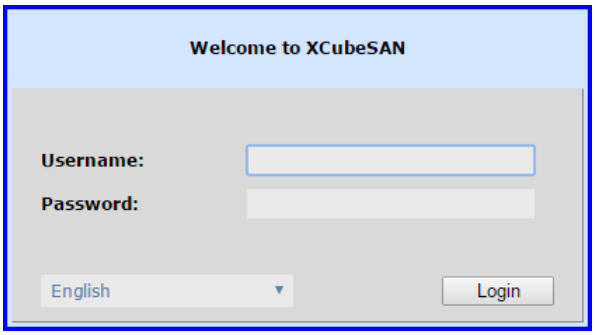

Figure 2-2 Login Page of web UI

To access the web user interface, you have to enter a username and password.

- Username: admin
- Password: <*Your Password*> (Default: 1234)

2.3. First-time Setup

This section describes how to perform a first-time system setup. Although the default management IP address is DHCP, we provide QFinder and QCentral utilities to search QSAN products on LAN. They are also highly portable software. You can search the management IP address via these utilities, and then login to execute the initial configuration wizard.

The latest utilities can be downloaded at the website:

- QFinder Utility:

  https://www.qsan.com/en/download.php?cid=&keywords=QFinder+utility&act=query

- QCentral Utility:

  https://www.qsan.com/en/download.php?cid=&keywords=QCentral+utility&act=query

QFinder and QCentral are java based programs. To execute this program, JRE (Java Runtime Environment) is required. You can visit the following websites to download and install JRE.

http://www.java.com/en/download/

2.3.1. QFinder Utility

After JRE is installed, run the **QFinder.jar** program. The main page will be displayed.

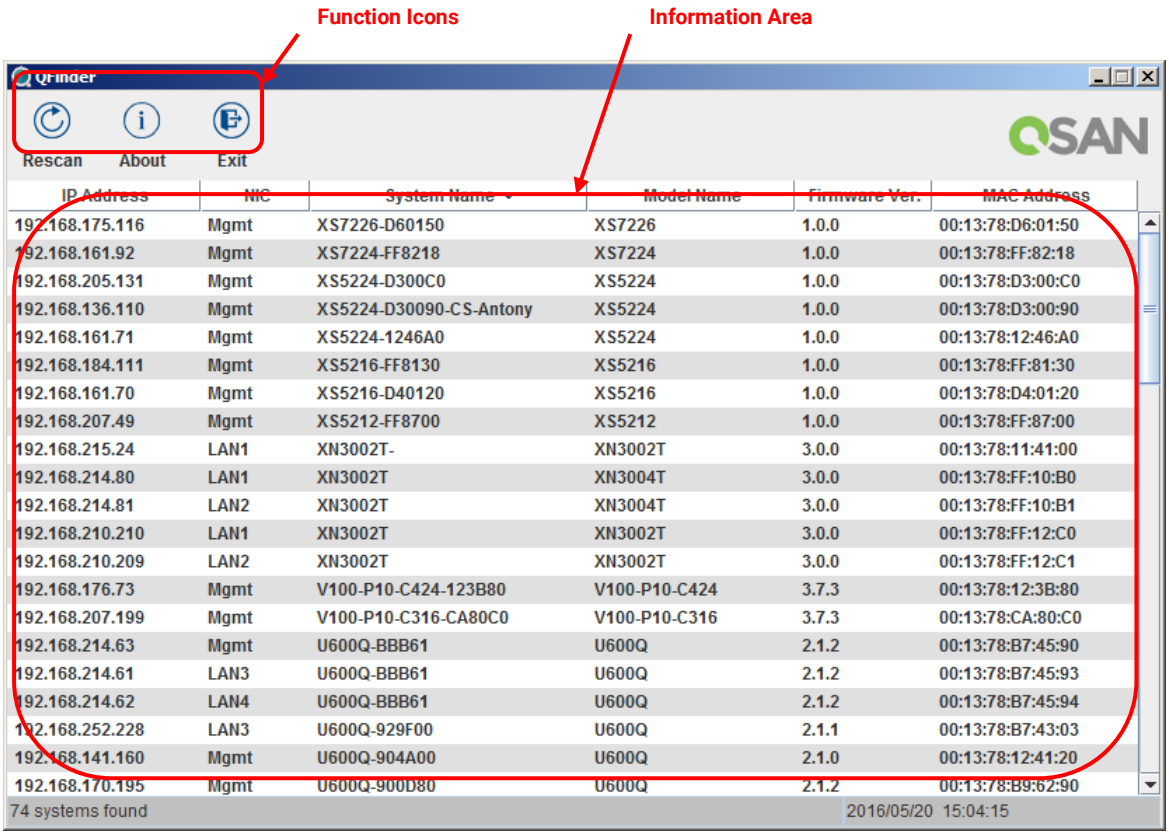

Figure 2-3 QFinder Utility

Select a system by **System Name**. Double click the column will open a browser to manage the system.

2.3.2. QCentral Utility

After JRE is installed, run the **QCentral.jar** program. The main page will be displayed.

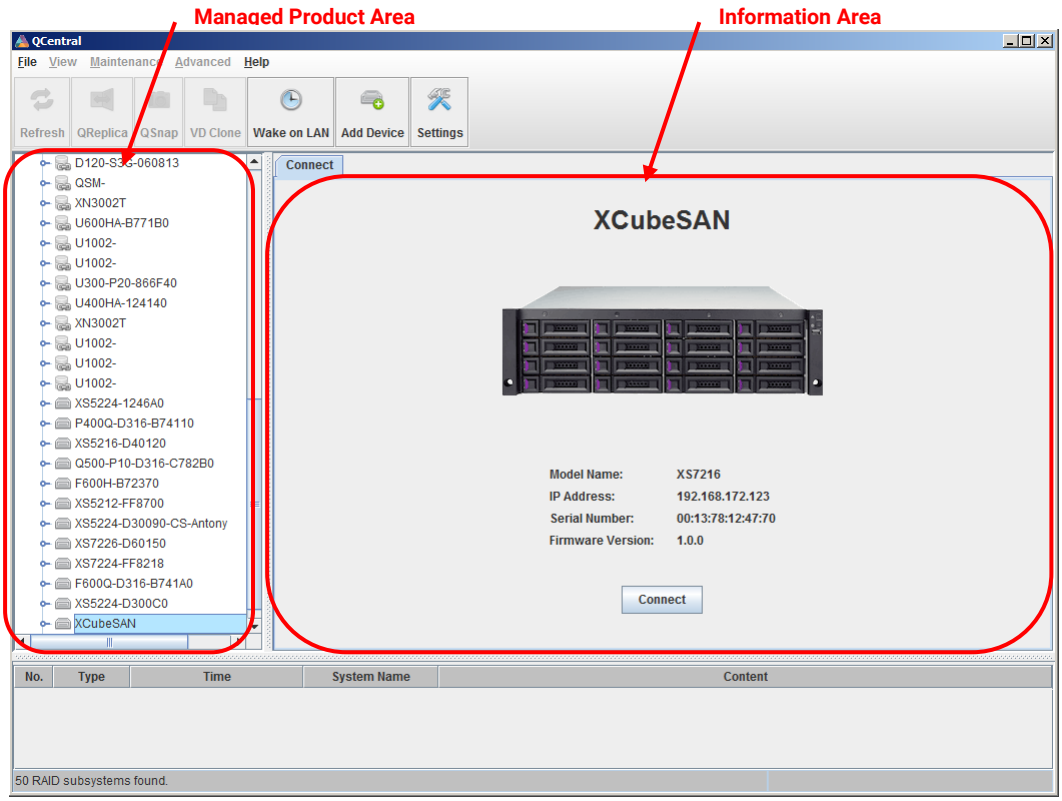

Figure 2-4 QCentral Utility

Select a system by **System Name**. Click **Connect** button will open a browser to manage the system.

2.3.3. Initial Configuration Wizard

This section describes how to complete the initial configuration, including the following tasks:

- System Setup
- Network Setup
- Initial Configuration Summary

If you just completed the initial configuration setup, that wizard automatically redirects to the SAN Storage and this wizard will not be shown again until reset to system default. Otherwise, complete the following steps to complete the initial configuration.

1. In QFinder or QCentral, find the managed system. Click the system will open a browser to system login page as shown in Figure 2-2.
2. To access the web user interface, you have to enter a username and password. The initial defaults for administrator login are:
    - Username: admin
    - Password: 1234

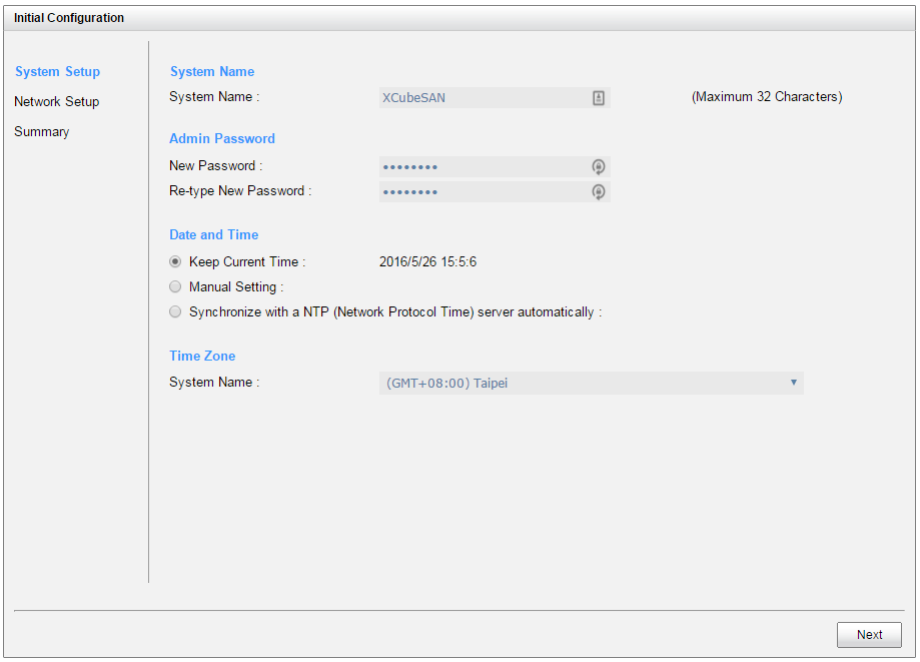

Figure 2-5 Initial Configuration Step 1

3. Enter a **System Name**, The maximum length of name is 32 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
4. Change **Admin Password**, The maximum length of name is 12 characters.
5. Set up the **Date and Time**. Date and time can be set by manually or synchronized with a NTP (Network Time Protocol) server.
6. Select a **Time Zone** depend on your location.
7. Click **Next** button to proceed.

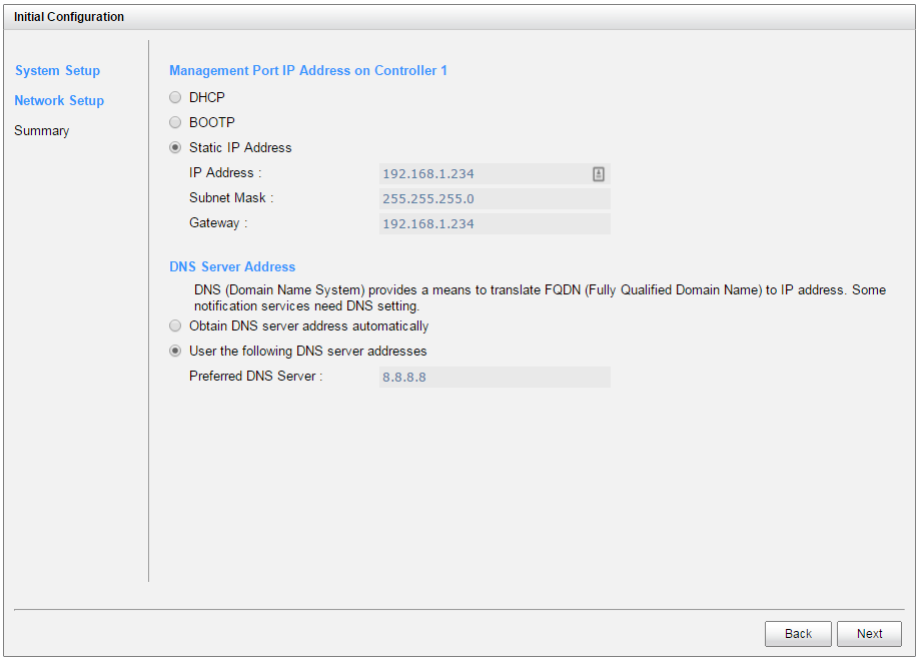

Figure 2-6 Initial Configuration Step 2

8. Confirm or change the management port IP address and DNS server. DNS (Domain Name System) provides a means to translate FQDN (Fully Qualified Domain Name) to IP address. Some notification services need DNS setting.
9. Click **Next** button to proceed.

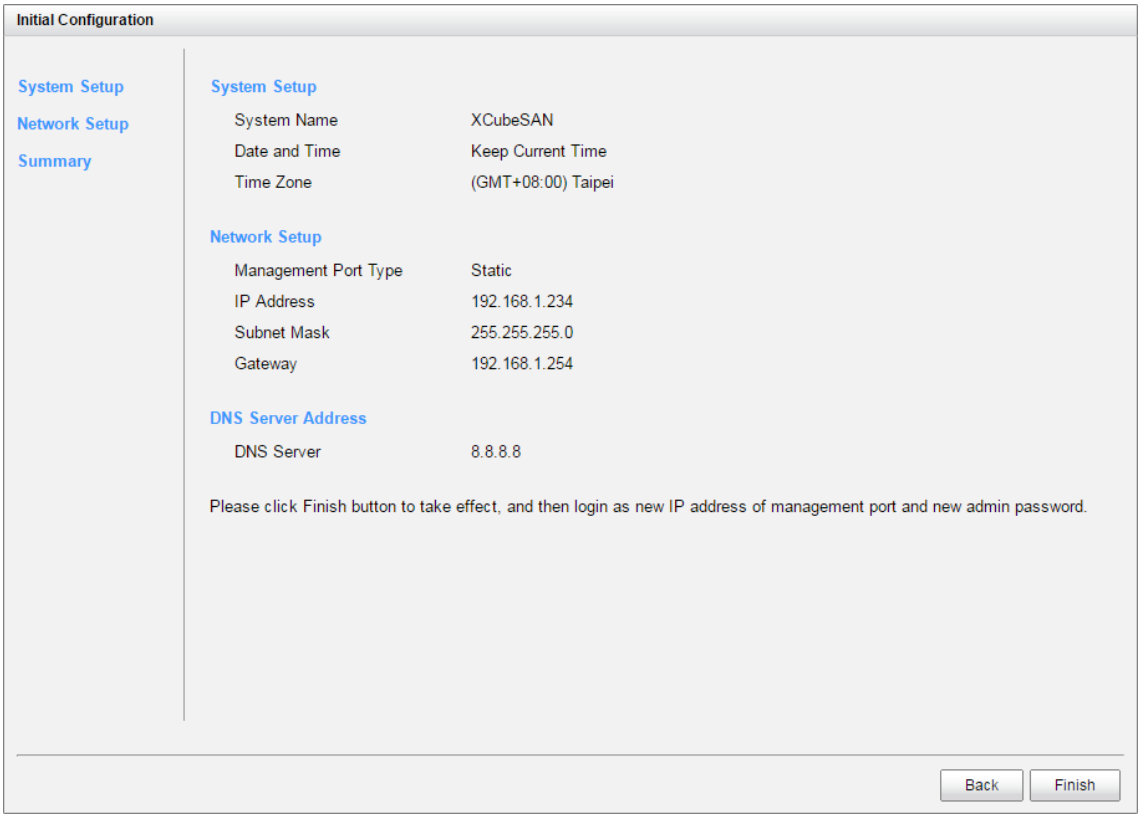

Figure 2-7 Initial Configuration Step 3

10. Verify all items, and then click **Finish** button to complete the initial configuration. You have to login as new IP address of management port and new admin password next time.

3. User Interface Overview

This chapter provides an overview of the web user interface.

3.1. Getting Started

This section provides information about accessing the management web user interface and describes the layout of the system panel.

3.1.1. Accessing the Management Web UI

To access the management web user interface, open a supported web browser and enter the management IP address or Hostname of the system. If you complete the initial configuration process, the browser will redirect to the new IP address of management port. The login panel is displayed, as shown in Figure 3-1.

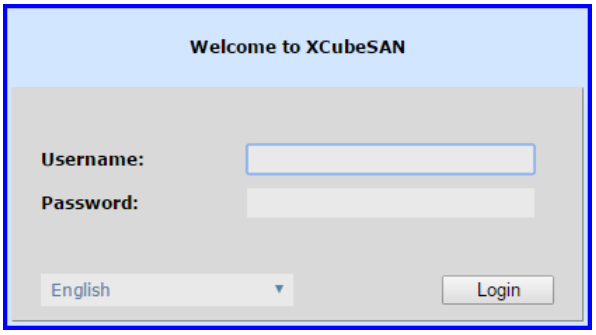

Figure 3-1 Login Page of web UI

To access the web user interface, you have to enter a username and password.

- Username: admin
- Password: <*Your Password*>

3.1.2. System Panel Layout

When the password has been verified, the home page is displayed. As shown in Figure 3-2, the System panel has five main sections:

- System Name
- Function Menu
- Function Tab
- Information Area
- Status Indicators

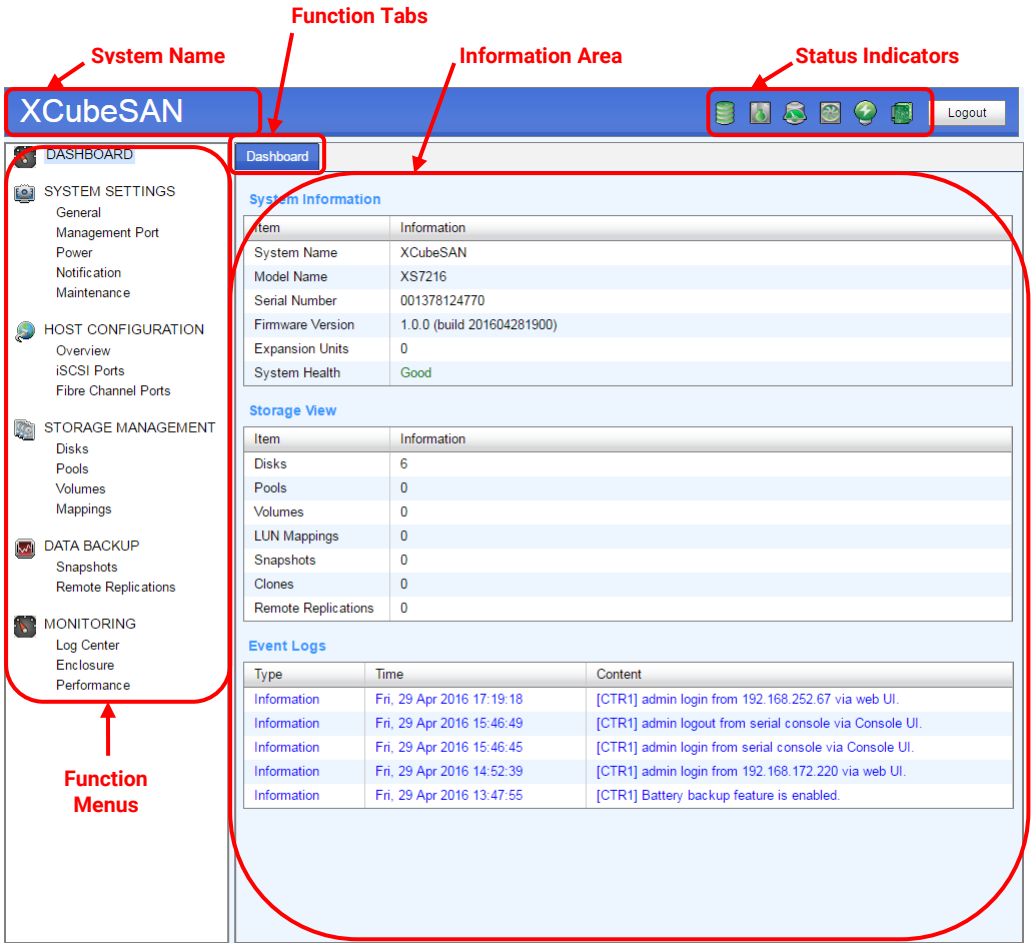

Figure 3-2 System Panel Layout

All Function Menus are listed at the left side of the window. Click anyone of them will change the related information in Information Area. Some main functions expand sub functions which are displayed at function tabs. The top-right corner displays Status Indicators which will be described in next section, and a **Logout** button.

3.1.3. Status Indicators

There are up to seven status indicators and three icons at the top-right corner. The last indicator (Dual-controller) is only visible when two controllers are installed.

Figure 3-3 Status Indicators

This table shows the descriptions.

Table 3-1 System Indicators

Icon	Description
	RAID indicator: • Green: All pools are functioning. • Red: A pool is degraded or has failed.
	Temperature indicator: • Green: Temperature is normal. • Red: Temperature is too high.
	Voltage indicator: • Green: Voltage values are normal. • Red: Voltage values are out of range.
	UPS indicator: • Green: UPS is functioning or no UPS is connected. • Red: UPS connection has failed.
	Fan indicator: • Green: Fan is working well. • Red: Fan failed.
	Power indicator: • Green: Power supplies are connected and working well. • Red: A power supply has failed or is no longer connected.
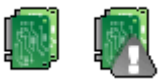	Dual controller indicator: • Green: Dual controllers are active and working well. • Amber: One of the dual controllers has failed.

NOTE:

If the status indicators in IE (Internet Explorer) are displayed in gray, but not in blinking red, please enable Tools -> Internet Options -> Advanced -> Play animations in WebPages options in IE. The default value is enabled, but some applications disable it.

3.2. Function Menus

The management web user interface provides function menus that are an efficient and quick mechanism for navigation. As described in System Panel Layout section, click the functions from the Function Menus on the left side of the window to make any configuration changes. Figure 3-5 shows all function menus and all their function tabs.

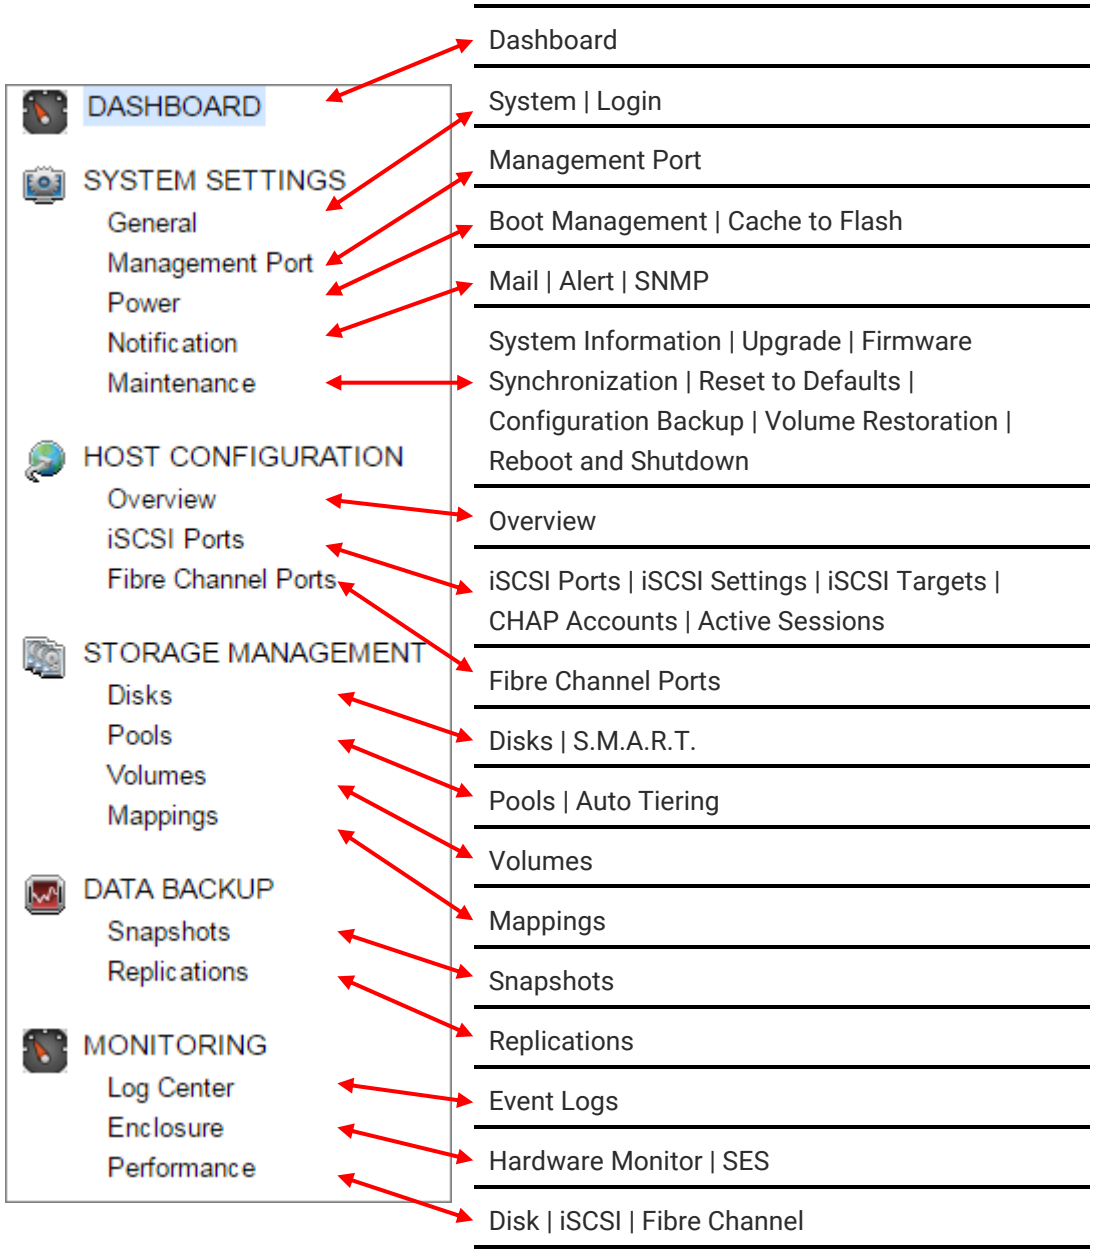

Figure 3-4 All Function Menus

3.2.1. Dashboard Menu

The **Dashboard** tab in **DASHBOARD** menu provides summary of system information. This page may help user to realize the basic information and system health.

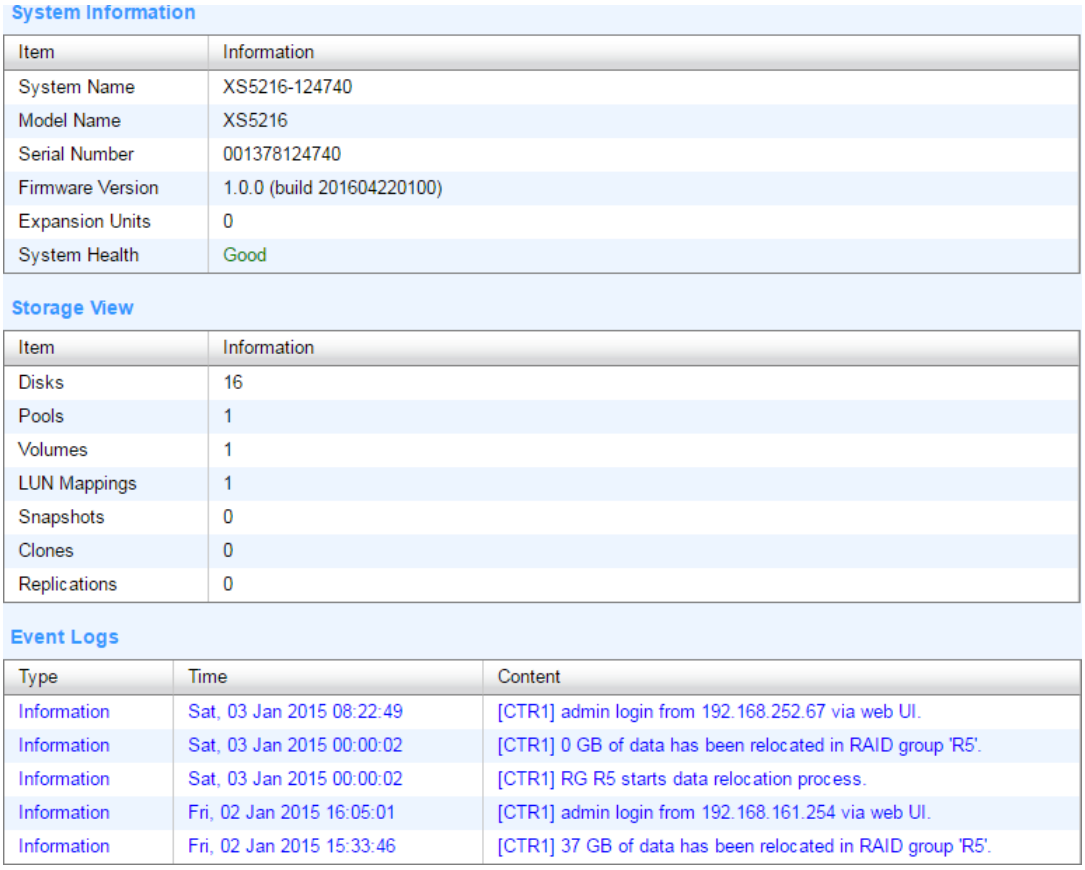

System Information

Item	Information
System Name	XS5216-124740
Model Name	XS5216
Serial Number	001378124740
Firmware Version	1.0.0 (build 201604220100)
Expansion Units	0
System Health	Good

Storage View

Item	Information
Disks	16
Pools	1
Volumes	1
LUN Mappings	1
Snapshots	0
Clones	0
Replications	0

Event Logs

Type	Time	Content
Information	Sat, 03 Jan 2015 08:22:49	[CTR1] admin login from 192.168.252.67 via web UI.
Information	Sat, 03 Jan 2015 00:00:02	[CTR1] 0 GB of data has been relocated in RAID group 'R5'.
Information	Sat, 03 Jan 2015 00:00:02	[CTR1] RG R5 starts data relocation process.
Information	Fri, 02 Jan 2015 16:05:01	[CTR1] admin login from 192.168.161.254 via web UI.
Information	Fri, 02 Jan 2015 15:33:46	[CTR1] 37 GB of data has been relocated in RAID group 'R5'.

Figure 3-5 Dashboard

There are three main sections:

- **System Information:** Include System Name, Model Name, Serial Number, Firmware Version, Expansion Units, and System Health.
- **Storage View:** Display the quantity of Disks, Pools, Volumes, LUN Mappings, Snapshots, Clones, and Replications.
- **Event Logs:** Display top 5 event logs.

3.2.2. System Settings Menu

The **SYSTEM SETTINGS** menu provides access to function menus of General, Management Port, Power, Notification, and Maintenance.

General

Select **General** function menu to setup system and login settings. For more information, see General System Settings section in System Settings chapter.

Management Port

Select **Management Port** function menu to setup IP address and DNS settings of management port. For more information, see Management Port Settings section in System Settings chapter.

Power

Select **Power** function menu to setup boot management which includes auto shutdown, wake on LAN, and wake on SAS …etc, and view the status of cache to flash. For more information, see Power Settings section in System Settings chapter.

Notification

Select **Notification** function menu to setup notification settings of mail, syslog, popup alert, LCM alert, and SNMP trap. For more information, see Notification Settings section in System Settings chapter.

Maintenance

Select **Maintenance** function menu to view system information, upgrade system firmware and check firmware synchronization, reset to system defaults, import and export configuration files, restore volumes if mis-delete occurs, reboot and shutdown. For more information, see Maintenance section in System Settings chapter.

3.2.3. Host Configuration Menu

The **HOST CONFIGURATION** menu provides access to function menus of Overview, iSCSI Port, and Fibre Channel Ports which is visible when fibre channel host cards are installed.

Overview

Select **Overview** function menu to view all host connectivity which includes status and settings of all host cards. For more information, see Overview for Host Connectivity section in Host Configuration chapter.

iSCSI Ports

Select **iSCSI Ports** function menu to setup iSCSI port settings, including iSCSI port IP address, link aggregation, VLAN ID, jumbo frame, iSNS server, iSCSI targets, CHAP accounts, and view iSCSI active sessions. For more information, see Configure for iSCSI Connectivity section in Host Configuration chapter.

Fibre Channel Ports

Select **Fibre Channel** function menu to setup fibre channel port settings, including link speed, typology, target configuration, and clear counters of fibre channel statistics. For more information, see Configure for Fibre Channel Connectivity section in Host Configuration chapter.

3.2.4. Storage Management Menu

The **STORAGE MANAGEMENT** menu provides access to function menus of Disks, Volumes, and Mappings.

Disks

Select **Disks** function menu to view status of disk drives and S.M.A.R.T. (Self-Monitoring Analysis and Reporting Technology) information. For more information, see Working with Disk Drives section in Storage Management chapter.

Pools

Select **Pools** function menu to configure storage pools. For more information, see Configuring Storage Pools section in Storage Management chapter.

Volumes

Select **Volumes** function menu to configure volumes and manage clone tasks. For more information, see Configuring Volume section in Storage Management chapter and Managing Local Clones section in Data Backup chapter.

Mappings

Select **Mappings** function menu to configure LUN mappings. For more information, see Configuring LUN Mappings section in Storage Management chapter.

3.2.5. Data Backup Menu

The **DATA BACKUP** menu provides access to function menus of Snapshots and Replications.

Snapshots

Select **Snapshots** function menu to manage snapshots, including setup snapshot space, take snapshot, rollback snapshot, and schedule snapshots. For more information, see Managing Snapshots section in Data Backup chapter.

Replications

Select **Replications** function menu to manage remote replications, including setup replication tasks, schedule replication tasks, configure shaping settings. For more information, see Managing remote replications section in Data Backup chapter.

3.2.6. Monitoring Menu

The **MONITORING** menu provides access to function menus of Log Center, Enclosure, and Performance.

Log Center

Select **Log Center** function menu to view event logs, download event logs, or clear event logs. Also mute buzzer if system happens alerts. For more information, see Event Logs section in Monitoring chapter.

Enclosure

Select **Enclosure** function menu to view hardware monitor status and enable SES (SCSI Enclosure Services). For more information, see Hardware Monitor section in Monitoring chapter.

Performance

Select **Performance** function menu to view system performance. For more information, see Performance Monitor section in Monitoring chapter.

4. System Settings

The **SYSTEM SETTINGS** function menu provides submenus of **General**, **Management Port**, **Power**, **Notification**, and **Maintenance**.

4.1. General System Settings

General provides **System** and **Login** tabs to show information and configuration setup.

4.1.1. System Settings

System tab shows information of system name, date and time and UID LED control mechanism. System name, date and time could be set as required. The system UID LED helps user to easily get system location on a rack.

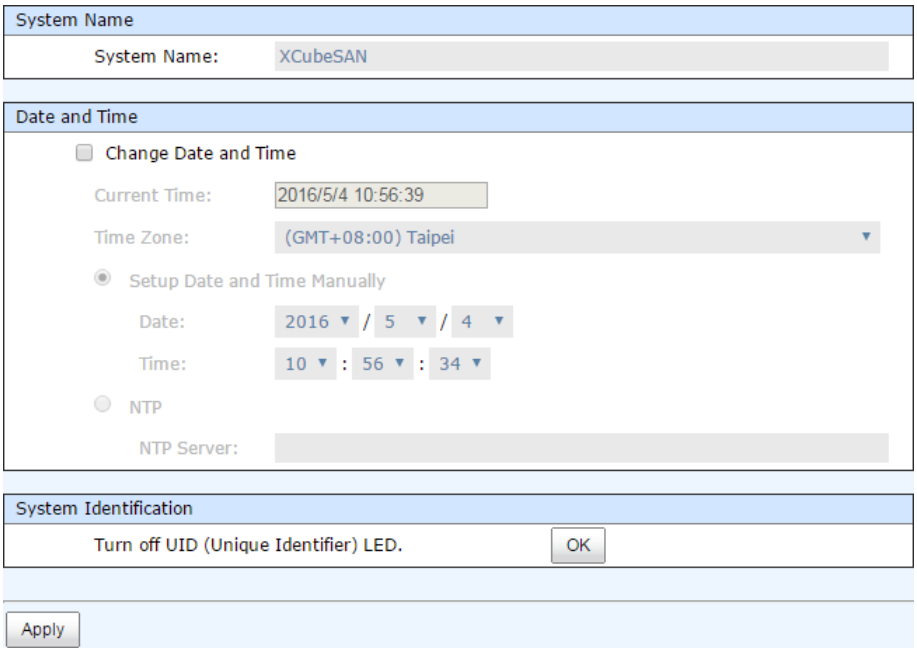

Figure 4-1 System Settings

Options of System Settings

The options are available on this tab:

- **System Name:** Change the system name; highlight the old name and type in a new one. Maximum length of name is 32 characters. Valid characters are [A~Z | a~z | 0~9 | -_].
- **Date and Time:** Check **Change Date and Time** button and change the current date, time and time zone as required. Date and time can be set by manually or synchronized from an NTP (Network Time Protocol) server.

- **System Identification:** Turn on/off system UID LEDs which are blue color, located on the right panel of front view and both controllers of rear view.

  When UID LEDs are turn off, click **OK** button to turn on UID LEDs. When UID LEDs are steady on, click **OK** button to turn off UID LEDs.

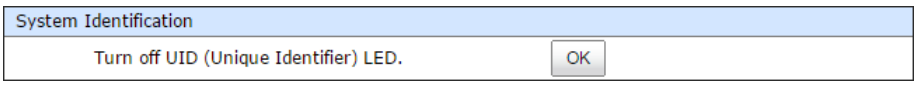

Figure 4-2 Turn off UID (Unique Identifier) LED

When above items are set done, click **Apply** button to take effect.

4.1.2. Login Settings

Login tab shows information of login related settings and administrator can set the login options, administrator password and user password. All the items could be set according to management purpose.

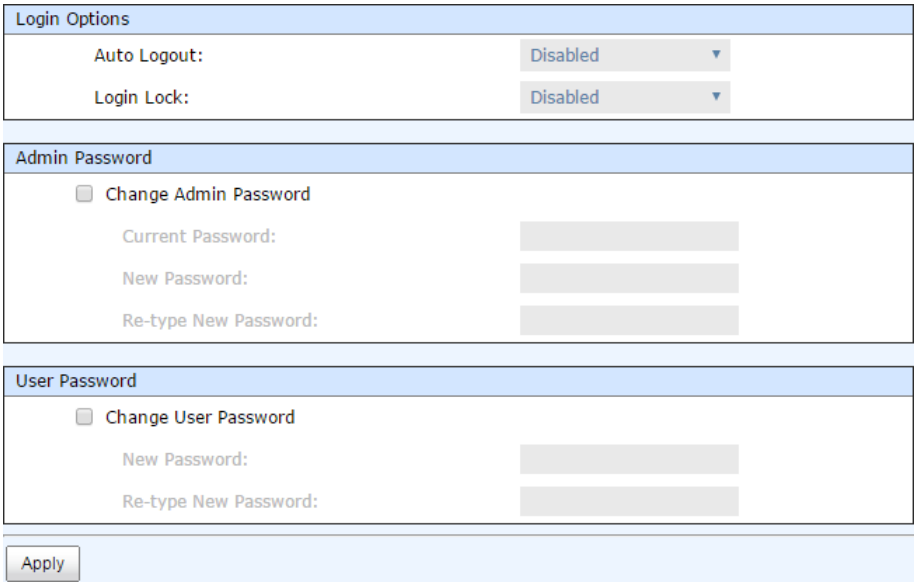

Figure 4-3 Login Settings

Options of Login Settings

The options are available on this tab:

- **Auto Logout:** When the auto logout option is enabled, you will be logged out of the admin interface after the time specified. There are Disabled (default), 5 minutes, 30 minutes and 1 hour options.
- **Login Lock:** When the login lock is enabled, the system allows only one user to login to the web UI at a time. There are Disabled (default) and Enabled options.

- **Change Admin Password:** Check it to change administrator password. The maximum length of password is 12 alphanumeric characters. Valid characters are [A~Z | a~z | 0~9].
- **Change User Password:** Check it to change user password. The maximum length of password is 12 alphanumeric characters. Valid characters are [A~Z | a~z | 0~9].

When it is done, click **Apply** button.

4.2. Management Port Settings

Management Port shows information of management ports: MAC address, IP address, DNS (Domain Name Server) server address and service ports. MAC address is fixed and could be changed. IP address, DNS server and service ports could be modified according to management purpose.

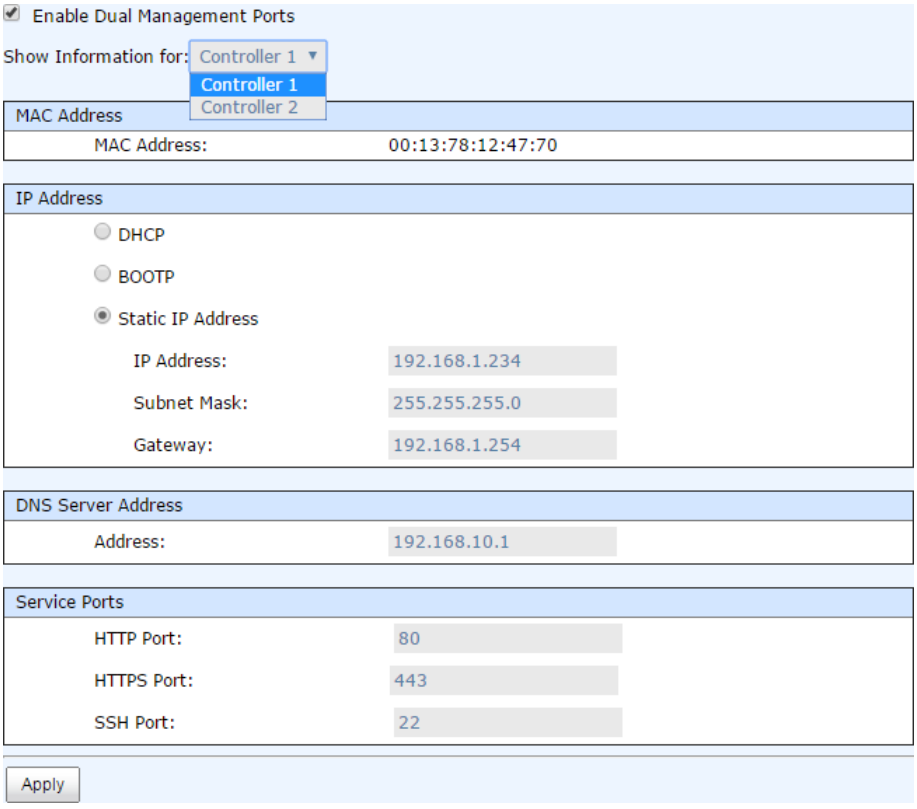

Figure 4-4 Management port settings

Options of Management Port Settings

The options are available on this tab:

- **Enable dual management ports:** This is for dual controller models. Check it to enable dual management ports. Both management ports of controller 1 and 2 have their own IP addresses, and both are active. Otherwise, only master controller is active, the other one is standby.
- **MAC Address:** Display the MAC address of the management port in the system.

- **IP Address:** The option can change IP address for remote administration usage. There are three options. They are **DHCP**, **BOOTP**, and Specify a **Static IP Address**.
- **DNS Server Address:** DNS (Domain Name System) provides a means to translate FQDN (Fully Qualified Domain Name) to IP address. Some notification services need DNS setting. Enter a IP address of DNS server here.
- **Service Ports:** If the default port numbers of **HTTP Port**, **HTTPS Port**, and **SSH Port** are not allowed on the network, they can be changed here.

When it is done, click **Apply** button.

4.3. Power Settings

Power provides **Boot Management** and **Cache to Flash** tabs to show information and configuration setup.

4.3.1. Boot Management Settings

Boot Management enables or disables options about boot.

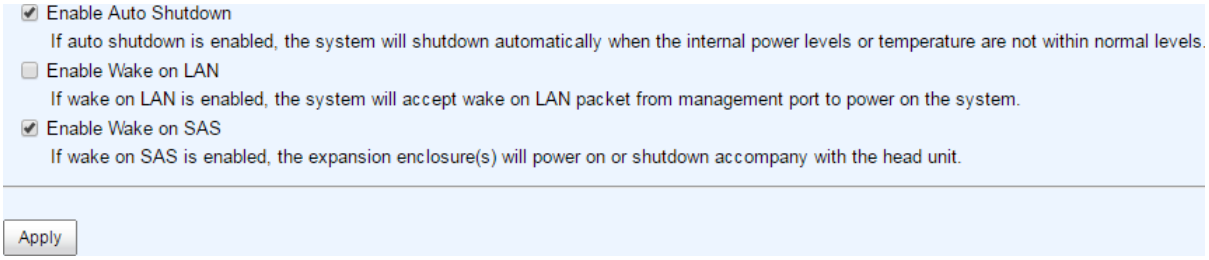

Figure 4-5 Boot Management Settings

Options of Boot Management Settings

The options are available on this tab:

- **Enable Auto Shutdown:** Check to enable auto shutdown, the system will shutdown automatically when the internal power levels or temperature are not within normal levels.
- **Enable Wake on LAN:** Check to enable wake on LAN, the system will accept wake on LAN packet from management port to power on the system.
- **Enable Wake on SAS:** Check to enable wake on SAS, the expansion enclosure(s) will power on or shutdown accompany with the head unit.

When it is done, click **Apply** button.

Auto Shutdown

For better protection and avoiding single short period of high temperature that could trigger an automatic shutdown, the system uses to gauge if a shutdown is needed. This is done using several sensors placed on key systems that the system checks every 30 seconds for present temperatures.

- The core processor temperature is greater than 95℃.
- The SAS Controller and SAS Expandor temperature is greater than 95℃.
- The host card temperature limit is 85℃.

When one of these sensors reports a temperature above the threshold for three contifuous minutes, the system shuts down automatically. For more information, see Hardware Monitor section in Monitoring chapter.

4.3.2. Cache to Flash

Cache to Flash shows information of power module and flash module for cache to flash protection. Power module indicates BBM (Battery Backup Module) or SCM (Super Capacitor Module)

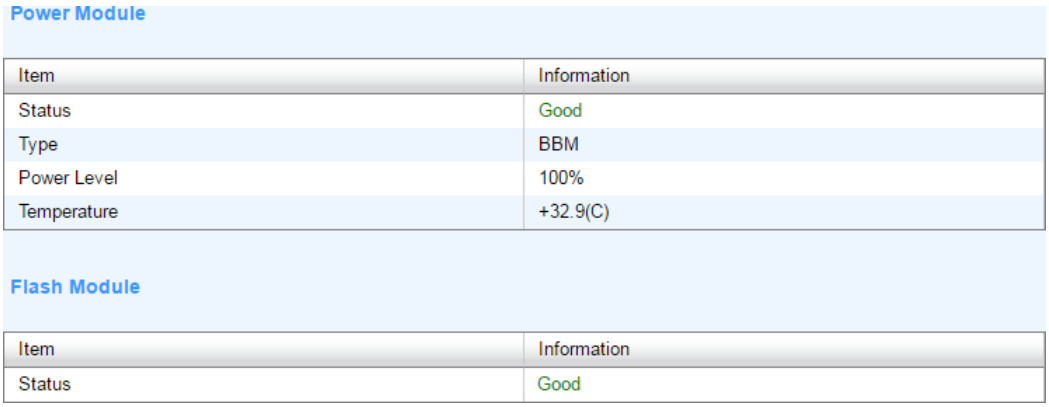

Figure 4-6 Cache to Flash

This table shows the column descriptions.

Table 4-1 Cache to Flash Column Descriptions

Column Name	Description
Status	The status of power module: • Good: The power module is good. • Failed: The power module is failure. • Absent: The power module is absent. • Charging: The power module is charging
Type	The type of power module: • BBM (Battery Backup Module) • SCM (Super Capacitor Module)
Power Level	The power level of the module.
Temperature	The temperature of the power module..
Status	The status of flash module: • Good: The flash module is good. • Failed: The flash module is failure. • Absent: The flash module is absent.

	• Detecting: The flash module is detecting

4.4. Notification Settings

Notification provides **Mail**, **Alert** and **SNMP** tabs to show information and configuration setup.

4.4.1. Mail Settings

The **Email Settings** tab is used to enter up to three email addresses for receiving the event notifications. Fill in the necessary fields and click **Send Test Email** button to test whether it is available. Some email servers will check the mail-from address and need the SMTP relay settings for authentication.

NOTE:

Please make sure the DNS server IP is well-setup in **Management Port** Settings. So the event notification emails can be sent successfully.

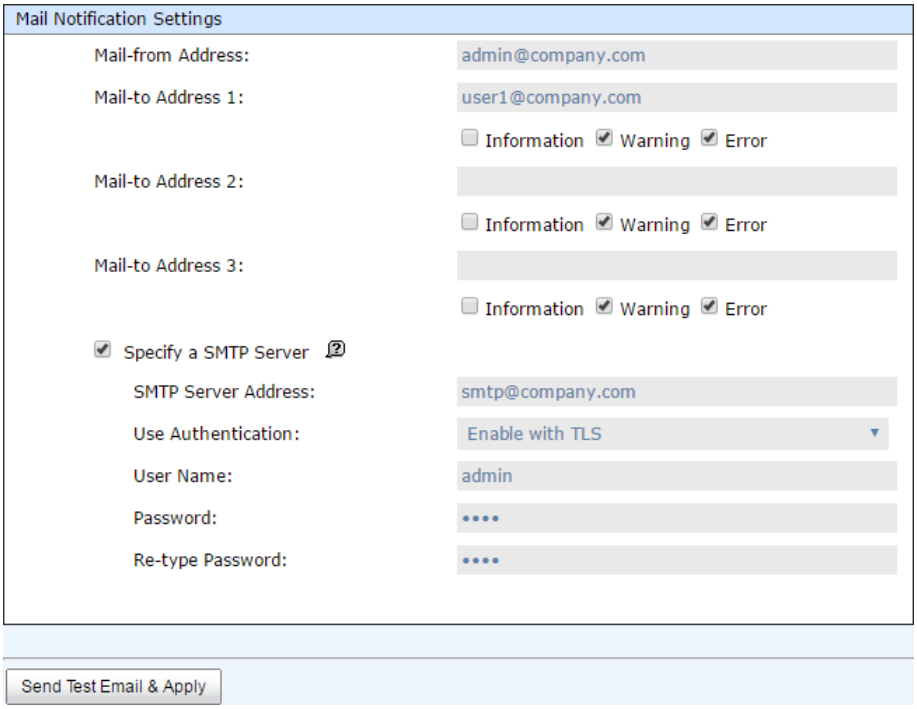

Figure 4-7 Mail Settings

You can also select which levels of event logs which you would like to receive. The default setting only includes **Warning** and **Error** event logs. When it is done, click **Apply** button.

4.4.2. Alert Settings

The **Alert Settings** tab is used to setup alerts via the **Syslog** protocol, the pop-up alerts and alerts on the front display. The device buzzer is also managed here.

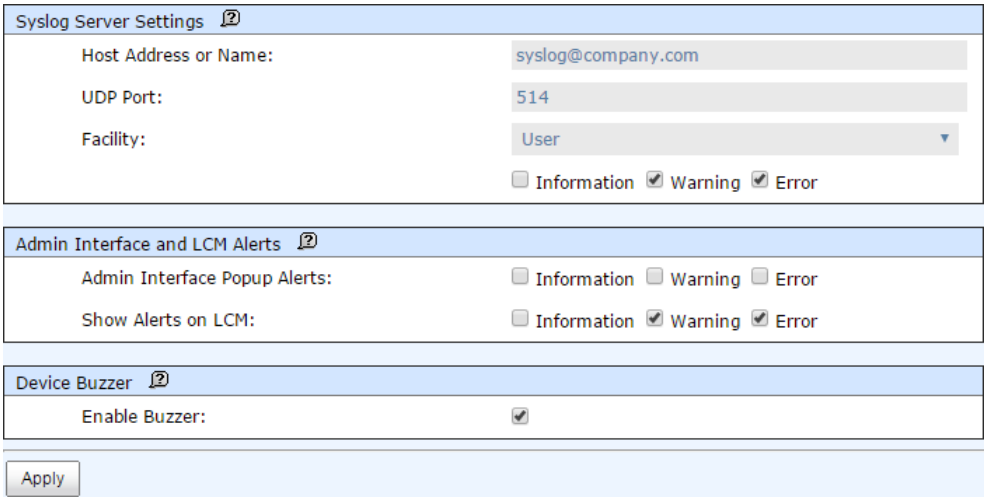

Figure 4-8 Alert Settings

Options of Alert Settings

The options are available on this tab:

- **System Server Settings:** Fill in the host address and the facility for syslog service. The default UDP port is 514. You can also check the alert levels here. Most LINUX/UNIX systems built in syslog daemon.
- **Admin Interface and LCM Alerts:** You can check or uncheck the alert levels which you would like to have pop-up message in the Web UI and show on LCM.
- **Device Buzzer:** Check it to enable the device buzzer. Uncheck it to disable device buzzer.

When it is done, click **Apply** button.

4.4.3. SNMP Settings

The **SNMP Settings** tab is used to setup SNMP traps (for alerting via SNMP).

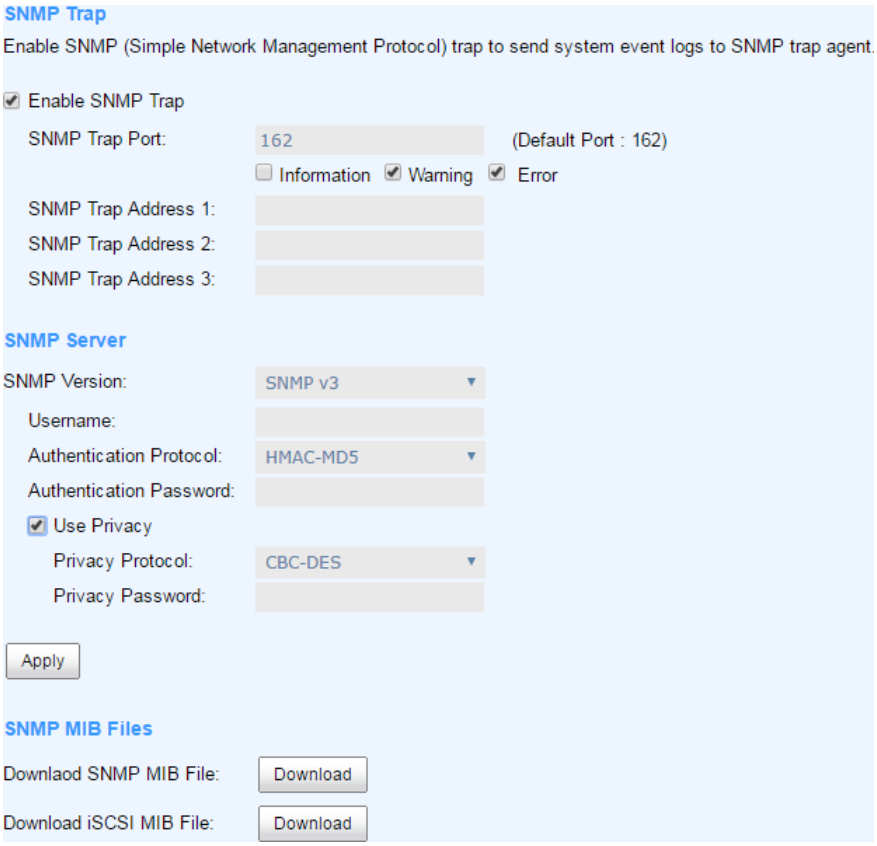

Figure 4-9 SNMP Settings

Options of SNMP Trap Settings

The options are available on this tab:

- **Enable SNMP Trap:** Check to enable SNMP (Simple Network Management Protocol) trap to send system event logs to SNMP trap agent. The default SNMP trap port is 162. You can check or uncheck the alert levels which you would like to receive. And then fill in up to three SNMP trap addresses for receiving the event notifications.

Options of SNMP Server Settings

The options are available on this tab:

- **SNMP Version:** Select SNMP supported versions, v1/v2 or v3. The default community setting of v1/v2 is public. If select v3, it needs more options for authentication. Follow the instructions to complete the settings.

When it is done, click **Apply** button.

Options of SNMP MIB Files

The options are available on this tab:

- **Download SNMP MIB File:** Click **Download** button to save the SNMP MIB file which can be imported to the SNMP client tool to get system information.

- **Download iSCSI MIB File:** Click **Download** button to save the iSCSI MIB file which can be imported to the SNMP client tool to get network information.

4.5. Maintenance

The **Maintenance** option is accessing the **System Information**, **Upgrade**, **Firmware Synchronization** (This option is only visible when dual controllers is installed.), **Reset to Defaults**, **Configuration Backup**, **Volume Restoration**, and **Reboot / Shutdown** option tabs. The **Volume Restoration** function will be described at Volume Restoration section in Advanced Volume Administration chapter; the other functions are described on the following sections.

4.5.1. System Information

The **System Information** provides to display system information.

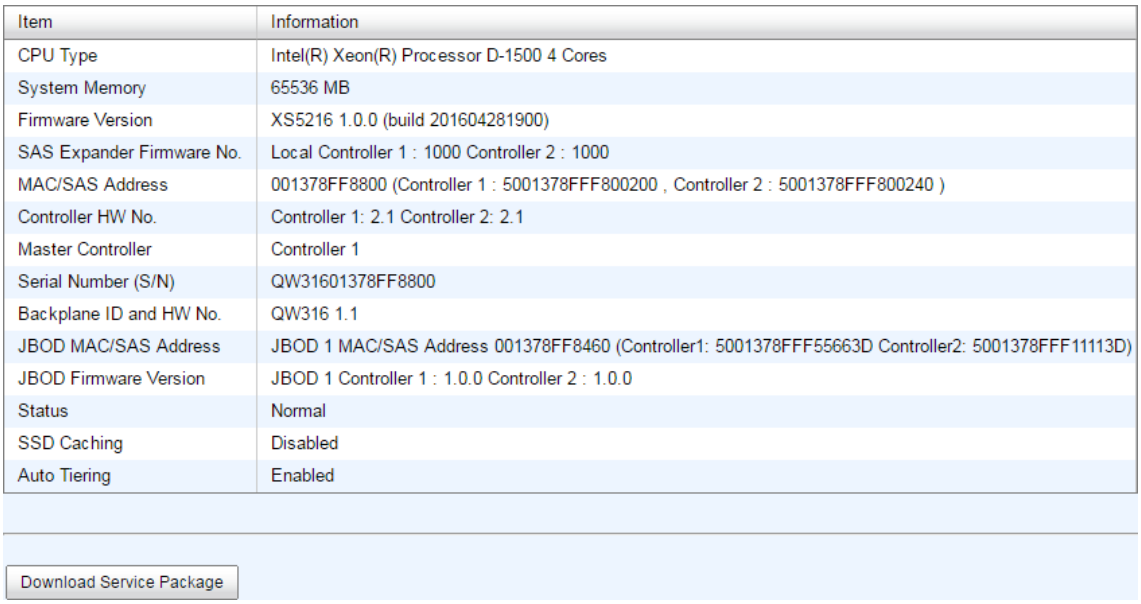

Figure 4-10 System Information

This table shows the descriptions.

Table 4-2 Status in System Information Column Descriptions

Column Name	Description
Status	The status of system: • Normal: Dual controllers and JBODs are in normal stage. • Degraded: One controller or JBOD fails or has been plugged out. • Lockdown: The firmware of two controllers is different or the size of memory of two controllers is different. • Single: Single controller mode.
Degraded	One controller or JBOD fails or has been plugged out.
Lockdown	The firmware of two controllers is different or the size of memory of two

	controllers is different.
Single	Single controller mode.

Options of System Information

The options are available on this tab:

- **Download Service Package:** Click button to download system information for service.

CAUTION:

If you try to increase the system memory and running in dual controller mode, please make sure both controllers have the same DIMM on each corresponding memory slot. Failing to do so will result in controller malfunction, which will not be covered by warranty.

4.5.2. Firmware Upgrade

The **Upgrade** tab is used to upgrade controller firmware, JBOD firmware, change operation mode, and active SSD caching license. Before upgrade, we recommend you to export your system configurations fist in **Configuration Backup** tab.

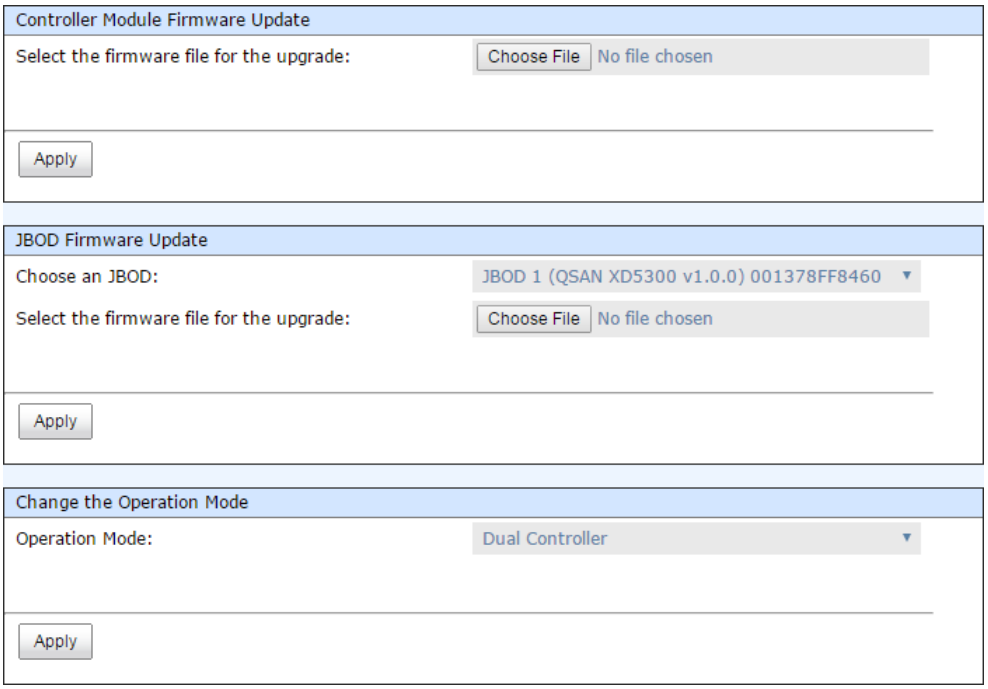

Figure 4-11 Firmware Upgrade

Options of Firmware Upgrade

The options are available on this tab:

- **Controller Module Firmware Update:** Please prepare new controller firmware file named “xxxx.bin” in local hard drive, click **Choose File** button to select the firmware file. Then Click **Apply** button, it will pop up a warning message, click **OK** button to start upgrading the firmware.

  When upgrading, there is a progress bar running. After finished upgrading, the system must reboot manually to make the new firmware take effect.

- **JBOD Firmware Update:** To upgrade JBOD firmware, choose a JBOD enclosure first, the other steps are the same as controller firmware update. After finished upgrading, the JBOD enclosure must reboot manually to make the new firmware take effect.

- **Change the Operation Mode:** This option can be modified to operate in dual-controller or single-controller mode here. If the system installed only one controller, switch this mode to **Single Controller**, and then click **Apply** button. After changing the operation mode, the system must reboot manually to take effect.

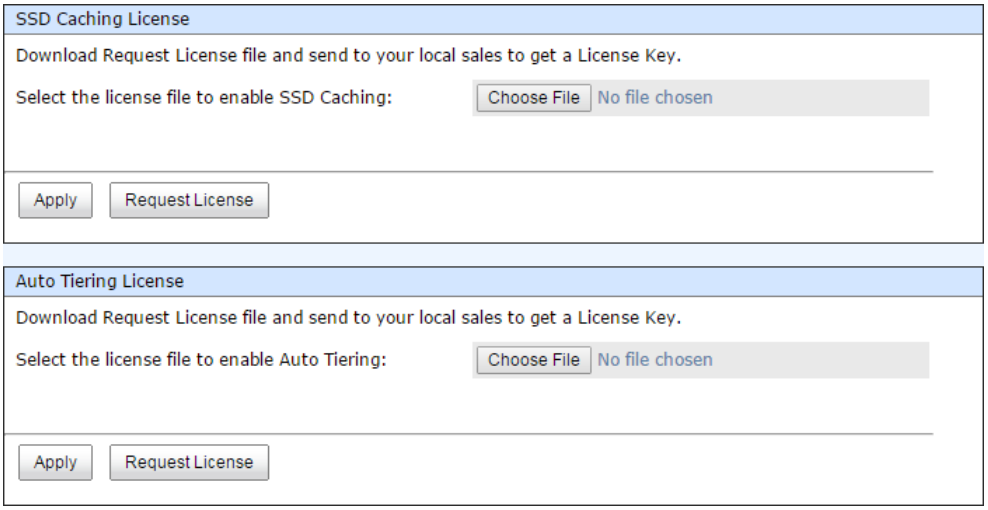

Figure 4-12 Enable Licenses

- **SSD Caching License:** This option enables SSD caching. Download **Request License** file and send to your local sales to obtain a License Key. After getting the license key, click Choose File button to select it, and then click **Apply** button to enable. Each license key is unique and dedicated to a specific system. If you have already enabled, this option will be invisible. After enabling the license, the system must reboot manually to take effect.

- **Auto Tiering License:** This option enables auto tiering. The operation is the same as SSC caching.

4.5.3. Firmware Synchronization

The **Firmware Synchronization** tab is used on dual controller systems to synchronize the controller firmware versions when the firmware of the master controller and the slave controller are different. The firmware of slave controller is always changed to match the firmware of the

master controller. It doesn't matter if the firmware version of slave controller is newer or older than that of the master. Normally, the firmware versions in both controllers are the same.

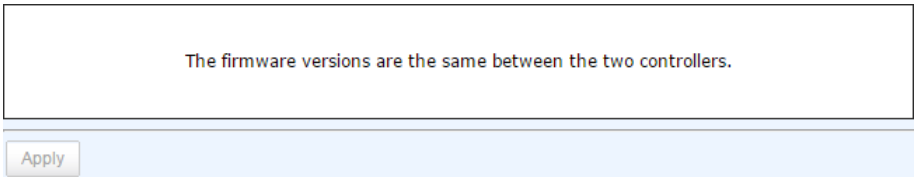

Figure 4-13 Both Firmware Versions are Synchronized

If the firmware versions between two controllers are different, it will display the warning message. Click **Apply** button to synchronize.

NOTE:

This tab is only visible when the dual controllers are installed. A single controller system does not have this option.

4.5.4. Reset to Factory Default Settings

The **Reset to Defaults** tab allows users to reset the system configurations back to the factory default settings.

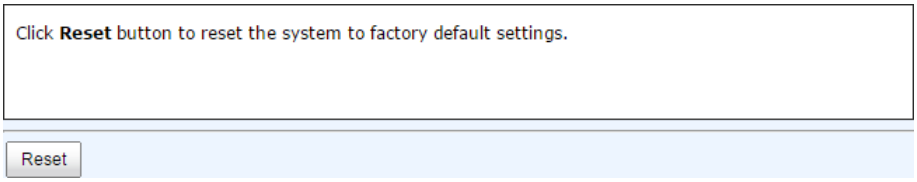

Figure 4-14 Reset to Defaults

The default values are:

- Management IP Address: DHCP
- User Name: admin
- Password: 1234

4.5.5. Configuration Backup

The **Configuration Backup** is used to either save system configuration (export) or apply a saved configuration (import).

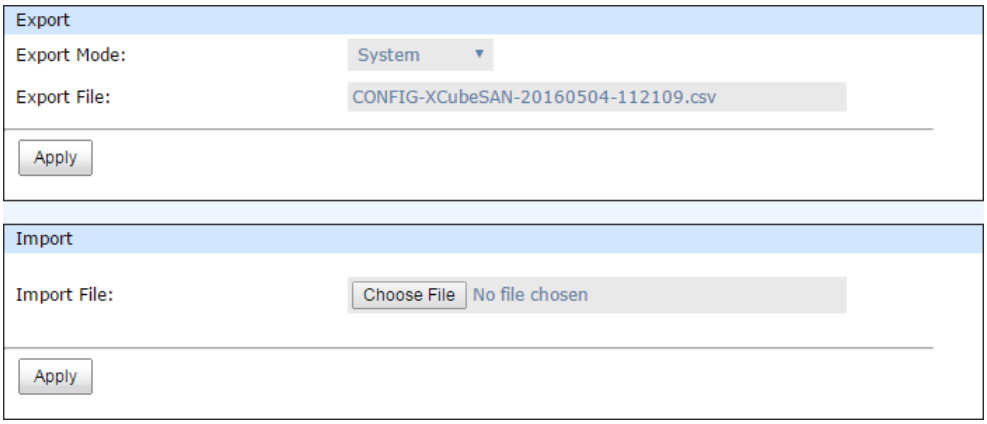

Figure 4-15 Configuration Backup

While the volume configuration settings are available for exporting, to prevent conflicts and overwriting existing data, they cannot be imported.

Options of Configuration Backup

The options are available on this tab:

- **Import:** Import all system configurations excluding volume configuration.
- **Export:** Export all configurations to a file.

CAUTION:

Import option will import all system configurations excluding volume configuration. The current system configurations will be replaced.

4.5.6. Reboot and Shutdown

The **Reboot / Shutdown** tab is used to reboot or shutdown the system. Before powering off the system, it is highly recommended to execute **Shutdown** function to flush all data from cache into the disk drives. The step is important for data protection.

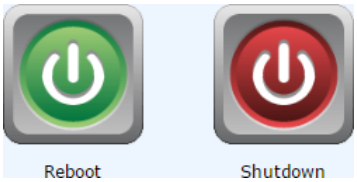

Figure 4-16 Reboot and Shutdown

The **Reboot** function has three options; reboot both controllers, controller 1 only or controller 2 only.

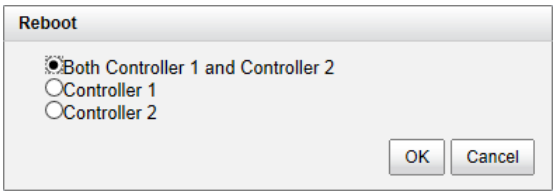

Figure 4-17 Reboot Options

5. Host Configuration

The **HOST CONFIGURATION** function menu provides submenus of **Overview**, **iSCSI Ports**, and **Fibre Channel Ports**.

5.1. Overview for Host Connectivity

XCubeSAN provides different type of host connectivity according to the system configuration, it could be base system or host cards installed system. The base system has onboard 2x 10GbE iSCSI port per controller. The host cards are installed the same type on both controllers. Currently host card has three types, 1GbE iSCSI (RJ45), 10GbE iSCSI (SFP+), and 16Gb FC (SFP+), for selection according to system infrastructure.

Overview page displays all the host connectivity in system. The status and information of all host ports are listed.

The following is an example to show onboard 10GbE LAN port, 1GbE iSCSI host card in slot 2 and 16Gb FC host cards in slot 1.

Controller 1

Location	Name	Status	MAC Address/WWPN
Onboard	LAN1 (10Gb)	Down	00:13:78:ff:88:02
Onboard	LAN2 (10Gb)	Down	00:13:78:ff:88:03
Slot2	LAN1 (1Gb)	1 Gb/s	00:13:78:ff:88:04
Slot2	LAN2 (1Gb)	1 Gb/s	00:13:78:ff:88:05
Slot2	LAN3 (1Gb)	1 Gb/s	00:13:78:ff:88:06
Slot2	LAN4 (1Gb)	1 Gb/s	00:13:78:ff:88:07
Slot1	FC1 (16Gb)	16Gb/s	2100001378FF8800
Slot1	FC2 (16Gb)	16Gb/s	2200001378FF8800
Slot1	FC3 (16Gb)	16Gb/s	2300001378FF8800
Slot1	FC4 (16Gb)	16Gb/s	2400001378FF8800

Controller 2

Location	Name	Status	MAC Address/WWPN
Onboard	LAN1 (10Gb)	Down	00:13:78:ff:88:0c
Onboard	LAN2 (10Gb)	Down	00:13:78:ff:88:0d
Slot2	LAN1 (1Gb)	1 Gb/s	00:13:78:ff:88:0e
Slot2	LAN2 (1Gb)	1 Gb/s	00:13:78:ff:88:0f
Slot2	LAN3 (1Gb)	1 Gb/s	00:13:78:ff:88:10
Slot2	LAN4 (1Gb)	1 Gb/s	00:13:78:ff:88:11
Slot1	FC1 (16Gb)	16Gb/s	2900001378FF8800
Slot1	FC2 (16Gb)	16Gb/s	2A00001378FF8800
Slot1	FC3 (16Gb)	16Gb/s	2B00001378FF8800
Slot1	FC4 (16Gb)	16Gb/s	2C00001378FF8800

Figure 5-1 Host Port Overview

5.2. Configure for iSCSI Connectivity

iSCSI Ports provides **iSCSI Ports**, **iSCSI Settings**, **iSCSI Targets**, **CHAP Accounts**, and **Sessions** tabs to configuration iSCSI ports.

5.2.1. Configure for iSCSI Ports

iSCSI Ports tab shows information of iSCSI ports where are located onboard or host cards. Each iSCSI port could be configured by clicking ▼ of the specific port.

Controller 1

	Location	Name	Status	LAG	VLAN ID	IP Address	Gateway	Jumbo Frame	MAC Address
▼	Onboard	LAN1 (10Gb)	Down	N/A	N/A	192.168.1.1	192.168.1.254	Disabled	00:13:78:ff:88:02
▼	Onboard	LAN2 (10Gb)	Down	N/A	N/A	192.168.2.1	192.168.2.254	Disabled	00:13:78:ff:88:03
▼	Slot2	LAN1 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.11		Disabled	00:13:78:ff:88:04
▼	Slot2	LAN2 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.12		Disabled	00:13:78:ff:88:05
▼	Slot2	LAN3 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.13		Disabled	00:13:78:ff:88:06
▼	Slot2	LAN4 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.14		Disabled	00:13:78:ff:88:07

Controller 2

	Location	Name	Status	LAG	VLAN ID	IP Address	Gateway	Jumbo Frame	MAC Address
▼	Onboard	LAN1 (10Gb)	Down	N/A	N/A	192.168.11.1	192.168.11.254	Disabled	00:13:78:ff:88:0c
▼	Onboard	LAN2 (10Gb)	Down	N/A	N/A	192.168.12.1	192.168.12.254	Disabled	00:13:78:ff:88:0d
▼	Slot2	LAN1 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.21		Disabled	00:13:78:ff:88:0e
▼	Slot2	LAN2 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.22		Disabled	00:13:78:ff:88:0f
▼	Slot2	LAN3 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.23		Disabled	00:13:78:ff:88:10
▼	Slot2	LAN4 (1Gb)	1 Gb/s	N/A	N/A	192.168.100.24		Disabled	00:13:78:ff:88:11

Figure 5-2 List iSCSI Ports

The columns display information of **Location**, **Name**, **Status**, **LAG** (Link Aggregation), **VLAN ID** (Virtual LAN ID), **IP address**, **Gateway** IP address, **Jumbo Frame** status, and **MAC address**.

Set IP Address

Click ▼ **-> Set IP Address** which can assign an iSCSI IP address of the iSCSI data port. There are two options: Use DHCP to acquire an IP address automatically or specify a Static IP Address to set the IP address manually.

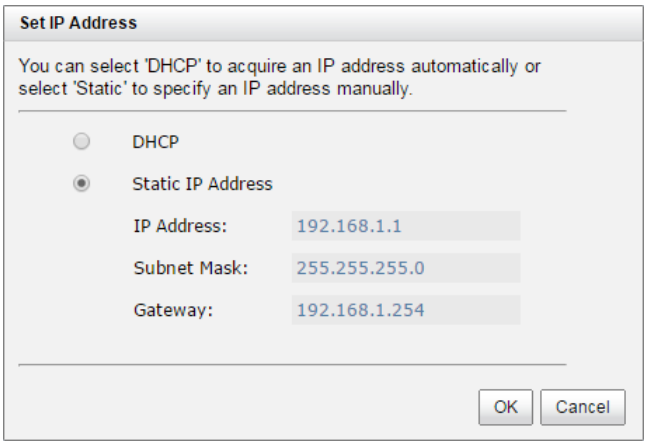

Figure 5-3 iSCSI IP Address Settings

Set Link Aggregation

Click ▼ **-> Set Link Aggregation**, the default mode of each iSCSI data port is connected without any bonding. Two bonding methods, Trunking and LACP (Link Aggregation Control Protocol), can be selected. At least two iSCSI data ports must be checked for iSCSI link aggregation.

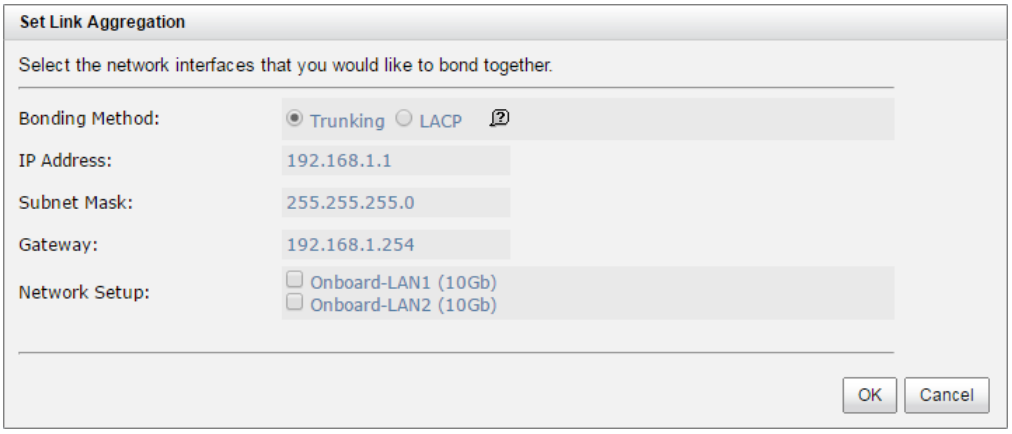

Figure 5-4 Set Link Aggregation

NOTE:

Trunking: Configures multiple iSCSI ports to be grouped together into one in order to increase the connection speed beyond the limit of a single iSCSI port.

LACP: The Link Aggregation Control Protocol is part of IEEE 802.3ad that allows bonding several physical ports together to form a single logical channel. LACP allows a network switch to negotiate an automatic bundle by sending LACP packets to the peer. LACP can increase bandwidth usage and automatically perform failover when the link status fails on a port.

Set VLAN ID

Click ▼ **-> Set VLAN ID**, VLAN (Virtual LAN) is a logical grouping mechanism implemented on switch device. VLANs are collections of switching ports that comprise a single broadcast domain. It allows network traffic to transfer more efficiently within these logical subgroups. Please consult your network switch user manual for VLAN setting instructions. Most of the work is done at the switch part. Please make sure that your VLAN ID of iSCSI port matches that of switch port. If your network environment supports VLAN, you can use this function to change the configurations. Fill in VLAN ID and Priority settings to enable VLAN.

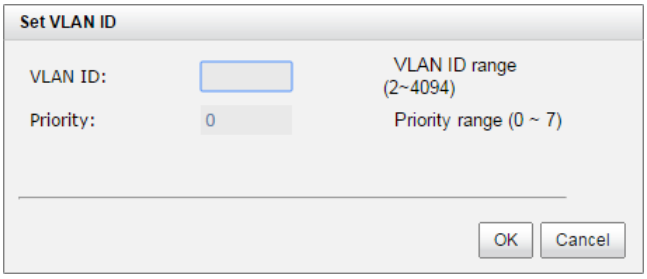

Figure 5-5 Set VLAN ID

NOTE:

VLAN ID: VLAN ID is a number ranges from 2 to 4094. Three numbers (0, 1, and 4095) are reserved for special purposes.

Priority: The PCP (Priority Code Point) is a number ranges from 0 to 7 and reserved for QoS (Quality of Service). The definition is compliant with IEEE 802.1p protocol and 0 is the default value. In normal cases, you don't need to set this value.

Set Default Gateway

Click ▼ **-> Set Default Gateway** to set the gateway of the IP address as default gateway. There can be only one default gateway.

Remove Default Gateway

To remove the default gateway, click ▼ **-> Remove Default Gateway**.

Set Jumbo Frames

Click ▼ **-> Set Jumbo Frames** to set the MTU (Maximum Transmission Unit) size. The jumbo frame size could be set as 4000 or 9000 bytes. Jumbo Frame is disabled by default.

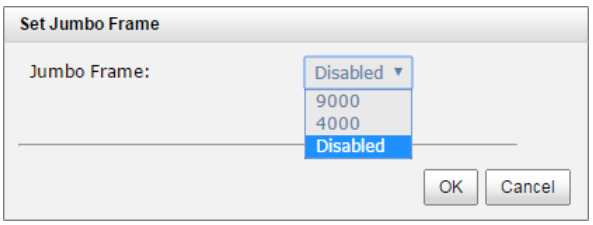

Figure 5-6 Set Jumbo Frame

CAUTION:

VLAN ID, Jumbo Frames for switching hub and HBA on host must be enabled. Otherwise, the LAN connection cannot work properly.

Ping Host

Click ▼ **-> Ping Host**, it can verify the port connection from a target to the corresponding host data port. Input the host's IP address and click **Start** button. The system will display the ping result. Click **Stop** button will stop ping activity.

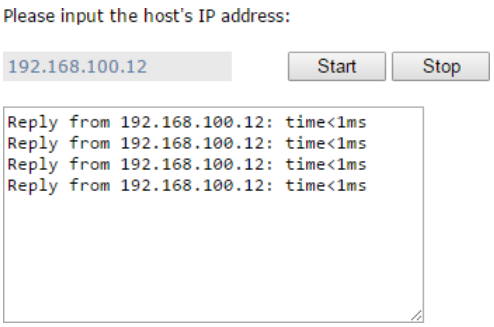

Figure 5-7 Ping Host

Reset Port

Click ▼ -> **Reset Port** which is generally used to recover from port malfunction.

5.2.2. Configure for iSCSI Settings

iSCSI Settings tab provides to setup entity name of the system and iSNS (Internet Storage Name Service) server. The entity name is default in IQN (iSCSI Qualified Name) format and could be modified for management purpose. The iSNS IP is used by iSNS protocol for automated discovery, management and configuration of iSCSI devices on a TCP/IP network. To use iSNS, an iSNS server must be added to the SAN. The iSNS server IP address must be added to the storage system for iSCSI initiator service to send queries.

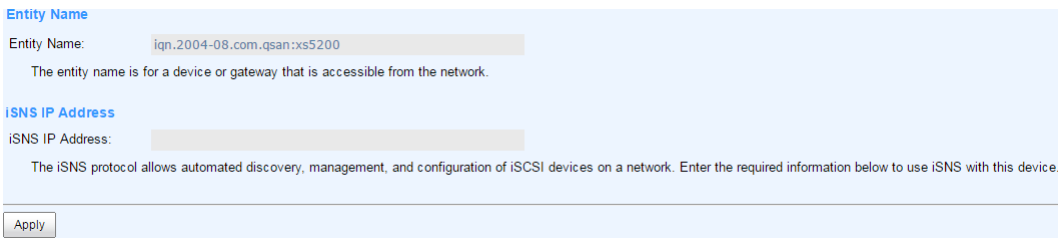

Figure 5-8 Entity Name and iSNS Settings

Options of iSCSI Settings

The options are available on this tab:

- **Entity Name:** Change the entity name; highlight the old name and type in a new one. Maximum length of name is 223 characters. Valid characters are [A~Z | a~z | 0~9 | -.:].
- **iSNS IP Address:** The option can change iSNS IP address for internet storage name service.

When it is done, click **Apply** button.

5.2.3. Configure for iSCSI Targets

iSCSI Target tab shows the iSCSI target information for iSCSI initiator. Each iSCSI target can be configured by clicking ▼ of the specific target.

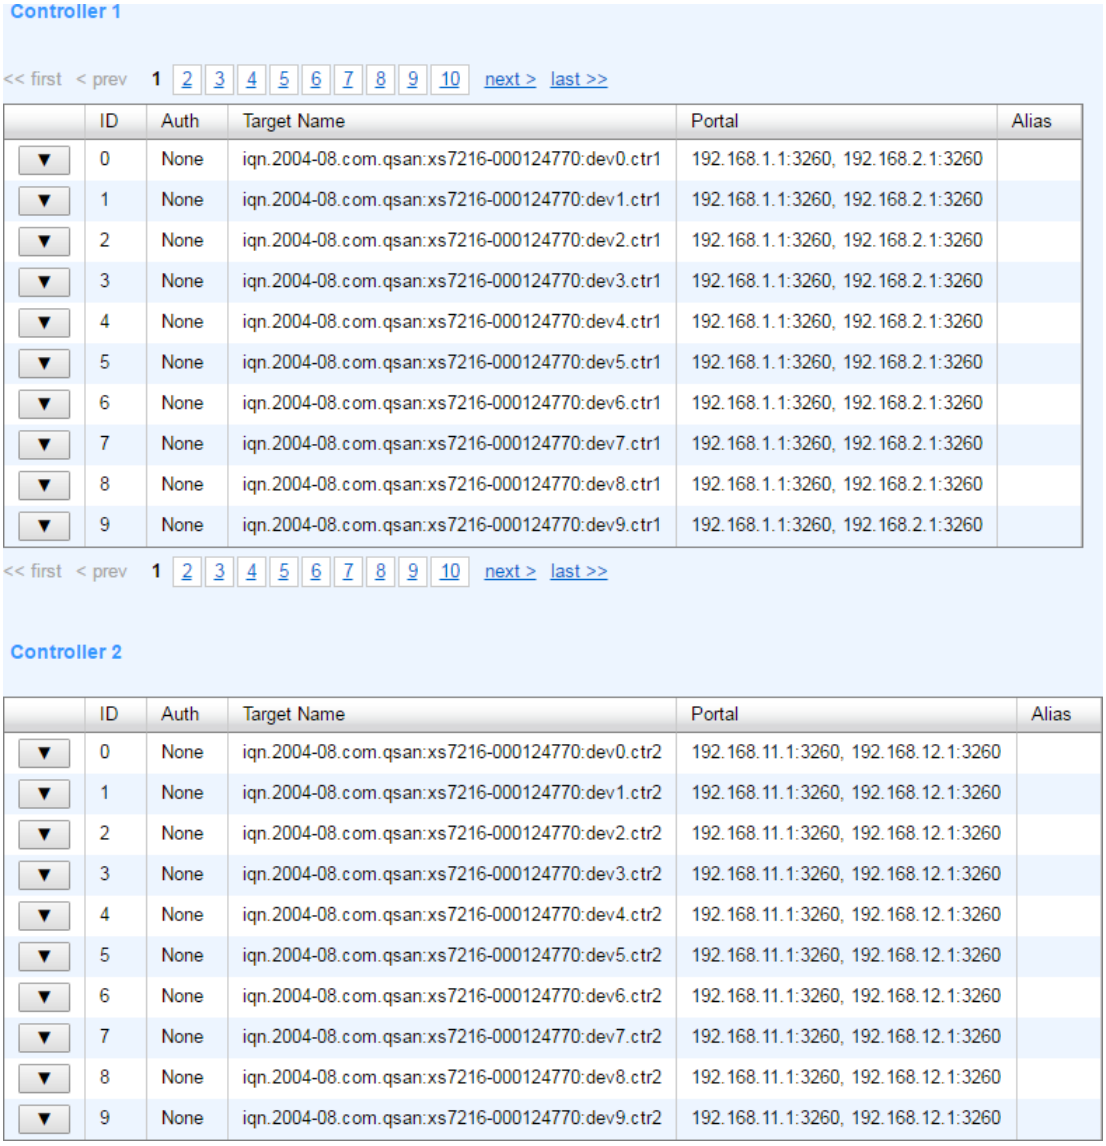

Figure 5-9 iSCSI Targets

Change Authentication Mode

Click ▼ **-> Authentication Method** to enable CHAP (Challenge Handshake Authentication Protocol) authentication method used in point-to-point for user login. It's a type of authentication in which the authentication server sends the client a key to be used for encrypting the username and password. CHAP enables the username and password to transmit in an encrypted form for protection.

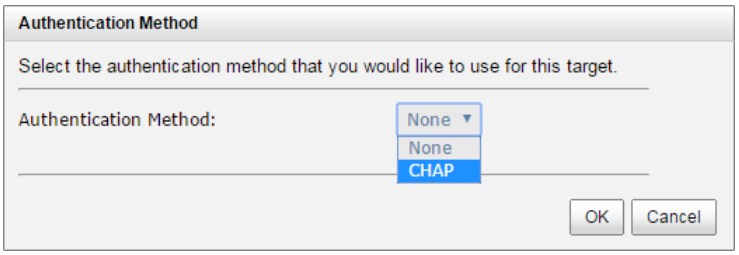

Figure 5-10 Authentication Method

NOTE:

A CHAP account must be added before you can use this authentication method. Please refer to CHAP Accounts section to create an account.

After enable CHAP authentication mode, the host initiator should be set with the same CHAP account. Otherwise, the host cannot connect to the volume.

Change CHAP Users

Click ▼ **-> Change CHAP Users**, and then select the CHAP users that you would like to have access to this target. CHAP users can be more than one, but it must be at least one CHAP to enable on the target.

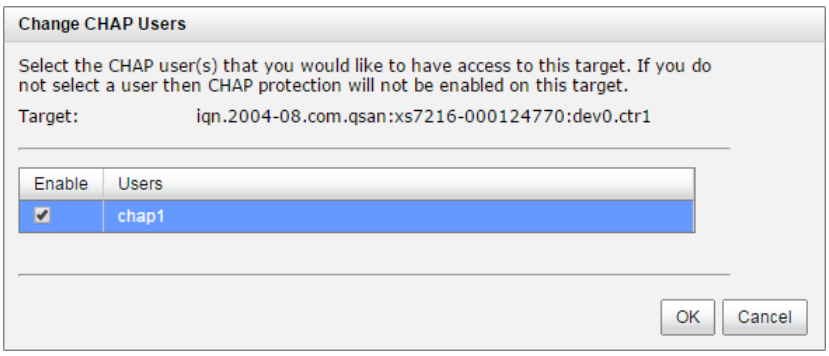

Figure 5-11 Change CHAP Users

Change Target Name

Click ▼ **-> Change Target Name** to change iSCSI target name if necessary. Maximum length of name is 223 characters. Valid characters are [A~Z | a~z | 0~9 | -.:].

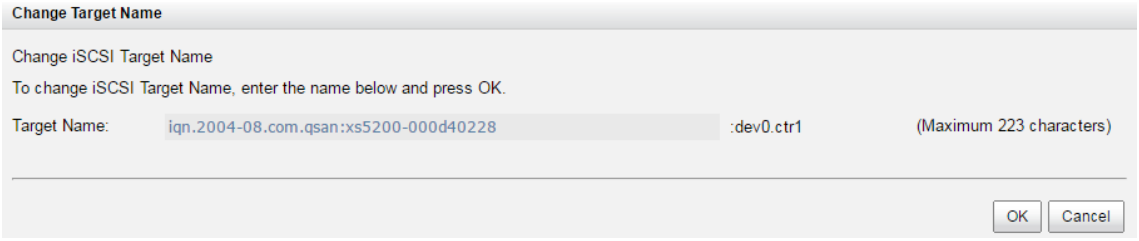

Figure 5-12 Change Target Name

Change Network Portal

Click ▼ **-> Change Network Portal**, and then select the network ports that you would like to be available for this iSCSI target.

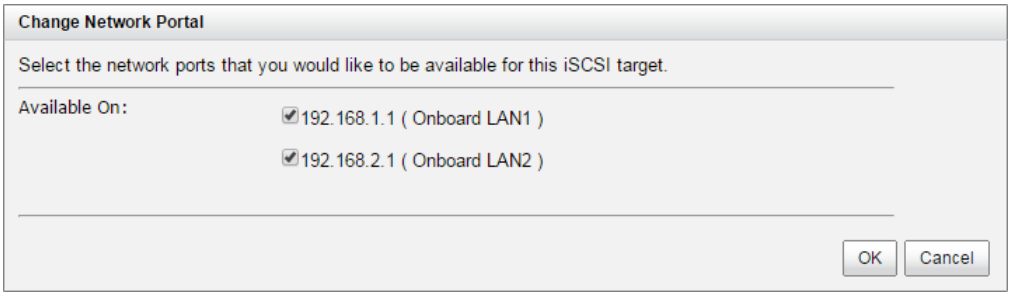

Figure 5-13 Change Network Portal

Change Alias

Click ▼ **-> Change Alias** to add or change the alias name. To remove an alias, clear out the current name. Maximum length of name is 223 characters. Valid characters are [A~Z | a~z | 0~9 | -.:].

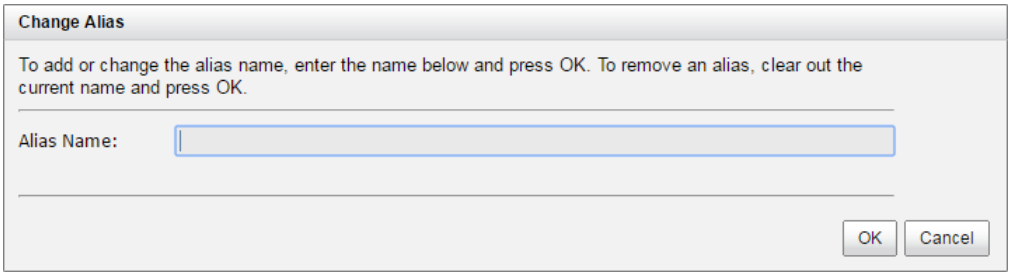

Figure 5-14 Change Alias

5.2.4. Configure for iSCSI CHAP Accounts

CHAP Accounts tab provides to manage the CHAP accounts on the system.

Create a CHAP user

Here is an example of creating a CHAP user.

1. In **CHAP Account** tab, click **Create CHAP User** button.

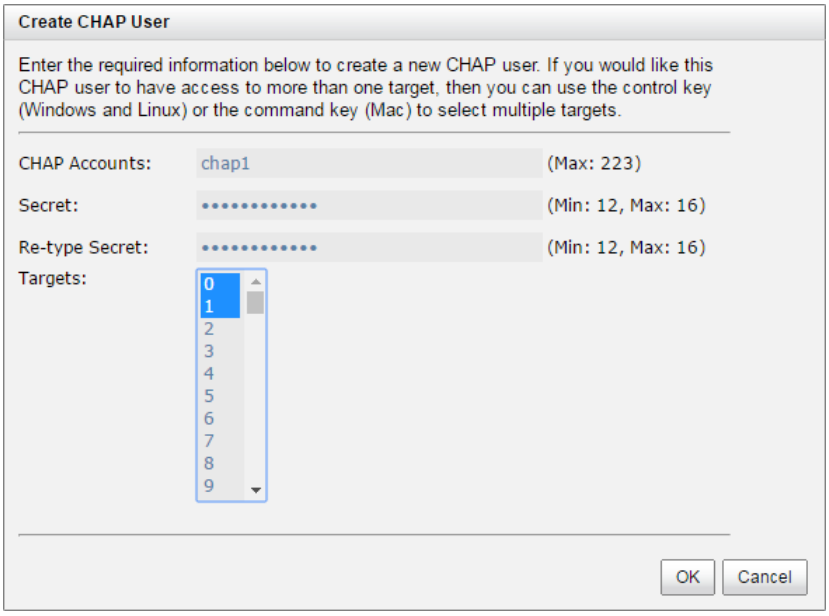

Figure 5-15 Create a CHAP User

2. Enter a name of **CHAP Accounts**. Maximum length of name is 223 characters. Valid characters are [A~Z | a~z | 0~9 | -.:].
3. Enter **Secret**, and **Re-type Secret**. The maximum length is 16 alphanumeric characters. Valid characters are [A~Z | a~z | 0~9].
4. Click **OK** button to create a CHAP user.

List CHAP Users

Here is an example of list CHAP users. Each CHAP account can be configured by clicking ▼ of the specific one.

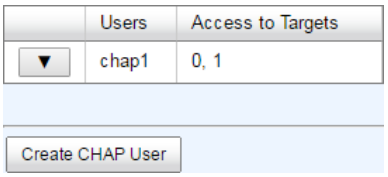

Figure 5-16 List CHAP Users

Modify CHAP User

Click ▼ -> **Modify CHAP User** to modify the selected CHAP user information. To change the targets that this user has access to, please go to **iSCSI Targets** tab, click on the option menu, and

select **Change CHAP Users**. For more information, please see Configure for iSCSI Targets section in this chapter.

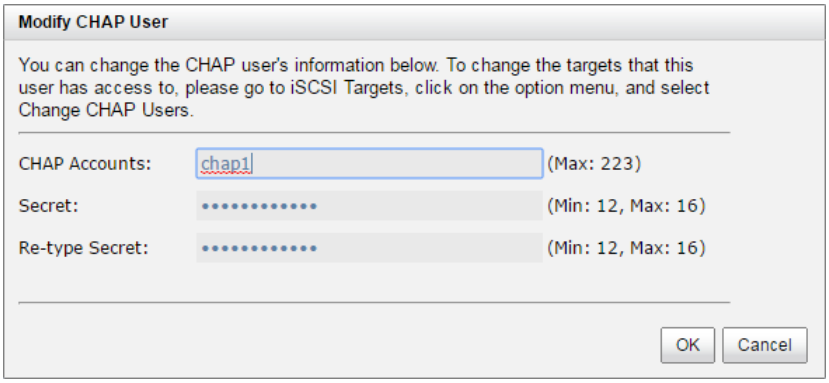

Figure 5-17 Modify a CHAP User

Delete CHAP User

Click ▼ -> **Delete CHAP User** to delete the selected CHAP user.

5.2.5. Active Sessions

Sessions tab shows all currently active iSCSI sessions and their connection information.

Figure 5-18 Active Sessions

This table shows the column descriptions.

Table 5-1 Active Sessions Column Descriptions

Column Name	Description
TSIH	TSIH (Target Session Identifying Handle) is used for this active session.
Initiator Name	It displays the host computer name.
Target Name	It displays the controller name.
InitialR2T	InitialR2T (Initial Ready to Transfer) is used to turn off either the use of a unidirectional R2T command or the output part of a bidirectional command. The default value is Yes.
Immed. data	Immed. data (Immediate Data) sets the support for immediate data between the initiator and the target. Both must be set to the same setting. The default value is Yes.
MaxDataOutR2T	MaxDataOutR2T (Maximum Data Outstanding Ready to Transfer) determines the maximum number of outstanding ready to transfer per

	task. The default value is 1.
MaxDataBurstLen	MaxDataBurstLen (Maximum Data Burst Length) determines the maximum SCSI data payload. The default value is 256kb.
DataSeginOrder	DataSeginOrder (Data Sequence in Order) determines if the PDU (Protocol Data Units) are transferred in continuously non-decreasing sequence offsets. The default value is Yes.
DataPDU InOrder	DataPDU InOrder (Data PDU in Order) determines if the data PDUs within sequences are to be in order and overlays forbidden. The default value is Yes.

List Connection Details

Click ▼ -> **Connection Details** which will list all connection(s) of the selected session.

Disconnect Session

Click ▼ -> **Disconnect** will disconnect the selected session, click **OK** button to confirm.

5.3. Configure for Fibre Channel Connectivity

Fibre Channel Ports tab shows information of fibre channel ports where are located on host card. Each fibre channel port can be configured by clicking ▼ of the specific port.

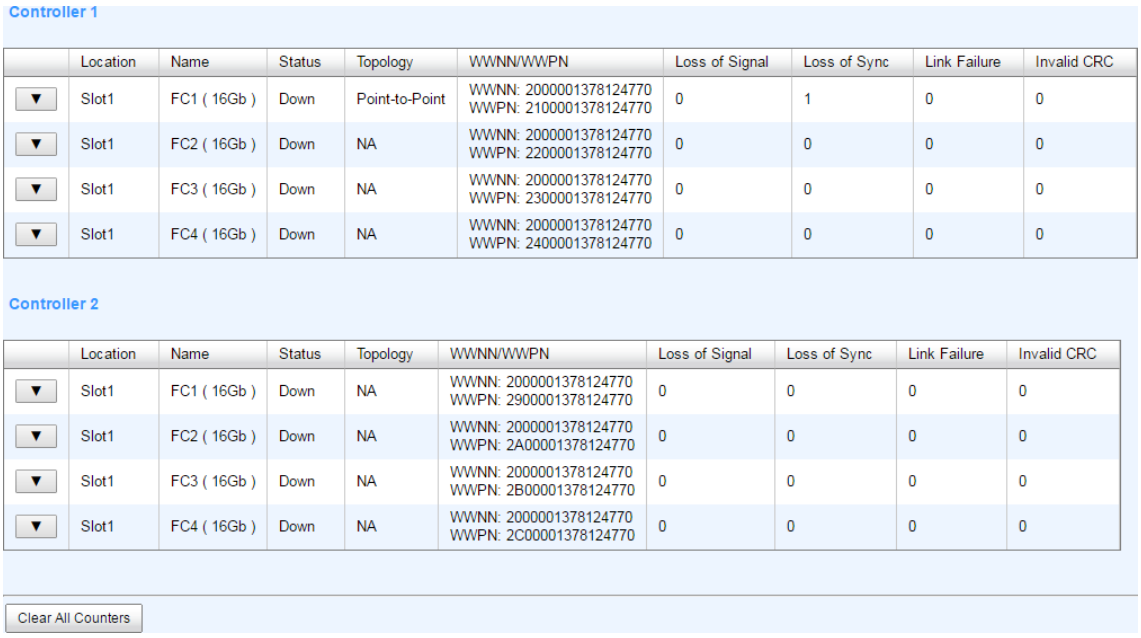

Figure 5-19 List Fibre Channel Ports

The columns display information of **Location**, **Name**, **Status**, **Topology**, **WWNN** (World Wide Node Name), **WWPN** (World Wide Port Name), and some statistical information.

5.3.1. Configure for FC Link Speed

Click ▼ -> **Change Link Speed** to change the link speed of fibre channel.

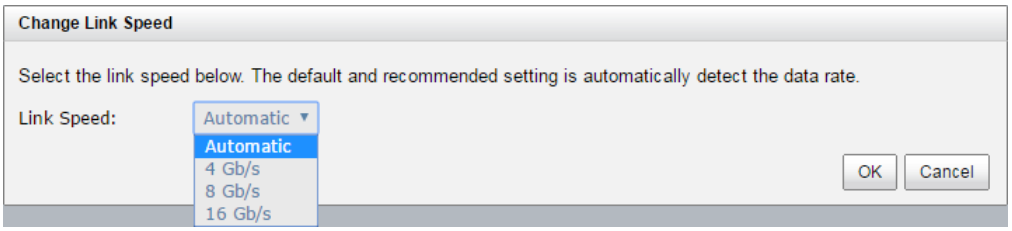

Figure 5-20 Change Link Speed

5.3.2. Configure for FC Topology

Click ▼ -> **Change Topology** to change the topology of fibre channel. It supports Fabric if link speed changes to 8 Gb/s.

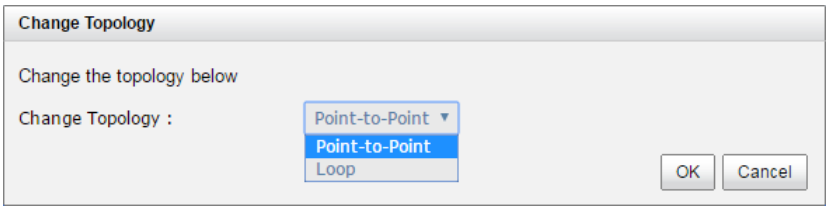

Figure 5-21 Change Topology

NOTE:

Point-to-Point (FC-P2P): Two devices are connected directly by FC interface. This is the simplest topology with limited connectivity.

Loop (FC-AL, Arbitrated Loop): All devices are connection in loop or ring, similar to token ring networking. Add or remove any device will affect activities on the loop. The failure of any device will cause ring broken. Fibre Channel hub connects multiple devices together and may bypass the failed ports. A loop may also be made by cabling each port to the next in a ring.

Fabric (Switched Fabric)(FC-SW): All devices or loops of devices are connected to Fibre Channel switches, similar conceptually to modern Ethernet implementations. Advantages of this topology over FC-P2P or FC-AL include.

5.3.3. Configure for FC Targets

Click ▼ -> **Target Configuration** to set multi-target configurations which are accessible by the host.

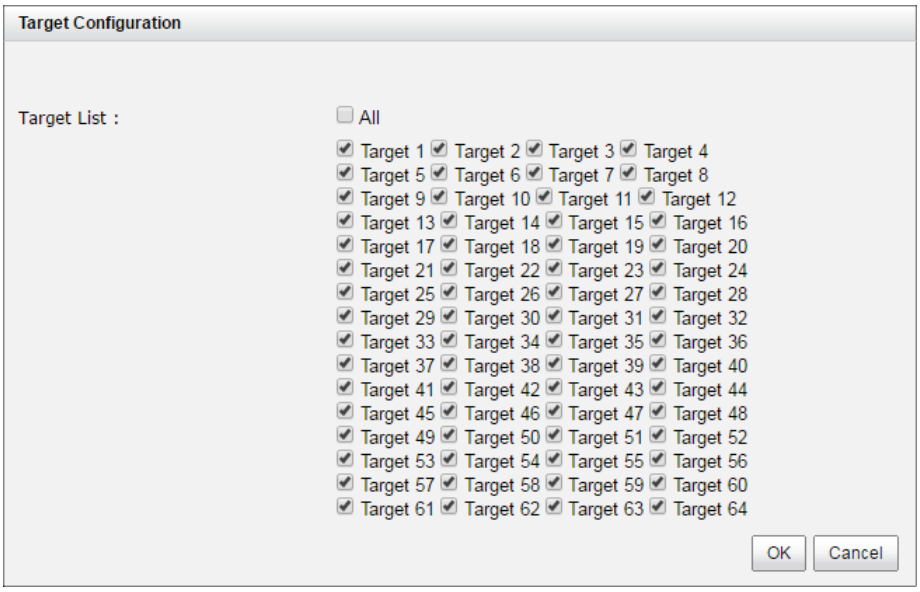

Figure 5-22 Target Configuration

CAUTION:

Point-to-Point connection mode does not support multi-target.

5.3.4. Clear FC Counters

Click **Clear All Counters** button to clear all counters of fibre channels.

Click ▼ -> **Clear Counters** to clear the counters of the selected fibre channel.

6. Storage Management

The **STROAGE MANAGEMENT** function menu provides submenus of **Disks**, **Pools**, **Volume**, and **Mappings**.

6.1. Working with Disk Drives

The **Disks** tab provides the status of the disk drives in the system. A system contains up to 446 disk drives with expansion units. For more information about system topology, please see QSAN XCubeSAN Hardware Manual. The document is available at this website:

https://www.qsan.com/en/download.php?cid=&keywords=XCubeSAN+Hardware+Manual&act=query

Table 6-1 Disk Parameters

Item	Value
Maximum Disk Quantity in a System	446

6.1.1. List Disks

The two drop-down lists at the top enable you to switch between the local system and any expansion JBOD systems attached. The other is to change the drive capacity units (MB or GB). Each disk drive can be configured by clicking ▼ of the specific one.

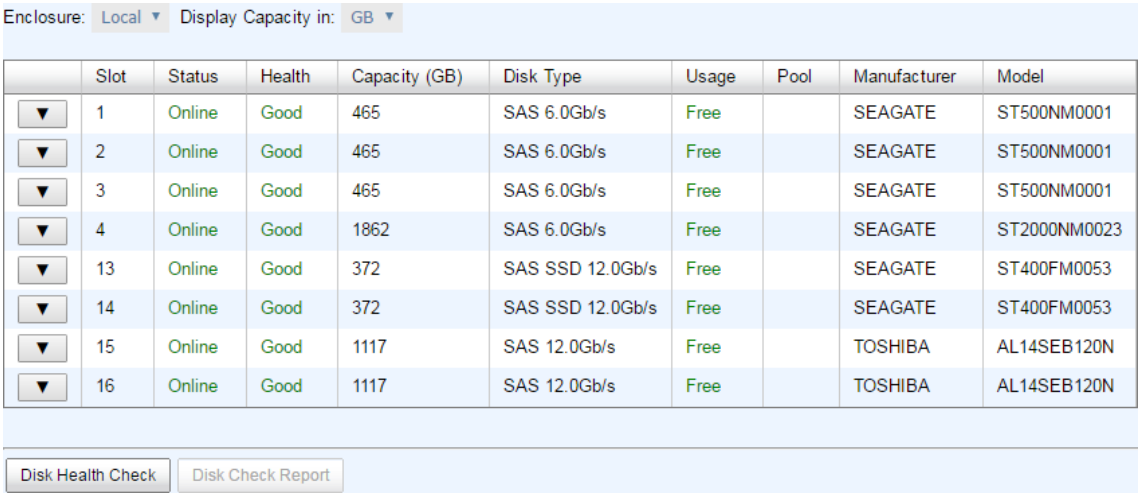

Figure 6-1 List Disks

This table shows the column descriptions.

Table 6-2 Disk Column Descriptions

Column Name	Description
Slot	The position of the disk drive.
Status	The status of the disk drive: • Online: The disk drive is online. • Rebuilding: The disk drive is being rebuilt. • Transitioning: The disk drive is being migrated or is replaced by another disk when rebuilding occurs. • Scrubbing: The disk drive is being scrubbed.
Health	The health of the disk drive: • Good: The disk drive is good. • Failed: The disk drive is failed. • Errors Detected: The disk drive has errors detected. • Error Alert: S.M.A.R.T. error alerts. • Read Errors: The disk drive has unrecoverable read errors.
Capacity (GB) or (MB)	The capacity of disk drive. The unit can be displayed in GB or MB.
Pool	The disk drive belongs to which pool.
Disk Type	The type of disk drive: • SAS 12.0Gb/s • SAS 6.0Gb/s • SAS 3.0Gb/s • SAS SSD 12.0Gb/s • SAS SSD 6.0Gb/s • SATA 6.0Gb/s • SATA 3.0Gb/s • SATA 1.5Gb/s • SATA SSD 6.0Gb/s
Usage	The usage of the disk drive: • RAID: This disk drive has been set to a pool. • Free: This disk drive is free for use. • Dedicated Spare: This disk drive has been set as dedicated spare of a pool. • Local Spare: This disk drive has been set as local spare of the enclosure. • Global Spare: This disk drive has been set as global spare of whole system.
Manufacturer	The manufacturer of the disk drive.
Model	The model name of disk drive.

6.1.2. Operations on Disks

The options are available on this tab:

Disk Health Check

Click **Disk Health Check** button to check the health of the selected disks. Also select the quantity of bad block to stop disk health check. And then click OK button to proceed. It cannot check the disks which are in used.

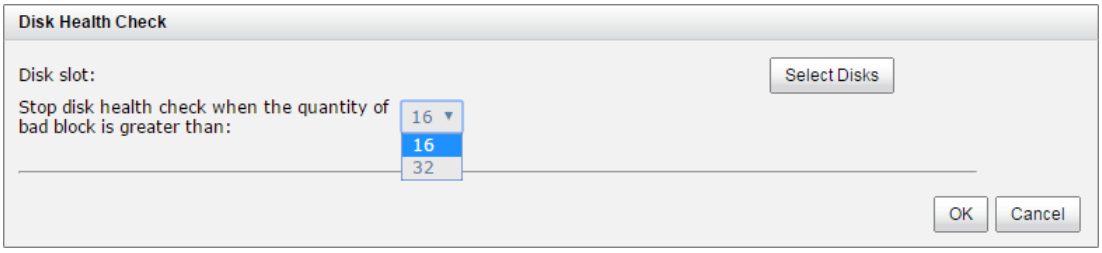

Figure 6-2 Disk Health Check

NOTE:

Disk Health Check function cannot check the disk drives which are in used.

Disk Check Report

Click **Disk Check Report** button to download the disk check report. It's available after executing **Disk Health Check**.

Set Free Disk

Click ▼ **-> Set Free Disk** to make the selected disk drive be free for use.

Set Spare Disk

Click ▼ **-> Set Global Spare** to set the selected disk drive to global spare of all pools.

Click ▼ **-> Set Local Spare** to set the selected disk drive to local spare of the pools which located in the same enclosure.

Click ▼ **-> Set Dedicated Spare** to set a disk drive to dedicated spare of the selected pool.

Take an example to set the disk to dedicated spare disk.

1. Check ▼ **-> Set Dedicated Spare** at one disk.

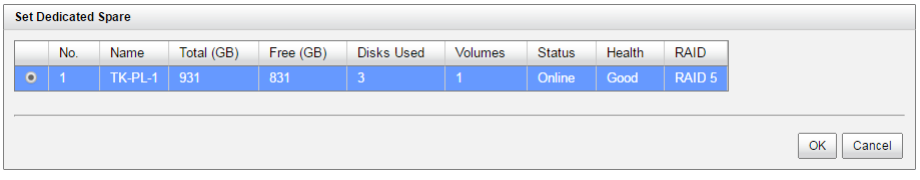

Figure 6-3 Set Dedicated Spare

2. If there is any pool which is in protected RAID level and can be set with dedicate spare disk, select one pool, and then click **OK** button.

Disk Scrub and Clear Disk Read Error

Click ▼ **-> Disk Scrub** which will scrub the disk drive. It's not available when the disk drive is in used.

Click ▼ **-> Clear Disk Read Error** to clean the read error of the disk drive.

Upgrade Disk Firmware

Click ▼ **-> Upgrade Disk Firmware** to upgrade the firmware of the disk drive.

Turn on Disk LED

Click ▼ **-> Turn on Disk LED** to turn on the indication LED of the disk drive.

Turn off Disk LED

Click ▼ **-> Turn off Disk LED** to turn off the indication LED of the disk drive.

More Information of the Disk

Click ▼ **-> More Information** to show the detail information of the disk drive.

6.1.3. S.M.A.R.T.

S.M.A.R.T. (Self-Monitoring Analysis and Reporting Technology) is a diagnostic tool for disk drives to deliver warning of drive failures in advance. It provides users a chance to take actions before a possible drive failure.

S.M.A.R.T. measures many attributes of disk drives all the time and inspects the properties of disk drives which are close to being out of tolerance. The advanced notice of possible disk drive failure will allow users to back up the data of disk drive or replace the disk drive. This is much better than disk drive crash when it is writing data or rebuilding a failed disk drive.

This tool displays S.M.A.R.T. information of disk drives. The number is the current value; the number in parenthesis is the threshold value. The threshold values from different disk drive vendors are different; please refer to the specification of disk drive for details.

S.M.A.R.T. supports SATA drives only. SAS drives have temperature parameter; the other columns will display N/A.

6.2. Configuring Storage Pools

The **Pools** tab provides to create, modify, delete, or view the status of the pools. We will describe thick provisioning pool on the following section, and keep thin provisioning in Thin Provisioning chapter and auto tiering in Auto Tiering chapter. A thick provisioning pool contains up to 64 disk drives.

Table 6-3 Thick Provisioning Pool Parameters

Item	Value
Maximum Disk Drive Quantity in a Pool (Include Dedicated Spares)	64

Maximum Pool Quantity per System	64
Maximum Dedicated Spare Quantity in a Pool	8

6.2.1. Create a Thick Provisioning Pool

Here is an example of creating a thick provisioning pool.

1. In **Pools** tab, click **Create Pool** button. It will scan available disks first.

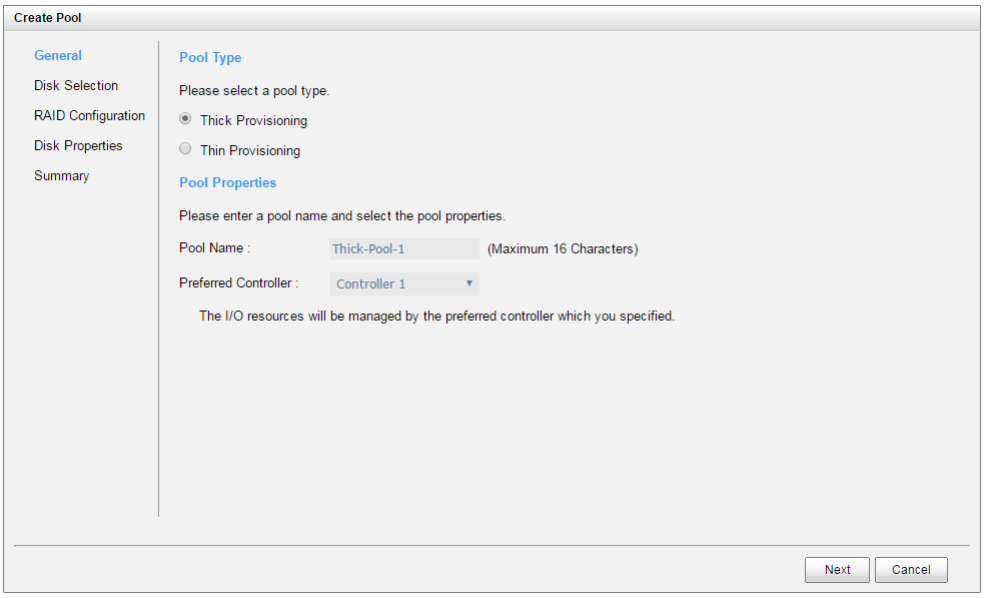

Figure 6-4 Create a Thick Provision Pool Step 1

2. Select **Pool Type** as **Thick Provisioning**.
3. Enter a **Pool Name** for the pool. Maximum length of name is 16 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
4. Select a **Preferred Controller** from the drop-down list. The I/O resources will be managed by the preferred controller which you specified. This option is available when dual controllers are installed.
5. Click **Next** button to continue.

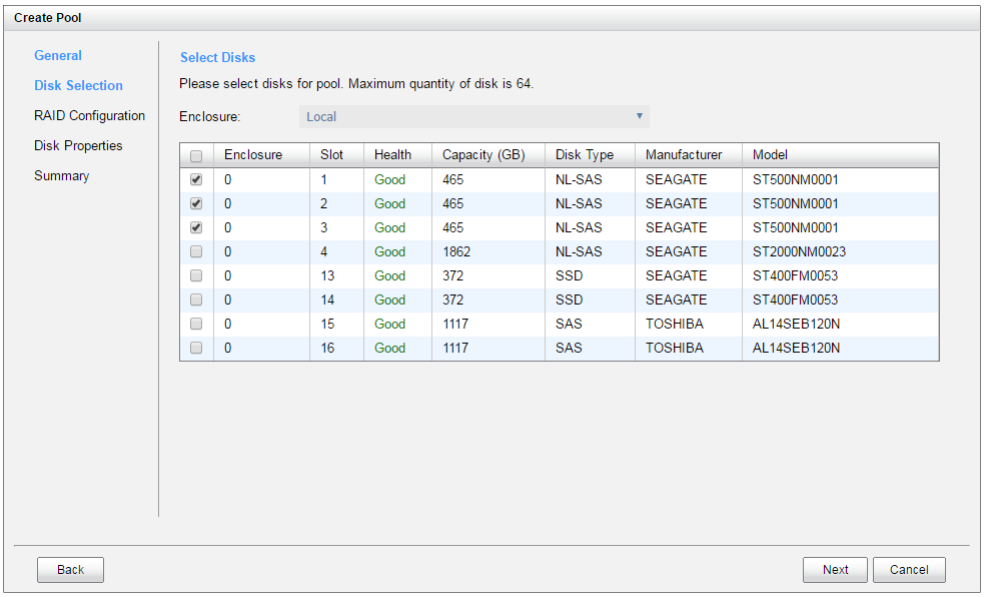

Figure 6-5 Create a Thick Provision Pool Step 2

6. Please select disks for pool. Maximum quantity of disk is 64. You can also select disks from expansion enclosures. Select an **Enclosure** from the drop-down list.
7. Click **Next** button to continue.

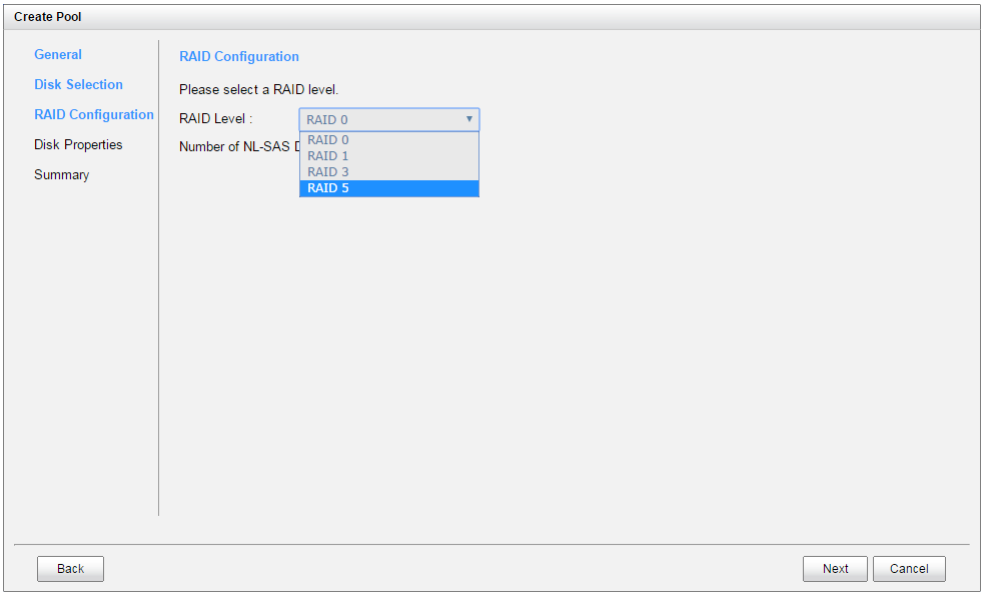

Figure 6-6 Create a Thick Provision Pool Step 3

8. Select a **RAID Level** from the drop-down list which lists available RAID level only according to the disk selection.
9. Click **Next** button to continue.

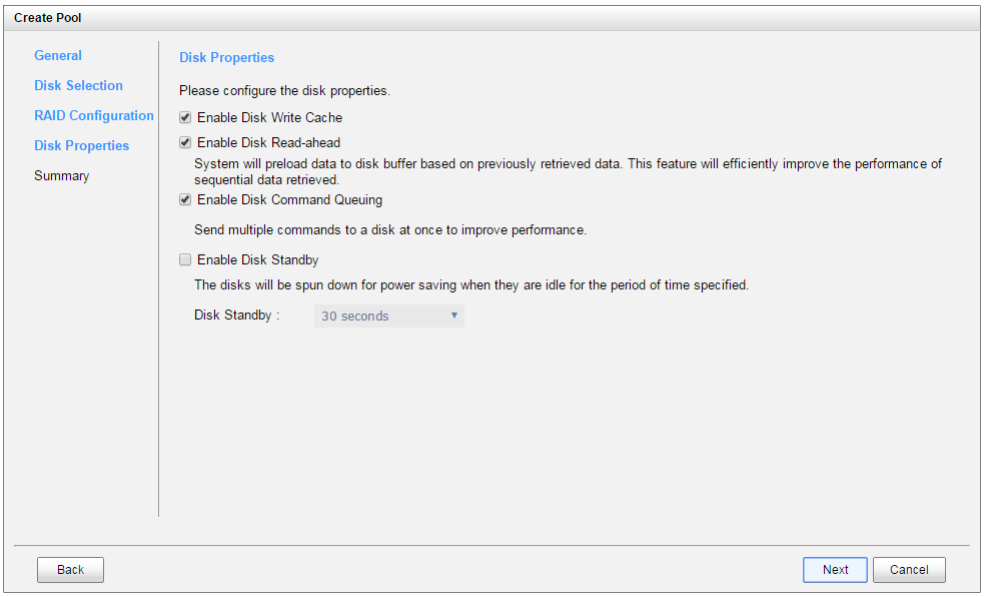

Figure 6-7 Create a Thick Provision Pool Step 4

10. Optionally, configure the disk properties:
    - **Enable Disk Write Cache:** Check to enable the write cache option of disks.
    - **Enable Disk Read-ahead:** Check to enable the read-ahead function of disks. System will preload data to disk buffer based on previously retrieved data. This feature will efficiently improve the performance of sequential data retrieved.
    - **Enable Disk Command Queuing:** Check to enable the command queue function of disks. Send multiple commands to a disk at once to improve performance.
    - **Enable Disk Standby:** Check to enable the auto spin down function of disks. The disks will be spun down for power saving when they are idle for the period of time specified.
11. Click **Next** button to continue.

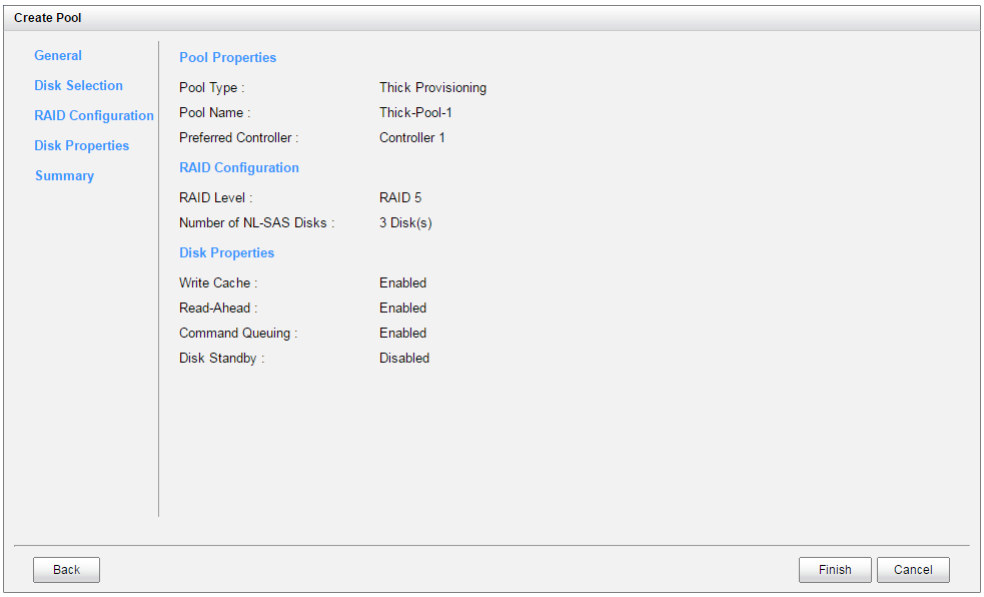

Figure 6-8 Create a Thick Provision Pool Step 5

12. After confirmation at summary page, click **Finish** button to create a pool.

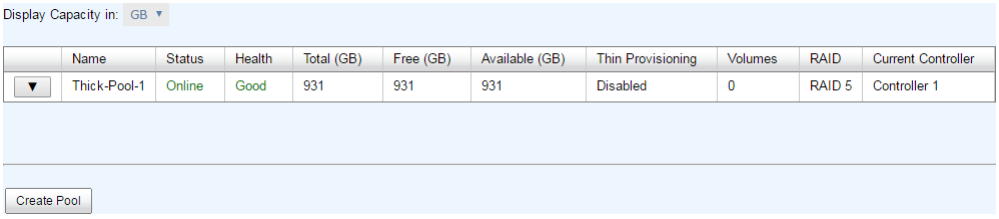

Figure 6-9 Thick Provisioning Pool is Created

13. The pool has been created. If necessary, click **Create Pool** button to create another.

6.2.2. List Thick Provisioning Pools

Use the drop-down list at the top to change the drive capacity units (MB or GB). Select one of the pools; it will display the related disk drives. Each pool can be configured by clicking ▼ of the specific one.

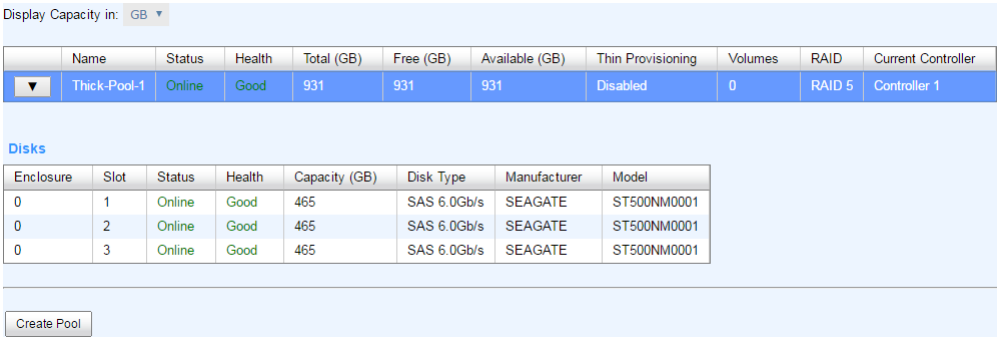

Figure 6-10 List Thick Provisioning Pools

This table shows the column descriptions.

Table 6-4 Pool Column Descriptions

Column Name	Description
Name	Pool name.
Status	The status of the pool: • Online: The pool is online. • Offline: The pool is offline. • Rebuilding: The pool is being rebuilt. • Migrating: The pool is being migrated. • Scrubbing: The pool is being scrubbed. • Relocation: The pool is being relocation.

Health	The health of the pool: • Good: The pool is good. • Failed: The pool fails. • Degraded: The pool is not healthy and not completed. The reason could be lack of disk(s) or have failed disk.
Total (GB) or (MB)	Total capacity of the pool. The unit can be displayed in GB or MB.
Free (GB) or (MB)	Free capacity of the pool. The unit can be displayed in GB or MB.
Available (GB) or (MB)	Available capacity of the pool. The unit can be displayed in GB or MB.
Thin Provisioning	The status of Thin provisioning: • Disabled. • Enabled.
Volumes	The quantity of volumes in the pool.
RAID	The RAID level of the pool.
Current Controller (*This option is only visible when dual controllers are installed.*)	The current running controller of the pool.

Table 6-5 Disk Column Descriptions

Column Name	Description
Enclosure	The number of enclosure.
Slot	The position of the disk drive.
Status	The status of the disk drive: • Online: The disk drive is online. • Missing: The disk drive is missing in the pool. • Rebuilding: The disk drive is being rebuilt. • Transitioning: The disk drive is being migrated or is replaced by another disk when rebuilding occurs. • Scrubbing: The disk drive is being scrubbed.
Health	The health of the disk drive: • Good: The disk drive is good. • Failed: The disk drive is failed. • Errors Detected: The disk drive has errors detected. • Error Alert: S.M.A.R.T. error alerts. • Read Errors: The disk drive has unrecoverable read errors.
Capacity (GB) or (MB)	The capacity of disk drive. The unit can be displayed in GB or MB.
Disk Type	The type of disk drive: • SAS 12.0Gb/s • SAS 6.0Gb/s • SAS 3.0Gb/s

	• SAS SSD 12.0Gb/s • SAS SSD 6.0Gb/s • SATA 6.0Gb/s • SATA 3.0Gb/s • SATA 1.5Gb/s • SATA SSD 6.0Gb/s
Manufacturer	The manufacturer of the disk drive.
Model	The model name of disk drive.

6.2.3. Operations on Thick Provisioning Pools

The options are available on this tab:

Activate and Deactivate the Pool

Click ▼ **-> Activate/Deactivate**, these options are for the purpose of online disk roaming. Deactivate can be executed when the status is online. Conversely, activate can be executed when the pool status is offline. For more information, please see Disk Roaming section in this chapter.

Change Disk Properties of the Pool

Click ▼ **-> Change Disk Properties** to change disk properties of the pool.

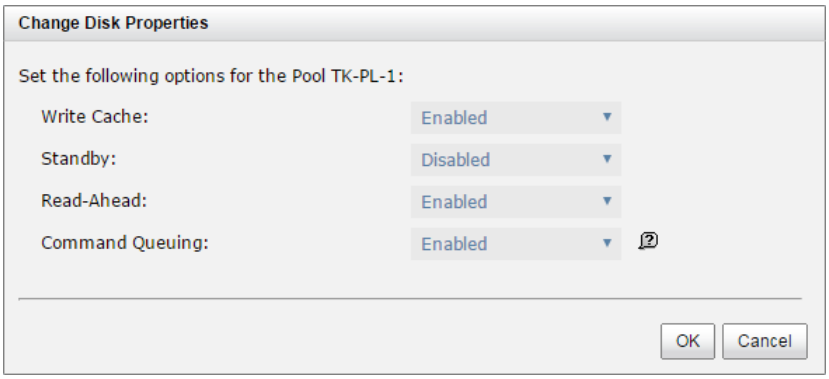

Figure 6-11 Change Disk Properties

Change Thin Provisioning Policy of the Pool

Click ▼ **-> Change Thin Provisioning Policy** to change policy of the thin provisioning pool. For more information, please see Thin Provisioning chapter.

Change Preferred Controller of the Pool

Click ▼ **-> Change Preferred Controller** to set the pool ownership to the other controller.

Verify Parity of the Pool

Click ▼ **-> Verify Parity** which will regenerate parity for the pool. It supports the RAID level 3 / 5 / 6 / 30 / 50 / 60.

Add a Disk Group into the Pool

Click ▼ **-> Add Disk Group** to add a disk group for the thin provisioning pool. For more information, please see Thin Provisioning chapter.

Migrate a Thick Provisioning Pool

Click ▼ **-> Migrate Pool** to change the RAID level of a pool or move the member disk drives of the pool to different disks. For more information, please see Migrate a Thick Provisioning Pool section in this chapter.

Delete a Pool

Click ▼ **-> Delete** to delete the pool. The pool cannot be deleted when there are volumes in the pool.

More Information of the Pool

Click ▼ **-> More Information** to show the detail information of the pool.

6.2.4. Migrate a Thick Provisioning Pool

Migrate Pool function changes the pool to different RAID level or adds the member disks of the pool for larger capacity. Usually, the pool migrates to higher RAID level for better protection. To do migration, the total size of pool must be larger than or equal to the original pool. The limitation is that it's not allowed expanding the same RAID level with the same physical disks of the original pool. If the RAID level doesn't change, the function can also move the member disks of the pool to totally different disk drives.

In addition, thin provisioning pool cannot execute migrate or move, it uses **Add Disk Group** to enlarge capacity. For more information, please see Thin Provisioning chapter.

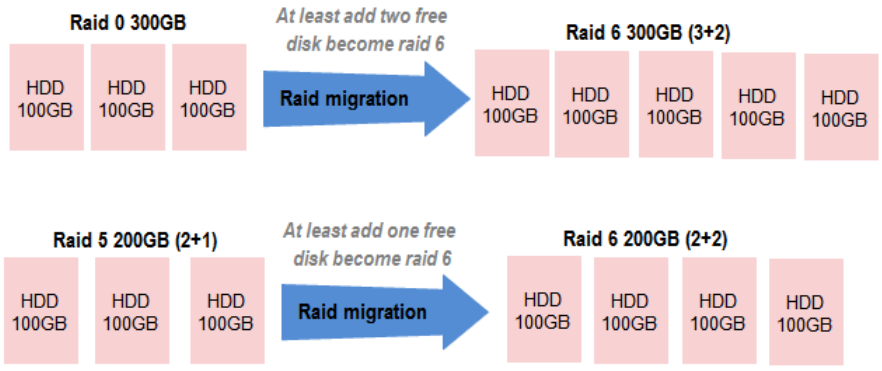

Figure 6-12 Migrate a Thick Provisioning Pool

Take an example of migrate a thick provisioning pool.

1. In **Pools** tab, Select a pool, and then click ▼ **-> Migrate Pool**
2. Select a **RAID Level** from the drop-down list.

3. Click the **Select Disks** button to select disks from either local or expansion JBOD systems, and click **OK** to complete the selection. The selected disks are displayed at **Disks Used**.

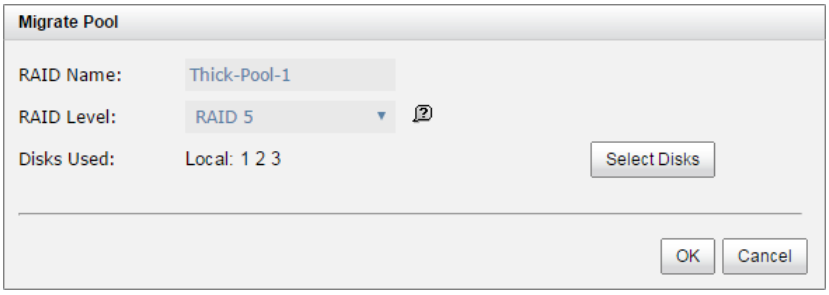

Figure 6-13 Migrate RAID Level

4. At the confirmation dialog, click **OK** button to execute migration.
5. Migration starts. The status of **Disks**, **Pools** and **Volumes** are changing. The complete percentage of migration is displayed in **Status**.
6. It's done when the complete percentage reaches 100%.

6.2.5. Migrate Pool Notices

There are some notices when a pool is being migrated or moved. System would reject these operations:

1. Add dedicated spare.
2. Remove a dedicated spare.
3. Create a new volume.
4. Delete a volume.
5. Extend a volume.
6. Scrub a volume.
7. Perform another migration operation.
8. Scrub entire pool.
9. Take a snapshot.
10. Delete a snapshot.
11. Expose a snapshot.
12. Rollback to a snapshot.

NOTE:

Migrate function will migrate the member disks of pool to the same disks but it should increase the number of disks or it should be different RAID level. Move function will move the member disks of pool to totally different physical

disks.

CAUTION:

Pool migration or moving cannot be executed during rebuilding or volume extension.

6.3. Configuring Volume

The **Volumes** tab provides to create, modify, delete, or view the status of the volumes. The same as pool, we will describe thick provisioning volume on the following section, and keep thin provisioning in Thin Provisioning chapter and auto tiering in Auto Tiering chapter. A pool contains up to 96 volumes and a system can contain up to 4,096 volumes including snapshot volumes.

Table 6-6 Volumes Parameters

Item	Value
Maximum Volume Quantity in a Pool	96
Maximum Volume Quantity Per System (include Snapshot Volumes)	4,096
Maximum Host Number Per Volume	16
Provisioning Granularity	1 GB
Maximum Thin Volume Capacity	128 TB

6.3.1. Create a Volume

Here is an example of creating a volume in think provisioning Pool.

1. In **Volume** tab, click **Create Volume** button.

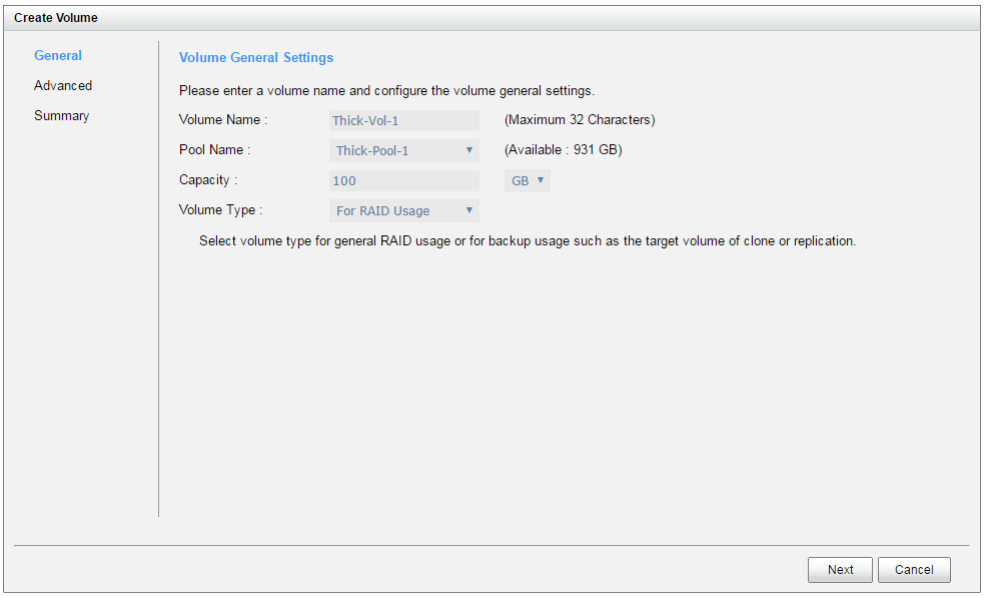

Figure 6-14 Create a Volume in Thick Provisioning Pool Step 1

2. Enter a **Volume Name** for the pool. Maximum length of name is 32 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
3. Select a **Pool Name** from the drop-down list. It will also display the available capacity of the pool.
4. Enter required **Capacity**. The unit can be selected from the drop-down list.
5. Select **Volume Type**. The options are **For RAID Usage** (for general usage) and **For Backup Usage** (for the target of local clone or remote replication).
6. Click **Next** button to continue.

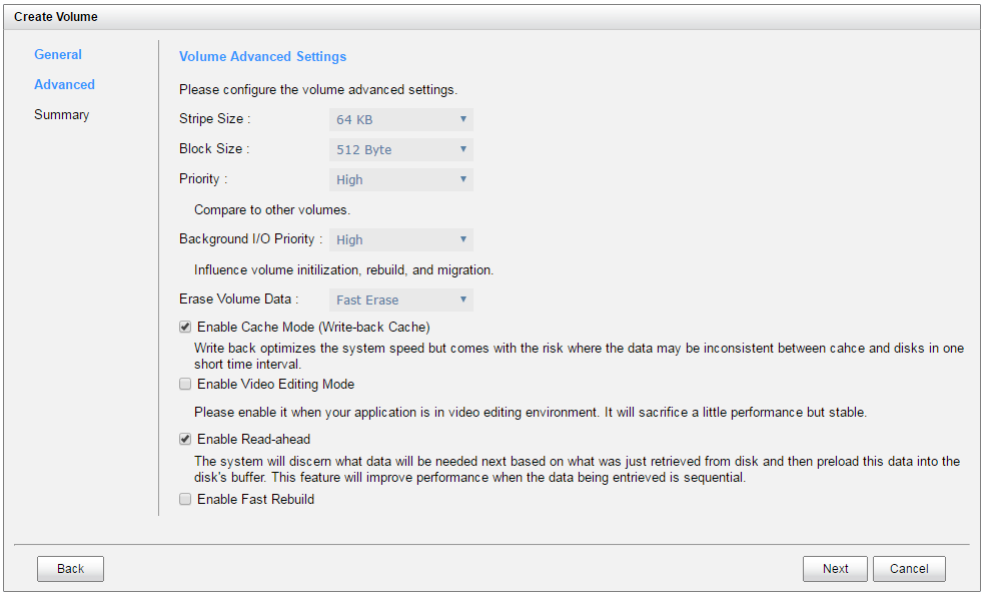

Figure 6-15 Create a Volume in Thick Provisioning Pool Step 2

7. Optionally, configure the volume advanced settings:
    - **Stripe Size:** The options are 4 KB to 64 KB.
    - **Block Size:** The options are 512 Bytes to 4,096 Bytes.
    - **Priority:** The options are High, Medium, and Low. The priority compares to other volumes. Set it as High if the volume has many I/O.
    - **Background I/O Priority:** The options are High, Medium, and Low. It will influence volume initialization, rebuild, and migration. Set it as High if the volume is important.
    - **Erase Volume Data:** This option is available when the pool is thick provisioning. This option will wipe out old data in volume to prevent that OS recognizes the old partition. The options are Do Not Erase, Fast Erase or Full Disk.
    - **Enable Cache Mode (Write-back Cache):** Check to enable cache mode function of volume. Write back optimizes the system speed but comes with the risk where the data may be inconsistent between cache and disks in one short time interval.
    - **Enable Video Editing Mode:** Check to enable video editing mode function. It is optimized for video editing usage. Please enable it when your application is in video editing environment. It will sacrifice a little performance but stable.
    - **Enable Read-ahead:** Check to enable the read ahead function of volume. The system will discern what data will be needed next based on what was just retrieved from disk and then preload this data into the disk's buffer. This feature will improve performance when the data being retrieved is sequential.
    - **Enable Fast Rebuild:** This option is available when the pool needs to rebuild. Let's say, RAID 1 / 3 / 5 / 6 / 0+1 / 10 / 30 / 50 / 60. For more information, please see Fast Reuild section in this chapter.
8. Click **Next** button to continue.

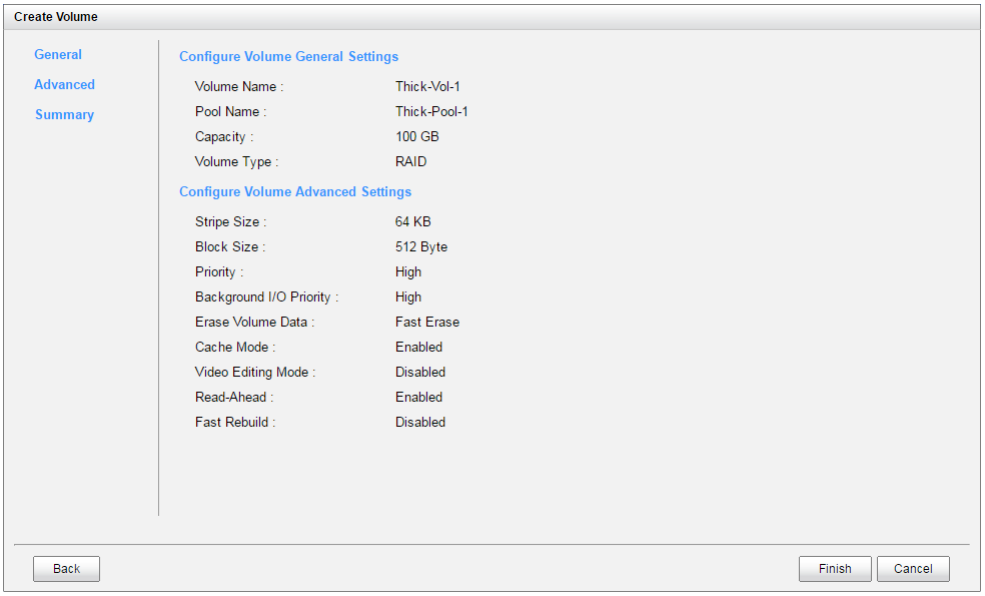

Figure 6-16 Create a Volume in Thick Provisioning Pool Step 3

9. After confirmation at summary page, click **Finish** button to create a volume.

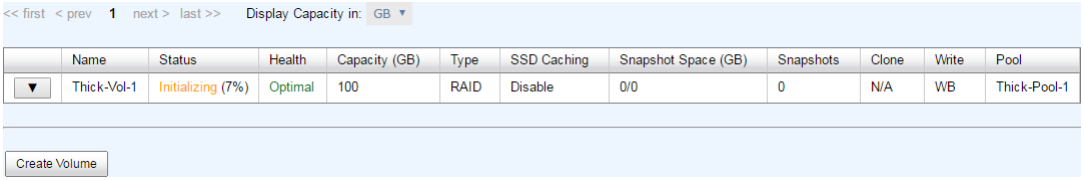

Figure 6-17 Volume in Thick Provisioning Pool is Created

10. The volume has been created. If necessary, click **Create Volume** button to create another.

6.3.2. List Volumes

Use the drop-down list at the top to change the drive size units (MB or GB). Select one of the volumes; it will display the related LUN Mappings if the volume is mapped. Each volume can be configured by clicking ▼ of the specific one.

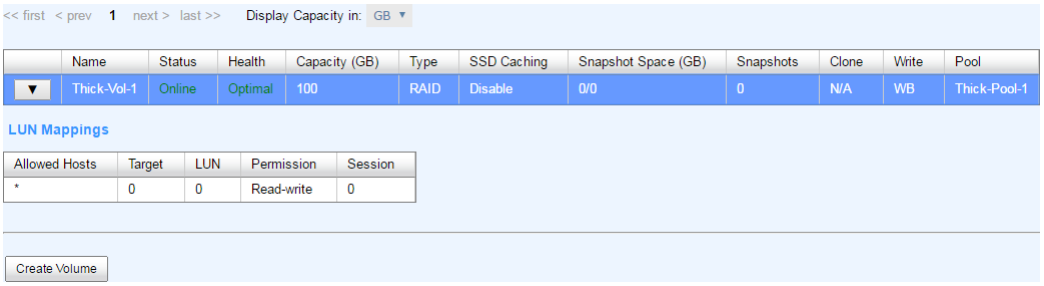

Figure 6-18 List Volumes

This table shows the column descriptions.

Table 6-7 Volume Column Descriptions

Column Name	Description
Name	Volume name.
Status	The status of the volume: • Online: The volume is online. • Offline: The volume is offline. • Erasing: The volume is being erased if **Erase Volume Data** is set. • Initiating: The volume is being initialized. • Rebuilding: The volume is being rebuilt. • Migrating: The volume is being migrated. • Rollback: The volume is being rolled back. • Parity Checking: The volume is being parity check. • Relocation: The volume is being relocation.
Health	The health of volume:

	• Optimal: the volume is working well and there is no failed disk in the RG. • Failed: the pool disk of the VD has single or multiple failed disks than its RAID level can recover from data loss. • Degraded: At least one disk from the RG of the Volume is failed or plugged out. • Partially optimal: the volume has experienced recoverable read errors. After passing parity check, the health will become Optimal.
Capacity (GB) or (MB)	Total capacity of the volume. The unit can be displayed in GB or MB.
SSD Caching	The SSD caching policy: • Disable: Disable SSD caching. • Database: Enable SSD caching and set it as database policy. • File System: Enable SSD caching and set it as file system policy. • Web Service: Enable SSD caching and set it as web service policy. • Custom: Enable SSD caching and set it as customization policy.
Snapshot space (GB) or (MB)	The volume size that is used for snapshot. The number means Used snapshot space / Total snapshot space. The unit can be displayed in GB or MB.
Clone	The target name of the clone volume.
Write	The right of volume: • WT: Write Through. • WB: Write Back. • RO: Read Only.
Pool	The volume belongs to which pool.

6.3.3. Operations on Volumes

The options are available on this tab:

Setup Cloning Options

Click **Cloning Options** button to set the clone options. For more information, please see Cloning Options section in Data Backup chapter.

LUN Mapping Operations

Click ▼ **-> Map LUN** to map a logical unit number to the volume. For more information, please see Configure LUN Mappings section in this chapter.

Click ▼ **-> Unmap LUNs** to unmap logical unit numbers from the volume.

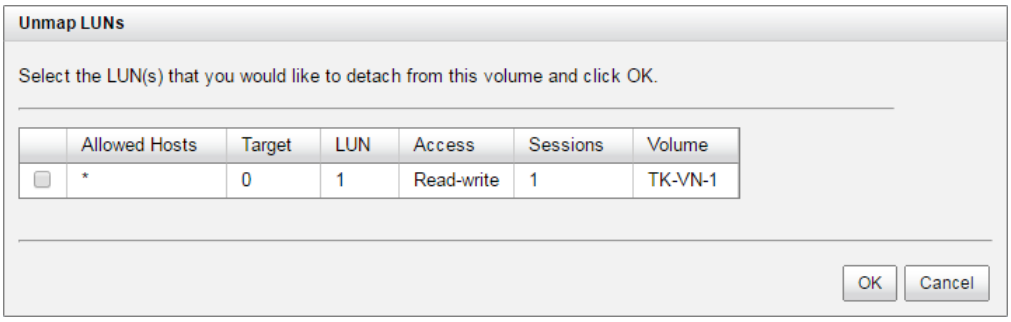

Figure 6-19 Unmap LUNs

Snapshot Operations

Click ▼ **-> Set Snapshot Space** to set snapshot space for preparing to take snapshots. For more information, please see Managing Snapshots section in Data Backup chapter.

Click ▼ **-> Take Snapshot** to take a snapshot on the volume.

Click ▼ **-> Schedule Snapshots** to set the snapshots by schedule.

Click ▼ **-> List Snapshots** to list all snapshots of the volume.

Click ▼ **-> Cleanup Snapshots** to clean all snapshots of the volume and release the snapshot space.

Setup SSD Caching

Click ▼ **-> Set SSD Caching** to set SSD caching for the volume. For more information, please see SSD Caching chapter.

Local Clone Operations

Click ▼ **-> Create Local Clone** to set the target volume for clone. For more information, please see Managing Local Clones section in Data Backup chapter.

Click ▼ **-> Clear Clone** to clear the clone.

Click ▼ **-> Start Clone** to start the clone.

Click ▼ **-> Stop Clone** to stop the clone.

Click ▼ **-> Schedule Clone** to set the clone function by schedule.

Click ▼ **-> Change Replication Options** to change the clone to Replication relationship. For more information, please see Managing Remote Replications section in Data Backup chapter.

Change Volume Properties

Click ▼ **-> Change Volume Properties** to change the volume properties of the volume.

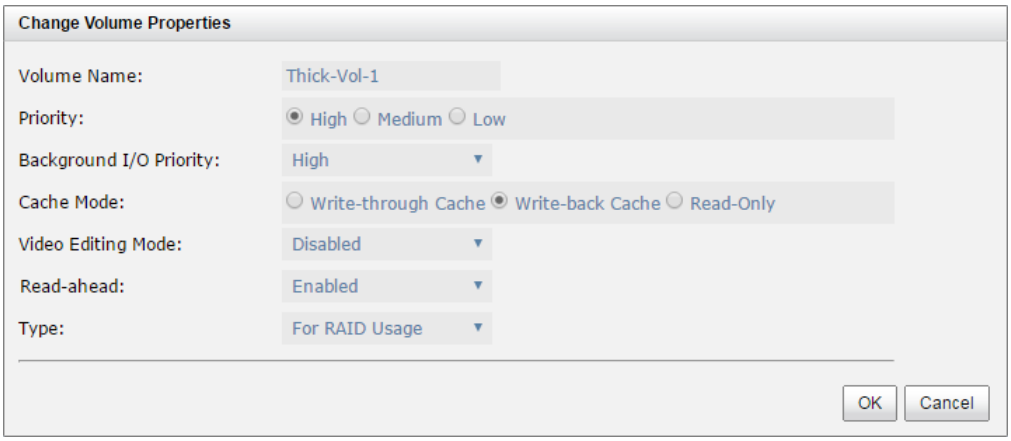

Figure 6-20 Change Volume Properties

Change Volume Properties

Click ▼ **-> Space Reclamation** to reclaim space for the volume. This option is available when the volume belongs to thin provisioning pool.

Verify Parity of the Volume

Click ▼ **-> Verify Parity** to execute parity check for the volume. It supports RAID 3 / 5 / 6 / 30 / 50 / 60. This volume can either be verified and repaired or only verified for data inconsistencies. This process takes longer.

Extent Volume Capacity

Click ▼ **-> Extend Volume** to extend the volume capacity. For more information, please see Extend Volume Capacity section in this chapter.

Delete Volume

Click ▼ **-> Delete** to delete the volume. The related LUN mappings will also be deleted.

More Information of the Volume

Click ▼ **-> More Information** to show the detail information of the volume.

6.3.4. Extend Volume Capacity

Extend Volume function extends the capacity of the volume if there is enough free space.

Take an example of extending the volume.

1. In **Volumes** tab, Select a volume, and then click ▼ **-> Extend Volume**.
2. Change the volume capacity. The capacity must be larger than the current, and then click **OK** button to start extension.

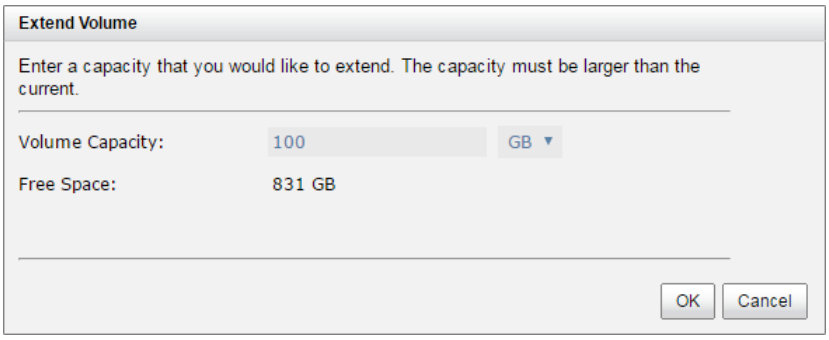

Figure 6-21 Extend Volume Capacity

3. If the volume needs initialization, it will display the status **Initiating** and the complete percentage of initialization in **Status**.
4. It's done when the complete percentage reaches 100%.

NOTE:

The extension capacity must be larger than the current capacity.

CAUTION:

Extension cannot be executed during rebuilding or migration.

6.3.5. Fast Rebuild

When executing rebuild, the Fast Rebuild feature skips any partition of the volume where no write changes have occurred, it will focus only on the parts that have changed. Enable the Fast Rebuild function when creating a volume. For more information, please refer to Fast Rebuild White Paper. The document is available at this website:

https://www.qsan.com/en/download.php?cid=&keywords=Fast+Rebuild&act=query

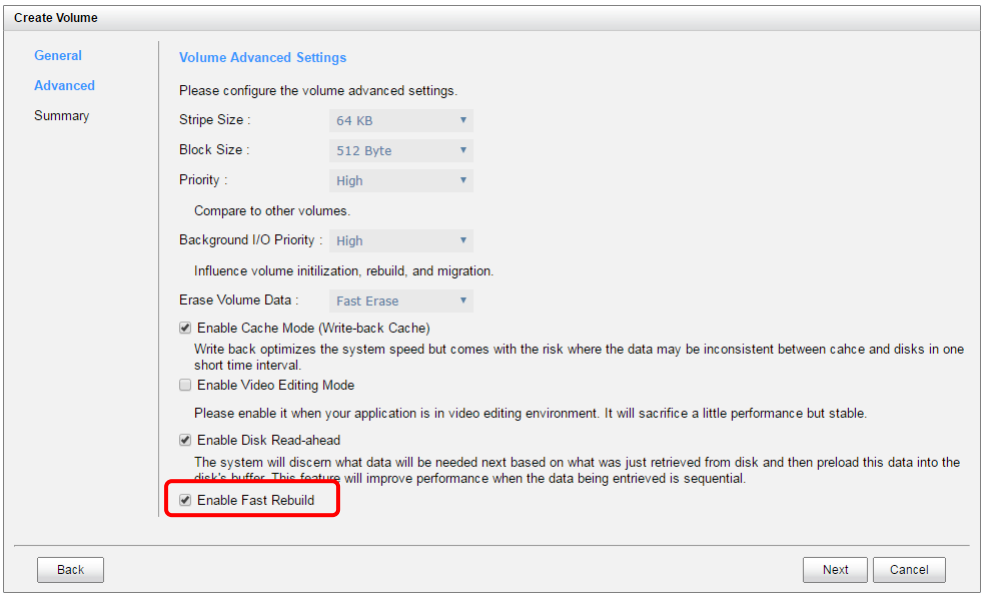

Figure 6-22 Enable Fast Rebuild When Creating a Volume

Fast Rebuild Notices

Here are some notices about Fast Rebuild.

- Only thick provisioning pool supports this feature. Thin provision pool already has this feature implement.
- When rebuild happened in a fast rebuild volume, clean partitions are not rebuilt since there are no data saved there. Though clean partitions are never rebuilt, their health status is good.
- If all partitions of the fast rebuild volume are clean, then no rebuild would happen and no event would be sent.
- The RAID stacks could not use optimize algorithm to compute parities of a partition which is not rebuilt. Thus, the performance of random write in a clean partition would be worse.

CAUTION:

Suggest not enabling Fast Rebuild function when the access pattern to the volume is random write.

6.4. Configure LUN Mappings

The **Mappings** tab provides to map, unmap or view the status of LUN (Logical Unit Number) for each volume. A system can contain up to 4,096 LUNs.

Table 6-8 LUN Parameters

Item	Value
Maximum Quantity of LUN	4,096

6.4.1. Map a LUN of iSCSI Connectivity

Take an example of mapping a LUN of iSCSI connectivity.

1. In **Mappings** tab, click **Map LUN** button.

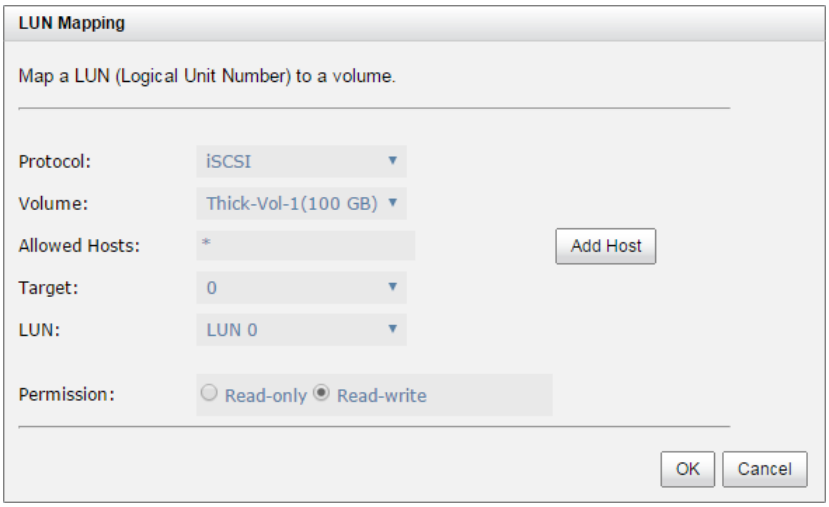

Figure 6-23 Map a LUN of iSCSI Connectivity

2. Select the **Protocol** as iSCSI.
3. Select a **Volume** from the drop-down list.
4. Enter the **Allowed Hosts** with semicolons (;) or click **Add Host** button to add one by one. Fill-in wildcard (*) for access by all hosts.
5. Select a **Target** from the drop-down list.
6. Select a **LUN** from the drop-down list.
7. Select a **Permission** level, normally set it as **Read-write**.
8. Click **OK** button to map a LUN.

The matching rules of access control are followed from created time of the LUNs. The earlier created LUN is prior to the matching rules. For example: there are two LUN rules which are set to the same volume:

1. Allow Hosts sets *, LUN is *LUN 0*
2. Allow Hosts sets *iqn.host1*, LUN *is LUN 1*

The host *iqn.host2* can login successfully because it matches the rule 1.

Wildcard * and ? are allowed in this field.

- Wildcard * can replace any word.
- Wildcard ?: can replace only one character.

For example:

- iqn.host? -> iqn.host1 and iqn.host2 are accepted.
- iqn.host* -> iqn.host1 and iqn.host12345 are accepted.

This field cannot accept comma, so iqn.host1, iqn.host2 stands a long string, not two iqns.

6.4.2. Map a LUN of FC Connectivity

Take an example of mapping a LUN of fibre channel connectivity.

1. In **Mappings** tab, click **Map LUN** button.

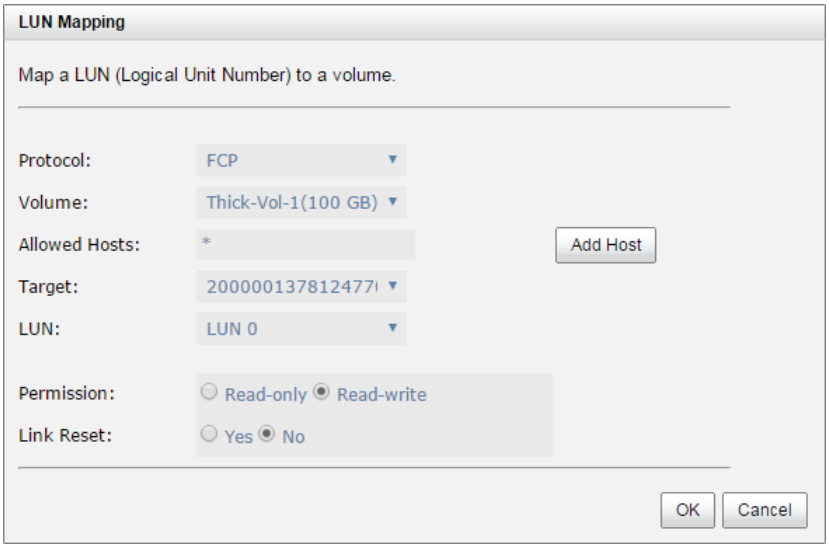

Figure 6-24 Map a LUN of FC Connectivity

2. Select the **Protocol** as FCP (Fibre Channel Protocol).
3. Select a **Volume** from the drop-down list.
4. Enter the **Allowed Hosts** with semicolons (;) or click **Add Host** button to add one by one. Fill-in wildcard (*) for access by all hosts.
5. Select a **Target** from the drop-down list.
6. Select a **LUN** from the drop-down list.
7. Select a **Permission** level, normally set it as **Read-write**.
8. Set **Link Reset** to **Yes** if fibre channel needs to reset when connected.
9. Click **OK** button to map a LUN.

6.4.3. List LUN Mappings

List all LUN mappings in this page. Each LUN can be configured by clicking ▼ of the specific one.

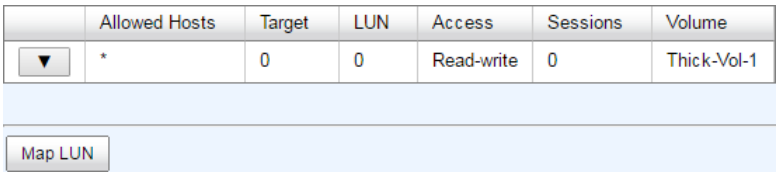

Figure 6-25 List LUN Mappings

This table shows the column descriptions.

Table 6-9 LUN Mapping Column Descriptions

Column Name	Description
Allowed Hosts	The target of FC / iSCSI for access control or a wildcard (*) for access by all hosts.
Target	The number of the target.
LUN	The number of the LUN assigned.
Permission	The permission level: • Read-write • Read-only
Sessions	The number of the active iSCSI connection linked to the logical unit. Show N/A if the protocol is FCP.
Volume	The name of the volume assigned to this LUN.

6.4.4. Operations on LUN Mappings

The options are available on this tab:

Unmap a LUN

Click ▼ **-> Unmap LUN** to unmap a logical unit number from the volume.

Active Sessions

Click ▼ **-> Active Sessions** to show the active sessions of iSCSI connection. This option is available when the protocol of LUN mapping is iSCSI.

6.5. Connect by Host Initiator

After map a LUN to a volume, the host can connect the volume by initiator program. We provide some documents of host connections for reference. The documents are available at the website:

- How to Configure iSCSI Initiator in Microsoft Windows

  https://www.qsan.com/en/download.php?cid=&keywords=iscsi&act=query
- How to Configure iSCSI Initiator in ESXi6.x

  https://www.qsan.com/en/download.php?cid=&keywords=iscsi&act=query
- Implement iSCSI Multipath in RHEL 6.5

https://www.qsan.com/en/download.php?cid=&keywords=iscsi&act=query

- Implement iSCSI in Mac OS X

  https://www.qsan.com/en/download.php?cid=&keywords=iscsi&act=query

7. Thin Provisioning

This chapter describes an overview and operations of thin provisioning.

7.1. Overview

Nowadays thin provisioning is a hot topic people talk about in IT management and storage industry. To make contrast to thin provisioning, it naturally brings to our minds with the opposite term – thick or fat provisioning, which is the traditional way IT administrators allocate storage space to each logical volume that is used by an application or a group of users. When it comes to the point to decide how much space a logical volume requires for three years or for the lifetime of an application, it's really hard to make the prediction correctly and precisely. To avoid the complexity of adding more space to the volumes frequently, IT administrators might as well allocate more storage space to each logical volume than it needs in the beginning. This is why it's called thick or fat provisioning. Usually it turns out that a lot of free space is sitting around idle. This stranded capacity is wasted, which equals to waste of investment and inefficiency. Various studies indicate that as much as 75% of the storage capacity in small and medium enterprises or large data centers is allocated but unused. And this is where thin provisioning kicks in.

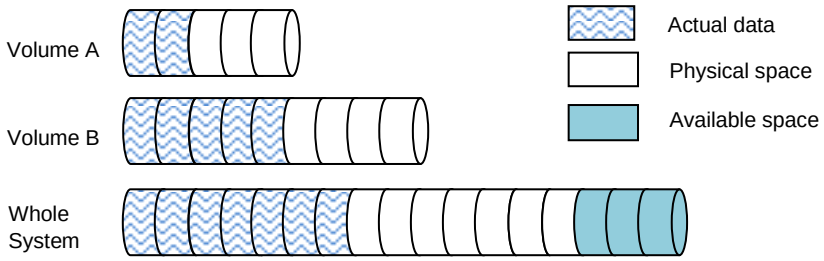

Figure 7-1 Traditional Thick Provisioning

Thin provisioning sometimes is known as just-in-time capacity or over allocation. As the term explains itself, it provides storage space by requests dynamically. Thin provisioning presents more storage space to the hosts or servers connecting to the storage system than is actually available on the storage system. Put it in another way, thin provisioning allocates storage space that may or may not exist. The whole idea is actually another way of virtualization. Virtualization is always about a logical pool of physical assets and provides better utilization over those assets. Here the virtualization mechanism behind thin provisioning is storage pool. The capacity of the storage pool is shared by all volumes. When write requests come in, the space will be drawn dynamically from this storage pool to meet the needs.

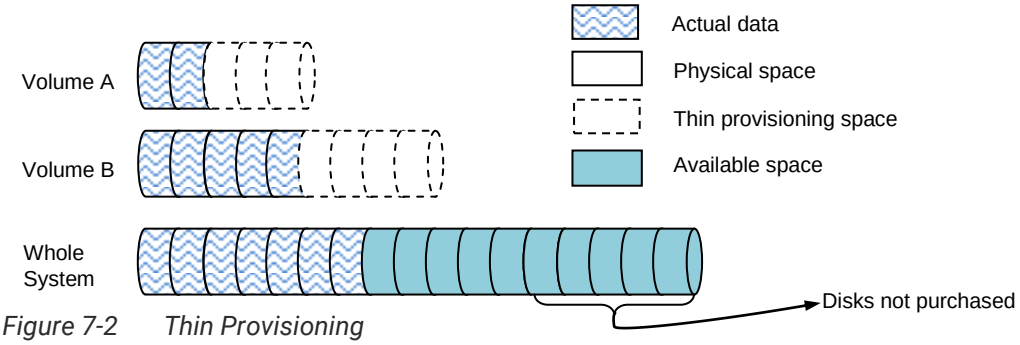

Figure 7-2 Thin Provisioning

7.1.1. Features Highlight

The following describes the features highlight of thin provisioning.

Write on Demand or Allocate on Demand

This is the most distinctive function in thin provisioning. Compare to the thick provisioning, it will occupy the space when there are data in. That's called write on demand or allocate on demand. And it shows the storage efficiency.

Expand Capacity on Demand without Downtime

Extra disk group can be added to the thin pool to increase the size of free storage pool. A thin pool can have up to 32 disk groups with each disk group containing up to 8 disk drives. The maximum size of each disk group is 64TB.

Allocation Unit (granularity) is 1GB

This is a number that demands careful balance between efficiency and performance. The smaller it is, the better the efficiency and the worse the performance becomes, and vice versa.

Thin Provisioning Snapshot Space and Snapshot Writable

Snapshot space sits at the same pool of the volume which the snapshot is taken against. Therefore when you expose the snapshot into a volume, it becomes a thin provisioning volume. It will only take up the just the right amount of space to store the data, not the full size of the volume.

Convert Traditional Volume to Thin Provisioning and Vice Versa

You can enjoy the benefits of thin provisioning right now and right this moment. Move all your existing thick provisioning volumes to thin provisioning ones. Volume clone function can be performed on both directions - thick-to-thin and thin-to-thick, depending on your application needs.

Automatic Space Reclamation to Recycle Unused Space

Automatic space reclamation will be automatically activated in pool initialization process or it can be set manually through capacity policy. When space usage ratio grows over the threshold set in the policy, space reclamation will be enabled automatically at the background with the lowest priority or when the I/O is low. The resource impact is reduced to minimum.

7.1.2. Benefits of Thin Provisioning

The benefits of thin provisioning are described on the following.

- Less disk purchase is needed initially when setting up a new storage system. You don't need to buy more capacity to meet your future data growth at present time. Usually hard drive price declines as time progresses. You can buy the same hard drives with cheaper price at a later time. Why not save money upfront while you can?
- No stranded storage capacity, better utilization efficiency and lower total cost of ownership. Thin provisioning can make full use of the stranded capacity that traditional provisioning can't. All free capacity can be made available to other hosts. A single storage system can serve more hosts and servers to achieve high consolidation ratio. Thin provisioning can help you achieve the same level of services with less hard drives purchased upfront, which can significantly reduce your total cost of ownership.

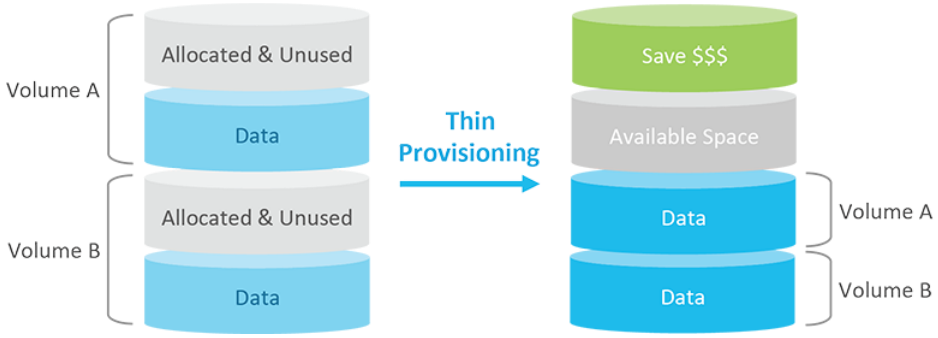

Figure 7-3 Benefits of Thin Provisioning

- Scalability: storage pool can grow on demand. When the storage pool has reached the threshold you set before. Up to 32 disk groups can be added to the pool to increase the capacity on demand without interrupting I/O. Each disk group can have up to 8 disks.
- Automatic space reclamation mechanism to recycle unused blocks. The technology used here is called zero reclamation. When a thin pool is created, the initialization process will try to fill out all the storage pool space with zero. This process will run in background with low priority in order not to impact the I/O performance. This is the reason why when there is no I/O traffic from the hosts, the hard drive LED will keep blinking as if there are I/O activities. The purpose of zero reclamation is that when the actual user data happens to have all zero in a basic allocation unit (granularity), the storage system will treat it as free space and recycle it. Until the next time there is data update to this reclaimed all zero basic unit, the storage system can swiftly return a basic unit from the free storage pool because it's already filled with zero.
- An eco-friendly green feature that helps to reduce energy consumption. Hard drive is the top power consumer in a storage system. Because you can use less hard drives to achieve the same amount of work, this translates directly to a huge reduction of power consumption and more green in your pocket.

7.2. Operations

Disk groups which contain disk drives can be added to a thin provisioning pool at any time to increase the capacity of the pool. For simplifying usage and better performance, every disk group must have the same quantity of disk drives. A thin provisioning pool can have up to 32 disk groups with each disk group containing up to 8 disk drives. The maximum capacity of each disk group is 64TB. So the maximum capacity in a system is 256TB.

Table 7-1 Thin Provisioning Pool Parameters

Item	Value
Maximum Disk Group Quantity in a Pool	32
Maximum Disk Drive Quantity in a Disk Group	8
Maximum Disk Drive Quantity in a Pool	256 (= 32 x 8)
Maximum Capacity of a Disk Group	64 TB
Maximum Thin Provisioning / Auto tiering Pool Capacity per System	256 TB

The section will describe all operations of thin provisioning.

7.2.1. Create a Thin Provisioning Pool

Here is an example of creating a thin provisioning pool.

1. In **Pools** tab, click **Create Pool** button. It will scan available disks first.

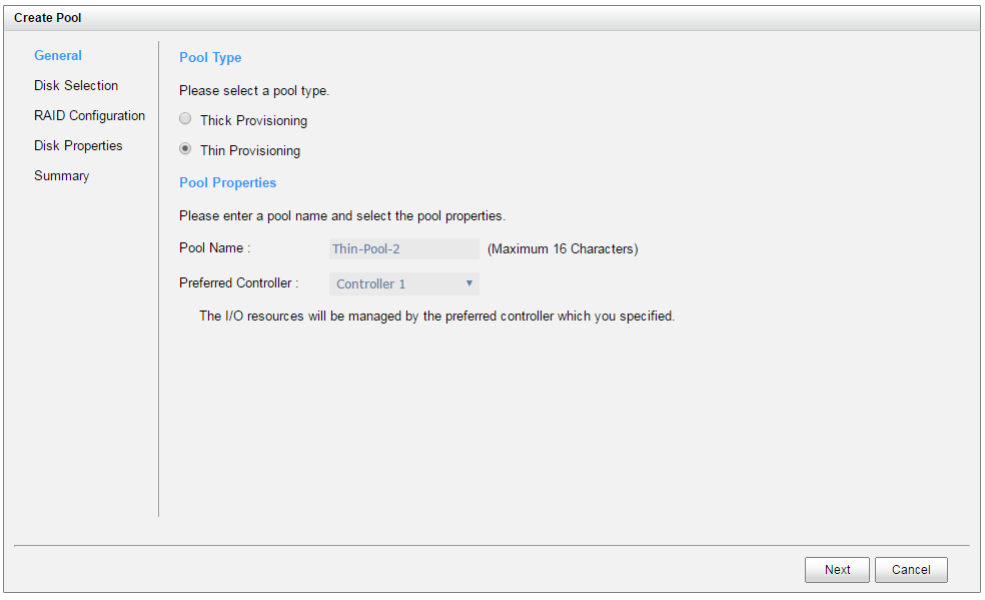

Figure 7-4 Create a Thin Provisioning Pool Step 1

2. Select **Pool Type** as **Thin Provisioning**.

3. Enter a **Pool Name** for the pool. Maximum length of name is 16 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
4. Select a **Preferred Controller** from the drop-down list. The I/O resources will be managed by the preferred controller which you specified. This option is available when dual controllers are installed.
5. Click **Next** button to continue.

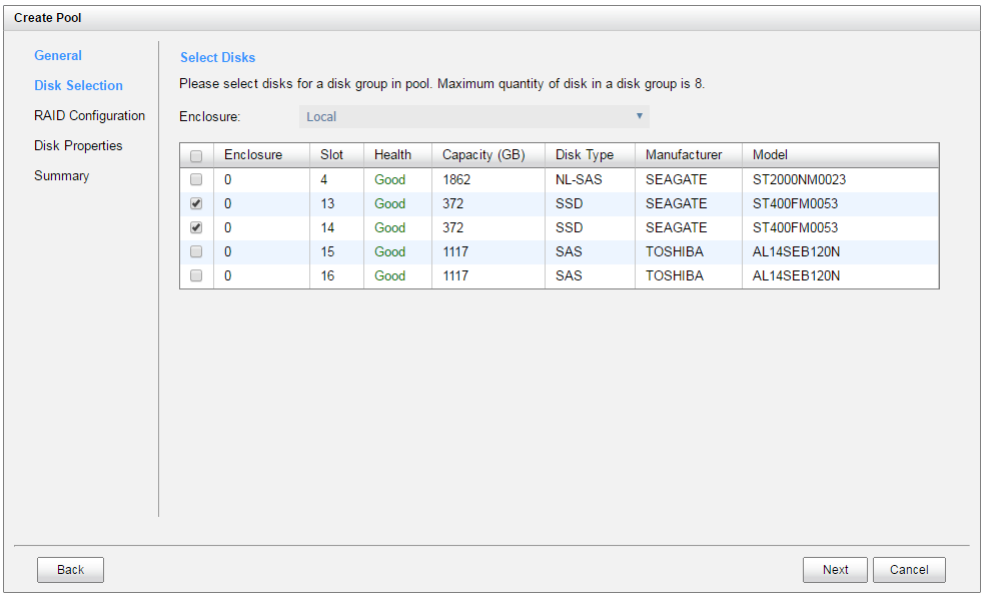

Figure 7-5 Create a Thin Provisioning Pool Step 2

6. Please select disks for a disk group in pool. Maximum quantity of disk in a disk group is 8. You can also select disks from expansion enclosures. Select an **Enclosure** from the drop-down list.
7. Click **Next** button to continue.

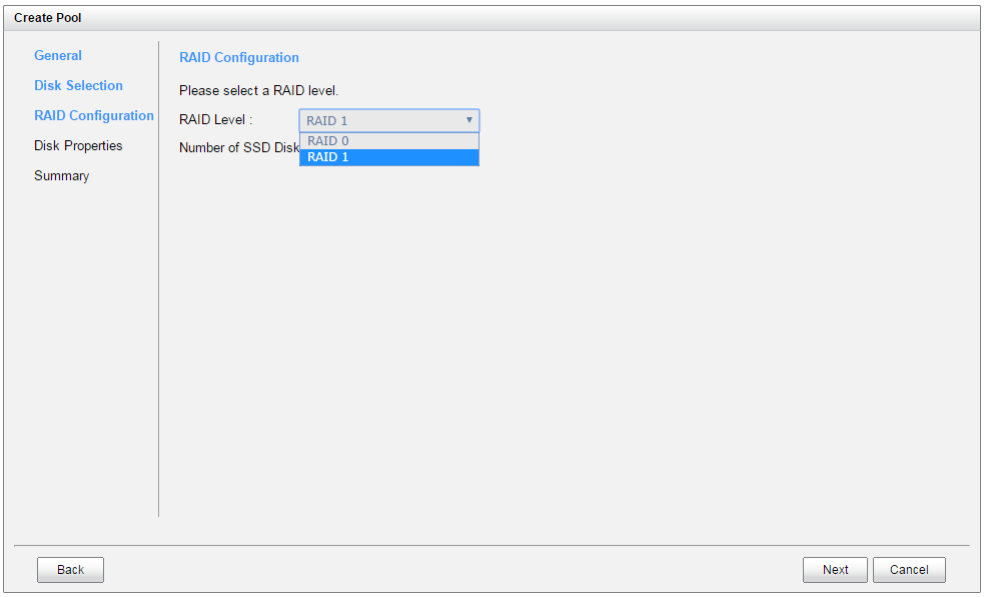

Figure 7-6 Create a Thin Provisioning Pool Step 3

8. Select a **RAID Level** from the drop-down list which lists available RAID level only according to the disk selection.
9. Click **Next** button to continue.

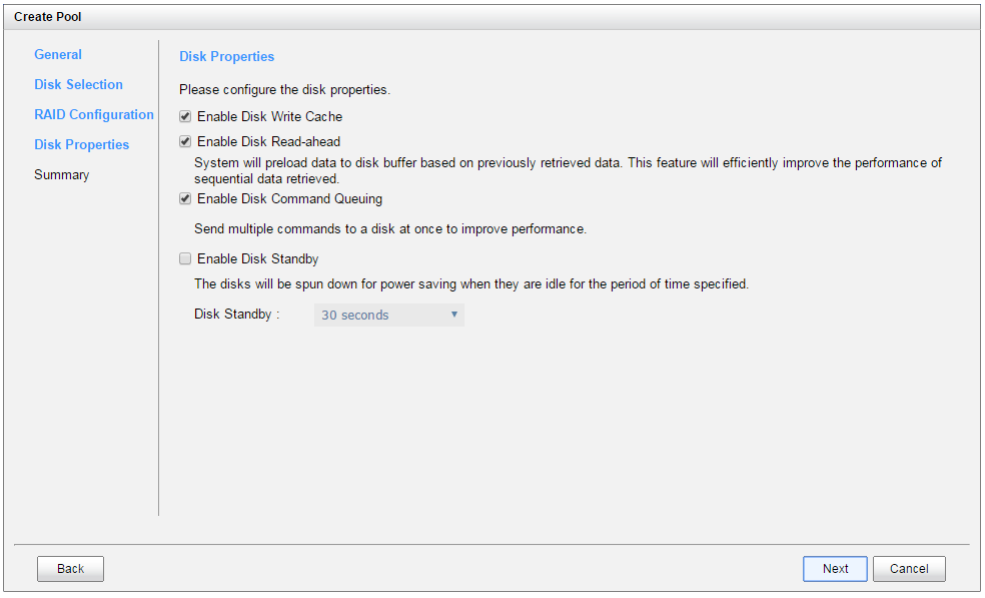

Figure 7-7 Create a Thin Provisioning Pool Step 4

10. Optionally, configure the disk properties:
    - **Enable Disk Write Cache:** Check to enable the write cache option of disks.
    - **Enable Disk Read-ahead:** Check to enable the read-ahead function of disks. System will preload data to disk buffer based on previously retrieved data. This feature will efficiently improve the performance of sequential data retrieved.

- **Enable Disk Command Queuing:** Check to enable the command queue function of disks. Send multiple commands to a disk at once to improve performance.
- **Enable Disk Standby:** Check to enable the auto spin down function of disks. The disks will be spun down for power saving when they are idle for the period of time specified.

11. Click **Next** button to continue.

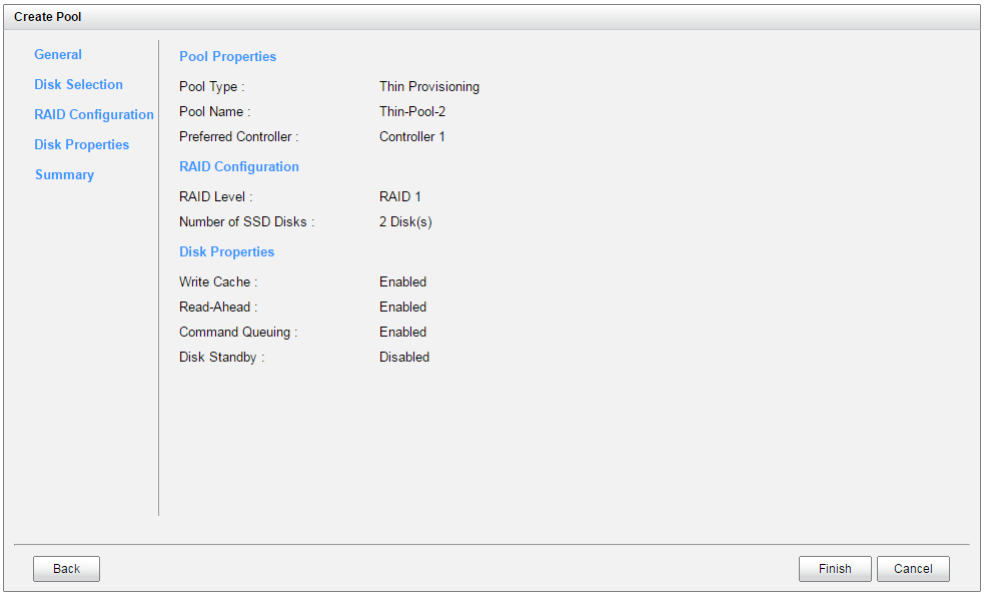

Figure 7-8 Create a Thin Provisioning Pool Step 5

12. After confirmation at summary page, click **Finish** button to create a pool.

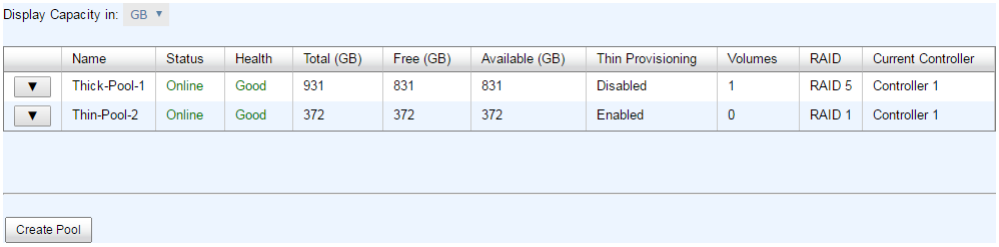

Figure 7-9 Thin Provisioning Pool is Created

13. The pool has been created. If necessary, click **Create Pool** button to create another.

7.2.2. List Thin Provisioning Pools

Use the drop-down list at the top to change the drive size units (MB or GB). According to the storage architecture, select one of the pools, it will display the related disk groups. The same, select one of the disk groups, it will display the related disk drives. Each pool can be configured by clicking ▼ of the specific one.

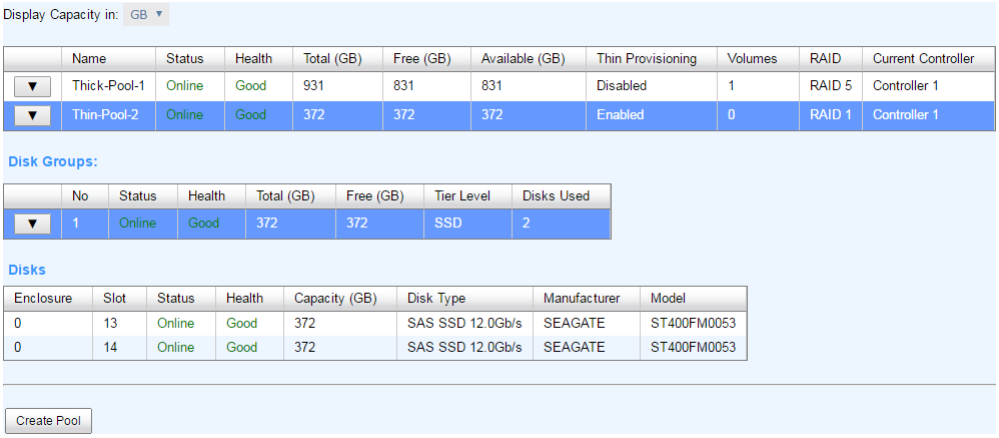

Figure 7-10 List Thin Provisioning Pools

This table shows the column descriptions.

Table 7-2 Pool Column Descriptions

Column Name	Description
Name	Pool name.
Status	The status of the pool: • Online: The pool is online. • Offline: The pool is offline. • Rebuilding: The pool is being rebuilt. • Migrating: The pool is being migrated. • Scrubbing: The pool is being scrubbed. • Relocation: The pool is being relocation.
Health	The health of the pool: • Good: The pool is good. • Failed: The pool fails. • Degraded: The pool is not healthy and not completed. The reason could be lack of disk(s) or have failed disk.
Total (GB) or (MB)	Total capacity of the pool. The unit can be displayed in GB or MB.
Free (GB) or (MB)	Free capacity of the pool. The unit can be displayed in GB or MB.
Available (GB) or (MB)	Available capacity of the pool. The unit can be displayed in GB or MB.
Thin Provisioning	The status of Thin provisioning: • Disabled. • Enabled.
Volumes	The quantity of volumes in the pool.
RAID	The RAID level of the pool.
Current Controller (*This option is only visible when dual controllers*	The current running controller of the pool.

are installed.)	

Table 7-3 Disk Group Column Descriptions

Column Name	Description
No	The number of disk group.
Status	The status of the disk group: • Online: The disk group is online. • Offline: The disk group is offline. • Rebuilding: The disk group is being rebuilt. • Migrating: The disk group is being migrated. • Scrubbing: The disk group is being scrubbed. • Relocation: The disk group is being relocation.
Health	The health of the disk group: • Good: The disk group is good. • Failed: The disk group fails. • Degraded: The disk group is not healthy and not completed. The reason could be lack of disk(s) or have failed disk.
Total (GB) or (MB)	Total capacity of the disk group. The unit can be displayed in GB or MB.
Free (GB) or (MB)	Free capacity of the disk group. The unit can be displayed in GB or MB.
Disks Used	The quantity of disk drives in the disk group.

Table 7-4 Disk Column Descriptions

Column Name	Description
Enclosure	The number of enclosure.
Slot	The position of the disk drive.
Status	The status of the disk drive: • Online: The disk drive is online. • Missing: The disk drive is missing in the pool. • Rebuilding: The disk drive is being rebuilt. • Transitioning: The disk drive is being migrated or is replaced by another disk when rebuilding occurs. • Scrubbing: The disk drive is being scrubbed.
Health	The health of the disk drive: • Good: The disk drive is good. • Failed: The disk drive is failed. • Errors Detected: The disk drive has errors detected. • Error Alert: S.M.A.R.T. error alerts. • Read Errors: The disk drive has unrecoverable read errors.
Capacity (GB) or (MB)	The capacity of disk drive. The unit can be displayed in GB or MB.
Disk Type	The type of disk drive: • SAS 12.0Gb/s • SAS 6.0Gb/s

	• SAS 3.0Gb/s • SAS SSD 12.0Gb/s • SAS SSD 6.0Gb/s • SATA 6.0Gb/s • SATA 3.0Gb/s • SATA 1.5Gb/s • SATA SSD 6.0Gb/s
Manufacturer	The manufacturer of the disk drive.
Model	The model name of disk drive.

7.2.3. Operations on Thin Provisioning Pools

Most operations are described in Configuring Storage Pools; please see Operations on Thick Provisioning Pools section. We describe the operations about thin provisioning on the following:

Change Thin Provisioning Policy of the Pool

Click ▼ **-> Change Thin Provisioning Policy** in pools to change the policies of the thin provisioning pool. It can change the action when the usage of the disk group reaches the threshold.

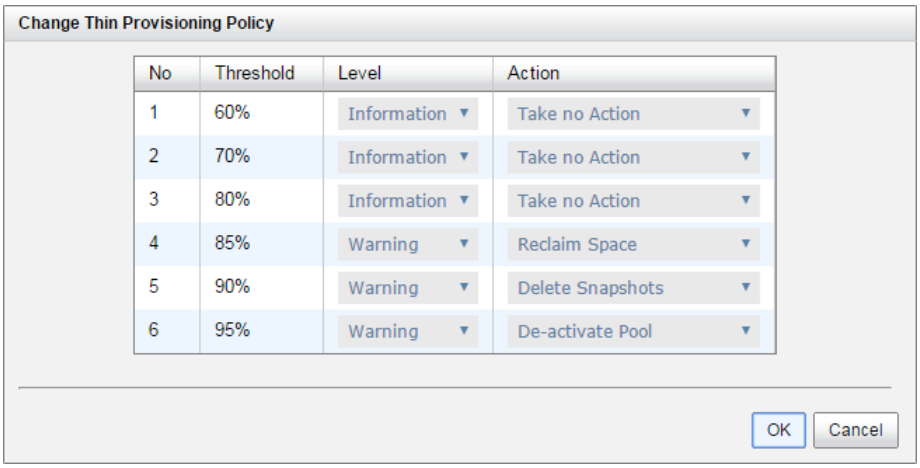

Figure 7-11 Change Thin Provisioning Policy

Table 7-5 Thin Provisioning Policy Column Descriptions

Column Name	Description
Threshold	The threshold of the disk group.
Level	Define the event log level when the usage of the disk group reaches the threshold. The options are: • Information • Warning • Error
Action	Take action of the system when the usage of the disk group reaches the threshold. The options are:

	• Take no Action • Reclaim Space • Delete Snapshots • De-activate Pool

Remove Disk Group

Click ▼ **-> Remove** in disk group to remove the selected disk group. The option is gray out if the there is only one disk group.

Move Disk Group Member Disks

Click ▼ **-> Move Disk Group** in disk group to move the member disks of disk group to other disk drives.

7.2.4. Add a Disk Group in Thin Provisioning Pool

Here is an example of adding a disk group in thin provisioning pool.

1. Click ▼ **-> Add Disk Group** in pools to add a disk group in thin provisioning pool.

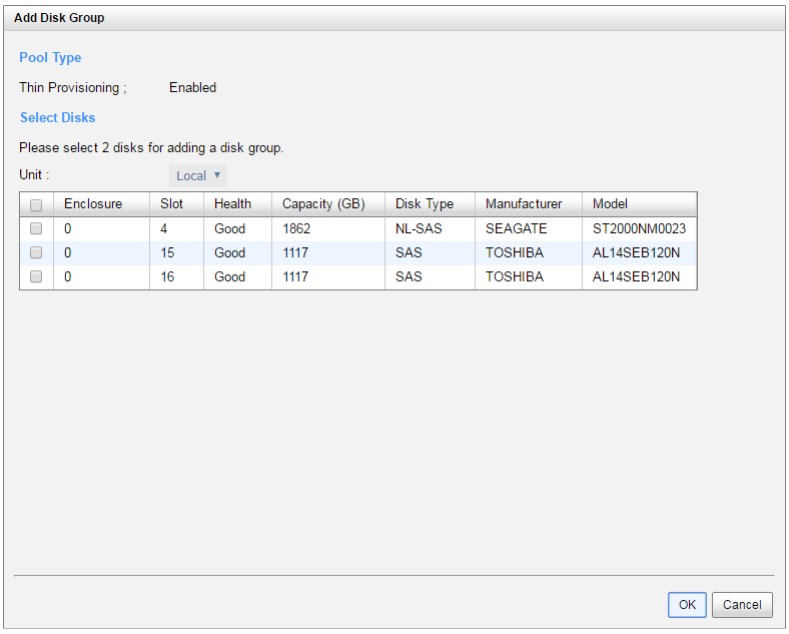

Figure 7-12 Add Disk Group

2. Please select disks for adding a disk group in pool. The quantity of selected disks must be the same as the quantity of current disk group. You can also select disks from expansion enclosures. Select an **Enclosure** from the drop-down list.
3. Click **OK** button to add a disk group.

7.2.5. Create a Volume in Thin Provisioning Pool

Here is an example of creating a volume of thin provisioning pool.

1. In **Volume** tab, click **Create Volume** button.

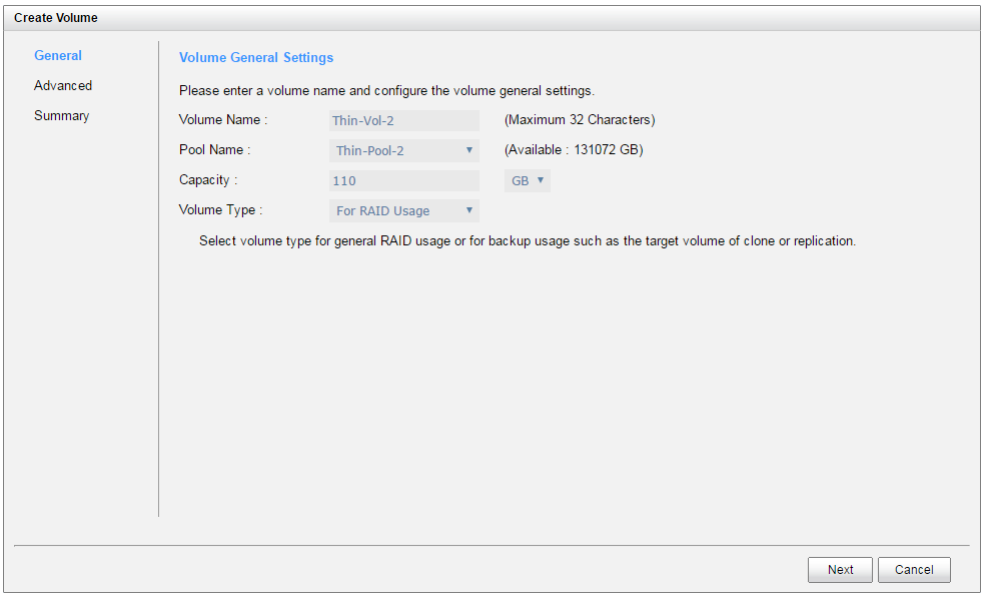

Figure 7-13 Create a Volume in Thin Provisioning Pool Step 1

2. Enter a **Volume Name** for the pool. Maximum length of name is 32 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
3. Select a **Pool Name** from the drop-down list. It will also display the available capacity of the pool.
4. Enter required **Capacity**. The unit can be selected from the drop-down list.
5. Select **Volume Type**. The options are **For RAID Usage** (for general usage) and **For Backup Usage** (for the target of local clone or remote replication).
6. Click **Next** button to continue.

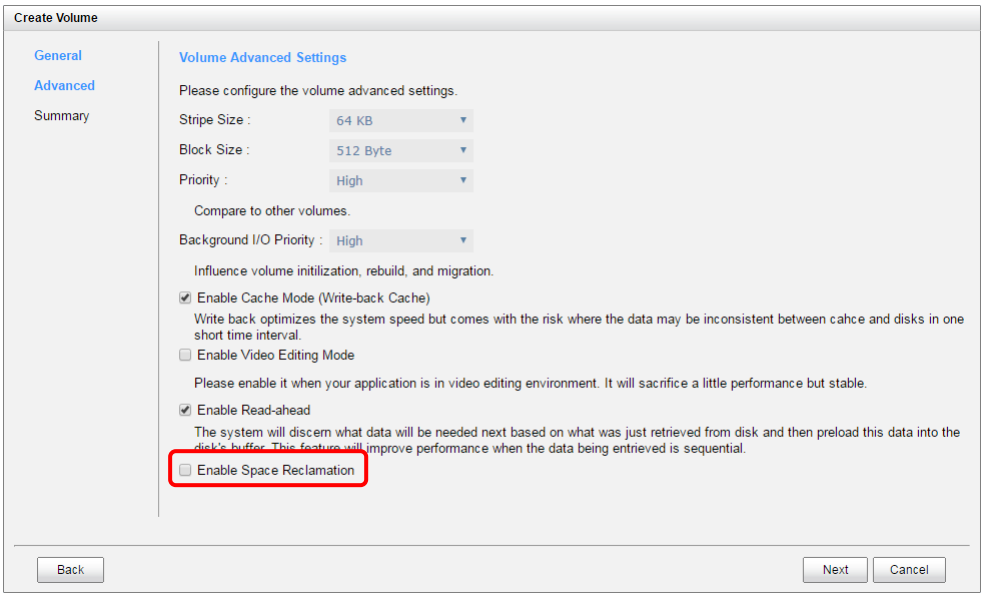

Figure 7-14 Create a Volume in Thin Provisioning Pool Step 2

7. Optionally, configure the volume advanced settings:
   - **Stripe Size:** The options are 4 KB to 64 KB.
   - **Block Size:** The options are 512 Bytes to 4,096 Bytes.
   - **Priority:** The options are High, Medium, and Low. The priority compares to other volumes. Set it as High if the volume has many I/O.
   - **Background I/O Priority:** The options are High, Medium, and Low. It will influence volume initialization, rebuild, and migration. Set it as High if the volume is important.
   - **Enable Cache Mode (Write-back Cache):** Check to enable cache mode function of volume. Write back optimizes the system speed but comes with the risk where the data may be inconsistent between cache and disks in one short time interval.
   - **Enable Video Editing Mode:** Check to enable video editing mode function. It is optimized for video editing usage. Please enable it when your application is in video editing environment. It will sacrifice a little performance but stable.
   - **Enable Read-ahead:** Check to enable the read ahead function of volume. The system will discern what data will be needed next based on what was just retrieved from disk and then preload this data into the disk's buffer. This feature will improve performance when the data being retrieved is sequential.
   - **Enable Space Reclamation:** Check to enable the space reclamation function of volume. This option is available when the pool is thin provisioning.
8. Click **Next** button to continue.

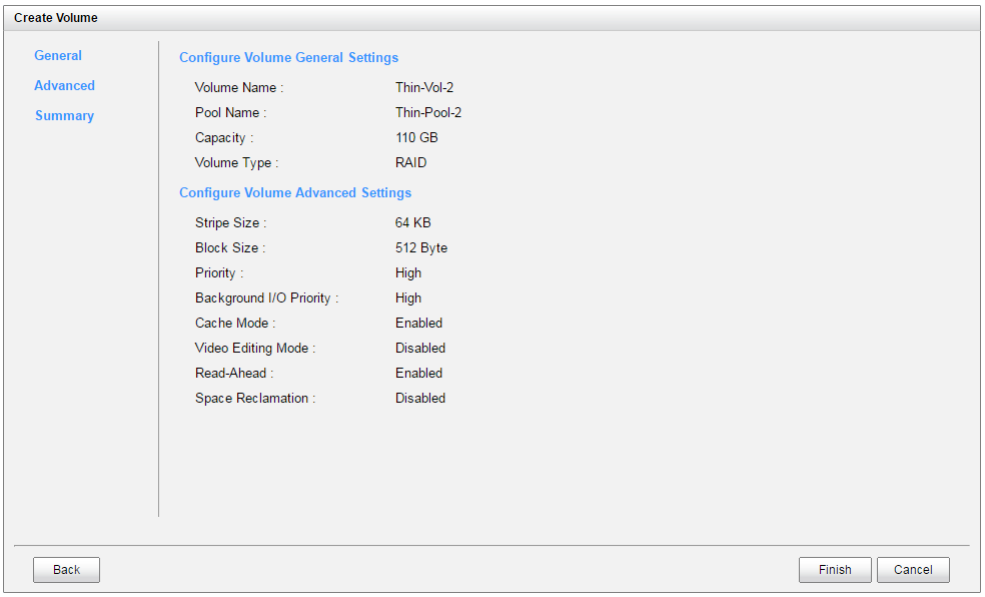

Figure 7-15 Create a Volume in Thin Provisioning Pool Step 3

9. After confirmation at summary page, click **Finish** button to create a volume.

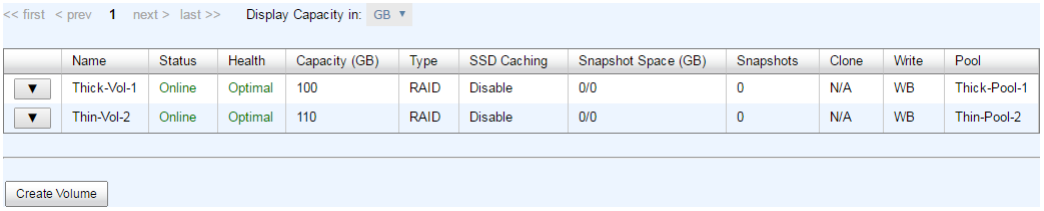

Figure 7-16 Volume in Thin Provisioning Pool is Created

10. The volume has been created. If necessary, click **Create Volume** button to create another

7.2.6. Configure LUN Mappings and Connect by Host Initiator

Next step you can configure LUN mapping and connect by host initiator. For more information about LUN mapping, please refer to Configure LUN Mappings section for detail. For more information about host initiator, please refer to Connect by Host Initiator section for detail.

8. Auto Tiering

This chapter describes an overview and operations of auto tiering.

8.1. Overview

From the perspective of storage features, the performance of SSD is high, but the cost is also high per capacity. Relatively speaking, the cost of traditional hard drive is low, so as performance is relatively poor. If we follow the 80/20 rule to configure storage systems, all-SSD configuration is unreasonable. In fact, SSD will be needed in only a small part due to the application of high-performance, regardless of whether or not a critical application, giving the SSD resources obviously substandard cost-effective. Although traditional hard disk is enough for general application which I/O requirements are not high, the traditional hard drive configuration is also gradually been inadequate.

In the other hand, the storage data has a lifecycle. Since the data in the development of its life cycle, it has experienced different levels of activity. In common usage, when creating the data, it is usually used. As the age of the data goes by, it is accessed less often.

The Solution

Therefore, in order to balance performance and cost, while using SSD and traditional hard drive hybrid storage architecture, is the most current IT environments reasonable approach. To mix in a small proportion of SSD storage devices (generally the total storage capacity of 10 to 15%) should be enough because a small part of the critical applications of high I / O requirements. The remaining applications are provided through traditional hard drive storage. Automated tiering pool is a simple solution to change data access frequency for dynamic matching storage requirements.

Tier Categories

As the name suggestion, auto tiering must have two tiers at least. Automated tiering pool segregated disk drives into three categories for dual controllers and four for single controller

- Tier 1: SSD drives for extreme performance tier
- Tier 2: SAS drives (15K or 10K SAS HDD) for performance tier
- Tier 3: Nearline SAS drives (7.2K SAS HDD) for capacity tier
- Tier 4: SATA drives for capacity tier (Only for single controller)

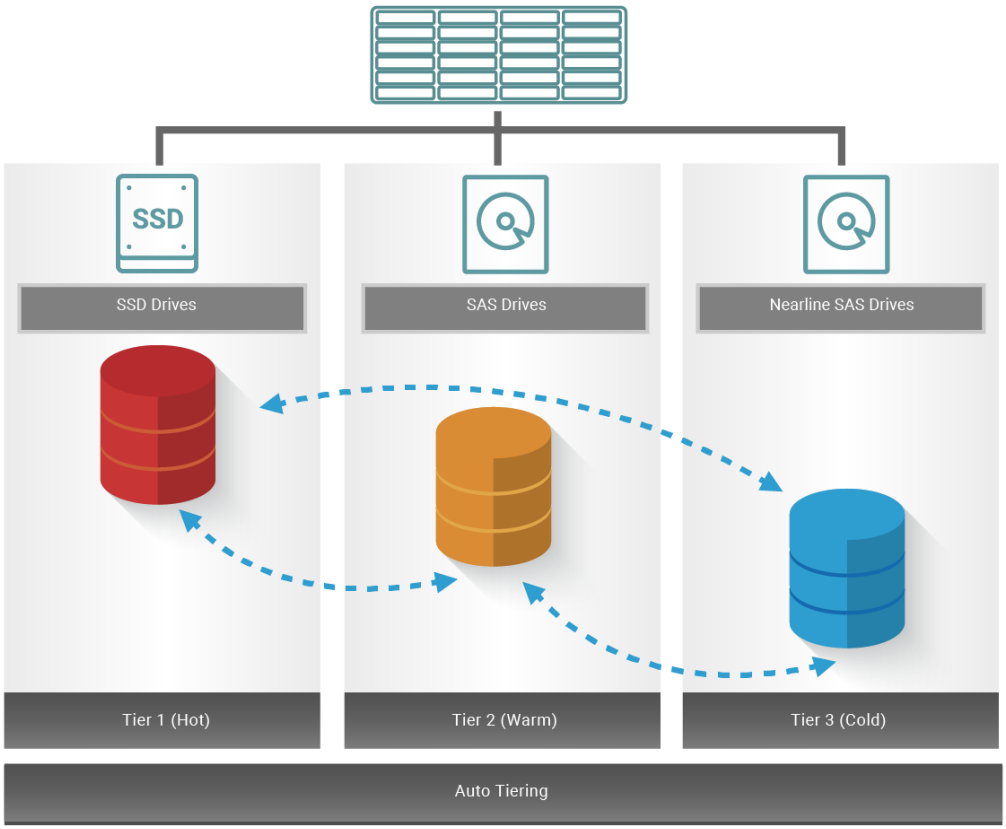

Figure 8-1 Auto Tiering Pool

8.1.1. Features Highlight

The following describes the features highlight of auto tiering.

Add tiers on demand

Base on the thin provisioning technology, each tier is a disk group. The storage pool can expand capacity on demand without downtime.

Accurate Calculation of Hotness Analysis

Data of volumes are relocated among tiers of a pool according to hotness of the data. An intelligent algorithm calculates the hot data accurately. This affects the data to move up to fast tier or down to slow tier. The more accurate data analysis, better performance it has.

Efficiency Relocation Data

Relocation is the process to move data between tiers in storage pool according to the ranking in hotness analysis process. We provide relocation data at background and also provide the function by schedule. It will not affect normal I/O by well planned.

Flexible Tiering Policy

Support up to five tiering policies for administrator to deploy the IT environment flexibly.

Convert Traditional Pool to Auto Tiering

Support both thin provisioning pool and thick provisioning pool to enable to auto tiering feature.

8.1.2. The Benefits of Auto Tiering

Automated tiering is aimed at a substantial reduction of ownership (TCO) total cost and maintaining performance features. Instead of using only one type of drive, mixing SSD, SAS and NL-SAS drives lets you keep the performance requirements while reducing the quantity of drives.

In summary, the benefits of auto tiering are listed on the following.

- Lower the Total Cost of Ownership (TCO)
- Increase performance by intelligently managing data placement
- Get the best Cost/Performance combination according to application

In some cases, the drive counts nearly two-thirds of the reduction can be achieved, while in other cases, the throughput performance by simply adding less than 10% of the pool twice a total capacity of SSD.

8.2. Design Philosophy

Auto tiering storage is the assignment of different categories of data to different disk types. It operates by relocating the most active data up to the highest available tier and the least active data down to the lowest tier. Auto tiering works at an allocation unit (granularity) of 1GB and relocates data by moving the entire unit to the appropriate tier, depending on the tiering policy selected for that particular volume.

In order to ensure sufficient space in the higher tiers, 10% of the space is maintained in each higher tier to absorb the data allocation for those tiering policies which would allocate initial space in highest available tiers. By reclaiming this 10% headroom, the least active units within each tier move to lower tiers. The whole mechanism of auto tiering contains three steps; they are statistic collection by accessed counts, ranking hotness data by the statistic collection, and then relocation data via ranking.

8.2.1. Hotness Analysis

The volume space is divided into units of equal size in which the hotness is collected and analyzed per hour. This is also called sub LUN. The analysis consists of following two phases:

Statistics Collection

Activity level of a sub LUN is determined by counting number of read and write access on the sub LUN. Logical volume manger maintains a cumulative I/O count and weights each I/O by how recently it arrived. The new coming I/O is given a full weight. After approximately 24 hours, the weight of this IO is nearly cut in half and continues to decrease. The reduction weight is processing per hour by our precision algorism. This statistics collection occurs continuously in the background for auto tiering pool.

Ranking

This analysis produces a rank ordering of each sub LUN within the pool. Note that the policies of volumes would affect how sub LUNs are ranked.

After analysis, the system would generate following information for each tier:

- The amount of data to be moved up
- The amount of data to be moved down
- The amount of data to be moved into a tier.

NOTE:

The hotness analysis process which includes statistics collection and ranking may take minutes to complete.

8.2.2. Relocation

According to the hotness analysis, relocation is processed during the user-defined relocation window; which is the number of seconds given to the relocation process. When the window closes, the relocation process would stop relocating data. The other parameter is relocation rate which controls speed of the relocation process. Valid range of relocation rate is fast, normal, and slow.

Auto tiering promotes sub LUNs according to the candidate list that it created in the analysis stage. During relocation, it prioritizes relocating sub LUNs to higher tiers. At the same time, sub LUNs are only relocated to lower tiers if the space they occupy is required for a higher priority. Using the mechanism, auto tiering makes sure that the higher performing drives are always used.

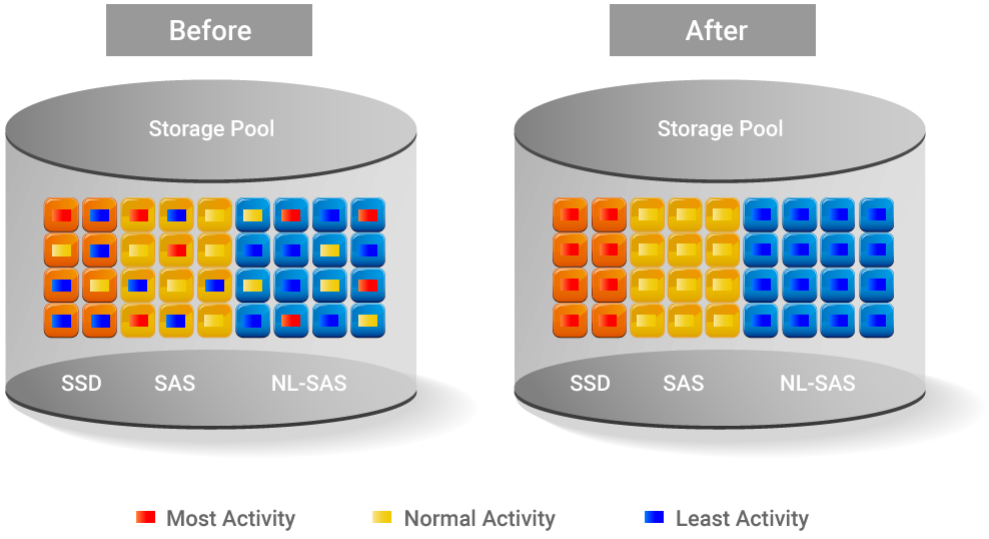

Figure 8-2 Auto Tiering Relocation

During I/O data is written to a pool, auto tiering attempts to move it to the higher tiers if space is available and the tiering policy allows for it. As we describe before, the relocation process will keep 10% of the free space in all tiers. This space is reserved for any new allocations of higher priority sub LUNs before the next relocation. Lower tiers are used for capacity when needed. The entire relocation process is complete automatically based on the user-defined relocation schedule, or manually if user triggers by himself. Figure 9-2 provides an illustration of how auto tiering can improve sub LUN placement in a pool.

8.2.3. Tiering Policies

For a best performance in various environments, auto tiering has a completely automated feature that implements a set of tiering polices. Tiering policies determine how new allocations and ongoing relocations should apply within a volume for those requirements. Auto tiering uses an algorithm to make data relocation decisions based on the activity level of each unit. It ranks the order of data relocation across all volumes within each separate pool. The system uses this information in combination with the tiering policy per volume to create a candidate list for data movement.

The following volume policies are available:

Auto Tiering

It allows for moving a small percentage of the “hot” data to higher tiers while maintaining the rest of the data in the lower tiers. This policy automatically relocates data to the most appropriate tier based on the activity level of each data. Sub LUNs are relocated based on the highest performance disk drives available and its hotness. Although this setting relocates data based on the performance statistics of the volume, the volume sets with “Highest available Tier” take precedence. Initial space is allocated in the tier which is healthier and has more free capacity than other tiers, then relocated according to hotness of the data.

Start Highest then Auto Tiering

This is the recommended policy and it is the default policy for each newly created volume, because it takes advantage of the both “Highest Available Tier” and “Auto Tiering” policies. “Start Highest then Auto Tiering” sets the preferred tier for initial data allocation to the highest performing disks with available space, and then it relocates the volume’s data based on the performance statistics and the auto-tiering algorithm. With this tiering policy, less active data is moved to lower tiers, making room for more active data in the higher tiers. Initial space is allocated in highest available tier first, then relocated according to hotness of the data.

Highest Available Tier

Use this policy when quick response times are a priority. This tier is effective for volumes which require high levels of performance whenever they are accessed. The policy starts with the “hottest” first and places them in the highest available tier until the tier’s capacity or performance capability limit is hit. Then it places the sub LUNs into the next highest tier. Initial space is allocated in highest available tier. Auto tiering would prioritize sub LUNs with highest available tier selected above all other settings.

Lowest Tier

Use this policy when cost effectiveness is the highest priority. With this policy, data is initially placed on the lowest available tier with capacity. Select this policy for volumes that are not performance sensitive or response-time sensitive. Regardless of their activity level, all sub volumes of these volumes will remain on the lowest storage tier available in their pool. Data of volumes with “Lowest tier” policy would always reside in the lowest tier. Changing policy of a volume with data in higher tiers to ‘”Lowest tier” would cause all its data in higher tier to be relocated down to the lowest tier.

No Data Movement

If a volume is configured with this policy, no sub LUN provisioned to the volumes is relocated across tiers. Data remains in its current position, but can still be relocated within the tier. The system still collects statistics on these sub LUNs after the tiering policy is changed. Initial space is allocated in the tier which is healthier and has more free capacity than other tiers. No relocation would be performed in a volume which selects “No data movement” tiering policy.

8.3. Operations

Disk groups which contain disk drives can be added to an auto tiering pool at any time to increase the capacity of the pool. For simplifying usage and better performance, every disk group must have the same quantity of disk drives. An auto tiering pool can have up to 32 disk groups with each disk group containing up to 8 disk drives. The maximum capacity of each disk group is 64TB. So the maximum capacity in a system is 256TB.

Table 8-1 Auto Tiering Pool Parameters

Item	Value
Maximum Disk Group Quantity in a Pool	32
Maximum Disk Drive Quantity in a Disk Group	8
Maximum Disk Drive Quantity in a Pool	256 (= 32 x 8)
Maximum Capacity of a Disk Group	64 TB
Maximum Thin Provisioning / Auto tiering Pool Capacity per System	256 TB

The section will describe all operations of auto tiering.

8.3.1. Enable Auto Tiering License

Download **Request License** file and send to your local sales to obtain a License Key. After getting the license key, click **Choose File** button to select it, and then click **Apply** button to enable. When the license is enabled, please reboot the system. Each license key is unique and dedicated to a specific system. If you have already enabled, this option will be invisible.

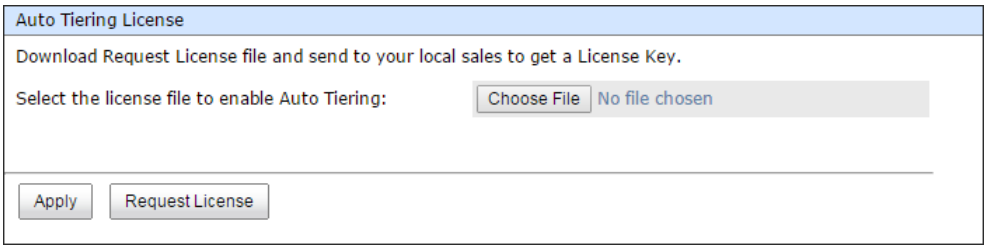

Figure 8-3 Enable Auto Tiering License

8.3.2. Create an Auto Tiering Pool

Here is an example of creating an auto tiering pool.

1. In **Pools** tab, click **Create Pool** button. It will scan available disks first.

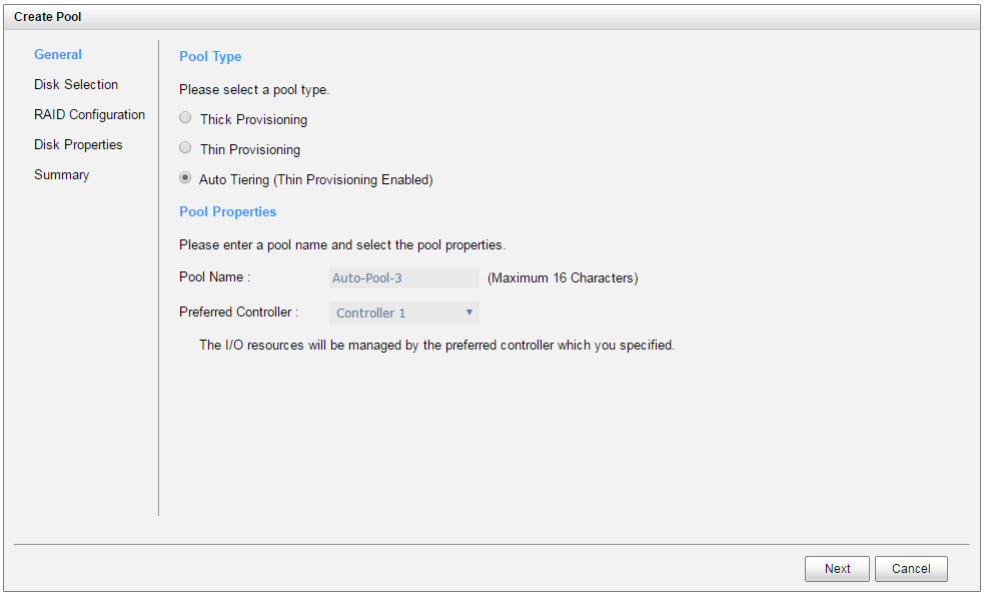

Figure 8-4 Create an Auto Tiering Pool Step 1

2. Select **Pool Type** as **Auto Tiering (Thin Provisioning enabled)**. This option is available when auto-tiering license is enabled.
3. Enter a **Pool Name** for the pool. Maximum length of name is 16 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
4. Select a **Preferred Controller** from the drop-down list. The I/O resources will be managed by the preferred controller which you specified. This option is available when dual controllers are installed.
5. Click **Next** button to continue.

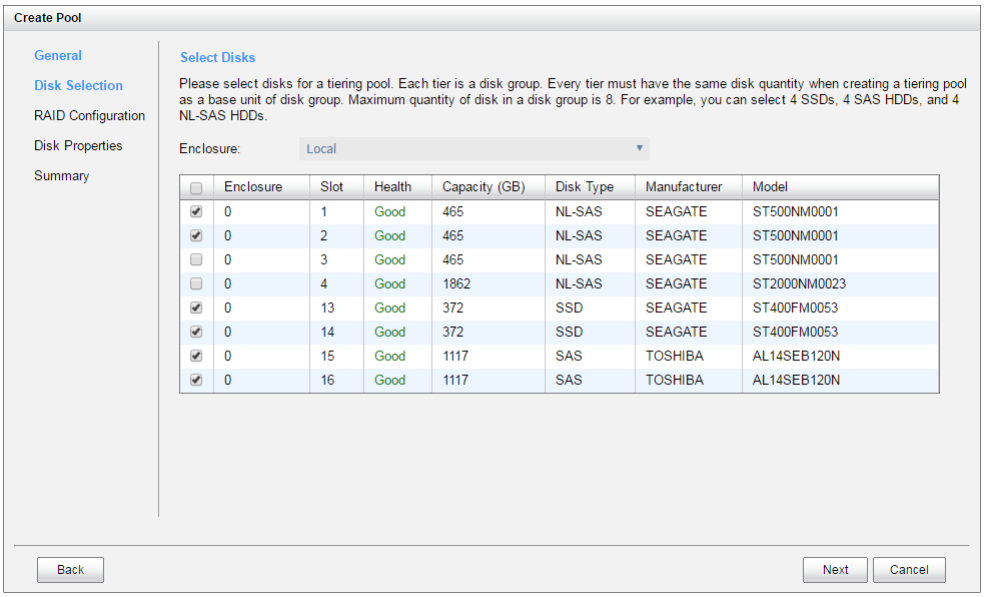

Figure 8-5 Create an Auto Tiering Pool Step 2

6. Please select disks for a tiering pool. Each tier is a disk group. Every tier must have the same disk quantity when creating a tiering pool as a base unit of disk group. Maximum quantity of disk in a disk group is 8. For example, you can select 4 SSDs, 4 SAS HDDs, and 4 NL-SAS HDDs.
7. Click **Next** button to continue.

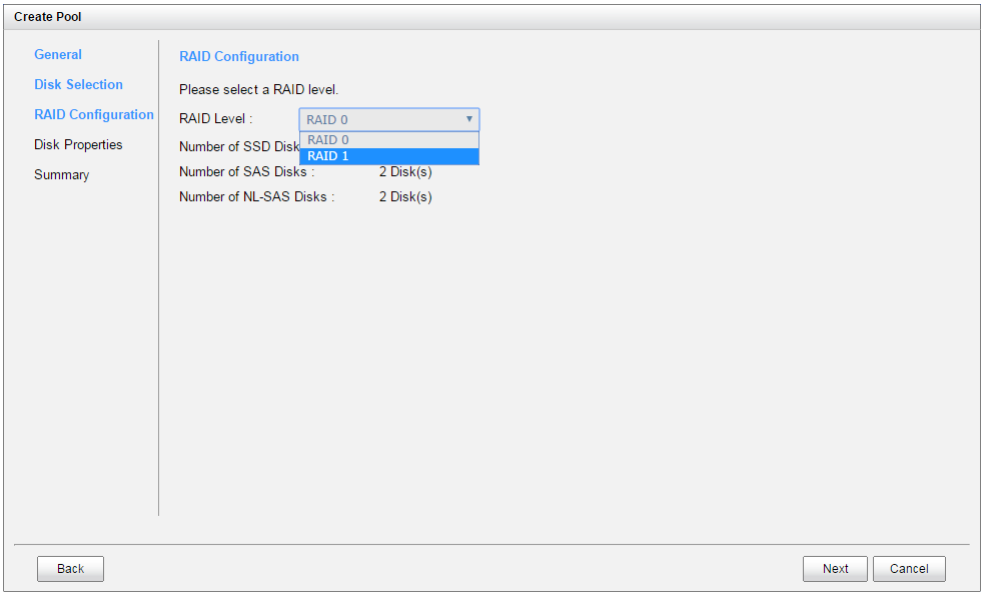

Figure 8-6 Create an Auto Tiering Pool Step 3

8. Select a **RAID Level** from the drop-down list which lists available RAID level only according to the disk selection.
9. Click **Next** button to continue.

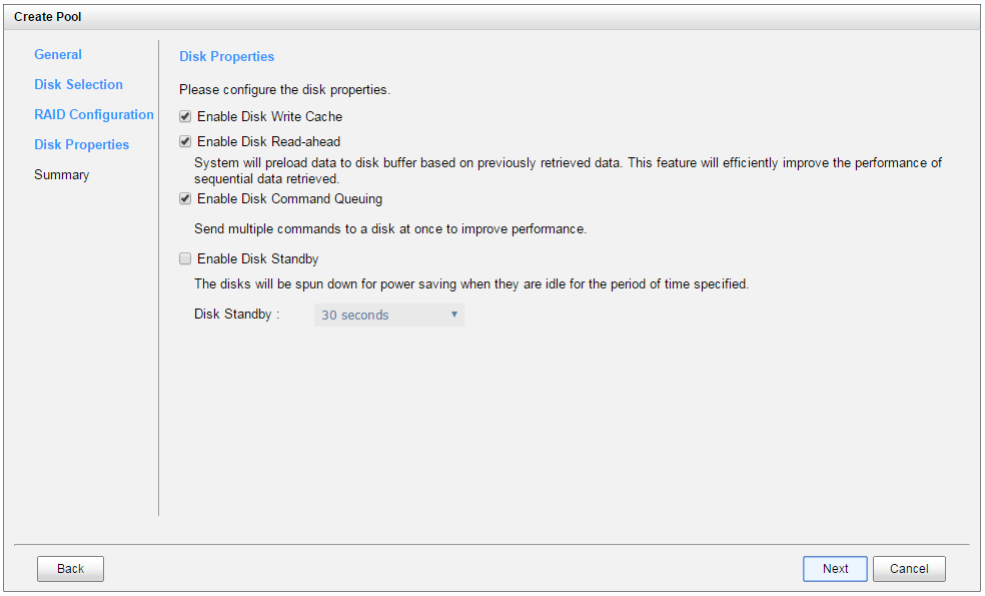

Figure 8-7 Create an Auto Tiering Pool Step 4

10. Optionally, configure the disk properties:
    - **Enable Disk Write Cache:** Check to enable the write cache option of disks.
    - **Enable Disk Read-ahead:** Check to enable the read-ahead function of disks. System will preload data to disk buffer based on previously retrieved data. This feature will efficiently improve the performance of sequential data retrieved.
    - **Enable Disk Command Queuing:** Check to enable the command queue function of disks. Send multiple commands to a disk at once to improve performance.
    - **Enable Disk Standby:** Check to enable the auto spin down function of disks. The disks will be spun down for power saving when they are idle for the period of time specified.
11. Click **Next** button to continue.

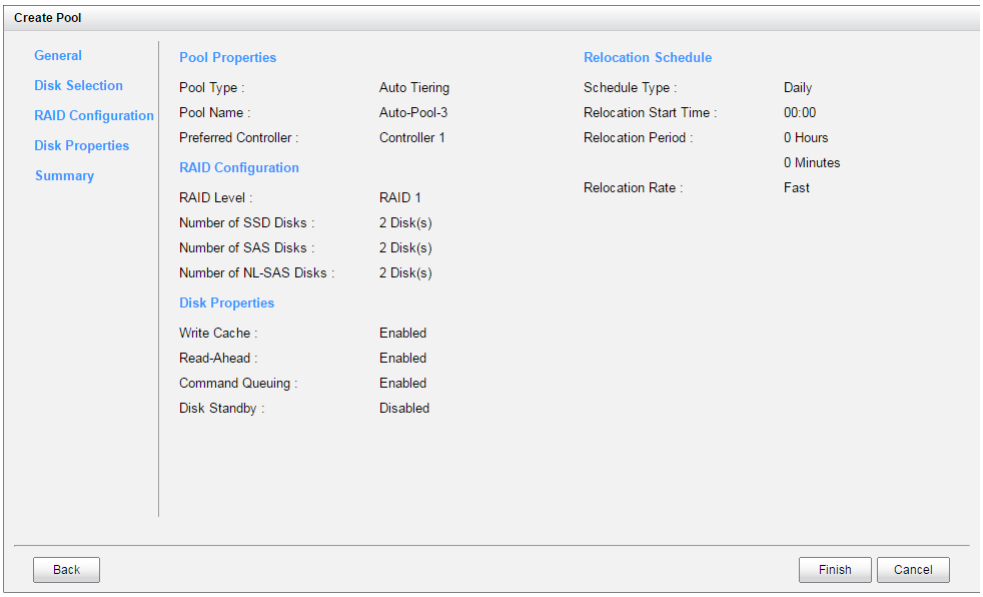

Figure 8-8 Create an Auto Tiering Pool Wizard Step 5

12. By default, we set relocation schedule at 00:00 daily, relocation period set to 00:00 which means let relocation process run until it finishes, and relocation rate to fast.
13. After confirmation at summary page, click **Finish** button to create a pool.

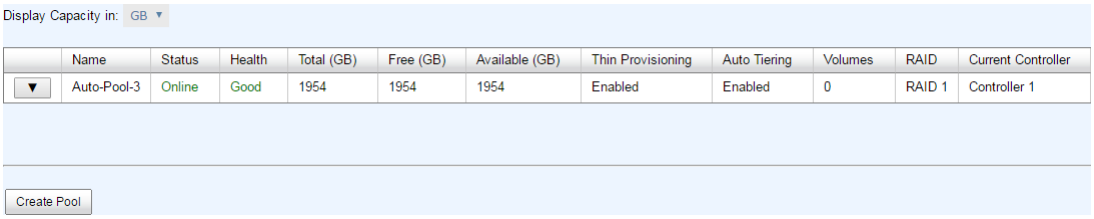

Figure 8-9 Auto Tiering Pool Has Been Created

14. The pool has been created. If necessary, click **Create Pool** button to create another.

8.3.3. List Auto Tiering Pools

Pool View

In **Pools** tab, use the drop-down list at the top to change the drive size units (MB or GB). According to the storage architecture, select one of the pools, it will display the related disk groups. The same, select one of the disk groups, it will display the related disk drives. Each pool can be configured by clicking ▼ of the specific one.

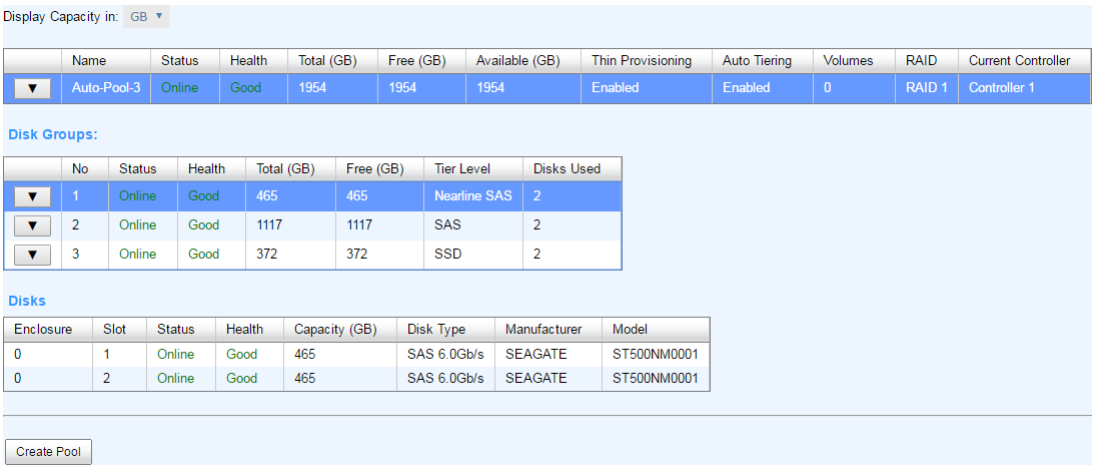

Figure 8-10 List Auto Tiering Pools

This table shows the column descriptions.

Table 8-2 Pool Column Descriptions

Column Name	Description
Name	Pool name.
Status	The status of the pool: • Online: The pool is online. • Offline: The pool is offline. • Rebuilding: The pool is being rebuilt. • Migrating: The pool is being migrated. • Scrubbing: The pool is being scrubbed. • Relocation: The pool is being relocation.
Health	The health of the pool: • Good: The pool is good. • Failed: The pool fails. • Degraded: The pool is not healthy and not completed. The reason could be lack of disk(s) or have failed disk.
Total (GB) or (MB)	Total capacity of the pool. The unit can be displayed in GB or MB.
Free (GB) or (MB)	Free capacity of the pool. The unit can be displayed in GB or MB.
Available (GB) or (MB)	Available capacity of the pool. The unit can be displayed in GB or MB.
Thin Provisioning	The status of Thin provisioning: • Disabled. • Enabled.
Auto Tiering	The status of Auto Tiering: • Disabled. • Enabled.
Volumes	The quantity of volumes in the pool.
RAID	The RAID level of the pool.

Current Controller (*This option is only visible when dual controllers are installed.*)	The current running controller of the pool.

Table 8-3 Disk Group Column Descriptions

Column Name	Description
No	The number of disk group.
Status	The status of the disk group: • Online: The disk group is online. • Offline: The disk group is offline. • Rebuilding: The disk group is being rebuilt. • Migrating: The disk group is being migrated. • Scrubbing: The disk group is being scrubbed. • Relocation: The disk group is being relocation.
Health	The health of the disk group: • Good: The disk group is good. • Failed: The disk group fails. • Degraded: The disk group is not healthy and not completed. The reason could be lack of disk(s) or have failed disk.
Total (GB) or (MB)	Total capacity of the disk group. The unit can be displayed in GB or MB.
Free (GB) or (MB)	Free capacity of the disk group. The unit can be displayed in GB or MB.
Disks Used	The quantity of disk drives in the disk group.

Table 8-4 Disk Column Descriptions

Column Name	Description
Enclosure	The number of enclosure.
Slot	The position of the disk drive.
Status	The status of the disk drive: • Online: The disk drive is online. • Missing: The disk drive is missing in the pool. • Rebuilding: The disk drive is being rebuilt. • Transitioning: The disk drive is being migrated or is replaced by another disk when rebuilding occurs. • Scrubbing: The disk drive is being scrubbed.
Health	The health of the disk drive: • Good: The disk drive is good. • Failed: The disk drive is failed. • Errors Detected: The disk drive has errors detected. • Error Alert: S.M.A.R.T. error alerts. • Read Errors: The disk drive has unrecoverable read errors.
Capacity (GB) or	The capacity of disk drive. The unit can be displayed in GB or MB.

(MB)	
Disk Type	The type of disk drive: • SAS 12.0Gb/s • SAS 6.0Gb/s • SAS 3.0Gb/s • SAS SSD 12.0Gb/s • SAS SSD 6.0Gb/s • SATA 6.0Gb/s • SATA 3.0Gb/s • SATA 1.5Gb/s • SATA SSD 6.0Gb/s
Manufacturer	The manufacturer of the disk drive.
Model	The model name of disk drive.

Auto Tiering View

Auto Tiering tab is only visible when auto tiering license is enabled. In **Auto Tiering** tab, use the drop-down list at the top to change the drive size units (MB or GB). Select one of the pools; it will display the related tiering status. The same, each pool can be configured by clicking ▼ of the specific one.

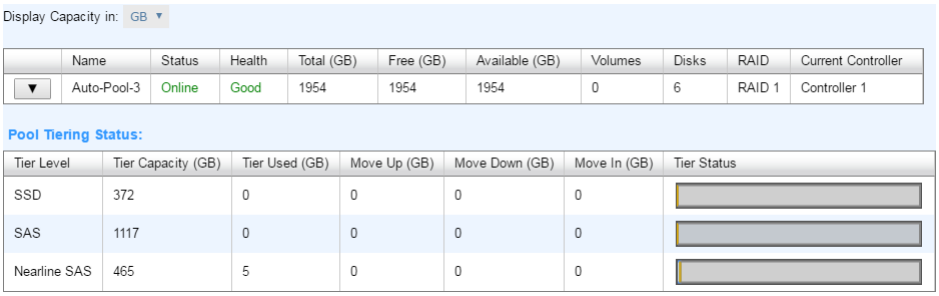

Figure 8-11 Auto Tiering Pools and Status

This table shows the column descriptions.

Table 8-5 Pool Tiering Status Column Descriptions

Column Name	Description
Tier Level	Tier categories, there are SSD, SAS, Nearline SAS, and SATA. The system will hide the tiers without any disk groups.
Tier Capacity (GB) or (MB)	Total capacity of the tier. The unit can be displayed in GB or MB.
Tier Used (GB) or (MB)	Used capacity of the tier. The unit can be displayed in GB or MB.
Move Up (GB) or (MB)	The capacity prepares to move up to higher tier. The unit can be displayed in GB or MB.

Move Down (GB) or (MB)	The capacity prepares to move down to lower tier. The unit can be displayed in GB or MB.
Move In (GB) or (MB)	The capacity prepares to move in from other tiers. The unit can be displayed in GB or MB.
Tier Status	Bar chart to show the tier status: • Light Blue: Used capacity. • Orange: The data will move in. • Gray: Unallocated.

8.3.4. Operations on Auto Tiering Pools

Most operations are described in Configuring Storage Pools; please see Operations on Thick Provisioning Pools section and Operations on Thin Provisioning Pools section. We describe the operations about auto tiering on the following:

Relocation Schedule

Click ▼ **-> Relocation Schedule** to setup relocation schedule in auto tiering pool. If **Relocation Period** sets as 00:00, it will let relocation process run until it finishes.

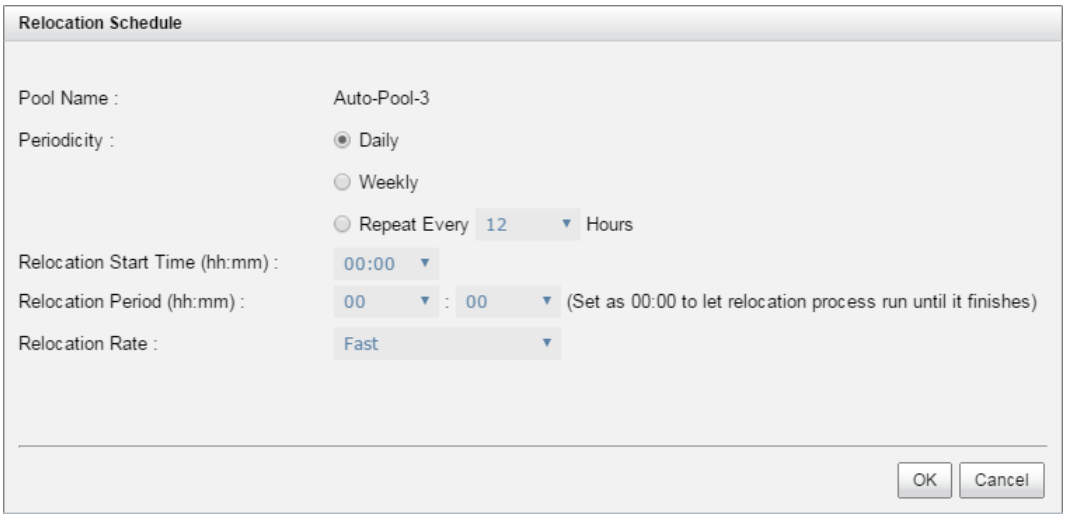

Figure 8-12 Relocation Schedule

Relocation Now

Click ▼ **-> Relocation Now** to execute relocation right now in auto tiering pool. The same as above, if **Relocation Period** sets as 00:00, it will let relocation process run until it finishes.

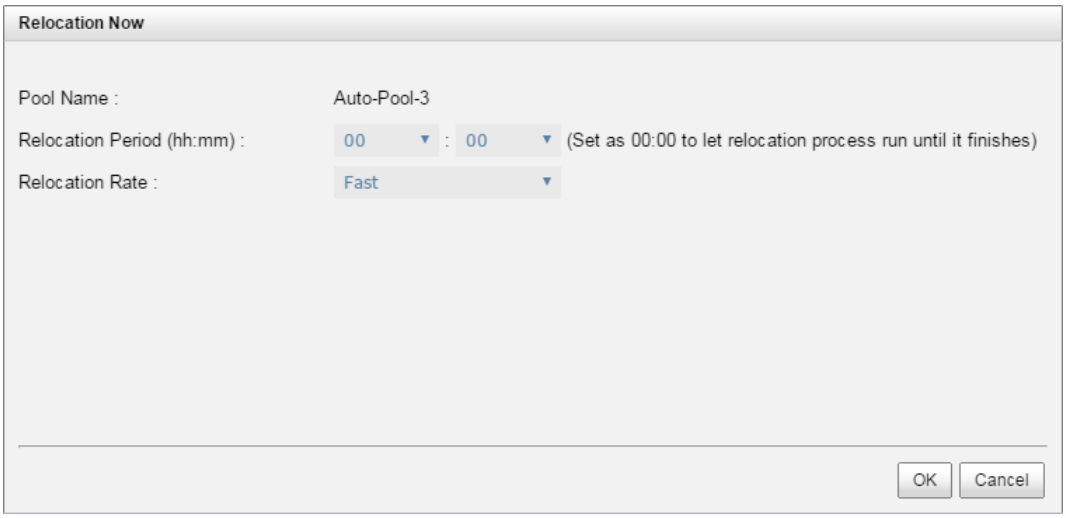

Figure 8-13 Relocation Now

8.3.5. Add a Tier (Disk Group) in Auto Tiering Pool

Here is an example of adding a disk group in thin provisioning pool.

1. In **Pools** tab, click ▼ **-> Add Disk Group** in pools to add a disk group in auto tiering pool.

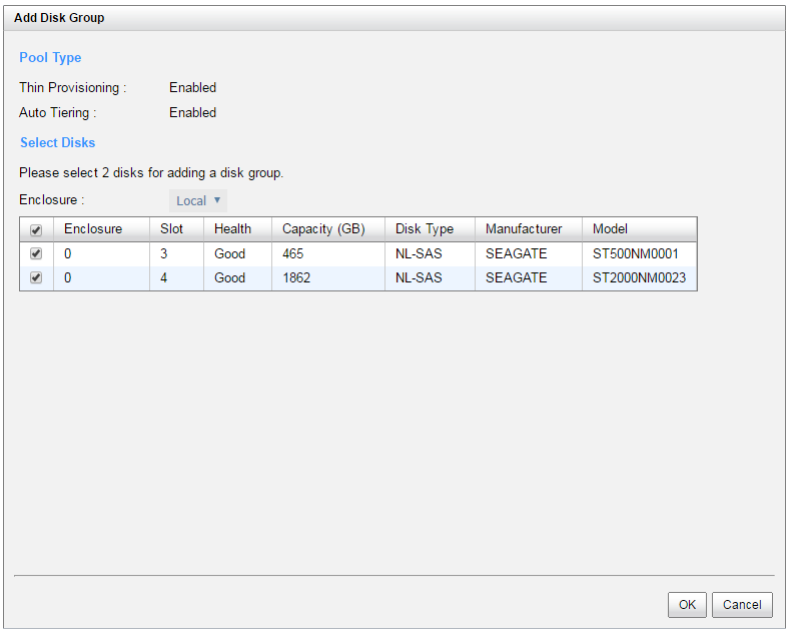

Figure 8-14 Add Disk Group

2. Please select disks for adding a disk group in pool. The quantity of selected disks must be the same as the quantity of current disk group. You can also select disks from expansion enclosures. Select an **Enclosure** from the drop-down list.
3. Click **OK** button to add a disk group.

8.3.6. Create a Volume in Auto Tiering Pool

Here is an example of creating a volume in auto tiering pool.

1. In **Volume** tab, click **Create Volume** button.

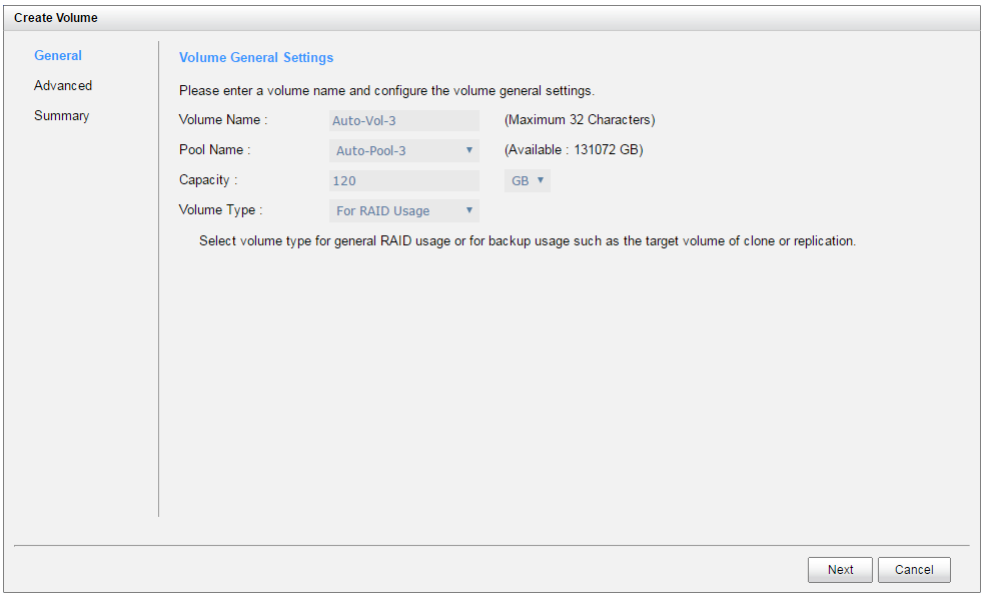

Figure 8-15 Create a Volume of Auto Tiering Pool Step 1

2. Enter a **Volume Name** for the pool. Maximum length of name is 32 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
3. Select a **Pool Name** from the drop-down list. It will also display the available capacity of the pool.
4. Enter required **Capacity**. The unit can be selected from the drop-down list.
5. Select **Volume Type**. The options are **For RAID Usage** (for general usage) and **For Backup Usage** (for the target of local clone or remote replication).
6. Click **Next** button to continue.

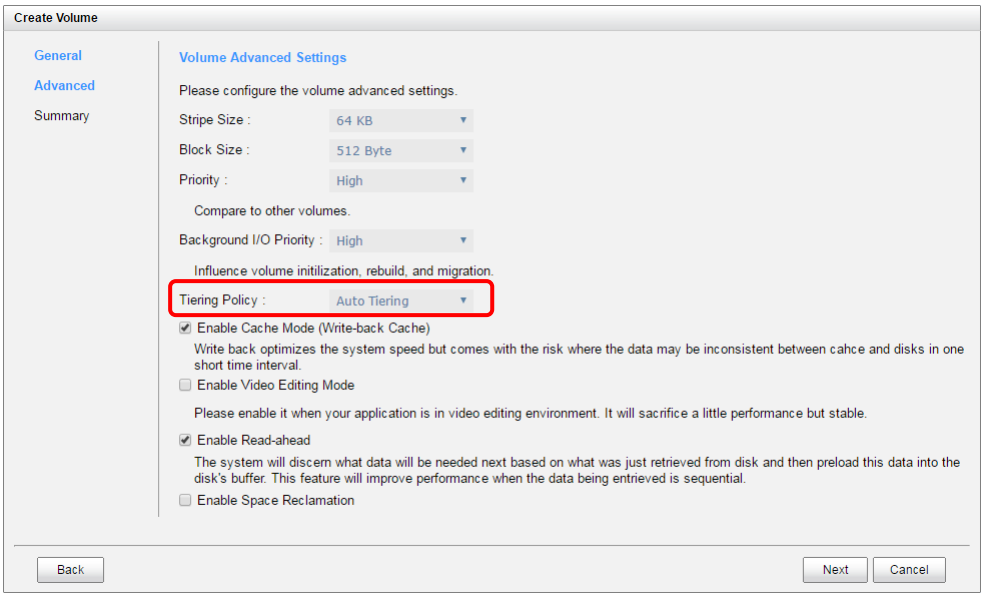

Figure 8-16 Create a Volume of Auto Tiering Pool Step 2

7. Optionally, configure the volume advanced settings:
   - **Stripe Size:** The options are 4 KB to 64 KB.
   - **Block Size:** The options are 512 Bytes to 4,096 Bytes.
   - **Priority:** The options are High, Medium, and Low. The priority compares to other volumes. Set it as High if the volume has many I/O.
   - **Background I/O Priority:** The options are High, Medium, and Low. It will influence volume initialization, rebuild, and migration. Set it as High if the volume is important.
   - **Tiering Policy:** The options are High Available Tier, Auto Tiering, Start Highest then Auto Tiering, and Lowest Tier. Please refer to Tiering Policies section for detail.
   - **Enable Cache Mode (Write-back Cache):** Check to enable cache mode function of volume. Write back optimizes the system speed but comes with the risk where the data may be inconsistent between cache and disks in one short time interval.
   - **Enable Video Editing Mode:** Check to enable video editing mode function. It is optimized for video editing usage. Please enable it when your application is in video editing environment. It will sacrifice a little performance but stable.
   - **Enable Read-ahead:** Check to enable the read ahead function of volume. The system will discern what data will be needed next based on what was just retrieved from disk and then preload this data into the disk's buffer. This feature will improve performance when the data being retrieved is sequential.
   - **Enable Space Reclamation:** Check to enable the space reclamation function of volume. This option is available when the pool is thin provisioning.
8. Click **Next** button to continue.

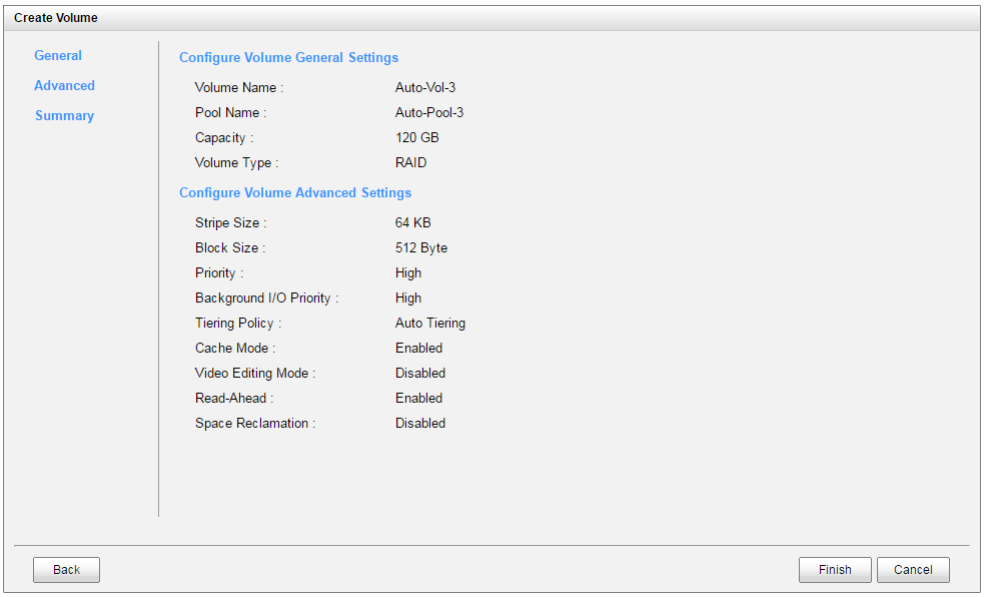

Figure 8-17 Create a Volume of Auto Tiering Pool Step 3

9. After confirmation at summary page, click **Finish** button to create a volume.

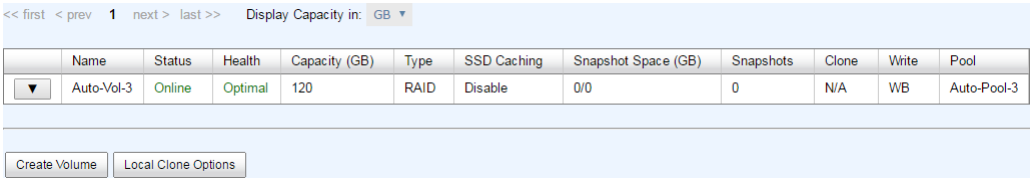

Figure 8-18 Volume in Auto Tiering Pool is Created

11. The volume has been created. If necessary, click **Create Volume** button to create another

8.3.7. Configure LUN Mappings and Connect by Host Initiator

Next step you can configure LUN mapping and connect by host initiator. For more information about LUN mapping, please refer to Configure LUN Mappings section for detail. For more information about host initiator, please refer to Connect by Host Initiator section for detail.

8.3.8. Auto Tiering Notices

There are some notices about auto tiering.

- The quantity of selected disks must be the same as the quantity of current disk group.
- The type of selected disks must be the same.
- Only disks faster than or equal to the lowest tier disk group can be added. The quantity of the disk group in the lowest tier must be greater or equal to the higher one.

9. SSD Caching

This chapter describes an overview and operations of SSD caching.

9.1. Overview

Traditionally, data are stored on the HDDs (Hard Disk Drives) and SSDs (Solid-State Drives) are mainly used for mission-critical applications that demand high-speed storage systems. In recent years, the capacity of HDDs has increased, but their random input/output (I/O) has not kept pace. For some applications such as web commerce, clouds, and virtualization that require both high capacity and performance, HDDs, though capacious, simply are not fast enough.

SSD caching technology leverages the strengths of both HDDs and SSDs, to cost-effectively meet the capacity and performance requirements of enterprise applications. Data are stored on HDDs while SSDs serve as an extended cache for many I/O operations. A single chassis, therefore, can provide both the capacity and economy of HDDs and the blistering performance of SSDs.

Generally, SSD caching is particularly effective when:

- Reads are far more common than writes in the production environment.
- The inferior speeds of HDD reads cause performance bottlenecks.
- The size of repeatedly accessed data is smaller than the capacity of the SSD cache.

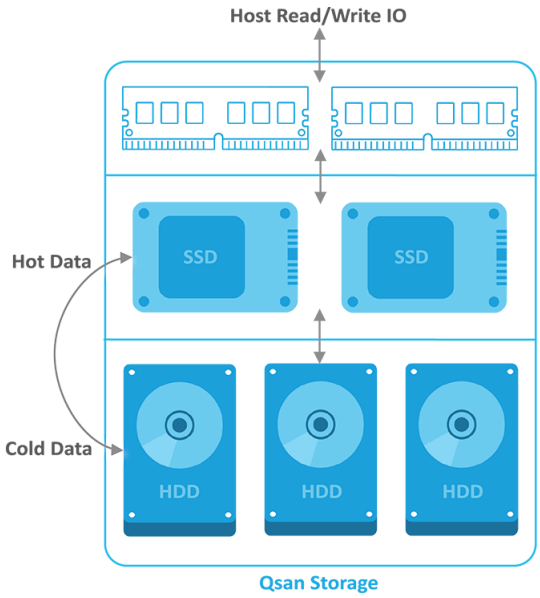

Figure 9-1 SSD Caching

The Solution

SSD caching is secondary cache that improves performance by keeping frequently accessed data on SSDs where they are read far more quickly than from HDDs. One or more SSDs can be

assigned to a single volume to provide the SSD cache. Note that the capacity allocated to the cache is not available for regular data storage. Currently, the maximum SSD cache capacity in a system is 3.6TB.

9.2. Design Philosophy

When reads or writes are performed, the data from the HDDs are copied into the SSD cache. Any subsequent reads of the same logical block addresses will be accessed directly from the SSD cache. Therefore, response times are much lower, increasing overall performance. In the very unlikely event the SSDs fail, there is no data loss because the cached data are copies of the original data residing on the HDDs.

An SSD cache is divided into a group of sectors of equal sizes. Each group is called a cache block; each block is divided into sub-blocks. The cache block size can be adjusted to suit a specific application - such as a database or web server.

9.2.1. Populating the Cache

The actions that read data from the HDD and then write to the SSD are called populating the cache. Typically, this is a background operation that immediately follows a host read or write operation. As the goal of the cache is to store frequently accessed data, not every I/O operation should trigger a cache population, but only ones that pass a certain threshold, implemented as a counter. There are both a populate-on-read threshold, and a populate-on-write threshold.

Populate-on-read Threshold

When the same data block to be read over the threshold, it is called hot data and populated to the SSD cache. The threshold must be greater than zero. If it is zero, no action is performed for a read cache.

Populate-on-write Threshold

When the same data block to be written over the threshold, it is called hot data and populated to the SSD cache. The threshold must be greater than zero. If it is zero, no action is performed for a write cache.

Each cache block on HDD volume has read and write counter associated. When a host requests to read data located in that cache block, the read count is increased. If the data is not found in the cache already and the read count is greater than or equal to the populate-on-read threshold, then a cache-populate operation is performed concurrently with the host read operation. If a cache hit occurs, the data is immediately returned from the SSD cache and a populate operation is not performed. If the read count is smaller than the threshold, a populate operation is not performed.

Write cases are the same scenario as read.

9.2.2. Read/Write Cache Cases

The following describes the read/write cache cases.

Read Data with Cache Miss

1. A host requests to read data. The system will check if the requested data is in memory cache or SSD cache. If not, it is called cache miss.
2. Data is read from the HDD volume because of cache miss.
3. The requested data is returned to the host. And the system will check whether the requested data is hot data.
4. If it is, the SSD cache is populated.

Read Data with Cache Hit

1. A host requests a read data. The system finds that the data is in SSD cache, so it is called cache hit.
2. Data is read from the SSD cache.
3. The requested data is returned to the host.
4. If there is an SSD cache error, data is read from the HDD volume.

Write Data

1. A host requests to write data.
2. Data is written to the HDD volume.
3. The status is returned to the host.
4. The SSD cache is populated if the write threshold is reached.

9.2.3. SSD Cache Tuning

The SSD cache can be tuned to maximize its efficiency base on application usage. Cache block size, populate-on-read threshold and populate-on-write-threshold are the main parameters.

Cache Block Size

A large cache block suits applications where frequently accessed data is close to each other, known as a high locality of reference. A large cache block will also fill up the SSD cache quickly - this is known as the warm-up time. After the cache is warmed up, the performance would be quite good for applications with high locality of reference. Such as the file system or web service usage, the frequently accessed data are based on some concentrated files which are usually in large block size. However large cache blocks will also generate larger I/O overhead, increasing response time, especially for cache misses.

A smaller cache block size suits applications with data that is less localized, meaning the data is accessed more randomly, such as database usage. The SSD cache will fill up slower, but with more cache blocks, there is greater chance of a cache hit, especially for data with less locality of reference. With a smaller cache block size, cache usage is usually less than with a larger cache block size, but overhead is less, so the penalty for cache misses is less severe.

Cache Sub-block Size

The cache block is divided into many sub-blocks. When any of the sub-blocks inside a cache block is accessed, the counter associated with the cache block will be increased. Once the read

or write counter reaches the population threshold, the system will copy the entire cache block into the SSD cache. The cache sub-block size is used to decide how many sub-blocks are divided from the cache block.

Population Threshold

The population threshold is the number of accesses at which point that cache block is copied to the SSD Cache. A higher number ensures that the cache only stores frequently accessed data so there will not be much cache turnover however it also means the cache will take longer to warm up and be fully effective. A lower number means the cache is warmed up quickly, but may cause excessive cache populations. A populate on read threshold of 2 is sufficient for many applications. Populate-on-write is useful when data that is written to is often read soon after. This is often the case in file systems. Other applications, such as database software, may not have this tendency so populate on write may sometimes even be disabled.

Table 9-1 I/O Type Table for SSD Caching

I/O Type	Block Size (Sectors)	Sub-block Size (Sectors)	Populate-on-Read Threshold	Populate-on-Write Threshold
Database	1MB (2,048)	8KB (16)	2	0
File System	2MB (4,096)	16KB (32)	2	2
Web Service	4MB (8,192)	64KB (128)	2	0
Customization	1MB/2MB/4MB	8KB/16KB/64KB	≥ 0	≥ 0

As you can see there are tradeoffs for increasing or decreasing each parameter. Understanding the data locality of the application is essential and it can be useful to do some field testing to see what works best.

For more information, please refer to SSD Caching White Paper. The document is available at this website:

https://www.qsan.com/en/download.php?cid=&keywords=SSD+Caching&act=query

9.3. Operations

The section will describe all operations of SSD Caching.

9.3.1. Enable SSD Caching License

Download **Request License** file and send to your local sales to obtain a License Key. After getting the license key, click **Choose File** button to select it, and then click **Apply** button to enable. When the license is enabled, please reboot the system. Each license key is unique and dedicated to a specific system. If you have already enabled, this option will be invisible.

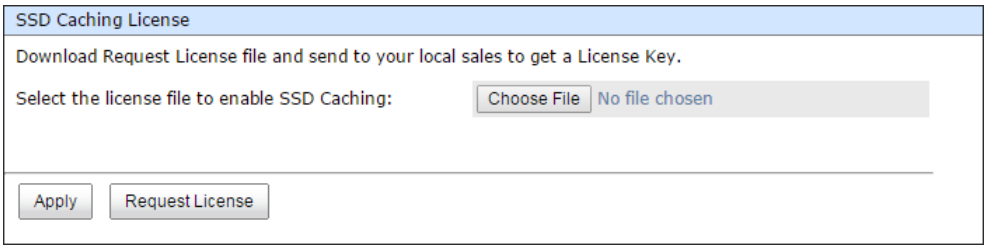

Figure 9-2 Enable SSD Caching License

9.3.2. Set SSD Caching on Volume

Here is an example of setting SSD caching on volume

1. Create a pool and then create a volume.
2. Click ▼ **-> Set SSD Caching** of the selected volume.

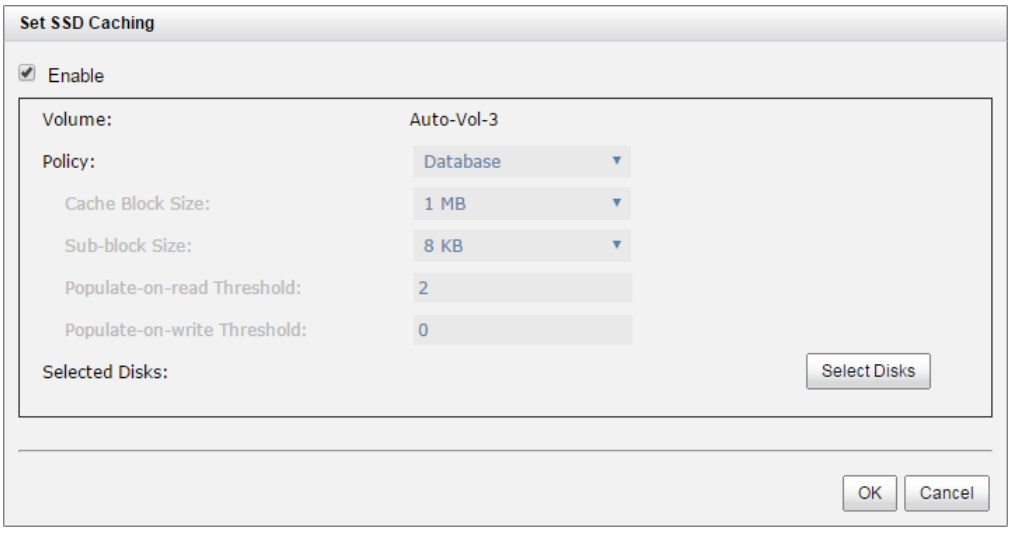

Figure 9-3 Set SSD Caching

3. Check **Enable** box.
4. Select the policy by drop down menu. Please refer to SSD Cache Tuning section for detail.
5. Click Select Disks, and then check the SSDs that are provided for SSD caching.
6. Click **OK** button to enable SSD caching.

9.3.3. SSD Caching Notices

The following are some notices about SSD caching.

- Only SSDs can be used for the SSD caching space of a volume.
- A SSD can be assigned to one and only one volume as its caching space.
- Up to 8 SSDs can be used as a SSD cache of a volume.
- There is support for up to 3.6TB of SSD caching space in one system.

10. Data Backup

The **DATA BACKUP** function menu provides submenus of **Snapshots** and **Replications**. Clone function is built in **Volumes** tab.

10.1. Overview

SAN product includes built-in data backup services for protecting the data from most unpredictable accidents. There are:

- Volume snapshot
- Local volume cloning
- Remote replication to have data duplicated to the remote sites

10.2. Managing Snapshots

Snapshot-on-the-box captures the instant state of data in the target volume in a logical sense. The underlying logic is copy-on-write, moving out the data which would be written to certain location where a write action occurs since the time of data capture. The certain location, named as “Snapshot volume”, is essentially a new volume which can be attached to a LUN provisioned to a host as a disk like other ordinary volumes in the system.

Rollback restores the data back to the state of any time which was previously captured in case for any unfortunate reason it might be (e.g. virus attack, data corruption, human errors and so on). Snapshot volume is allocated within the same pool in which the snapshot is taken; we suggest reserving 20% of the pool size or more for snapshot space.

The **Snapshots** tab provides to setup snapshot space, take snapshots, list snapshots, setup snapshot schedule, and delete snapshots. The maximum snapshot quantity per volume is 64, and the maximum volume quantity for snapshot is also 64. So a system can have 4,096 snapshots.

Table 10-1 Snapshot Parameters

Item	Value
Maximum Snapshot Quantity per Volume	64
Maximum Volume Quantity for Snapshot	64
Maximum Snapshot Quantity per System	4,096 (= 64 x 64)
Maximum Snapshot Space Capacity of a Thin Provisioning Volume	128 TB

10.2.1. Set Snapshot Space

Take an example of setting snapshot space.

1. Before taking a snapshot, it must reserve some storage space for saving variant data. There are two methods to set snapshot space. In **Snapshots** tab, click **Set Snapshot Space** button. Or in **Volumes** tab, selects a volume, then click ▼ **-> Set Snapshot Space**.

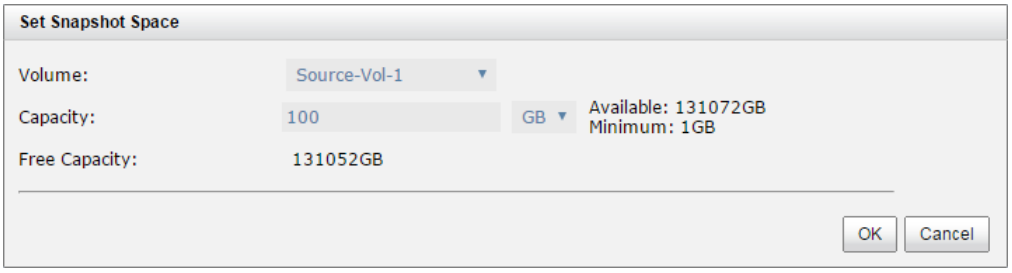

Figure 10-1 Set Snapshot Space

2. Enter a **Size** which is reserved for the snapshot space, and then click **OK** button. The minimum capacity is suggested to be 20% of the volume. Now there are two numbers in Snapshot Space (GB) column in **Volumes** tab. They mean used snapshot space and total snapshot space.

10.2.2. Take a Snapshot

Take an example of taking a snapshot.

1. There are two methods to take snapshot. In **Snapshots** tab, click **Take Snapshot** button. Or in **Volumes** tab, selects a volume, then click ▼ **-> Take Snapshot**.

Figure 10-2 Take Snapshot

2. Enter a **Snapshot Name**. Maximum length of name is 32 characters. Valid characters are [A~Z | a~z | 0~9 | ~!@#%^&*()-_+={}[]:;<>.?|/].
3. Click **OK** button. The snapshot is taken.

10.2.3. List Snapshots

In **Snapshots** tab, the two drop-down lists at the top enable you to switch the volume s. The other is to change the drive size units (MB or GB).

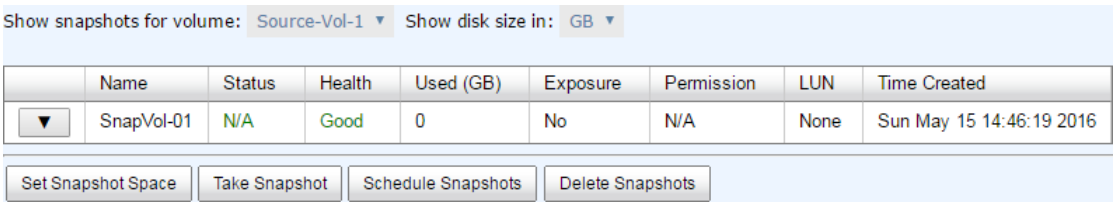

Figure 10-3 List Snapshots

This table shows the column descriptions.

Table 10-2 Snapshot Column Descriptions

Column Name	Description
Name	Snapshot name.
Status	The status of the snapshot: • N/A: The snapshot is normal. • Replicated: The snapshot is for clone or replication usage. • Abort: The snapshot is over space and abort.
Health	The health of the snapshot: • Good: The snapshot is good. • Failed: The snapshot fails.
Used (GB) or (MB)	The amount of the snapshot space that has been used. The unit can be displayed in GB or MB.
Exposure	The snapshot is exposed or not.
Permission	The permission of the snapshot: • N/A: Unknown when the snapshot is unexposed. • Read-write: The snapshot can be read / write. • Read-only: The snapshot is read only.
LUN	Number of LUN(s) that snapshot is mapped.
Time Created	The created time of the snapshot.

In **Volumes** tab, selects a volume, then click ▼ **-> List Snapshots**.

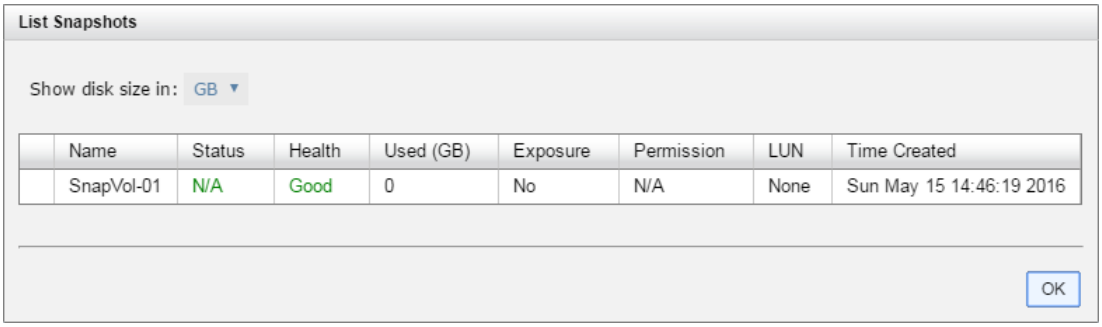

Figure 10-4 List Snapshots

10.2.4. Expose and Unexpose Snapshot

Expose Snapshot

Take an example of exposing a snapshot.

1. In **Snapshots** tab, click ▼ **-> Expose Snapshot** to set writable snapshot capacity to expose the snapshot.

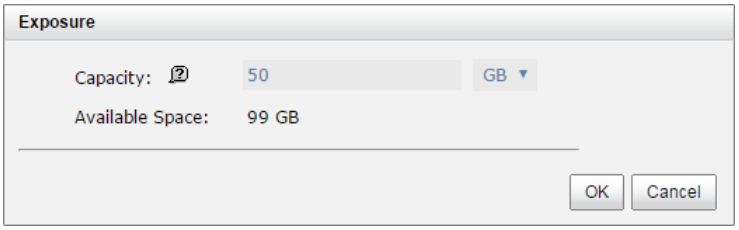

Figure 10-5 Set Writable Snapshot Capacity to Expose Snapshot

2. Enter a **Capacity** which is reserved for the snapshot. If the size is zero, the exposed snapshot will be read only. Otherwise, the exposed snapshot can be read / written, and the size will be the maximum capacity for writing.
3. Click **OK** button.
4. Click ▼ **-> Map LUN** to map a LUN to the snapshot. Please refer to Configure LUN Mappings section for detail.
5. Done. The Snapshot can be used as a volume.

Unexpose Snapshot

Take an example of unexposing a snapshot.

1. Click ▼ **-> Unexpose Snapshot** to unexposed snapshot
2. Click **OK** button.

10.2.5. Rollback Snapshot

The data in snapshot can rollback to the original volume. Please follow the procedures.

1. In **Snapshots** tab, selects a snapshot. And then click ▼ **-> Rollback Snapshot** to rollback the snapshot to the volume.
2. Click **OK** to apply.

CAUTION:

Before executing rollback, it is better that the disk is unmounted on the host computer for flushing data from cache.

When a snapshot has been rollbacked, the related snapshots which are earlier than it will also be removed. But the rest snapshots will be kept after rollback.

10.2.6. Delete Snapshots

Delete One Snapshot

Click ▼ **-> Delete** to delete the snapshot. All snapshots after deleted one will be deleted.

CAUTION:

If a snapshot has been deleted, the other snapshots which are earlier than it will also be deleted. The space occupied by these snapshots will be released after deleting.

Delete All Snapshots

To cleanup all the snapshots, please follow the procedures.

1. There are two methods to cleanup snapshots. In **Volumes** tab, selects a volume. And then click ▼ **-> Delete Snapshots** or in **Snapshots** tab, click **Delete Snapshots** button.
2. Click **OK** to apply. It will delete all snapshots of the volume and release the snapshot space.

10.2.7. Operations on Snapshots

Configure LUN Mappings

Click ▼ **-> Map LUN** to map a LUN to the snapshot. Please refer to Configure LUN Mappings section for detail.

Click ▼ **-> List LUNs** to list all LUNs of the snapshot.

Click ▼ **-> Unmap LUNs** to unmap LUNs from the snapshot.

10.2.8. Schedule Snapshots

The snapshots can be taken by schedule such as hourly or daily. Please follow the procedures.

1. There are two methods to set schedule snapshots. In **Volumes** tab, selects a volume. And then click ▼ **-> Schedule Snapshots** or in **Snapshots** tab, click **Schedule Snapshots** button.

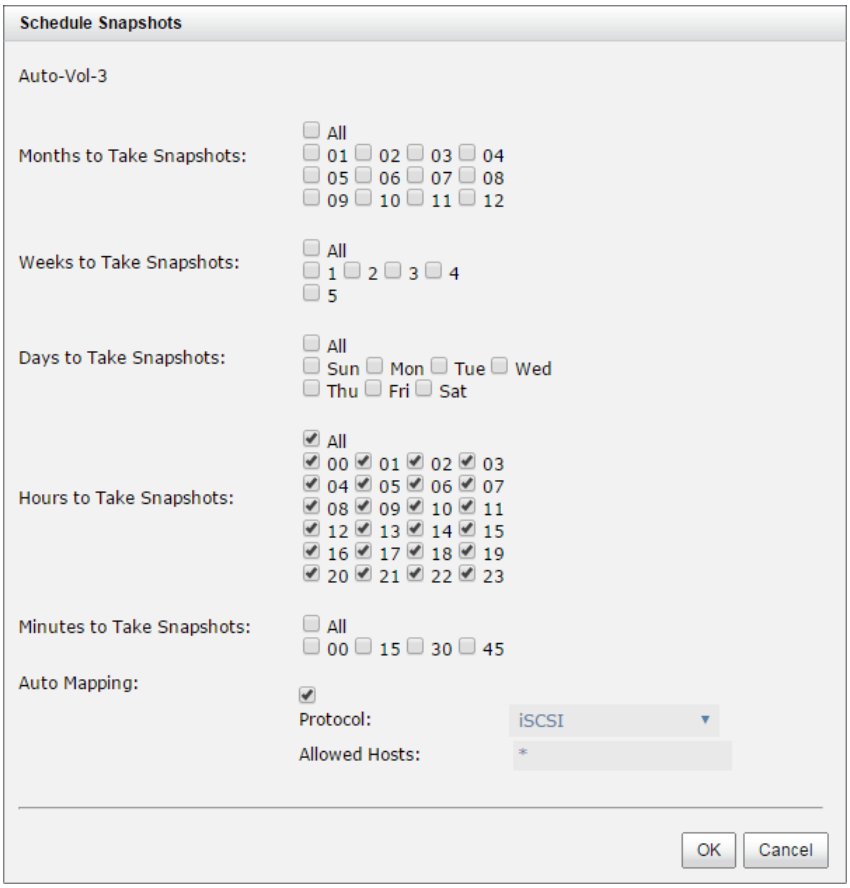

Figure 10-6 Schedule Snapshots

2. Check the schedules which you want. They can be set by monthly, weekly, daily, or hourly. Check **Auto Mapping** to map a LUN automatically when the snapshot is taken. And the LUN is allowed to access by **Allowed Hosts**.
3. Click **OK** to apply.

NOTE:

Daily snapshot will be taken at every 00:00.

Weekly snapshot will be taken every Sunday 00:00.

Monthly snapshot will be taken every first day of month 00:00.

10.2.9. Snapshot Notices

Snapshot function applies copy-on-write technique on volume and provides a quick and efficient backup methodology. When taking a snapshot, it does not copy any data at first time until a request of data modification comes in. The snapshot copies the original data to snapshot space and then overwrites the original data with new changes. With this technique, snapshot only copies the changed data instead of copying whole data. It will save a lot of disk space.

Data Consistent Snapshot

Before using snapshot, user has to know why sometimes the data corrupts after rollback of snapshot. Please refer to the following diagram.

When user modifies the data from host, the data will pass through file system and memory of the host (write caching). Then the host will flush the data from memory to physical disks, no matter the disk is local disk (IDE or SATA), DAS (SCSI or SAS), or SAN (fibre or iSCSI). From the viewpoint of storage device, it cannot control the behavior of host side. This case maybe happens. If a snapshot is taken, some data is still in memory and not flush to disk. Then the snapshot may have an incomplete image of original data. The problem does not belong to the storage device. To avoid this data inconsistent issue between snapshot and original data, user has to make the operating system flush the data from memory of host (write caching) into disk before taking snapshot.

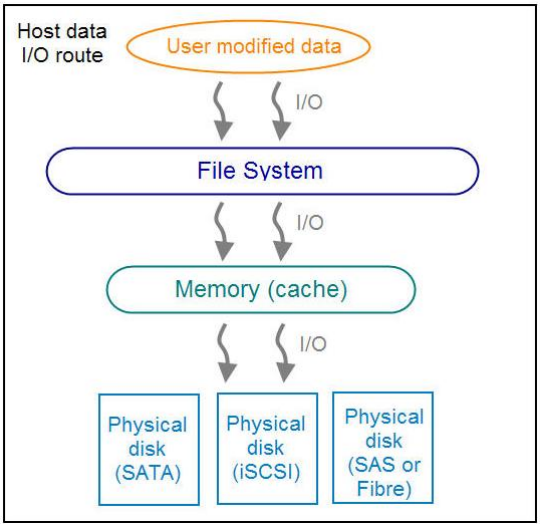

Figure 10-7 Host data I/O Route

On Linux and UNIX platform, a command named **sync** can be used to make the operating system flush data from write caching into disk. For Windows platform, Microsoft also provides a tool – **sync**, which can do exactly the same thing as the **sync** command in Linux/UNIX. It will tell the OS to flush the data on demand. For more detail about **sync** tool, please refer to http://technet.microsoft.com/en-us/sysinternals/bb897438.aspx

Besides the **sync** tool, Microsoft develops **VSS** (volume shadow copy service) to prevent this issue. VSS is a mechanism for creating consistent point-in-time copies of data known as shadow copies. It is a coordinator between backup software, application (SQL or Exchange…) and storages to make sure the snapshot without the problem of data-inconsistent. For more detail about the VSS, please refer to http://technet.microsoft.com/en-us/library/cc785914.aspx. QSAN storage system can support Microsoft VSS.

Run Out of Snapshot Space

Before using snapshot, a snapshot space is needed from pool capacity. After a period of working snapshot, what if the snapshot size over the snapshot space of user defined? There are two different situations:

1. If there are two or more snapshots existed, the system will try to remove the oldest snapshots (to release more space for the latest snapshot) until enough space is released.
2. If there is only one snapshot existed, the snapshot will fail. Because the snapshot space is run out.

For example, there are two or more snapshots existed on a volume and the latest snapshot keeps growing. When it comes to the moment that the snapshot space is run out, the system will try to remove the oldest snapshot to release more space for the latest snapshot usage. As the latest snapshot is growing, the system keeps removing the old snapshots. When it comes that the latest snapshot is the only one in system, there is no more snapshot space which can be released for incoming changes, then snapshot will fail.

Maximum Snapshot Quantity per Volume

There are up to 64 snapshots can be created per volume. What if the 65th snapshot has been taken? There are two different situations:

1. If the snapshot is configured as schedule snapshot, the latest one (the 65th snapshot) will replace the oldest one (the first snapshot) and so on.
2. If the snapshot is taken manually, when taking the 65th snapshot will fail and a warning message will be showed on web user interface.

Rollback and Delete Snapshot

When a snapshot has been rollbacked, the related snapshots which are earlier than it will also be removed. But the rest snapshots will be kept after rollback. If a snapshot has been deleted, the other snapshots which are earlier than it will also be deleted. The space occupied by these snapshots will be released after deleting.

10.3. Managing Local Clones

Clone function has another physical data copy as the original volume. At the beginning, copy all data from the source volume to target. It is also called full copy. Afterwards, use snapshot technology to perform the incremental copy. Please be fully aware that the incremental copy needs to use snapshot to compare the data difference. Therefore, the enough snapshot space for the volume is very important. Of course, clone task can also be set as schedule.

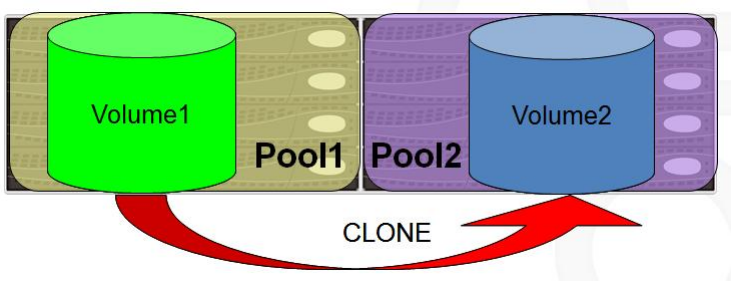

Figure 10-8 Local Clones

10.3.1. Create a Local Clone Task

Take an example of clone the volume.

1. Before cloning, it must have a backup target volume. In **Volumes** tab, click **Create Volume** button. And then select **Volume Type** to **For Backup Usage**. Please refer to Create a Volume section for detail.
2. Select the source volume, and then click ▼ **-> Create Local Clone**.
3. Select a target volume, and then click **OK** button.

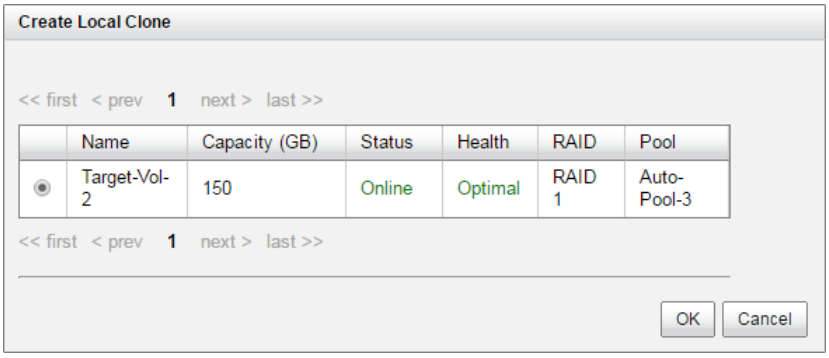

Figure 10-9 Create Local Clone

4. At this time, if the source volume has no snapshot space, it will be allocated snapshot space for clone usage automatically. The capacity will depend on the parameter of **Cloning Options**.

10.3.2. Start and Stop Local Clone Task

Start Local Clone Task

To start clone, please follow the procedures.

1. Select the source volume, and then click ▼ **-> Start Local Clone**.
2. Click **OK** button. The source volume will take a snapshot, and then start cloning.

Stop Local Clone Task

To stop clone, please follow the procedures.

1. Select the source volume, and then click ▼ **-> Stop Local Clone**.
2. Click **OK** button to stop cloning.

10.3.3. Schedule Local Clone Tasks

The clone task can be set by schedule such as hourly or daily. Please follow the procedures.

1. Select the source volume, and then click ▼ **-> Schedule Local Clone**.

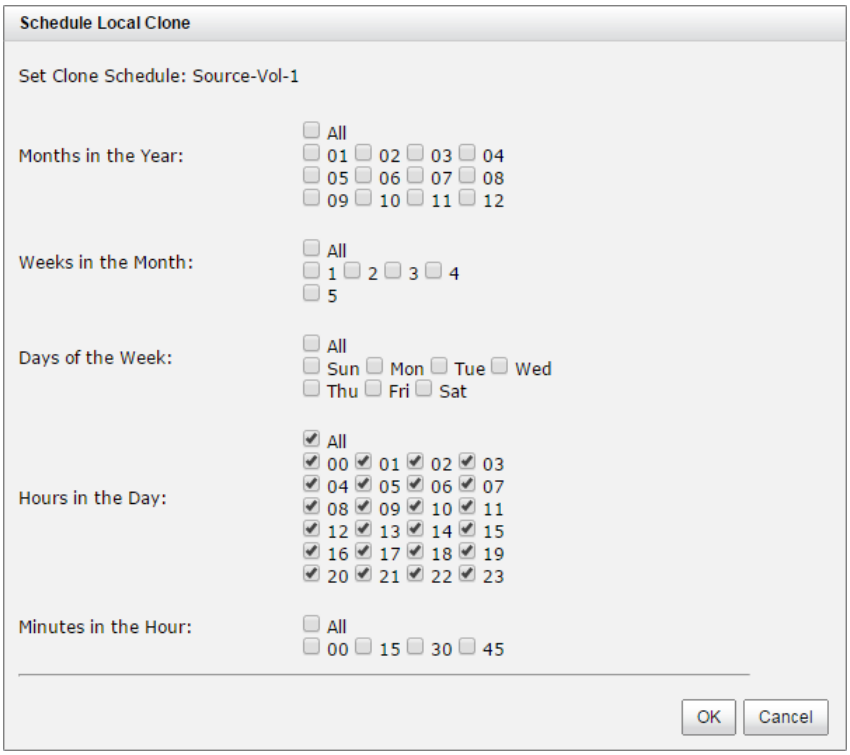

Figure 10-10 Schedule Local Clone

2. Check the schedules which you want. They can be set by monthly, weekly, daily, or hourly. Click **OK** button to apply.

10.3.4. Delete Local Clone Task

To clear the clone task, please follow the procedures.

1. Select the source volume, and then click ▼ **-> Delete Local Clone**.
2. Click **OK** button to clear clone task.

10.3.5. Cloning Options

There are three clone options, described on the following.

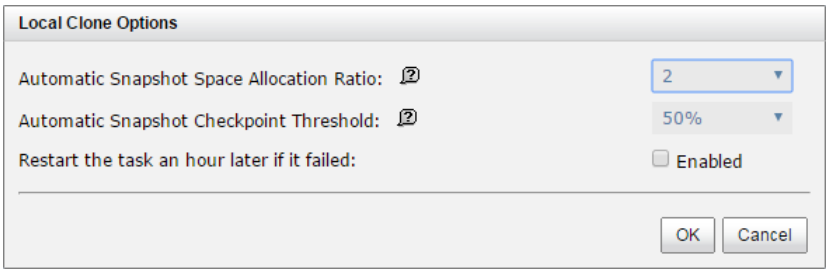

Figure 10-11 Cloning Options

- **Automatic Snapshot Space Allocation Ratio:** This setting is the ratio of the source volume and snapshot space. If the ratio sets to 2, it means when the clone process is starting, the system will book the free pool space to set as the snapshot space which capacity is double the source volume automatically. The options are 0.5 ~ 3.
- **Automatic Snapshot Checkpoint Threshold:** The setting will be effective after enabling schedule clone. The threshold will monitor the usage amount of the snapshot space. When the used snapshot space achieves the threshold, system will take a snapshot and start clone process automatically. The purpose of threshold could prevent the incremental copy failure immediately when running out of the snapshot space. For example, the default threshold is 50%. The system will check the snapshot space every hour. When the snapshot space is used over 50%, the system will start clone task automatically. And then continue monitoring the snapshot space. When the rest snapshot space has been used 50%, in other words, the total snapshot space has been used 75%, the system will start clone task again.
- **Restart the task an hour later if failed** The setting will be effective after enabling schedule clone. When running out of the snapshot space, the volume clone process will be stopped because there is no more available snapshot space. If this option is checked, the system will clear the snapshots of clone in order to release snapshot space automatically, and the clone task will be restarted after an hour. This task will start a full copy.

CAUTION:

The default snapshot space allocated by the system is two times the capacity of source volume. That is the best value of our suggestion. If user sets snapshot space by manually and lower than the default value, user should take the risk if the snapshot space is not enough and the clone task will fail.

10.3.6. Local Clone Notices

While the clone is processing manually, the increment data of the volume is over the snapshot space. The clone will complete the task, but the clone snapshot will fail. At the next time, when trying to start clone, it will get a warning message “This is not enough of snapshot space for the operation”. The user needs to clean up the snapshot space in order to operate the clone process. Each time the clone snapshot failed, it means that the system loses the reference value of incremental data. So it will start a full copy at the next clone process.

10.4. Managing Remote Replications

Remote Replication function can replicate data easily through LAN or WAN from one system to another. Here is the replication operation. At the beginning, copy all data from the source volume to target. It is also called full copy. Afterwards, use snapshot technology to perform the incremental copy. Please be fully aware that the incremental copy needs to use snapshot to compare the data difference. Therefore, the enough snapshot space for the volume is very important. Of course, replication task can also be set as schedule.

10.4.1. Create a Remote Replication Task

Take an example of creating a remote replication task.

1. Before replication, it must have a backup target volume. In **Volumes** tab of the target side, click **Create Volume** button. And then select **Volume Type** to **For Backup Usage**. Please refer to Create a Volume section for detail.
2. After creating a target volume, please also setup snapshot space at target side. So the snapshot of the source volume can replicate to the target volume. Please refer to Set Snapshot Space section for detail
3. Map a LUN of the source and target volume separately. Please refer to Configure LUN Mappings section for detail.
4. In **Replications** tab of the source side, click **Create** button.
5. Select a target volume, and then click **Next** button.

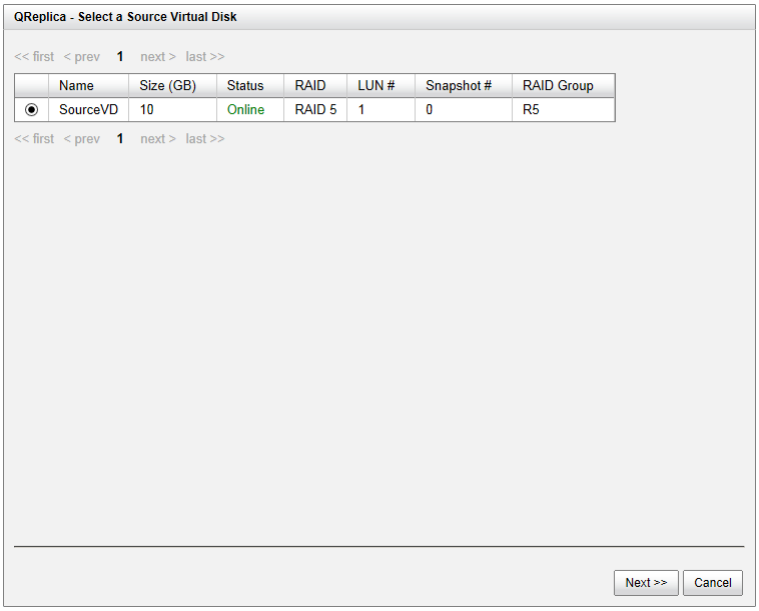

Figure 10-12 Create a Remote Replication Task Step 1

6. Select the **Source Port** and input the **Target IP**, and then click **Next** button.

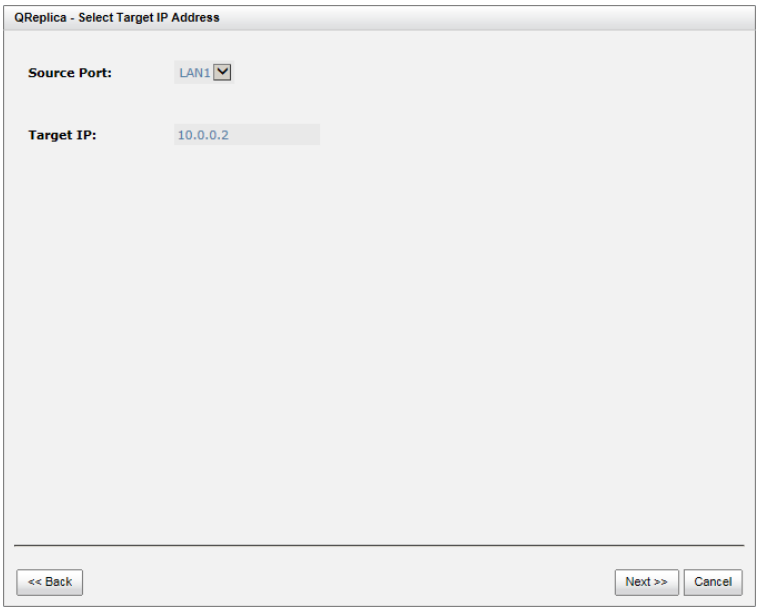

Figure 10-13 Create a Remote Replication Task Step 2

7. Select **Authentication Method** and input the CHAP user if needed. Select a **Target Node**, and then click **Next** button.

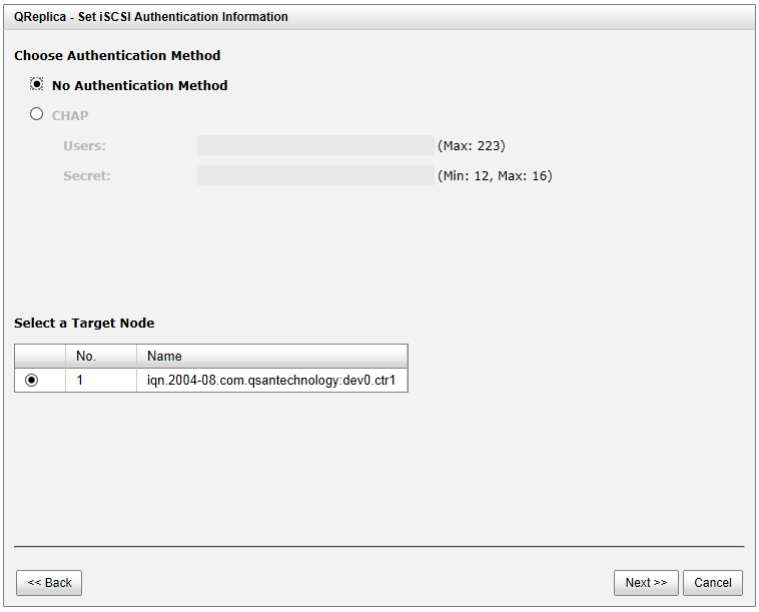

Figure 10-14 Create a Remote Replication Task Step 3

8. Select a Target LUN. When a replication task completes, it will take a snapshot on its target volume. Please make sure the snapshot space of the backup volume on the target side is properly configured. Finally, click **Finish** button.

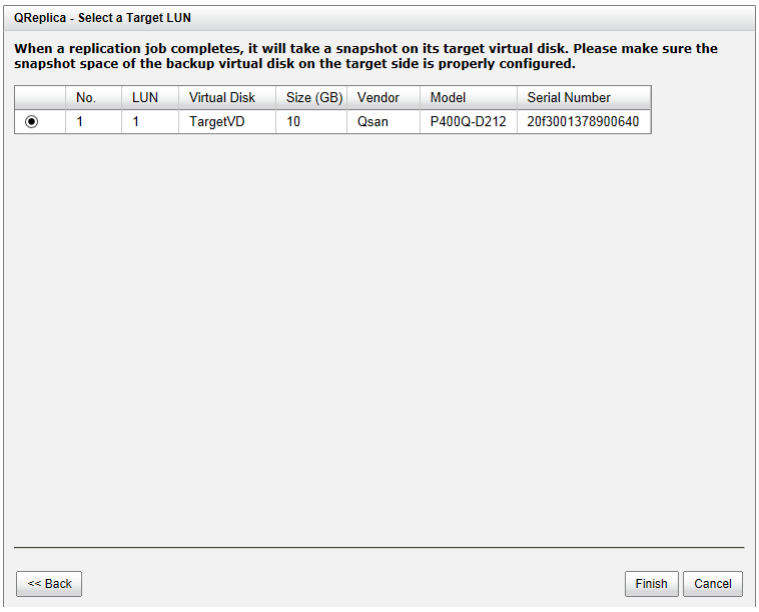

Figure 10-15 Create a Remote Replication Task Step 4

9. The replication task is created.

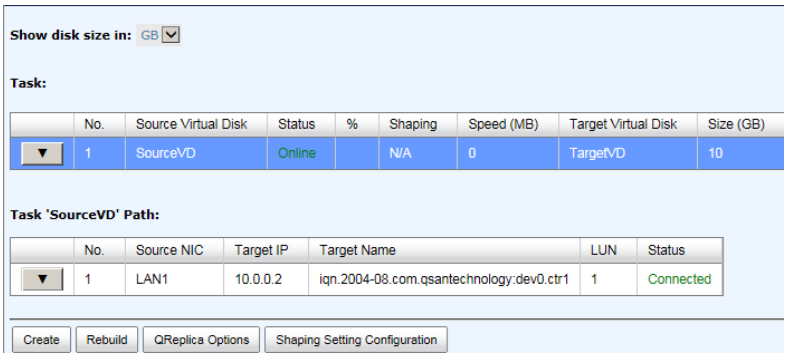

Figure 10-16 Create a Remote Replication Task Step 5

10. At this time, if the source volume has no snapshot space, it will be allocated snapshot space for replication usage automatically. The size will depend on the parameter of replication options.

10.4.2. Start and Stop Remote Replication Task

Start Remote Replication Task

To start replication task, please follow the procedures.

1. In **Replications** tab of the source side, select the task, and then click ▼ **-> Start**.
2. Click **OK** button. The source and target volumes will take snapshots, and then start replication.

Stop Remote Replication Task

To stop replication task, please follow the procedures.

1. In **Replications** tab of the source side, select the task, and then click ▼ **-> Stop**.
2. Click **OK** button to stop replication.

10.4.3. MPIO and MC/S of Replication Task

Setup MPIO of Replication Task

To setup MPIO (Multi Path Input/Ouput) of the replication task, please follow the procedures.

1. In **Replications** tab of the source side, select the task, and then click ▼ **-> Add Path**.
2. Next steps are the same as the procedure of creating a new replication task.

Delete MPIO of Replication Task

To delete multi path of the replication task, please follow the procedures.

1. In **Replications** tab of the source side, select the task, and then click ▼ **-> Delete Path**.
2. Select the path(s) which want to be deleted, and then click **OK** button.
3. The multi path(s) are deleted.

Setup MC/S of Replication Task

To setup MC/S (Multiple Connections per Session) of the replication task path, please follow the procedures.

1. In **Replications** tab of the source side, select the task path, and then click ▼ **-> Add Connection**.
2. Select the **Source Port** and input the **Target IP**, and then click **OK** button.
3. The connection is added.

Delete MC/S of Replication Task

To delete multi connections per session of the replication task path, please follow the procedures.

1. In **Replications** tab of the source side, select the task path, and then click ▼ **-> Delete Connection**.
2. Select the connection(s) which want to be deleted, and then click **OK** button.
3. The multi connection(s) are deleted.

10.4.4. Task Shaping

If the replication traffic affects the normal usage, we provide a method to limit it. There are eight shaping groups which can be set. In each shaping group, we also provide peak and off-peak time slot for different bandwidth. The following take an example of setting shaping group.

1. In **Replications** tab of the source side, click **Shaping Setting Configuration** button.

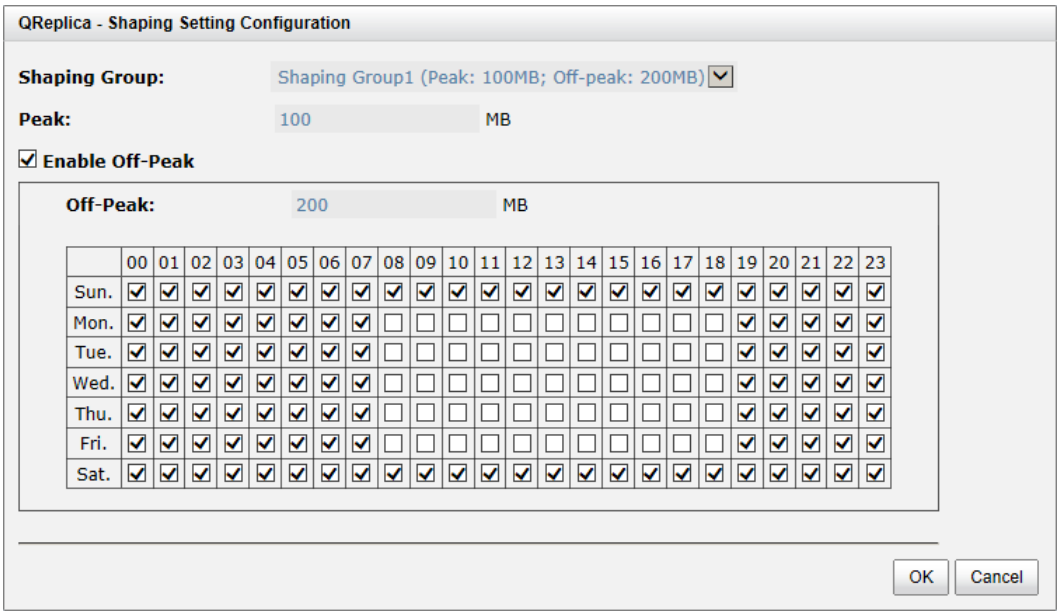

Figure 10-17 Shaping Setting Configuration

2. Select a **Shaping Group** to setup.
3. Input the bandwidth (MB) at the **Peak** time.
4. If needed, check **Enable Off-Peak**, and then input the bandwidth (MB) at **Off-Peak** time. And define the off-peak hour.
5. Click **OK** button.
6. In **Replications** tab, select the task, and then click ▼ **-> Set Task Shaping**.

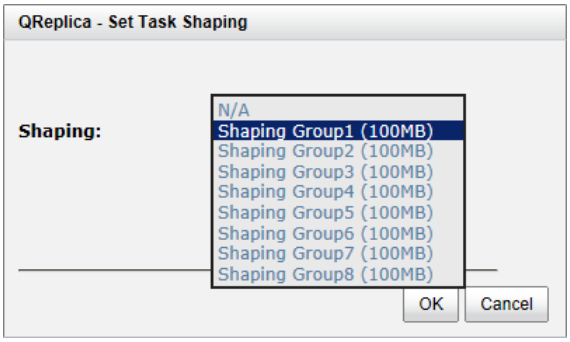

Figure 10-18 Set Task Shaping

7. Select a Shaping Group from the drop down list. And then click **OK** button.
8. The shaping group is applied to the replication task.

10.4.5. Schedule Remote Replication Tasks

The replication task can be set by schedule such as hourly or daily. Please follow the procedures.

1. In **Replications** tab, select the task, and then click ▼ **-> Schedule**.

2. Check the schedules which you want. They can be set by monthly, weekly, daily, or hourly. Click **OK** button to apply. Please refer to Schedule Snapshots section for detail.

10.4.6. Delete Replication Task

To delele the replication task, please follow the procedures.

1. Select the task in **Replications** tab, and then click ▼ **-> Delete**.
2. Click **OK** button to delete the replication task.

10.4.7. Replication Options

There are three **Replica Options**, described on the following.

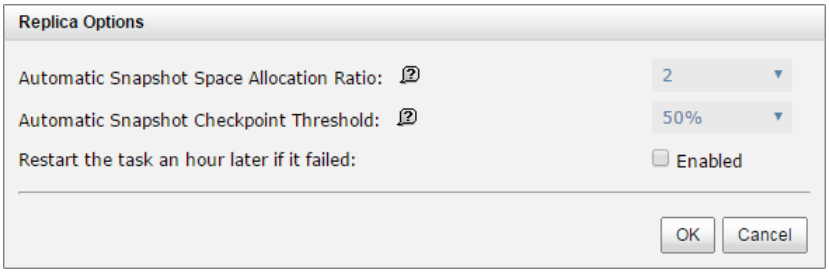

Figure 10-19 Replication Options

- **Automatic Snapshot Space Allocation Ratio:** This setting is the ratio of the source volume and snapshot space. If the ratio sets to 2, it means when the replication process is starting, the system will book the free pool space to set as the snapshot space which capacity is double the source volume automatically. The options are 0.5 ~ 3.
- **Automatic Snapshot Checkpoint Threshold:** The setting will be effective after enabling schedule replication. The threshold will monitor the usage amount of the snapshot space. When the used snapshot space achieves the threshold, system will take a snapshot and start replication process automatically. The purpose of threshold could prevent the incremental copy failure immediately when running out of the snapshot space. For example, the default threshold is 50%. The system will check the snapshot space every hour. When the snapshot space is used over 50%, the system will start replication task automatically. And then continue monitoring the snapshot space. When the rest snapshot space has been used 50%, in other words, the total snapshot space has been used 75%, the system will start replication task again.
- **Restart the task an hour later if failed:** The setting will be effective after enabling schedule replication. When running out of the snapshot space, the volume replication process will be stopped because there is no more available snapshot space. If this option is checked, the system will clear the snapshots of replication in order to release snapshot space automatically, and the replication task will be restarted after an hour. This task will start a full copy.

CAUTION:

The default snapshot space allocated by the system is two times the capacity of source volume. That is the best value of our suggestion. If user sets snapshot space by manually and lower than the default value, user should take the risk if the snapshot space is not enough and the replication task will fail.

10.4.8. Local Clone Transfers to Remote Replication

It is always being a problem that to do full copy over LAN or WAN when the replication task is executed at the first time. It may take days or weeks to replicate data from source to target within limited network bandwidth. We provide two methods to help user shorten the time of executing full copy.

1. One is to skip full copy on a new, clean volume. The term “clean” means that the volume has never been written data since created. For a new created volume which has not been accessed, the system will recognize it and skip full copy automatically when the replication task is created on this volume at the first time.

NOTE:

Any I/O access to the new created volume will make it as “not clean”, even though executing “Erase” function when a volume is created. The full copy will take place in such a case.

2. The other way is to use volume clone function, which is a local data copy function between volumes to execute full copy at the first time. Then move all the physical drives of the target volume to the target system and then turn the cloning task into replication task with differential copy afterward.

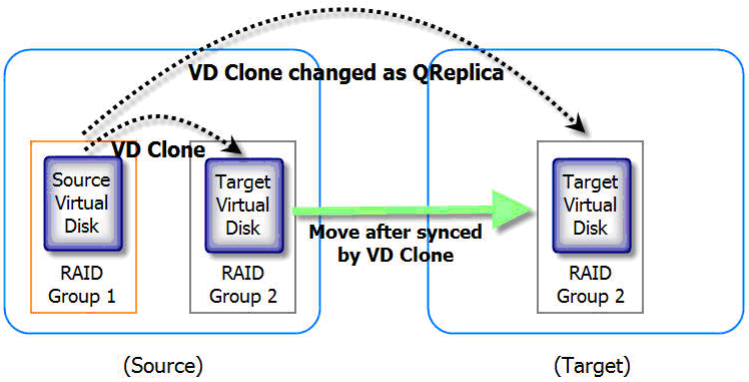

Figure 10-20 Local Clone Transfer to Remote Replication

To do that volume disk local clone transfers to remote replication, please follow the procedures.

1. Create a clone task on an existing volume with data stored already.
2. It is better that there is no host connected to the source volume. Then run **Set Clone**, **Start Clone** to synchronize the data between source and target volumes.
3. After the data is synchronized, change the cloning task to a Replication task. Select the source volume, and then click ▼ **-> Covert to Remote Replication**.
4. The **Clone** column of the source volume will be changed from the name of the target volume into **Replication**.

CAUTION:

Changing a cloning task to a replication task is only available when the cloning task has been finished. This change is irreversible.

5. Deactivate the pool which the target volume resides in and move all physical disks of the pool to the target system. Then activate the pool in the target system. Remember to set snapshot space for the target volume. And then attach the target volume to a LUN ID.
6. In **Replications** tab of the source side, click Rebuild button to rebuild the replication task which is changed from a cloning task formerly.
7. Rebuild the clone relationship, select a source volume.

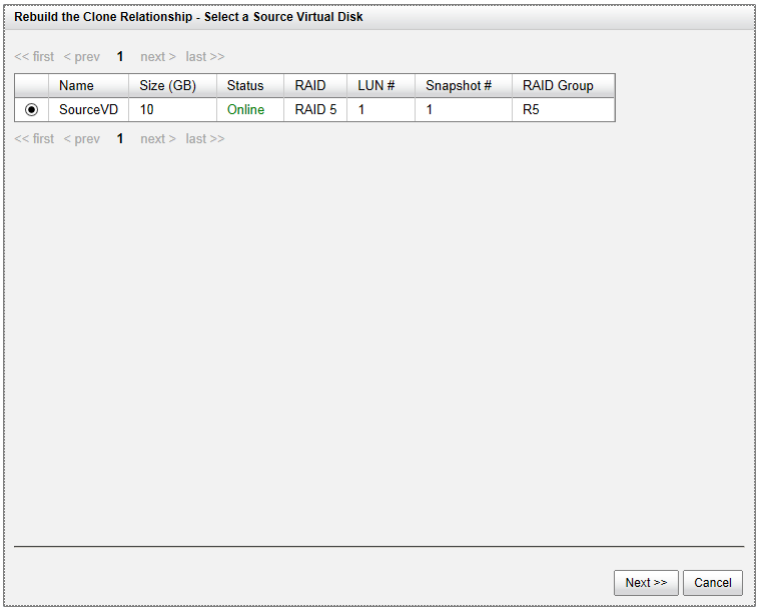

8. Next steps are the same as the procedure of creating a new replication task.
9. If a wrong target volume is selected when rebuilding the replication task, there will be an alert and the system stops the creation.

11. Monitoring

The **MONITORING** function menu provides submenus of **Log Center**, **Enclosure**, and **Performance**.

11.1. Log Center

The **Event Logs** tab provides event messages. Select or unselect the buttons of **Info**, **Warning**, or **Error** levels to show or hide those particular events.

Figure 11-1 Event Logs

The event logs are displayed in reverse order which means the latest event log is on the first / top page. They are actually saved in the first four hard drives; each hard drive has one copy of event log. For one system, there are four copies of event logs to make sure users can check event log any time when there are failed disks.

The event logs record all system events. Each event has time frame that identifies the type of event that occurred, and has one of the following severities:

- **Error:** A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
- **Warning:** A problem occurred that may affect system stability, but not data integrity. Evaluate the problem and correct it if necessary.
- **Information:** An operation recorded that may help to debug.

11.1.1. Operations on Event Logs

The options are available on this tab:

Mute Buzzer

Click **Mute Buzzer** button to stop alarm if the system alerts.

Download Event Logs

Click **Download** button to save the event log as a file. It will pop up a filter dialog as the following. The default it “Download all event logs”.

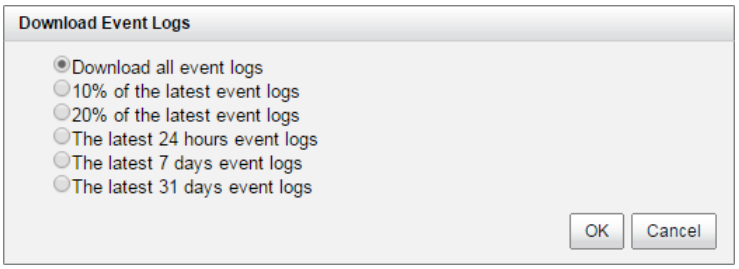

Figure 11-2 Download Event Logs

Clear Event Logs

Click **Clear** button to clear all event logs.

NOTE:

Please plug-in any of the first four hard drives, then event logs can be saved and displayed in next system boot up. Otherwise, the event logs cannot be saved and would be disappeared.

11.2. Monitoring Enclosure

Enclosure provides **Hardware Monitor** and **SES** tabs to show and monitor enclosure information.

11.2.1. Hardware Monitor

The **Hardware Monitor** tab shows the information of current voltages, temperatures and status of power supply and cooling.

Enclosure: Local Temperature: (C)

Controller 1 Monitors

Type	Item	Value	Low Critical	Low Warning	High Warning	High Critical	Status
Voltage	+3.3V	+3.33 V	+3.04 V	+3.14 V	+3.47 V	+3.56 V	OK
Voltage	Standby +3.3V	+3.34 V	+3.04 V	+3.14 V	+3.47 V	+3.56 V	OK
Voltage	CPU 10GbE MAC	+1.06 V	+0.98 V	+1.02 V	+1.07 V	+1.12 V	OK
Voltage	CPU South Complex	+1.00 V	+0.90 V	+0.95 V	+1.05 V	+1.10 V	OK
Voltage	CPU KR/KX PHY	+1.29 V	+1.23 V	+1.24 V	+1.35 V	+1.37 V	OK
Voltage	SAS Expander	+0.90 V	+0.81 V	+0.85 V	+0.95 V	+0.99 V	OK
Voltage	SAS IOC	+0.96 V	+0.87 V	+0.92 V	+1.03 V	+1.08 V	OK
Voltage	Battery	+3.17 V	+2.00 V	+2.30 V	+3.47 V	+3.56 V	OK
Voltage	+1.2V	+1.26 V	+1.08 V	+1.13 V	+1.28 V	+1.32 V	OK
Voltage	+0.82V	+0.82 V	+0.73 V	+0.77 V	+0.86 V	+0.90 V	OK
Voltage	VCore	+1.81 V	+1.40 V	+1.45 V	+1.97 V	+2.00 V	OK
Temperature	CPU	+52.0 (C)	+0.0 (C)	+5.0 (C)	+82.0 (C)	+95.0 (C)	OK
Temperature	SLOT1	+53.2 (C)	-20.0 (C)	+0.0 (C)	+80.0 (C)	+85.0 (C)	OK
Temperature	SAS Expander	+81.3 (C)	+0.0 (C)	+5.0 (C)	+90.0 (C)	+95.0 (C)	OK

Controller 2 Monitors

Type	Item	Value	Low Critical	Low Warning	High Warning	High Critical	Status
Voltage	+3.3V	+3.30 V	+3.04 V	+3.14 V	+3.47 V	+3.56 V	OK
Voltage	Standby +3.3V	+3.33 V	+3.04 V	+3.14 V	+3.47 V	+3.56 V	OK
Voltage	CPU 10GbE MAC	+1.05 V	+0.98 V	+1.02 V	+1.07 V	+1.12 V	OK
Voltage	CPU South Complex	+1.00 V	+0.90 V	+0.95 V	+1.05 V	+1.10 V	OK
Voltage	CPU KR/KX PHY	+1.29 V	+1.23 V	+1.24 V	+1.35 V	+1.37 V	OK
Voltage	SAS Expander	+0.90 V	+0.81 V	+0.85 V	+0.95 V	+0.99 V	OK
Voltage	SAS IOC	+0.95 V	+0.87 V	+0.92 V	+1.03 V	+1.08 V	OK
Voltage	Battery	+3.07 V	+2.00 V	+2.30 V	+3.47 V	+3.56 V	OK
Voltage	+1.2V	+1.26 V	+1.08 V	+1.13 V	+1.28 V	+1.32 V	OK
Voltage	+0.82V	+0.82 V	+0.73 V	+0.77 V	+0.86 V	+0.90 V	OK
Voltage	VCore	+1.80 V	+1.40 V	+1.45 V	+1.97 V	+2.00 V	OK
Temperature	CPU	+56.0 (C)	+0.0 (C)	+5.0 (C)	+82.0 (C)	+95.0 (C)	OK
Temperature	SLOT1	+53.7 (C)	-20.0 (C)	+0.0 (C)	+80.0 (C)	+85.0 (C)	OK
Temperature	SAS Expander	+74.0 (C)	+0.0 (C)	+5.0 (C)	+90.0 (C)	+95.0 (C)	OK

Backplane

Type	Item	Value	Low Critical	Low Warning	High Warning	High Critical	Status
Voltage	+5V	+4.98 V	+4.50 V	+4.60 V	+5.40 V	+5.50 V	OK
Voltage	PSU +12V	+12.07 V	+10.80 V	+11.04 V	+12.96 V	+13.20 V	OK
Voltage	+3.3V	+3.25 V	+2.97 V	+3.04 V	+3.56 V	+3.63 V	OK
Voltage	Standby +3.3V	+3.26 V	+2.97 V	+3.04 V	+3.56 V	+3.63 V	OK
Temperature	Location Left	+33.5 (C)	-20.0 (C)	+0.0 (C)	+55.0 (C)	+70.0 (C)	OK
Temperature	Location Middle	+34.5 (C)	-20.0 (C)	+0.0 (C)	+55.0 (C)	+70.0 (C)	OK
Temperature	Location Right	+32.0 (C)	-20.0 (C)	+0.0 (C)	+55.0 (C)	+70.0 (C)	OK
Power Supply	PSU1	N/A					OK
Power Supply	PSU2	N/A					OK
Cooling	FAN1	7377 RPM					OK
Cooling	FAN2	7458 RPM					OK
Cooling	FAN3	7458 RPM					OK
Cooling	FAN4	7417 RPM					OK

Figure 11-3 Hardware Monitor

Monitoring Notification

The voltage and temperature are normal if their values are between **Low Warning** and **High Warning**. If the value is lower than **Low Warning** or higher than **High Warning**, the system will send a warning event. If lower than **Low Critical** or higher than **High Critical**, error event will be sent.

Smart Cooling Mechanism

Fan speed will adjust automatically according to system thermal and the status of power and fan. For more information, please refer to XCubeSAN Hardware Manual. The document is available at the website:

https://www.qsan.com/en/download.php?cid=&keywords=XCubeSAN+Hardware+Manual&act=query

11.2.2. Configuring SES

SES (SCSI Enclosure Services) is an enclosure management standard. The host can communicate with the enclosure using a LUN and a specialized set of SCSI commands to monitor hardware characteristics. The **SES** tab is used to enable or disable the management of SES. Enable SES will map an iSCSI LUN or a FC LUN.

Enable SES

Take an example of enabling SES.

1. In **SES** tab, click **Enable SES** button.

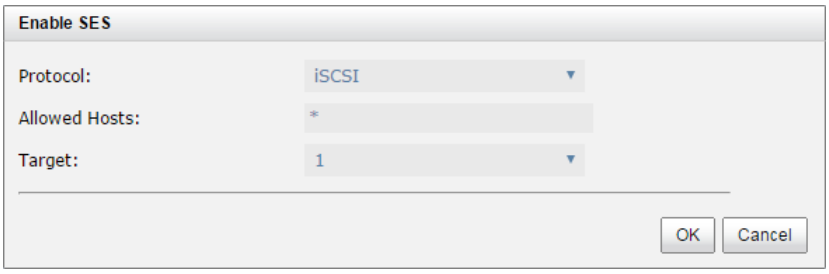

Figure 11-4 Enable SES

2. Select the **Protocol** as iSCSI.
3. Enter the **Allowed Hosts** with semicolons (;) or click **Add Host** button to add one by one. Fill-in wildcard (*) for access by all hosts.
4. Select a **Target** from the drop-down list.
5. Click **OK** button to enable SES.

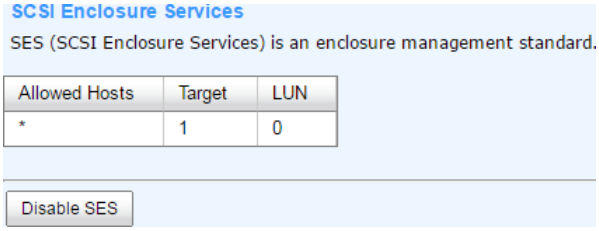

Figure 11-5 SES Status

Disable SES

Click the **Disable** button to disable SES.

SES Client Tool

The SES client software is available at the following web site:

SANtools: http://www.santools.com/

11.3. Monitoring Performance

Performance provides **Disk**, **iSCSI**, and **Fibre Channel** tabs to monitor performance.

11.3.1. Monitoring Disk Performance

The **Disk** tab provides to display the throughput and latency of the physical disk. Check the slots which you want to monitor.

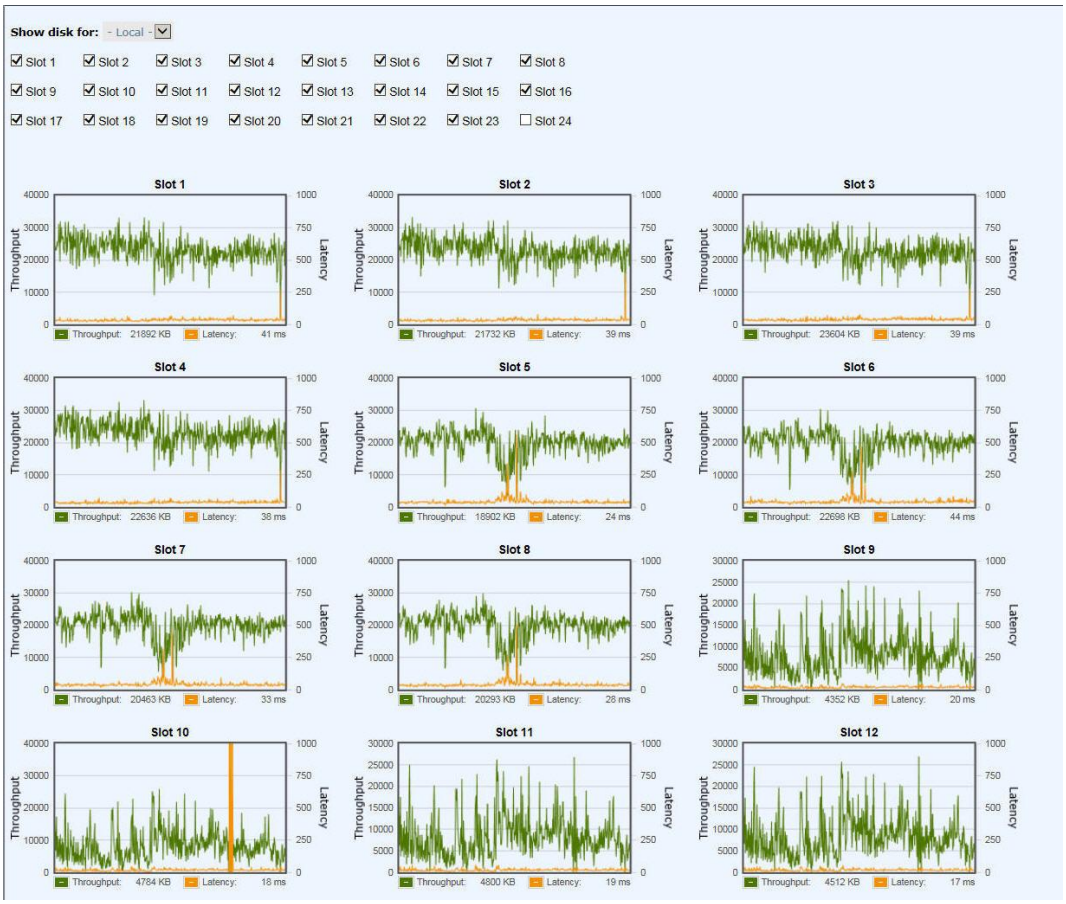

Figure 11-6 Monitoring Disk Performance

11.3.2. Monitoring iSCSI Performance

The **iSCSI** provides to display TX (Transmission) and RX (Reception) of the iSCSI ports. Check the interfaces which you want to monitor.

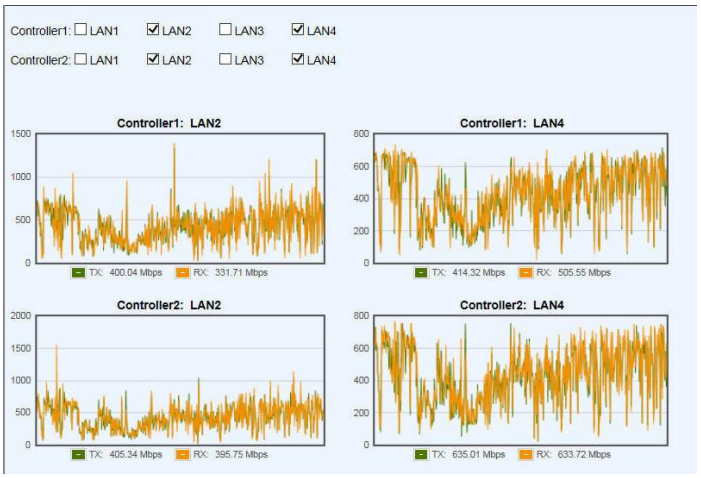

Figure 11-7 Monitoring iSCSI Performance

11.3.3. Fibre Channel

The **Fibre Channel** provides to display TX (Transmission) and RX (Reception) of the fibre channels. Check the interfaces which you want to monitor.

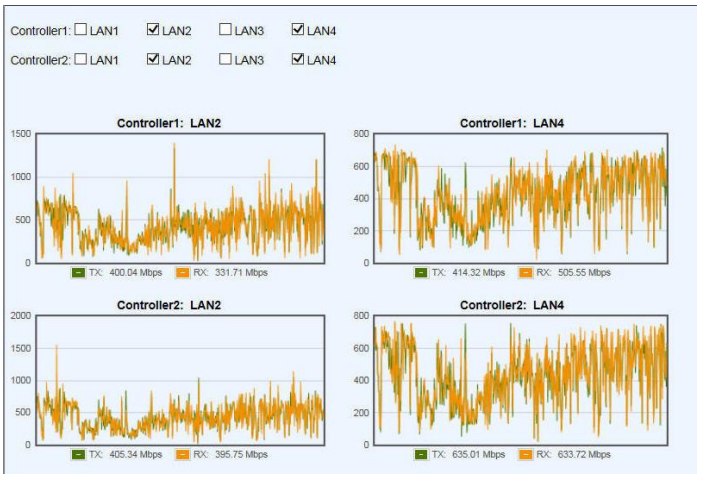

Figure 11-8 Monitoring Fibre Channel Performance

12. Troubleshooting

This chapter describes how to troubleshooting of software. Hardware troubleshooting is described in XCubeSAN Hardware Manual. The document is available at the website:

https://www.qsan.com/en/download.php?cid=&keywords=XCubeSAN+Hardware+Manual&act=query

12.1. Fault Isolation Methodology

This section presents the basic methodology used to locate faults within a storage system.

12.1.1. Basic Steps

The basic fault isolation steps are listed below:

- Gather fault information, including using system LEDs. Please refer to LED Descriptions section in XCubeSAN Hardware Manual for detail.
- Review event logs. Please refer Log Center section in Monitoring chapter.
- Determine where in the system the fault is occurring. If the fault is occurring from hardware, please refer to Troubleshooting section in XCubeSAN Hardware Manual for detail.
- If required, isolate the fault to disk drive, configuration, or system; please refer to Diagnostic Steps section.
- If not above all, please call support for help.

12.1.2. Stopping I/O

When troubleshooting disk drive and connectivity faults, stop I/O to the affected volume data from all hosts and remote systems as a data protection precaution. As an additional data protection precaution, it is helpful to conduct regularly scheduled backups of your data.

12.1.3. Diagnostic Steps

The diagnostic steps are listed below:

- If the fault is occurring from disk drive with event log of disk drive errors, we suggest replacing a healthy one. For more information, please refer to Rebuild section.
- If the fault is occurring from volume configuration. For example, the pool or volume configuration is deleted by accident; please refer to Volume Restoration section for disaster recovery.
- If the fault is occurring from components of the system, it needs to replace a device. The data can be move to new one by disk roaming. For more information, please refer to Disk Roaming section.

12.2. Rebuild

If one disk of the pool which is set as protected RAID level (e.g.: RAID 5, or RAID 6) fails or has been removed, then the status of pool will be changed to degraded mode. At the same time, the system will search the spare disk to execute volume rebuild the degraded pool into complete one.

There are three types of spare disks which can be set in **Disks**:

- **Dedicated Spare**: The disk drive has been set as dedicated spare of a pool.
- **Local Spare**: The disk drive has been set as local spare of the enclosure.
- **Global Spare**: The disk drive has been set as global spare of whole system.

The detection sequence is the dedicated spare disk as the rebuild disk first, then local spare disk and global spare disk.

The following examples are scenarios for a RAID 6.

1. When there is no global spare disk or dedicated spare disk in the system, The pool will be in degraded mode and wait until there is one disk assigned as spare disk, or the failed disk is removed and replaced with new clean disk, and then the Auto-Rebuild starts.
2. When there are spare disks for the degraded array, system starts Auto-Rebuild immediately. In RAID 6, if there is another disk failure occurs during rebuilding, system will start the above Auto-Rebuild process as well. Auto-Rebuild feature only works at that the status of pool is **Online**. Thus, it will not conflict with the online roaming feature.
3. In degraded mode, the health of the pool is **Degraded**. When rebuilding, the status of pool and volume disk will display **Rebuilding**, the column **R%** in volume will display the ratio in percentage. After complete rebuilding, the status will become **Online**.

NOTE:

The dedicated spare cannot be set if there is no pool or only pool s with RAID 0 level.

Sometimes, rebuild is called recover; they are the same meaning. This table describes the relationship between RAID levels and recovery.

Table 12-1 RAID Rebuild

Operation	Description
RAID 0	Disk striping. No protection for data. Pool fails if any disk drive fails or unplugs.
RAID 1	Disk mirroring over 2 disks. RAID 1 allows one disk drive fails or unplugging. Need one new disk drive to insert to the system and rebuild to be completed.
N-way mirror	Extension to RAID 1 level. It has N copies of the disk. N-way mirror allows N-1 disk drives failure or unplugging.
RAID 3	Striping with parity on the dedicated disk. RAID 3 allows one disk drive

	failure or unplugging.
RAID 5	Striping with interspersed parity over the member disks. RAID 5 allows one disk drive failure or unplugging.
RAID 6	2-dimensional parity protection over the member disks. RAID 6 allows two disk drives failure or unplugging. If it needs to rebuild two disk drives at the same time, it will rebuild the first one, then the other in sequence.
RAID 0+1	Mirroring of RAID 0 volumes. RAID 0+1 allows two disk drive failures or unplugging, but at the same array.
RAID 10	Striping over the member of RAID 1 volumes. RAID 10 allows two disk drive failure or unplugging, but in different arrays.
RAID 30	Striping over the member of RAID 3 volumes. RAID 30 allows two disk drive failure or unplugging, but in different arrays.
RAID 50	Striping over the member of RAID 5 volumes. RAID 50 allows two disk drive failures or unplugging, but in different arrays.
RAID 60	Striping over the member of RAID 6 volumes. RAID 60 allows four disk drive failures or unplugging, every two in different arrays.

12.3. Volume Restoration

The **Volume Restoration** can restore the volume configuration from the volume creation history. It is used for pool corrupt and tries to recreate the volume. When trying to do data recovery, the same volume configurations as original must be set and all member disks must be installed by the same sequence as original. Otherwise, data recovery will fail. The volume restoration does not guarantee that the lost data can be restored. Please get help from the expert before executing the function.

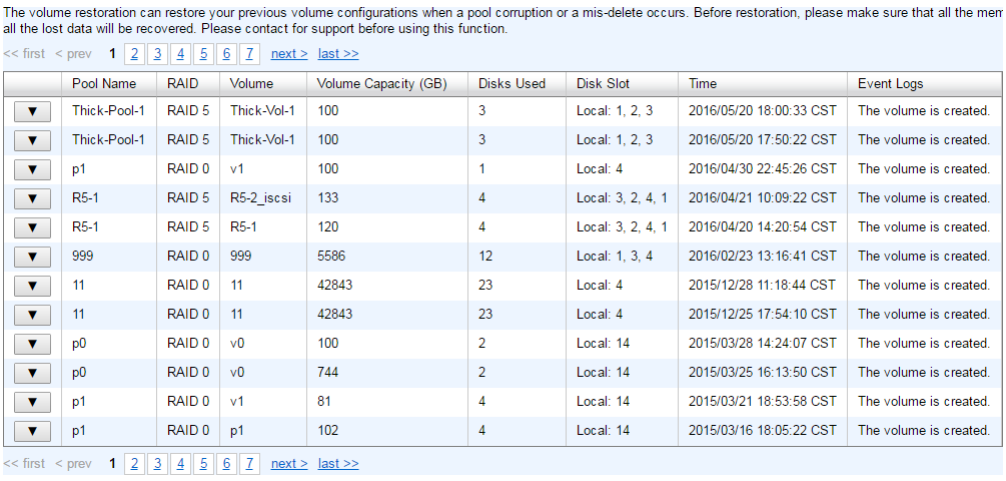
The volume restoration can restore your previous volume configurations when a pool corruption or a mis-delete occurs. Before restoration, please make sure that all the mem
all the lost data will be recovered. Please contact for support before using this function.

<< first < prev 1 2 3 4 5 6 7 next > last >>

	Pool Name	RAID	Volume	Volume Capacity (GB)	Disks Used	Disk Slot	Time	Event Logs
▼	Thick-Pool-1	RAID 5	Thick-Vol-1	100	3	Local: 1, 2, 3	2016/05/20 18:00:33 CST	The volume is created.
▼	Thick-Pool-1	RAID 5	Thick-Vol-1	100	3	Local: 1, 2, 3	2016/05/20 17:50:22 CST	The volume is created.
▼	p1	RAID 0	v1	100	1	Local: 4	2016/04/30 22:45:26 CST	The volume is created.
▼	R5-1	RAID 5	R5-2_iscsi	133	4	Local: 3, 2, 4, 1	2016/04/21 10:09:22 CST	The volume is created.
▼	R5-1	RAID 5	R5-1	120	4	Local: 3, 2, 4, 1	2016/04/20 14:20:54 CST	The volume is created.
▼	999	RAID 0	999	5586	12	Local: 1, 3, 4	2016/02/23 13:16:41 CST	The volume is created.
▼	11	RAID 0	11	42843	23	Local: 4	2015/12/28 11:18:44 CST	The volume is created.
▼	11	RAID 0	11	42843	23	Local: 4	2015/12/25 17:54:10 CST	The volume is created.
▼	p0	RAID 0	v0	100	2	Local: 14	2015/03/28 14:24:07 CST	The volume is created.
▼	p0	RAID 0	v0	744	2	Local: 14	2015/03/25 16:13:50 CST	The volume is created.
▼	p1	RAID 0	v1	81	4	Local: 14	2015/03/21 18:53:58 CST	The volume is created.
▼	p1	RAID 0	p1	102	4	Local: 14	2015/03/16 18:05:22 CST	The volume is created.

<< first < prev 1 2 3 4 5 6 7 next > last >>

This table shows the column descriptions.

Table 12-2 Volume Restoration Column Descriptions

Column Name	Description
Pool Name	The original pool name.
RAID	The original RAID level.
Volume	The original volume name.
Volume Size (GB)	The original capacity of the volume.
Disks Used	The original physical disk number of the pool.
Disk slot	The original physical disk locations.
Time	The last action time of the volume.
Event Log	The last event of the volume.

NOTE:

When trying to do data recovery, the same volume configurations as original must be set and all member disks must be installed by the same sequence as original. Otherwise, data recovery will fail.

CAUTION:

The data recovery does not guarantee that the lost data can be restored 100%. It depends on the real operation and the degree of physical damages on disks. Users will take their own risk to do these procedures.

12.3.1. Volume Restoration Options

Options on the Volume

The options are available on the volume:

- **Restore:** Restore the volume in the pool.

12.4. Disk Roaming

Disks can be re-sequenced in the same system or move all member disks in the same pool from system-1 to system-2. This is called disk roaming. System can execute disk roaming online. Please follow the procedures.

1. In **Pools** tab, selects a pool. And then click ▼ **-> Deactivate**.
2. Click **OK** button to apply. The **Status** changes to **Offline**.
3. Move all member disks of the pool to another system.
4. In **Volumes** tab, selects the pool. And then click ▼ **-> Activate**.
5. Click **OK** to apply. The **Status** changes to **Online**.

Disk roaming has some constraints as described in the followings:

1. Check the firmware version of two systems first. It is better that either systems have the same firmware version or the firmware version of the system-2 is newer.
2. All physical disks of the pool should be moved from system-1 to system-2 together. The configuration of both pool and volume will be kept but LUN configuration will be cleared in order to avoid conflict with the current setting of the system-2.

13. Support and Other Resources

13.1. Accessing Support

After installing your device, locate the serial number on the side sticker of the chassis and use it to register your product at https://partner.qsan.com/ (End-User Registration). We recommend registering your product in QSAN partner website for firmware updates, document download, and latest news in eDM. To contact QSAN Support, please use the following information.

- Website: http://www.qsan.com/en/contact_support.php
- Email: support@qsan.com (09:00 GMT+8 ~ 18:00 GMT+8, 09:00 GMT ~ 18:00 GMT)
- Skype ID: qsan.support (09:00 GMT+8 ~ 18:00 GMT+8, 09:00 GMT ~ 18:00 GMT)

Information to collect

- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages or capture screenshots
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Please refer System Information section to download system information for service by clicking **Download Service Package** button.

13.2. Accessing Update

To download product updates, please go to the following website:

https://www.qsan.com/en/download.php

13.3. Documentation Feedback

QSAN is committed to providing documentation that meets your needs. To help us improve the documentation, email any errors, suggestions, or comments to docsfeedback@qsan.com.

When submitting your feedback, include the document title, part number, revision, and publication date located on the front cover of the document.

Appendix

Glossary and Acronym List

Common Terminology

Item	Description
RAID	Redundant Array of Independent Disks. There are different RAID levels with different degree of data protection, data availability, and performance to host environment.
Disk	The Physical Disk belongs to the member disk of one specific RAID group.
Pool	A collection of removable media. One pool consists of a set of volumes and owns one RAID level attribute.
Volume	Each pool could be divided into several volumes. The volumes from one pool have the same RAID level, but may have different volume capacity.
LUN	Logical Unit Number. A logical unit number (LUN) is a unique identifier which enables it to differentiate among separate devices (each one is a logical unit).
WebUI	Web User Interface.
WT	Write-Through cache-write policy. A caching technique in which the completion of a write request is not signaled until data is safely stored in non-volatile media. Each data is synchronized in both data cache and accessed physical disks.
WB	Write-Back cache-write policy. A caching technique in which the completion of a write request is signaled as soon as the data is in cache and actual writing to non-volatile media occurs at a later time. It speeds up system write performance but needs to bear the risk where data may be inconsistent between data cache and the physical disks in one short time interval.
RO	Set the volume to be Read-Only.
DS	Dedicated Spare disks. The spare disks are only used by one specific RAID group. Others could not use these dedicated spare disks for any rebuilding purpose.
LS	Local Spare disks. The spare disks are only used by the RAID groups of the local enclosure. Other enclosure could not use these local spare disks for any rebuilding purpose.
GS	Global Spare disks. It is shared for rebuilding purpose. If some RAID groups need to use the global spare disks for rebuilding, they could get the spare disks out from the common spare disks pool for such requirement.

DG	DeGraded mode. Not all of the array's member disks are functioning, but the array is able to respond to application read and write requests to its virtual disks.
SCSI	Small Computer Systems Interface
SAS	Serial Attached SCSI
S.M.A.R.T.	Self-Monitoring Analysis and Reporting Technology
WWN	World Wide Name
HBA	Host Bus Adapter
SES	SCSI Enclosure Services
NIC	Network Interface Card
BBM	Battery Backup Module
SCM	Super Capacitor Module

FC / iSCSI / SAS Terminology

Item	Description
FC	Fibre Channel
FC-P2P	Point-to-Point
FC-AL	Arbitrated Loop
FC-SW	Switched Fabric
iSCSI	Internet Small Computer Systems Interface
LACP	Link Aggregation Control Protocol
MPIO	Multi-Path Input/Output
MC/S	Multiple Connections per Session
MTU	Maximum Transmission Unit
CHAP	Challenge Handshake Authentication Protocol. An optional security mechanism to control access to an iSCSI storage system over the iSCSI data ports.
iSNS	Internet Storage Name Service
SAS	Serial Attached SCSI

Dual Controller Terminology

Item	Description
SBB	Storage Bridge Bay. The objective of the Storage Bridge Bay Working Group (SBB) is to create a specification that defines mechanical, electrical and low-level enclosure management requirements for an enclosure controller slot that will support a variety of storage controllers from a variety of independent hardware vendors ("IHVs") and system vendors.

6G MUX	Bridge board is for SATA II disk to support dual controller mode.

End-User License Agreement (EULA)

Please read this document carefully before you use our product or open the package containing our product.

YOU AGREE TO ACCEPT TERMS OF THIS EULA BY USING OUR PRODUCT, OPENING THE PACKAGE CONTAINING OUR PRODUCT OR INSTALLING THE SOFTWARE INTO OUR PRODUCT. IF YOU DO NOT AGREE TO TERMS OF THIS EULA, YOU MAY RETURN THE PRODUCT TO THE RESELLER WHERE YOU PURCHASED IT FOR A REFUND IN ACCORDANCE WITH THE RESELLER'S APPLICABLE RETURN POLICY.

General

QSAN Technology, Inc. ("QSAN") is willing to grant you (“User”) a license of software, firmware and/or other product sold, manufactured or offered by QSAN (“the Product”) pursuant to this EULA.

License Grant

QSAN grants to User a personal, non-exclusive, non-transferable, non-distributable, non-assignable, non-sub-licensable license to install and use the Product pursuant to the terms of this EULA. Any right beyond this EULA will not be granted.

Intellectual Property Right

Intellectual property rights relative to the Product are the property of QSAN or its licensor(s). User will not acquire any intellectual property by this EULA.

License Limitations

User may not, and may not authorize or permit any third party to: (a) use the Product for any purpose other than in connection with the Product or in a manner inconsistent with the design or documentations of the Product; (b) license, distribute, lease, rent, lend, transfer, assign or otherwise dispose of the Product or use the Product in any commercial hosted or service bureau environment; (c) reverse engineer, decompile, disassemble or attempt to discover the source code for or any trade secrets related to the Product, except and only to the extent that such activity is expressly permitted by applicable law notwithstanding this limitation; (d) adapt, modify, alter, translate or create any derivative works of the Licensed Software; (e) remove, alter or obscure any copyright notice or other proprietary rights notice on the Product; or (f) circumvent or attempt to circumvent any methods employed by QSAN to control access to the components, features or functions of the Product.

Open Source

The Product may contain open source components licensed to QSAN. User may visit QSAN’s website to learn specifics of the open source components and the respective license terms thereof (“Open Source License”). The terms of the Open Source License will control solely with

respect to the open source components to the extent that this EULA conflicts with the requirements of the open source with respect to your use of the open source components, and, in such event, you agree to be bound by the Open Source License with respect to your use of such components.

Disclaimer

QSAN DISCLAIMS ALL WARRANTIES OF PRODUCT, INCLUDING BUT NOT LIMITED TO ANY MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, WORKMANLIKE EFFORT, TITLE, AND NON-INFRINGEMENT. ALL PRODUCTS ARE PROVIDE “AS IS” WITHOUT WARRANTY OF ANY KIND. QSAN MAKES NO WARRANTY THAT THE PRODUCT WILL BE FREE OF BUGS, ERRORS, VIRUSES OR OTHER DEFECTS.

IN NO EVENT WILL QSAN BE LIABLE FOR THE COST OF COVER OR FOR ANY DIRECT, INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, CONSEQUENTIAL OR SIMILAR DAMAGES OR LIABILITIES WHATSOEVER (INCLUDING, BUT NOT LIMITED TO LOSS OF DATA, INFORMATION, REVENUE, PROFIT OR BUSINESS) ARISING OUT OF OR RELATING TO THE USE OR INABILITY TO USE THE PRODUCT OR OTHERWISE UNDER OR IN CONNECTION WITH THIS EULA OR THE PRODUCT, WHETHER BASED ON CONTRACT, TORT (INCLUDING NEGLIGENCE), STRICT LIABILITY OR OTHER THEORY EVEN IF QSAN HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Limitation of Liability

IN ANY CASE, QSAN’S LIABILITY ARISING OUT OF OR IN CONNECTION WITH THIS EULA OR THE PRODUCT WILL BE LIMITED TO THE TOTAL AMOUNT ACTUALLY AND ORIGINALLY PAID BY CUSTOMER FOR THE PRODUCT. The foregoing Disclaimer and Limitation of Liability will apply to the maximum extent permitted by applicable law. Some jurisdictions do not allow the exclusion or limitation of incidental or consequential damages, so the exclusions and limitations set forth above may not apply.

Termination

If User breach any of its obligations under this EULA, QSAN may terminate this EULA and take remedies available to QSAN immediately.

Miscellaneous

1. QSAN reserves the right to modify this EULA.
2. QSAN reserves the right to renew the software or firmware anytime.
3. QSAN may assign its rights and obligations under this EULA to any third party without condition.
4. This EULA will be binding upon and will inure to User’s successors and permitted assigns.
5. This EULA shall be governed by and constructed according to the laws of R.O.C. Any disputes arising from or in connection with this EULA, User agree to submit to the jurisdiction of Taiwan Shilin district court as first instance trial.

Index