



# XCubeSAN Series White Paper

## Auto Tiering2.0



QSAN Technology, Inc.  
[www.QSAN.com](http://www.QSAN.com)



## Copyright

© Copyright 2018 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

## January 2018

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

## Trademarks

QSAN, the QSAN logo, XCubeSAN, and QSAN.com are trademarks or registered trademarks of QSAN Technology, Inc.

Microsoft, Windows, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Linux is a trademark of Linus Torvalds in the United States and/or other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Mac and OS X are trademarks of Apple Inc., registered in the U.S. and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

VMware, ESXi, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other countries.

Citrix and Xen are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries.

Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names or their products are the property of their respective owners.

## Notices

---

This XCubeSAN series white papers applicable to the following XCubeSAN models:

### XCubeSAN Storage System 4U 19" Rack Mount Models

Model Name	Controller Type	Form Factor, Bay Count, and Rack Unit
XS5224D	Dual Controller	LFF 24disk 4U Chassis
XS3224D	Dual Controller	LFF 24disk 4U Chassis
XS3224S	Single Controller	LFF 24disk 4U Chassis
XS1224D	Dual Controller	LFF 24disk 4U Chassis
XS1224S	Single Controller	LFF 24disk 4U Chassis

### XCubeSAN Storage System 3U 19" Rack Mount Models

Model Name	Controller Type	Form Factor, Bay Count, and Rack Unit
XS5216D	Dual Controller	LFF 16disk 3U Chassis
XS3216D	Dual Controller	LFF 16disk 3U Chassis
XS3216S	Single Controller	LFF 16disk 3U Chassis
XS1216D	Dual Controller	LFF 16disk 3U Chassis
XS1216S	Single Controller	LFF 16disk 3U Chassis

### XCubeSAN Storage System 2U 19" Rack Mount Models

Model Name	Controller Type	Form Factor, Bay Count, and Rack Unit
XS5212D	Dual Controller	LFF 12disk 2U Chassis
XS5212S	Single Controller	LFF 12disk 2U Chassis
XS3212D	Dual Controller	LFF 12disk 2U Chassis
XS3212S	Single Controller	LFF 12disk 2U Chassis
XS1212D	Dual Controller	LFF 12disk 2U Chassis
XS1212S	Single Controller	LFF 12disk 2U Chassis
XS5226D	Dual Controller	SFF 26disk 2U Chassis
XS5226S	Single Controller	SFF 26disk 2U Chassis
XS3226D	Dual Controller	SFF 26disk 2U Chassis
XS3226S	Single Controller	SFF 26disk 2U Chassis
XS1226D	Dual Controller	SFF 26disk 2U Chassis

XS1226S	Single Controller	SFF 26disk 2U Chassis
---------	-------------------	-----------------------

Information contained in this document has been reviewed for accuracy. It 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 product's features, functionality, and product specifications, which are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

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 refer to 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.

# Table of Contents

---

<b>Notices .....</b>	<b>i</b>
<b>Auto Tiering 2.0 .....</b>	<b>1</b>
ExecutiveSummary.....	1
Audience.....	2
Overview.....	2
Tier Categories.....	3
Flexible RAID and Disk Configurations.....	5
Theory of Operation.....	5
Auto Tiering Architecture.....	6
Intelligent Auto Tiering Mechanism.....	7
Tiering Policies.....	9
Configure Auto Tiering Pools.....	11
Enable Auto Tiering License.....	12
Create an Auto Tiering Pool.....	12
List Auto Tiering Pools.....	18
Operations on Auto Tiering Pools.....	22
Add a Tier (Disk Group) in an Auto Tiering Pool.....	23
Hot Spares in an Auto Tiering Pool.....	26
Configure Volumes.....	26
Create a Volume in an Auto Tiering Pool.....	26
List Volumes and Operations on Volumes.....	30
Configure LUN Mappings and Connect by Host Initiator.....	31
Transfer to Auto Tiering Pool.....	31
Transfer from Thick Provisioning Pool to Auto Tiering.....	32
Transfer from Thin Provisioning Pool to Auto Tiering.....	34
SSD Cache vs. Auto Tiering.....	37
Best Practice.....	39
Configuration Planning Advice.....	39
Case 1: Video Editing.....	43
Case 2: VMware.....	50
Case 3: Sudden Reaction.....	57
Auto Tiering Notices.....	62
Conclusion.....	63
Apply To.....	63
Reference.....	63
<b>Appendix.....</b>	<b>64</b>



Related Documents..... 64  
Technical Support..... 64

# Auto Tiering 2.0

---

## Executive Summary

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

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

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

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

Auto tiering is a feature available on XCubeSAN series and requires license to activate. This document discusses the Auto tiering technology and describes its features, functions, management and best practice



### INFORMATION:

Auto tiering 2.0 with flexible RAID and disk configurations is available in SANOS firmware 1.2.0

---

## Audience

This document is applicable for QSAN customers and partners who are familiar with QSAN products and considering using auto tiering function. Any settings which are configured with basic operations will not be detailed in this document. If there is any question, please refer to the user manuals of products, or contact QSAN support for further assistance.


## Overview

From the perspective of storage performance, the performance of SSDs are high, but the cost is also higher per GB. Relatively speaking, the cost of a traditional hard drive is low, so as performance is relatively poor. If we follow the 80/20 rule to configure storage systems, all SSD configurations are unreasonable for all but the most intensive applications. SSD will be needed in only a small part for most typical applications, regardless of whether or not a critical application, thus giving SSD resources for general storage needs is hugely cost prohibitive. Although traditional hard disk performance is enough for general applications which I/O requirements are not high, the traditional hard drive configuration is also gradually been inadequate.

On the other hand, the data has a life cycle. Since the data in the course of its life cycle, it has experienced different levels of activity. Common usage, when creating the data, it is usually used. As the age of the data increases, it is accessed less often.

### The Solution

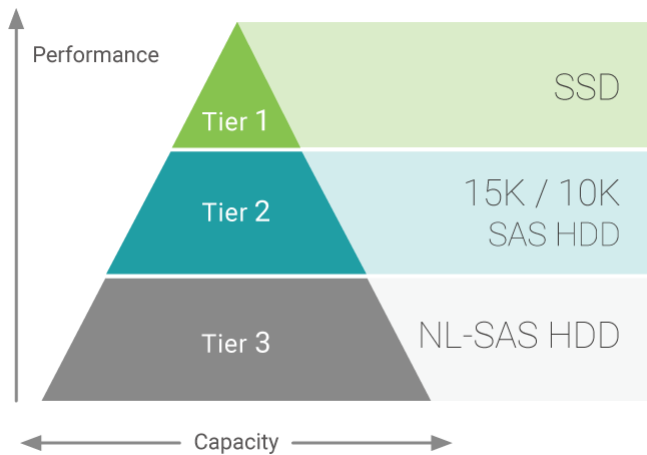
Therefore, to balance performance and cost factors, adapting hybrid storage architecture with a mixture of SSD and traditional HDDs seem to be the most reasonable approach for modern IT environments. Generally, SSD-based storage capacity in 10 to 15% of the total storage capacity should be enough to fulfill the requirements of critical high I/O applications. An automated tiering policy is a simple and elegant solution for dynamically matching storage requirements with changes in the frequency of data access.



## Tier Categories

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

- Tier 1:SSD drives for extreme performance tier
- Tier 2:SAS drives (15K or 10K RPM SAS HDD) for performance tier
- Tier 3:Nearline SAS drives (7.2K or lower RPM SAS HDD) for capacity tier
- Tier 4: SATA drives for capacity tier (for single controller only, not recommended)



### Tier 1 / SSD Tier / Extreme Performance Tier

Use the SSD tier when response time and performance are the most important criteria for storage. This tier uses flash technology that does not contain moving parts. This revolutionary technology eliminates the rotational latency and can improve performance and save energy significantly.

Compared to traditional spinning drives, SSD drives have higher cost per byte, but lower per IO cost. For the best practice use the SSD drive to get data that requires fast response time and high IOPS. Auto tiering enables you to optimize the use of these high performance resources because it automatically relocates "hot" data to the SSD.

### Tier 2 / SAS HDD Tier / Performance Tier

Use the SAS HDD tier to achieve a combination of performance and capacity. The SAS HDD tier provides high levels of performance, reliability, and capacity. SAS HDD stores data on a series of fast rotating disks based on mechanical hard disk drive technology.

This tier includes 15K and 10K RPM spinning drives, which are valuable because it provides a high level performance with consistent response time, high throughput and good bandwidth at moderate price.

### Tier 3 / NL-SAS HDD Tier / Capacity Tier

Use the NL-SAS HDD tier to reduce the cost per GB of data. This consists of 7.2K or lower RPM SAS HDD which is designed to achieve the maximum capacity at an appropriate performance level. While SAS HDDs have slower speeds than SAS HDDs, NL-SAS HDDs

can significantly reduce power consumption and extend capacity in more expensive and higher-performance storage tiers.

In a typical system, most of the application data has very low I/O activity. Because NL-SAS HDDs cost less per GB, they are the most appropriate media type for "cold" data. NL-SAS HDDs consume less power than SAS HDDs and provide a total cost of ownership improvement that takes into account purchase cost.

## Flexible RAID and Disk Configurations

Auto Tiering 2.0 supports flexible RAID and disk configurations. You can create each tier (disk group) with different RAID levels and a different quantity of disk. For example, SSD tier uses 4 disks with RAID 10 for extreme performance, SAS tier uses 6 disks with RAID 6, and NL-SAS tier uses 8 disks with RAID 5 for capacity. This feature is very important for IT administrators to arrange storage plans flexibly.

**RAID Configuration**

Please select RAID levels.

**SSD Tier**

RAID Level : RAID 10 ▼

Quantity of Subgroups : 2 ▼

Quantity of SSD Disks : 4 Disk(s)

**SAS Tier**

RAID Level : RAID 6 ▼

Quantity of SAS Disks : 6 Disk(s)

**NL-SAS Tier**

RAID Level : RAID 5 ▼

Quantity of NL-SAS Disks : 8 Disk(s)

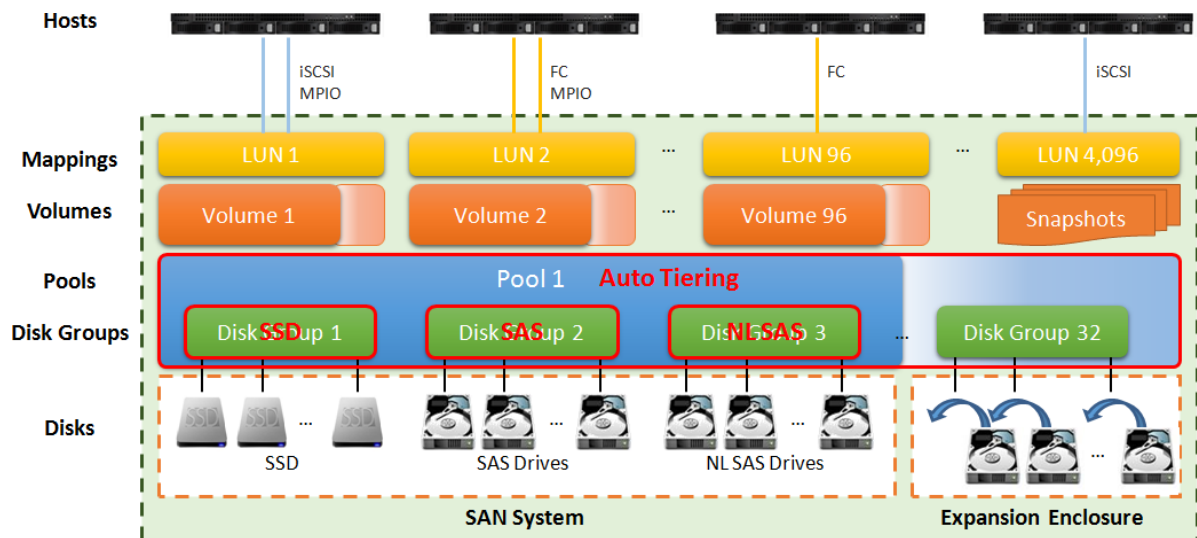
## Theory of Operation

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 storage 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 [cases](#) provides the ability to pin data to tiers permanently or for specific periods of time.

## Auto Tiering Architecture

A newly created auto tiering pool is based on provisioning technology. Each tier works based on one or more disk groups. The following is the storage architecture of an auto tiering pool.



To increase the capacity of an auto tiering pool, any tier (disk group) which contains either one tier of SSDs, SAS HDDs, or NL SAS HDDs can be added to the pool any time. An auto tiering pool can have up to 32 disk groups with each disk group contains up to 64 disk drives. And the maximum disk drive quantity in a pool is 256. The maximum addressable capacity of each disk group is 64TB. So the maximum capacity in a system is 2560TB. more information about pool operation, please refer to the [Configuring Auto Tiering Pools](#) section.

Item	Value
Maximum disk group quantity in a pool	32
Maximum disk drive quantity in a disk group (include dedicated spares)	64

Maximum disk drive quantity in a pool (include dedicated spares)	256
Maximum pool quantity per system	64
Maximum dedicated spare quantity in a pool	8
Maximum tiers (include SSD, SAS HDD, NSAS HDD)	3
Maximum addressable capacity of a disk group	64TB
Maximum addressable capacity of an auto tiering pool	256TB
Maximum addressable capacity of total auto tiering pools (include thin provisioning pools)	1,024TB
Provisioning granularity	1GB

By design, the auto tiering feature allows selecting policies that define how data are moved between different tiers and in many cases provides the ability to pin data to tiers permanently or for specific periods of time.

Auto tiering storage is the assignment of different categories of data to different tiers. It operates based on relocating the most active data up to the highest available tier and the least active data down to the lowest tier. Auto tiering works based on 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 reserved in each higher tier to prepare for the data allocation for those tiering policies which would allocate initial space in highest available tier. 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, statistic collection by accessed counts, ranking hotness data by the statistic collection, and then relocation data via ranking.

## Intelligent Auto Tiering Mechanism

Auto tiering storage management system manages the data relocation and monitors the data hotness ratio using half coefficient and advanced ranking algorithm. It operates on three major functions.

## Statistics Collection

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. Activity level of a sub-LUN is determined by counting the quantity of read and write access on the sub-LUN. Logical volume manager maintains a cumulative I/O count and weights each 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 algorithm. 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.



### TIP:

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


---

## Relocation

According to the hotness analysis, relocation is processing the used defined relocation window which is the number of minutes 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. Value of relocation rate is Fast, Medium, 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 higher tiers if the space they occupy is required for a higher priority. Using this mechanism, auto tiering makes sure that the higher performing drives are always used.

During I/O as 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 LUNs before the next location. 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 triggered by himself. The following figure provides an illustration of how auto tiering can improve sub-LUN placement in a pool.



## Tiering Policies

For the best performance in various environments, auto tiering is a completely automated feature that implements a set of tiering policies. 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 to create a candidate list for data movement. The following volume policies are available:

### Auto Tiering (Default)

It allows 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 sub-LUN. Sub-LUNs are relocated based on

the highest performance disk drives available as 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. This is the recommended policy and it is the default policy for each newly created volume.

## Start Highest then Auto Tiering

This takes advantage of 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 disk** 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 highest 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 sub LUNs into the second higher tier. Initial space is allocated in highest available tier. Auto tiering would prioritize sub LUNs with highest available tier selected above 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 LUN 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 **“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, sub LUN provisioned to the volume is not relocated across tiers. Data remains in its current position, but can still be relocated within

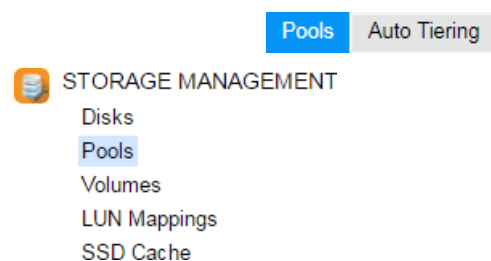
the tier. The system still collects statistics on these LUNs after the tiering policy is changed. Initial space is allocated in the tier which is healthier and has more capacity than other tiers. No relocation would be performed in a volume which selects “No data movement” tiering policy.

The following table summarizes the tiering policies.

Tiering Policy	Description
Auto Tiering	Sets the initial data placement to the optimized tier (disk group) and then relocates the data based on the statistics such that data is relocated among tiers according to the I/O activity.
Start Highest then Auto Tiering	First sets the preferred tier for the initial data placement to the highest tiers with available space, then relocates the data based on the statistics and the auto tiering algorithm.
Highest Available Tier	Sets the preferred tier for the initial data placement to the highest tiers with available space, and so as the succeeding data relocation.
Lowest Tier	Sets the preferred tier for the initial data placement to the lowest tiers with available space, and so as the succeeding data relocation.
No Data Movement	Sets the preferred tier for the initial data to the optimized tier and retains the data without movement.

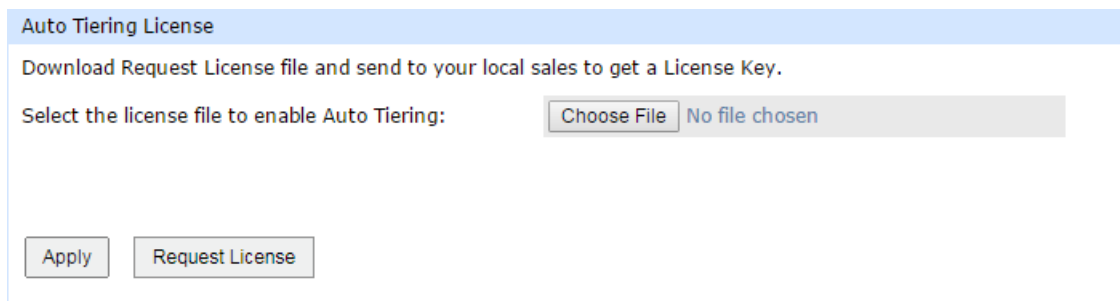
## Configure Auto Tiering Pools

This section will describe the operation of configuring auto tiering pool.



## Enable Auto Tiering License

The auto tiering function is optional. Before using it, you have to enable auto tiering license. Select the **Update** function tab in the **Maintenance** function submenu, download **Request License** file and send to your local sales to obtain a License Key. After getting the license key, click the **Choose File** button to select it, and then click the **Apply** button. 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.



## Create an Auto Tiering Pool

Here is an example of creating an auto tiering pool with 3 tiers, each tier has 3 disks configured in RAID 5. At the first time of creating an auto tiering pool, it may contain at least 2 tiers (disk groups) and the maximum quantity of disk in a tier (disk group) is 8.

1. Select the **Pools** function submenu, click the **Create Pool** button. It will scan available disks first.



**TIP:**

It may take 20 ~ 30 seconds to scan disks if your system has more than 200 disk drives. Please wait patiently.

**Create Pool**

**General**

- Disk Selection
- RAID Configuration
- Disk Properties
- Summary

**Pool Type**

Please select a pool type.

Thick Provisioning  
 Thin Provisioning  
 Auto Tiering (Thin Provisioning Enabled)

**Pool Properties**

Please enter a pool name and select preferred controller setting.

Pool Name :  ⓘ

Preferred Controller :  ▾

The I/O resources will be managed by the preferred controller which you specified.

Next Cancel

2. Select the **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. The maximum length of the pool name is 16 characters. Valid characters are [A-Z | a-z | 0-9-\_|<> ].
4. Select a **Preferred Controller** from the dropdown list. The backend I/O resources in this pool will be processed by the preferred controller which you specified. This option is available when dual controllers are installed.
5. Click the **Next** button to continue.

**Create Pool**

**General**

[Disk Selection](#)

RAID Configuration

Disk Properties

Summary

**Select Disks**

Please select at least two disk types of disks to add tiers an auto tiering pool. Each tier is a disk group. The maximum quantity of disk in a disk group is 64.

Enclosure ID: 0 (Head Unit: XS5216)

<input type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input checked="" type="checkbox"/>	0	1	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input checked="" type="checkbox"/>	0	2	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	3	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	4	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	5	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	6	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	7	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	8	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input checked="" type="checkbox"/>	0	9	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
<input checked="" type="checkbox"/>	0	10	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
<input checked="" type="checkbox"/>	0	11	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
<input checked="" type="checkbox"/>	0	12	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088

Back
Next
Cancel

6. Please select disks for pool and select at least two disk types of disks to add tiers an auto tiering pool. Each tier is a disk group. The maximum quantity of disk in a disk group is 64. Select an **Enclosure ID** from the dropdown list to select disks from expansion enclosures.
7. Click the **Next** button to continue

**Create Pool**

General  
 Disk Selection  
**RAID Configuration**  
 Disk Properties  
 Summary

**RAID Configuration**

Please select RAID levels.

SSD Tier  
 RAID Level : RAID 1  
 Quantity of SSD Disks : 2 Disk(s)

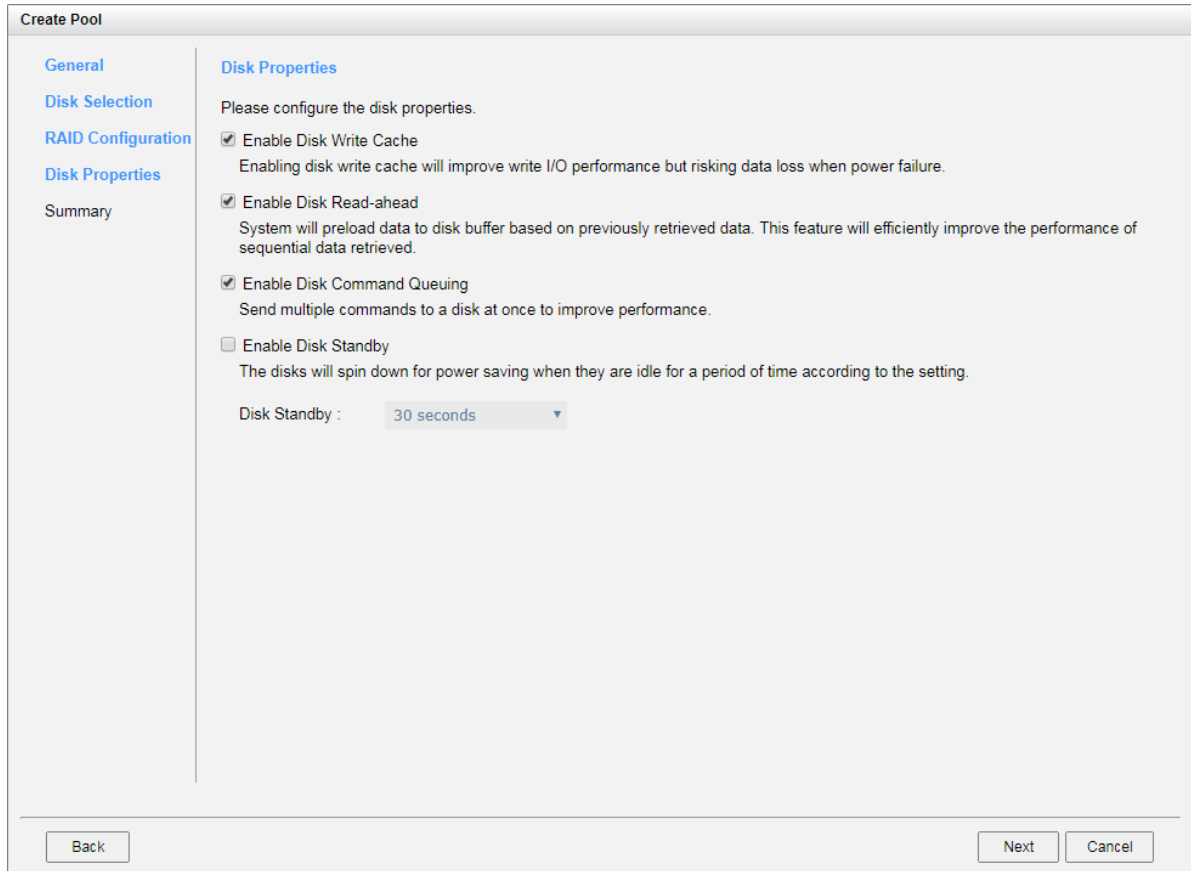
SAS Tier  
 RAID Level : RAID 6  
 Quantity of SAS Disks : 4 Disk(s)

NL-SAS Tier  
 RAID Level : RAID 5  
 Quantity of NL-SAS Disks :

RAID 5  
 RAID 0  
 RAID 1  
 RAID 3  
 RAID 5

Back Next Cancel

8. Select a **RAID Level** from the dropdown list which lists available RAID level only according to the disk selection for each tier. And also select **Quantity of Subgroups** if the combination RAID level is selected.
9. Click the **Next** button to continue



10. Disk properties can also be configured optionally in this step

- **Enable Disk Write Cache:** Check to enable the write cache option of disks. Enabling disk write cache will improve write I/O performance but have a risk of losing data when power failure.
- **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 the **Next** button to continue

**Create Pool**

**General**

[Disk Selection](#)

[RAID Configuration](#)

[Disk Properties](#)

[Summary](#)

**Pool Properties**

Pool Type : Auto Tiering

Pool Name : Pool-3

Preferred Controller : Controller 1

**RAID Configuration**

**SSD Tier**

RAID Level : RAID 1

Quantity of SSD Disks : 2 Disk(s)

**SAS Tier**

RAID Level : RAID 6

Quantity of SAS Disks : 4 Disk(s)

**NL-SAS Tier**

RAID Level : RAID 5

Quantity of NL-SAS Disks : 3 Disk(s)

**Disk Properties**

Disk Write Cache : Enabled

Disk Read-ahead : Enabled

Disk Command Queuing : Enabled

Disk Standby : Disabled

**Schedule Relocation**

Schedule Type : Daily

Relocation Start Time : 00:00

Relocation Period : 0 Hours  
0 Minutes

Relocation Rate : Fast

Back
Finish
Cancel

12. By default, we set relocation schedule at 00:00 daily, relocation period set to 00:00 which means set relocation process run until it finishes and relocation rate to fast.

13. After confirmation at summary page click the **Finish** button to create pool.

	Pool Name	Status	Health	Total	Free	Available	Thin Provisioning	Auto Tiering	Volumes	Current Controller
▼	Pool-3	Online	Good	18.92 TB	18.92 TB	18.92 TB	Enabled	Enabled	0	Controller 1

Create Pool

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



**TIP:**

Auto Tiering 2.0 supports flexible RAID and disk configurations. You can create each tier (disk group) with different RAID level and different quantity of disk. For example, SSD tier uses 4 disks with RAID 10 for extreme performance, SAS tier uses 6 disks with RAID 6, and NL-SAS tier uses 8 disks with RAID 5 for capacity.



**CAUTION:**

Because the auto tiering pool is based on thin provisioning technology, please always watch the system logs of thin provisioning pool. If the capacity of the thin provisioning pool reaches 95% (default thin provisioning policy), the system will deactivate the pool to avoid data loss. So the host cannot access the pool at this time. You have to expand the pool capacity, and then activate the pool to resolve the issue.

## List Auto Tiering Pools

### Pool View

Click a pool; it will display the related disk groups. Similarly, click a disk; it will display the related disk drives. The pool properties can be configured by clicking the functions button ▼ to the left side of the specific pool.

	Pool Name	Status	Health	Total	Free	Available	Thin Provisioning	Auto Tiering	Volumes	Current Controller
▼	Pool-3	Online	Good	13.46 TB	13.46 TB	13.46 TB	Enabled	Enabled	0	Controller 1

**Disk Groups**

	No.	Status	Health	Total	Free	Tier Level	Disks Used	RAID
▼	1	Online	Good	10.92 TB	10.92 TB	NL-SAS	3	RAID 5
▼	2	Online	Good	2.18 TB	2.18 TB	SAS	4	RAID 6
▼	3	Online	Good	372.00 GB	372.00 GB	SSD	2	RAID 1

**Disks**

Enclosure ID	Slot	Status	Health	Capacity	Disk Type	Manufacturer	Model
0	1	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
0	2	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053

Create Pool

This table shows the column descriptions.

Column Name	Description
Pool Name	The poolname.
Status	The status of the pool <ul style="list-style-type: none"> <li>• <b>Online</b>: The pool is online.</li> <li>• <b>Offline</b>: The pool is offline.</li> <li>• <b>Rebuilding</b>: The pool is being rebuilt.</li> <li>• <b>Migrating</b>: The pool is being migrated.</li> <li>• <b>Relocating</b>: The pool is being relocated.</li> </ul>
Health	The health of the pool <ul style="list-style-type: none"> <li>• <b>Good</b>: The pool is good.</li> <li>• <b>Failed</b>: The pool is failed.</li> <li>• <b>Degraded</b>: The pool is not healthy and not complete. The reason could be missing or failed disks.</li> </ul>
Total	Total capacity of the pool.
Free	Free capacity of the pool.
Available	Available capacity of the pool.
Thin Provisioning	The status of Thin provisioning: <ul style="list-style-type: none"> <li>• Disabled.</li> <li>• Enabled.</li> </ul>
Auto Tiering	The status of Auto Tiering: <ul style="list-style-type: none"> <li>• Disabled.</li> <li>• Enabled.</li> <li>• Not Supported: The pool contains disk groups with mixed disk type.</li> </ul>
Volumes	The quantity of volumes in the pool.
Current Controller ( )	The current running controller of the pool.

Column Name	Description
No	The number of disk group.
Status	The status of the disk group <ul style="list-style-type: none"> <li>• <b>Online</b> The disk group is online.</li> <li>• <b>Offline</b> The disk group is offline.</li> <li>• <b>Rebuilding</b> The disk group is being rebuilt.</li> <li>• <b>Migrating</b>: The disk group is being migrated.</li> <li>• <b>Relocating</b> The disk group is being relocated</li> </ul>
Health	The health of the disk group <ul style="list-style-type: none"> <li>• <b>Good</b>: The disk group is good.</li> <li>• <b>Failed</b>: The disk group fails.</li> <li>• <b>Degraded</b> The disk group is not healthy and not completed. The reason could be lack of disk(s) or have failed disk</li> </ul>
Total	Total capacity of the disk group
Free	Free capacity of the disk group
Disks Used	The quantity of disk drives in the disk group.
RAID	The RAID level of the disk group.

Column Name	Description
Enclosure ID	The enclosure ID.
Slot	The position of the disk drive
Status	The status of the disk drive <ul style="list-style-type: none"> <li>• <b>Online</b> The disk drive is online.</li> <li>• <b>Missing</b> The disk drive is missing in the pool</li> <li>• <b>Rebuilding</b> The disk drive is being rebuilt.</li> <li>• <b>Transitioning</b>: The disk drive is being migrated or is replaced by another disk when rebuilding occurs.</li> <li>• <b>Scrubbing</b> The disk drive is being scrubbed.</li> <li>• <b>Check Done</b> The disk drive has been checked the disk health</li> </ul>
Health	The health of the disk drive <ul style="list-style-type: none"> <li>• <b>Good</b>: The disk drive is good.</li> <li>• <b>Failed</b>: The disk drive is failed.</li> <li>• <b>Error Alert</b>: S.M.A.R.T. error alert</li> <li>• <b>Read Errors</b> The disk drive has unrecoverable read errors.</li> </ul>
Capacity	The capacity of the disk drive

Disk Type	The type of the disk drive <ul style="list-style-type: none"> <li>[ SAS HDD   NL-SAS HDD   SAS SSD   SATA SSD ]</li> <li>[ 12.0Gb/s   6.0Gb/s   3.0Gb/s   1.5Gb/s ]</li> </ul>
Manufacturer	The manufacturer of the disk drive
Model	The model name of disk drive

### Auto Tiering View

The **Auto Tiering** function tab in the **Pools** function submen is only visible when auto tiering license is enabled. Click pool; it will display the related tiering status. The properties can be configured by clicking the functions button on the left side of the specific pool.

	Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	Pool-3	Online	Good	13.46 TB	13.46 TB	13.46 TB	0	9	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	372	0	0	0	0	
SAS	2235	0	0	0	0	
NL-SAS	11177	0	0	0	0	

This table shows the column descriptions.

Column Name	Description
Tier Level	Tier categories there are SSD, SAS, Nearline SAS, and SATA. A SATA system will hide the tiers without any disk groups.
Tier Capacity	Total capacity of the tier.
Tier Used	Used capacity of the tier.
Move Up	The capacity prepares to move up to higher tier.
Move Down	The capacity prepares to move down to lower tier.
Move In	The capacity prepares to move in from other tiers.
Tier Status	Bar chart to show the tier status: <ul style="list-style-type: none"> <li>Light Blue Used capacity.</li> </ul>

	<ul style="list-style-type: none"> <li>• <b>Orange</b> The data will move in.</li> <li>• <b>Gray</b>. Unallocated.</li> </ul>
--	---

## Operations on Auto Tiering Pools

Most operations are described in the [Configuring Storage Pools](#) section. For more information, please refer to the [chapter 8.4.3 Operations on Thick Provisioning Pools](#) section and the [chapter 9.3.3 Operations on Thin Provisioning Pools](#) section in the [XCubeSAN SANOS 4.0 User's Manual](#). We describe the operations about auto tiering in the following.

### Schedule Relocation

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

**Schedule Relocation**

Pool Name : Pool-3

Frequency :  Daily  
 Weekly  
 Repeat Every 12 Hours

Relocation Start Time (hh:mm) : 00:00

Relocation Period (hh:mm) : 00 : 00 (Set as 00:00 to let relocation process run until it finishes.)

Relocation Rate : Fast

OK Cancel

## Relocate Now

Click ▼ -> **Relocate Now** to perform relocation right now in an auto tiering pool. If **Relocation Period** sets as 00:00 it will let relocation process run until it finishes.



Relocate Now

Pool Name : Pool-3

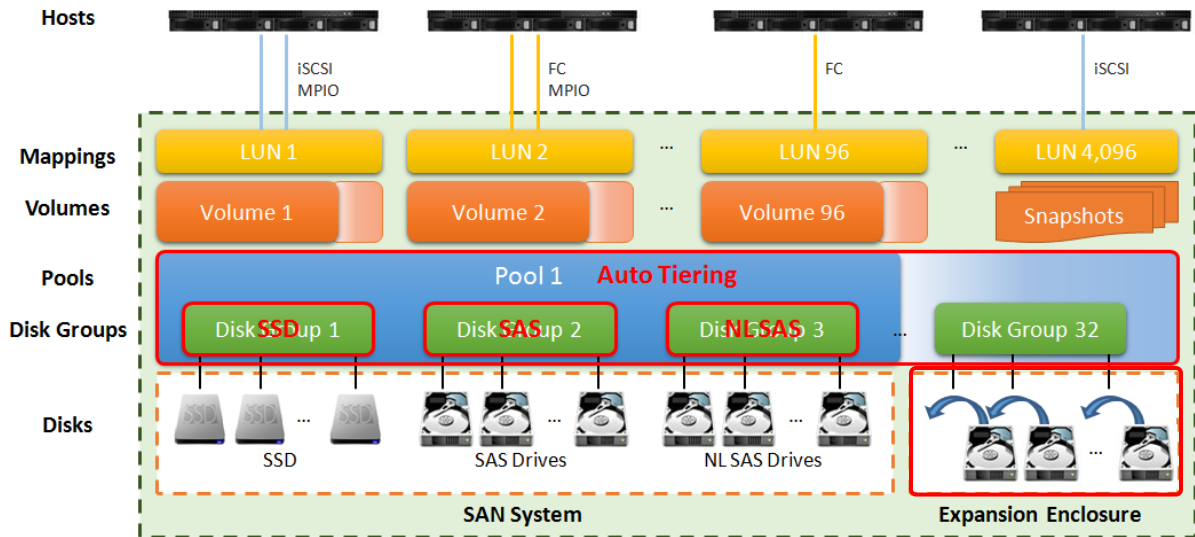
Relocation Period (hh:mm) : 00 : 00 (Set as 00:00 to let relocation process run until it finishes.)

Relocation Rate : Fast

OK Cancel

## Add a Tier (Disk Group) in an Auto Tiering Pool

The **Add Disk Group** function adds a disk group to a pool to increase the capacity.



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

1. Select a pool click ▼ -> **Add Disk Group** to add a disk group in the auto tiering pool

**Add Disk Group**

**Pool Type**

Thin Provisioning : Enabled  
 Auto Tiering : Enabled

**RAID Level**

Please select a RAID level.

RAID Level : RAID 10 ▼  
 Quantity of Subgroups : 2 ▼

**Select Disks**

Please select disks to add a disk group. The maximum quantity of disk in a disk group is 64.

Enclosure ID : 0 (Head Unit: XS5216) ▼

<input type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input type="checkbox"/>	0	3	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	4	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input checked="" type="checkbox"/>	0	5	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input checked="" type="checkbox"/>	0	6	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input checked="" type="checkbox"/>	0	7	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input checked="" type="checkbox"/>	0	8	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	16	Good	5.46 TB	NL-SAS HDD 12.0Gb/s	SEAGATE	ST6000NM0034

2. Select a **RAID Level** from the dropdownlist and also select **Quantity of Subgroups** if the combination RAID level is selected.
3. Please select disks to add a disk group. The maximum quantity of disk in a disk group is 64. Select an **Enclosure** from the dropdown list to select disks from the expansion enclosures.
4. Click the **OK** button to add a disk group

## Hot Spares in an Auto Tiering Pool

In an auto tiering pool, hot spare drives can only replace drives of the same disk type. For example, a SSD tier can only be assigned SSD type drives as hot spares drives.

Enclosure ID: 0 (Head Unit: XS5216) ▾

<< first < prev 1 next > last >>

	Slot	Status	Health	Capacity	Disk Type	Usage	Pool Name	Manufacturer	Model		
▾	1	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	RAID	Pool-3	SEAGATE	ST400FM0053		
▾	2	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	RAID	Pool-3	SEAGATE	ST400FM0053		
▾	3	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	Dedicated Spare	Pool-3	SEAGATE	ST400FM0053		
▾	4	Online	Good	372.36 GB	SAS SSD 12.0Gb/s	Free		SEAGATE	ST400FM0053		
▾	5	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	Free		MICRON	S630DC-800		
			MICRON	S630DC-800	▾	6	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	Free
			MICRON	S630DC-800	▾	7	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	Free
			MICRON	S630DC-800	▾	8	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	Free
	Pool-3	SEAGATE	ST1200MM0088	▾	9	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	RAID	
	Pool-3	SEAGATE	ST1200MM0088	▾	10	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	RAID	
	Pool-3	SEAGATE	ST1200MM0088	▾	11	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	RAID	
	Pool-3	SEAGATE	ST1200MM0088	▾	12	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	RAID	
	Pool-3	SEAGATE	ST6000NM0014	▾	13	Online	Good	5.46 TB	NL-SAS HDD 12.0Gb/s	RAID	
	Pool-3	SEAGATE	ST6000NM0014	▾	14	Online	Good	5.46 TB	NL-SAS HDD 12.0Gb/s	RAID	
	15	Online	Good	5.46 TB	NL-SAS HDD 12.0Gb/s	RAID	Pool-3	SEAGATE	ST6000NM0014	▾	
	16	Online	Good	5.46 TB	NL-SAS HDD 12.0Gb/s	Dedicated Spare	Pool-3	SEAGATE	ST6000NM0014	▾	

## Configure Volumes

This section will describe the operations of configuring volume in auto tiering pool.

### Create a Volume in an Auto Tiering Pool

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

1. Select the **Volumes** function sub menu and click the **Create Volume** button.

**Create Volume**

**General**

Advanced

Summary

**Volume General Settings**

Please enter a volume name and configure the volume general settings.

Volume Name :  ⓘ

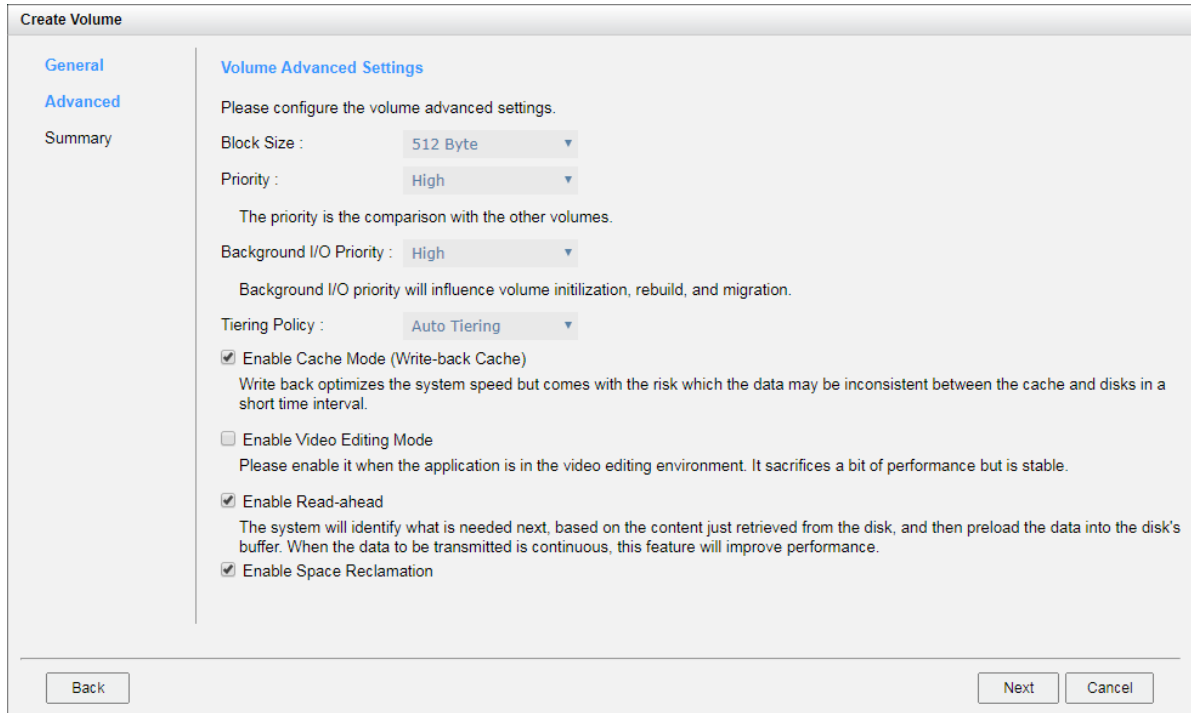
Pool Name :  (Available : 262144 GB)

Capacity :

Volume Type :

Select RAID Volume for general RAID usage or Backup Volume for backup usage such as the target volume of local clone or remote replication.

2. Enter a **Volume Name** for the pool. The maximum length of the volume name is 32 characters. Valid characters are [A-Z | a-z | 0-9 | < > ].
3. Select a **Pool Name** from the dropdown list. It will also display the available capacity of the pool.
4. Enter required **Capacity**. The unit can be selected from the dropdown list.
5. Select **Volume Type**. The options are **RAID Volume** (for general RAID usage) and **Backup Volume** (for the target volume of local clone or remote replication).
6. Click the **Next** button to continue.



7. Volume advanced settings can also be configured optionally in this step:

- **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.
- **Tiering Policy:** The options are Auto Tiering, Start Highest then Auto Tiering, High Available Tier, Lowest Tier, and No Data Movement. Please refer to the [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. This option provides a more stable performance figure without high and low peaks but slower in average.
- **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 the volume when the pool is auto tiering

8. Click the **Next** button to continue

**Create Volume**

[General](#)

[Advanced](#)

[Summary](#)

**Configure Volume General Settings**

Volume Name : Vol-3

Pool Name : Pool-3

Capacity : 100 GB

Volume Type : RAID Volume

**Configure Volume Advanced Settings**

Block Size : 512 Byte

Priority : High

Background I/O Priority : High

Tiering Policy : Auto Tiering

Cache Mode : Enabled

Video Editing Mode : Disabled

Read-ahead : Enabled

Space Reclamation : Enabled

Back
Finish
Cancel

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

10. The volume has been created. It will be initialized in protection RAID level (e.g., RAID 1, 3, 5, 6, 0+1, 10, 30, 50, and 60)

<< first < prev **1** next > last >>

	Volume Name	Status	Health	Capacity	Volume Type	SSD Cache	Snapshot Space	Snapshots	Clone	Write	Pool Name
▼	Vol-3	Online	Optimal	100.00 GB	RAID Volume	Disabled	0 MB / 0 MB	0	N/A	WB	Pool-3

<< first < prev **1** next > last >>

Create Volume
Local Clone Options

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



**TIP:**

SANOS supports instant RAID volume availability. The volume can be used immediately when it is initializing or rebuilding.



**TIP:**

If the pool contains some disk drives of 4Kn type, it is not available to set 512, 1024 or 2048 block size. When the case happens, it will pop up a warning message and suggest changing the block size to 4096.

## List Volumes and Operations on Volumes

Most operations are described in [chapter 8.5 Configuring Volumes](#) section in the [XCubeSAN SANOS 4.0 User Manual](#). We describe auto tiering operations below.

### Change Volume Properties

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

**Change Volume Properties**

Volume Name: Vol-3 ⓘ

Priority:  High  Medium  Low

Background I/O Priority: High ▼

Tiering Policy: Auto Tiering ▼

Cache Mode:  Write-through Cache  Write-back Cache  Read-Only ⓘ

Read-ahead: Enabled ⓘ

Space Reclamation: Enabled ⓘ

RAID Volume: RAID Volume ⓘ

Read-ahead: Read-ahead: [ ]

Space Reclamation: Space Reclamation: [ ]

Volume Type: Volume Type: [ ]

OK Cancel





- Click **▼ -> Add Disk Group** to transfer from a thick provisioning pool to an auto tiering pool. Select **Enabled** from the Auto Tiering dropdown list. The tier (disk group) must be added one at a time. Select **RAID Level** and **Select Disks**, and then click the **OK** button.

**Add Disk Group**

**Pool Type**

Thin Provisioning : Disabled

Auto Tiering : Disabled ▼ i

Disabled

Enabled

**RAID Level**

Please select a RAID level.

RAID Level : RAID 1 ▼

**Select Disks**

Please select disks to add a disk group. The maximum quantity of disk in a disk group is 64.

Enclosure ID : 0 (Head Unit: XS5216) ▼

<input type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input type="checkbox"/>	0	1	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	2	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	3	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	4	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input checked="" type="checkbox"/>	0	5	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input checked="" type="checkbox"/>	0	6	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	7	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	8	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
<input type="checkbox"/>	0	12	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
<input type="checkbox"/>	0	16	Good	5.46 TB	NL-SAS HDD 12.0Gb/s	SEAGATE	ST6000NM0034

OK
Cancel

- Use the same procedure to add another tier if necessary.

	Pool Name	Status	Health	Total	Free	Available	Thin Provisioning	Auto Tiering	Volumes	Current Controller
▼	Pool-1	Online	Good	2.18 TB	2.18 TB	2.18 TB	Disabled	Enabled	0	Controller 1
▼	Pool-2	Online	Good	10.92 TB	10.92 TB	10.92 TB	Enabled	Disabled	0	Controller 1

**Disk Groups**

	No.	Status	Health	Total	Free	Tier Level	Disks Used	RAID
▼	1	Online	Good	2.18 TB	2.18 TB	SAS	3	RAID 5
▼	2	Online	Good	744.00 GB	744.00 GB	SSD	2	RAID 1

**Disks**

Enclosure ID	Slot	Status	Health	Capacity	Disk Type	Manufacturer	Model
0	5	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
0	6	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800

Create Pool

- Auto Tiering** status is **Enabled**. The thick provisioning pool has been transferred to auto tiering.



**TIP:**

The total capacity of the pool does not change even adding the capacity of higher tiers



**CAUTION:**

The action of transferring from the thick provisioning pool to auto tiering is irreversible. Please consider carefully all possible consequences before taking this step.

## Transfer from Thin Provisioning Pool to Auto Tiering

First of all, make sure the auto tiering license is enabled. For more information about enabling license operation, please refer to the [Enable Auto Tiering License](#) section. And then use **Add Disk Group** function to add another tier (disk group). This is an example of transfer from thin provisioning pool to auto tiering one.

- Create a thin provisioning pool with NL-SAS disk drives. **Auto Tiering** status is **Disabled**.



2. Click **▼ -> Add Disk Group** to transfer from a thin provisioning pool to an auto tiering pool. Select **Enabled** from the Auto Tiering dropdown list. The tier (disk group) must be added one at a time. Select **RAID Level** and **Select Disks**, and then click the **OK** button.



	Pool Name	Status	Health	Total	Free	Available	Thin Provisioning	Auto Tiering	Volumes	Current Controller
	Pool-1	Online	Good	2.18 TB	2.18 TB	2.18 TB	Disabled	Enabled	0	Controller 1

**Disk Groups**

	No.	Status	Health	Total	Free	Tier Level	Disks Used	RAID
	1	Online	Good	10.92 TB	10.92 TB	NL-SAS	3	RAID 5
	2	Online	Good	744.00 GB	744.00 GB	SSD	2	RAID 1

**Disks**

Enclosure ID	Slot	Status	Health	Capacity	Disk Type	Manufacturer	Model
0	7	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800
0	8	Online	Good	744.96 GB	SAS SSD 12.0Gb/s	MICRON	S630DC-800

- Auto Tiering** status is **Enabled**. The thin provisioning pool has been transferred to auto tiering.



**TIP:**

The total capacity of the pool is the sum of all tiers.



**CAUTION:**

The action of transferring from the thin provisioning pool to auto tiering is irreversible. Please consider carefully all possible consequences before taking this step.

## SSD Cache vs. Auto Tiering

The SSD cache and auto tiering solutions work together and complement each other. A key difference between tiering and cache is that tiering moves data to SSD instead of simply caching it. Tiering can also move data both from slower storage to faster storage and vice versa. However, SSD cache 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. The important difference between moves and copies is that a cache does not need to have the



## Best Practice

Auto tiering technology provides a solution to achieve optimal storage efficiency and improved performance, making it the most cost effective storage solution for data center environments with dynamic workload changes.

If your applications belong to sequential I/O from beginning to end, such as surveillance or backup, or their access profiles are very random in the large address range, a homogeneous pool is recommended for your applications. In a homogeneous pool, only one drive type (SSD, SAS, or NL-SAS) is selected during pool creation. If using auto tiering technology in these applications, the data will move up and down frequently without any benefit.



**TIP:**

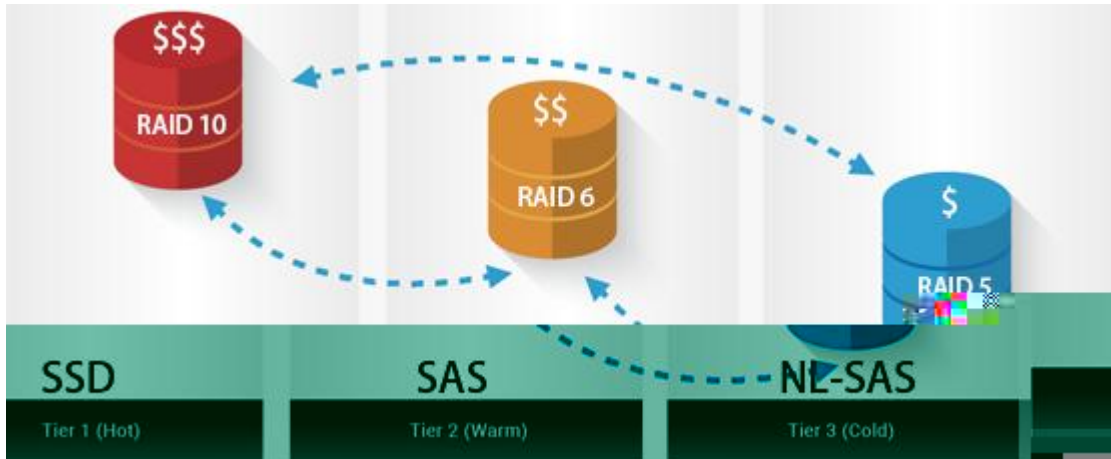
Homogeneous pool is suitable for the application of sequential I/O from beginning to end or very random in the large address range. In addition, auto tiering is suitable for the data which has a life cycle.

## Configuration Planning Advice

### SSD / SAS / NL-SAS Tier RAID Level and Capacity Ratio

The following is a general guide to the auto tiering pool planning. The user can fine-tune according to the actual situation.

- **SSD Tier (\$\$\$)**  
Suggest SSD tier using at least 4 disks with RAID 10 or 2 disks with RAID 0 for extreme performance. Prepare SSD storage capacity in 10% to 15% of the total pool capacity to fulfill the requirements of critical high I/O applications.
- **SAS Tier (\$\$)**  
Suggest SAS HDD tier configuring with RAID 10 or RAID 5. Prepare about 30% of the total storage capacity.
- **NL-SAS Tier (\$)**  
For capacity tier, suggest SAS HDD using RAID 5 level to store cold data. This tier occupies the rest of the storage capacity.



Take an example for reference. First, you can estimate the total capacity used, and estimate how much hot data or high I/O your application uses everyday. Assuming 666GB per day, the recommended SSD tier capacity is at least 1.5 times  $1.5 \times 666\text{GB} = 1\text{TB}$  as a conservative estimate. Then, calculate the SAS HDD tier capacity about 3 times of the SSD tier capacity,  $3 \times 1\text{TB} = 3\text{TB}$ , if the SSD tier full of the buffer, so that the performance does not drop too much. This tier is optional. The remaining space is left for NL-SAS HDD tier. The following table is the summary for reference.

Tier	Capacity per Drive	Quantity	RAID Level	Capacity per Tier	Capacity Ratio
SAS SSD Tier	500GB	4	RAID 10	$(4/2) \times 500\text{GB} = 1\text{TB}$	10%
SAS HDD Tier	1TB	5	RAID 6	$(5-2) \times 1\text{TB} = 3\text{TB}$	30%
NL-SAS HDD Tier	3TB	3	RAID 5	$(3-1) \times 3\text{TB} = 6\text{TB}$	60%

This is a rough planning proposal. Whether to meet customer requirements also requires users to calculate the performance and necessary capacity. Of course if more capacity is needed, you can also add a disk group to any tier.

### Relocation and Its Effect

In the [Intelligent Auto Tiering Mechanism](#) section, we introduced there are three major functions in auto tiering technology: statistics collection and ranking operate automatically but relocation can be configurable manually. We would like to suggest that users can set the schedule relocation at midnight everyday (**Daily 00:00**), the relocation period sets to









- Copy a 100GB file into the volume spends 2 minutes to complete and transmission speeds around 780 ~ 800 MB/s. The figure shows that the SSD tier is being used.

Pools **Auto Tiering**

	Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	Test	Online	Good	7.55 TB	7.45 TB	7.45 TB	1	13	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	81	0	0	0	
SAS	558	0	0	0	0	
NL-SAS	1393	0	0	0	0	

- The first coming file is located in SSD tier because the tiering policy is set as **Start Highest then Auto Tiering**.

Pools **Auto Tiering**

	Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	Test	Online	Good	2.09 TB	1.99 TB	1.99 TB	1	12	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	103	0	0	0	
SAS	557	0	0	0	0	
NL-SAS	1395	0	0	0	0	

- Copy another 100GB file to the volume. Since the capacity of SSD tier is full, the system will save the data at the next tier. It spends 2 minutes and 20 seconds to complete. The transmission speeds around 460 ~ 830 MB/s.

Pools **Auto Tiering**

Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
Test	Online	Good	7.55 TB	7.35 TB	7.35 TB	1	13	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	185	0	0	0	
SAS	558	3	0	0	0	
NL-SAS	1393	0	0	0	0	

- The second file is distributed in the SSD tier and SAS tier.


Pools **Auto Tiering**

Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
Test	Online	Good	2.09 TB	1.89 TB	1.89 TB	1	12	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	185	0	0	0	
SAS	557	17	0	0	0	
NL-SAS	1395	0	0	0	0	

7. After an hour the system analyzes the data automatically and the data will be relocated at midnight. The figure shows that 18GB data in SSD tier will be moved down to the SAS tier.



8. At the next day, 18GB data in SSD tier has been moved down to the SAS tier. The event log records how much data is moved. You can see that SSD tier reserved about 10% of the capacity for incoming data.



9. Continue copying the third 100GB file to the volume. It spends 3 minutes and 8 seconds to complete. The transmission speed is around 460~500 MB/s. The file is copied to the SAS tier.





14. Last, copy this file back to the source volume and observe the transmission speed. You can also compare the performance monitor of disks to the web UI and observe





3. Create eight VMs and save their data stories to the volume. When they are ready, run IOmeter on each VM to observe the performance. Because the tiering policy is set as **Auto Tiering**, the initial space is allocated in the tier which is healthier and has more free capacity than other tiers. The data is located in the NL-SAS tier with RAID 5.


Pools **Auto Tiering**

	Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	test	Online	Good	3.45 TB	3.25 TB	3.25 TB	1	12	Controller 1

**Pool Tiering Status:**

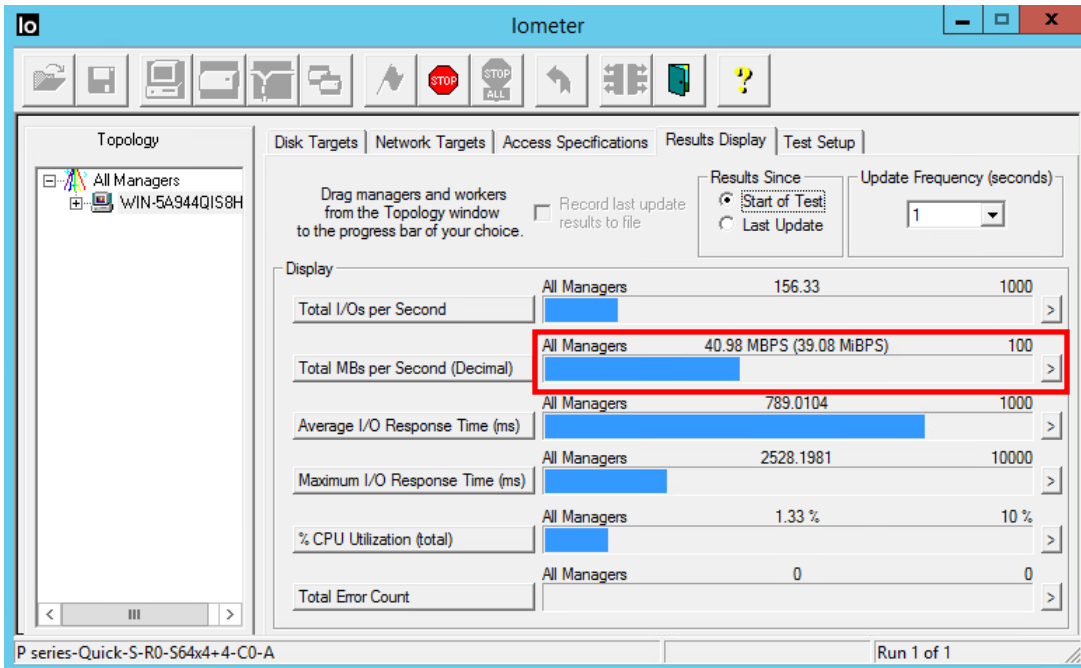
Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	0	0	0	0	
SAS	557	0	0	0	0	
NL-SAS	2793	210	0	0	0	

4. The followings are the throughput of VMs running by IOmeter at the beginning.
  - VM1: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 9.96 MB/s

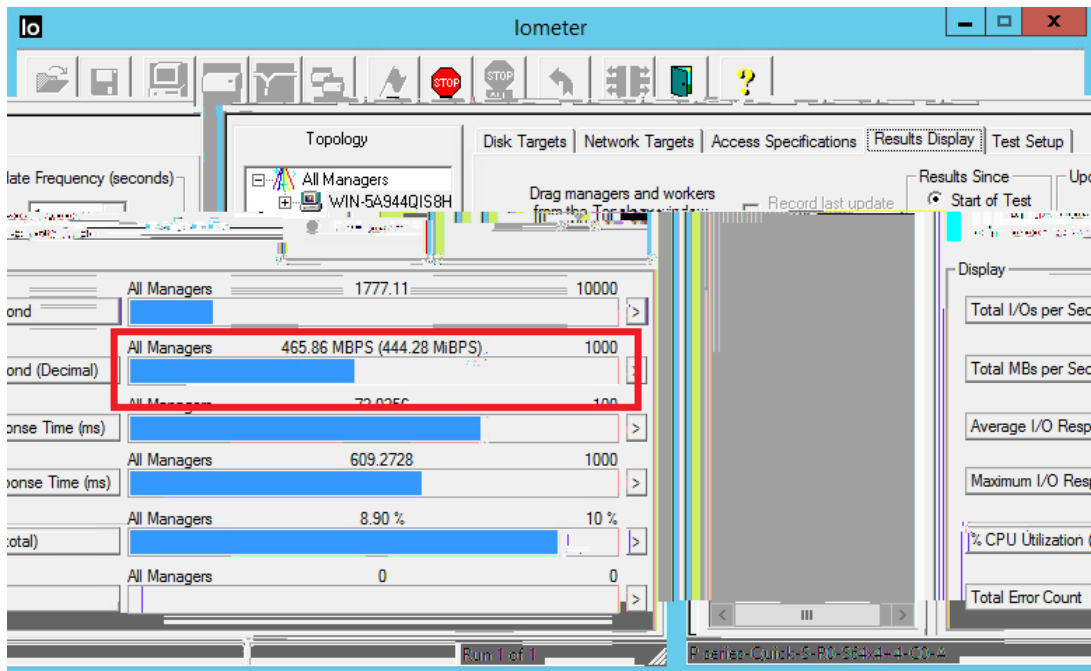


- VM2: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 4.78 MB/s
- VM3: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 4.41 MB/s
- VM4: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 4.13 MB/s
- VM5: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 3.98 MB/s
- VM6: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 3.79 MB/s
- VM7: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 3.70 MB/s
- VM8: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 3.61 MB/s

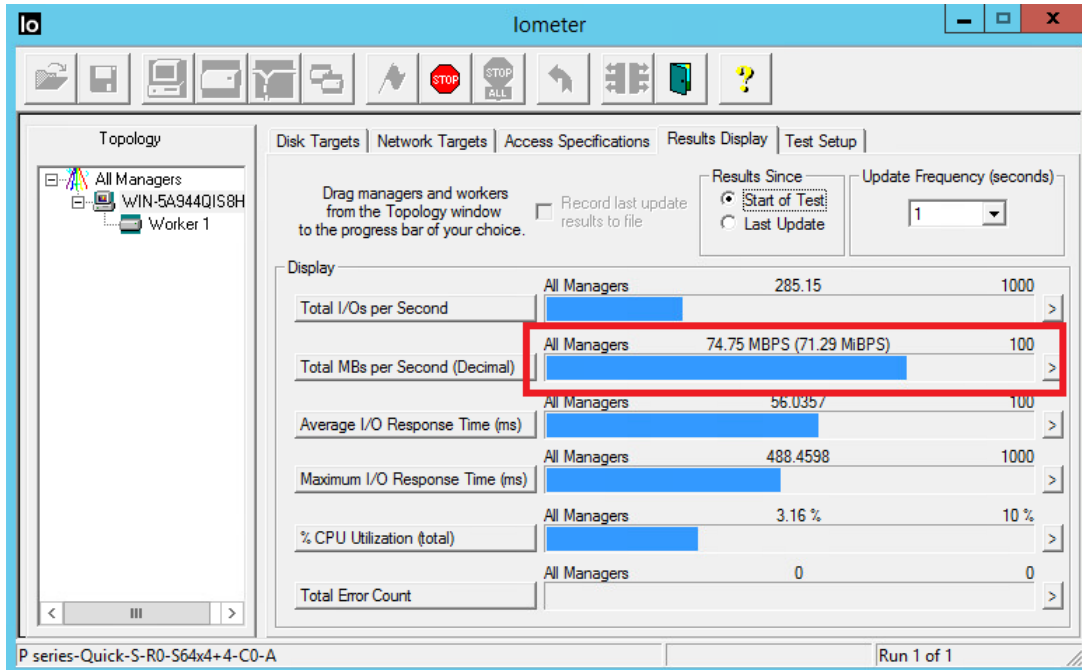
5. Stop VM2~VM8/O but keep VM1 running I/O the throughput of VM1 is up to 40.98 MB/s.



- Because VM1 keeps I/O, the data in VM1 will be accessed frequently than others. After analysis and relocation by auto tiering mechanism, the data in VM1 has been moved to a higher tier. We check the performance of VM1 again; the throughput is up to 46586MB/s.



- Run VM2~VM8 I/O again, check performance. The followings are the throughput of VMs running by Iometer.
  - VM2: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB, the throughput is 74.75 MB/s



- VM3: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 32 MB/s
- VM4: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 48 MB/s
- VM5: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 64 MB/s
- VM6: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 80 MB/s
- VM7: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 96 MB/s
- VM8: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 112 MB/s

## Summary

In case 2, although the autotiering policy is set to **Auto Tiering**, the data is allocated in the tier which is healthier and has more capacity than other tiers at the beginning. Then the data with frequently accessed I/O will be relocated to the higher tier for better performance. The following table summarizes the throughput before and after relocation and an improvement percentage calculation as a reference. This verifies the scenario meets the expectations of VMware.

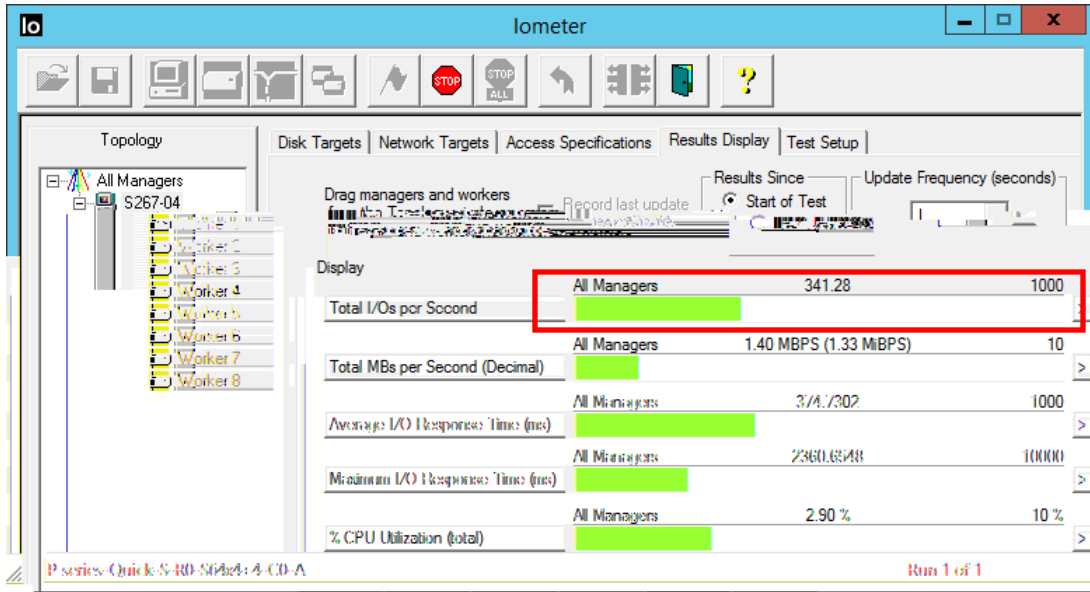
VM Name	Throughput Before Relocation	Throughput After Relocation	Improved
VM1	9.96 MB/s	465.86 MB/s	4,577%
VM2	4.78 MB/s	74.75 MB/s	1,464%
VM3	4.41 MB/s	68.78 MB/s	1,460%
VM4	4.13 MB/s	63.59 MB/s	1,440%
VM5	3.98 MB/s	60.03 MB/s	1,408%
VM6	3.79 MB/s	57.12 MB/s	1,407%
VM7	3.70 MB/s	54.90 MB/s	1,384%
VM8	3.61 MB/s	54.18 MB/s	1,401%

### Case 3: Sudden Reaction

In order to cope with an expected sudden event, administrators can move the required data to the SSD tier in advance. In general, w







- Assume that the data in this volume will be used frequently by your users. You will manually change the tiering policy to **Highest Available Tier**.



- After an hour the system analyzes the data automatically and it will be relocated at midnight. You can also manually execute relocation via the function **Relocation Now**. You can also set the relocation rate **Medium** or **Slow** to eliminate the possible performance impact. The figure shows that 5GB data in NL-SAS tier will be moved up to the SD tier.

Pools **Auto Tiering**

	Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	Teat	Online	Good	2.09 TB	2.04 TB	2.04 TB	1	12	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	0	0	0	52	
SAS	557	0	0	0	0	
NL-SAS	1395	52	52	0	0	

**Relocate Now**

Pool Name : Teat

Relocation Period (hh:mm) : 00 : 00 (Set as 00:00 to let relocation process run until it finishes.)

Relocation Rate : **Medium**

- Fast
- Medium**
- Slow

OK Cancel

6. The relocation completes. The data has been moved to the SSD tier.

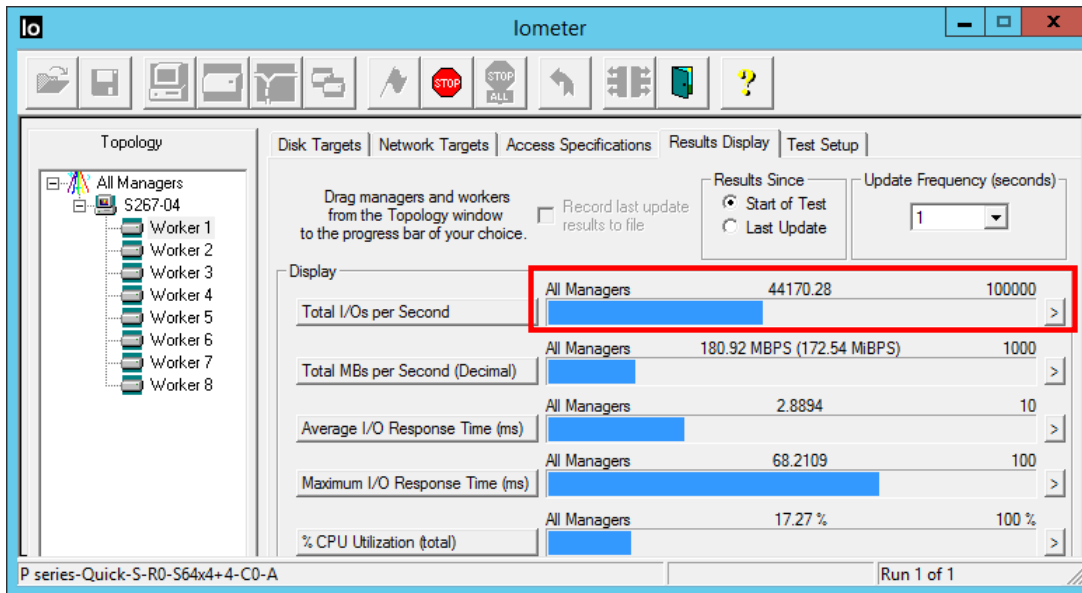
Pools **Auto Tiering**

	Pool Name	Status	Health	Total	Free	Available	Volumes	Disks	Current Controller
▼	Teat	Online	Good	2.09 TB	2.04 TB	2.04 TB	1	12	Controller 1

**Pool Tiering Status:**

Tier Level	Tier Capacity (GB)	Tier Used (GB)	Move Up (GB)	Move Down (GB)	Move In (GB)	Tier Status
SSD	185	52	0	0	0	
SAS	557	0	0	0	0	
NL-SAS	1395	0	0	0	0	

7. The IOPS of this volume increases to 44170.28.



## Summary

In case 3, IT administrator can manually control the data into the higher or lower tier in advance. The scenario meets the expectations of expected sudden event

## Auto Tiering Notices

There are some notices about auto tiering.

- In our design, the snapshot data will be located at the lowest tier in order to obtain economic benefit, and retain the highest space for performance usage. If an auto tiering pool enables snapshots, the performance may be limited to the HDDs at the lowest tier.
- If using SATA SSDs in dual controller systems, the performance of each SSD is limited to 270MB/s per SSD due to the MUX board
- In the [SSD Cache vs. Auto Tiering](#) section, we know that the effectiveness of SSD cache can be seen in a short term, and auto tiering is effective in a long term. Both functions can be used at the same time and achieve complementary effects. Be notice that the quantity and the capacity of SSDs which SSD cache and auto tiering include administrator should adjust via the performance monitor at any time to get better.

## Conclusion

With auto tiering technology, the XCubeSAN series can help you put the right data in the right place at the right time for optimal use of all storage tiers and allow you to reduce storage costs and management overhead while increasing performance and capacity.

Intelligent algorithm behind auto tiering manages the data relocation and monitors the data hotness ratio using half-life coefficient and advanced ranking mathematics. Relocations can occur on the user-defined relocation schedule, making auto tiering a truly automated offering.

## Apply To

- XCubeSAN XS5200 / XS3200 / XS1200 12.0 and later

## Reference

### SSD Cache 2.0 White Paper

- [SSD Cache 2.0 White Paper](#)

## Appendix

---

### Related Documents

There are related documents which can be downloaded from the website.

- [All XCubeSAN Documents](#)
- [XCubeSAN QIG \(Quick Installation Guide\)](#)
- [XCubeSAN Hardware Owner's Manual](#)
- [XCubeSAN Configuration Worksheet](#)
- [XCubeSAN SANOS 4.0 User's Manual](#)
- [Compatibility Matrix](#)
- [White Papers](#)
- [Application Notes](#)

### Technical Support

Do you have any questions or need help troubleshooting a problem? Please contact QSAN Support we will reply to you as soon as possible.

- Via the Web: <https://qsan.com/support>
- Via Telephone: +86-77202118 extension 136  
(Service hours: 09:30-18:00, Monday-Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support  
(Service hours: 09:30-2:00, Monday-Friday, UTC+8, Summer time: 09:30-1:00)
- Via Email: [support@qsan.com](mailto:support@qsan.com)