



XCubeSAN Series White Paper

Auto Tiering2.0



QSAN Technology, Inc.
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 paper is 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 24 -disk 4U Chassis |
| XS3224D | Dual Controller | LFF 24 -disk 4U Chassis |
| XS3224S | Single Controller | LFF 24 -disk 4U Chassis |
| XS1 224D | Dual Controller | LFF 24 -disk 4U Chassis |
| XS1 224S | Single Controller | LFF 24 -disk 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 16 -disk 3U Chassis |
| XS3216D | Dual Controller | LFF 16 -disk 3U Chassis |
| XS3216S | Single Controller | LFF 16 -disk 3U Chassis |
| XS1 216D | Dual Controller | LFF 16 -disk 3U Chassis |
| XS1 216S | Single Controller | LFF 16 -disk 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 12 -disk 2U Chassis |
| XS5212S | Single Controller | LFF 12 -disk 2U Chassis |
| XS3212D | Dual Controller | LFF 12 -disk 2U Chassis |
| XS3212S | Single Controller | LFF 12 -disk 2U Chassis |
| XS1 212D | Dual Controller | LFF 12 -disk 2U Chassis |
| XS1 212S | Single Controller | LFF 12 -disk 2U Chassis |
| XS5226D | Dual Controller | SFF 26 -disk 2U Chassis |
| XS5226S | Single Controller | SFF 26 -disk 2U Chassis |
| XS3226D | Dual Controller | SFF 26 -disk 2U Chassis |
| XS3226S | Single Controller | SFF 26 -disk 2U Chassis |
| XS1 226D | Dual Controller | SFF 26 -disk 2U Chassis |

| | | |
|----------|-------------------|-------------------------|
| XS1 226S | Single Controller | SFF 26 -disk 2U Chassis |
|----------|-------------------|-------------------------|

Information contained in this document has been reviewed for accuracy but 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. All features, functionality and product specifications 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 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.

Table of Contents

| | |
|--|----|
| Notices | i |
| Auto Tiering 2.0 | 1 |
| Executive Summary..... | 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 |
| Appendix..... | 64 |



Related Documents..... 64

Technical Support..... 64

Auto Tiering 2.0

Executive Summary

QSAN autotiering costeffectively 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 storage should be assigned based on this information. Then at the scheduled time, the most accessed blocks that have been marked **as “hot” data will be migrated into the highest performing tier, the least accessed or “cold”** data will be migrated into the lowest-cost-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.

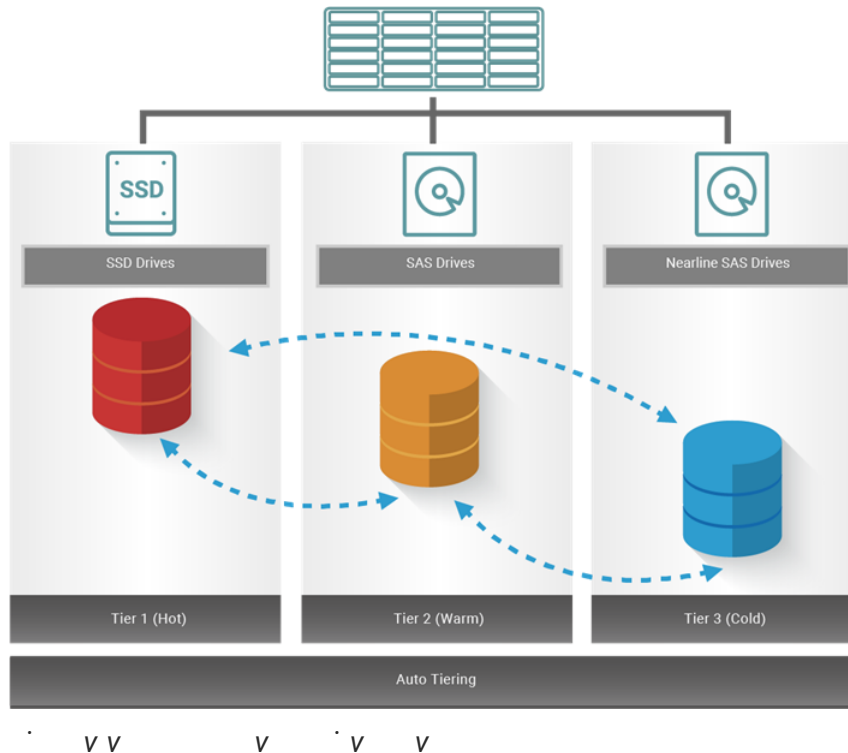
Overview

From the perspective of storage features, 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. In fact, 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 all-hard-drive configuration is also gradually been inadequate.

On the other hand, the data itself has a lifecycle. Since the data in the course 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 increases, it is accessed less often.

The Solution

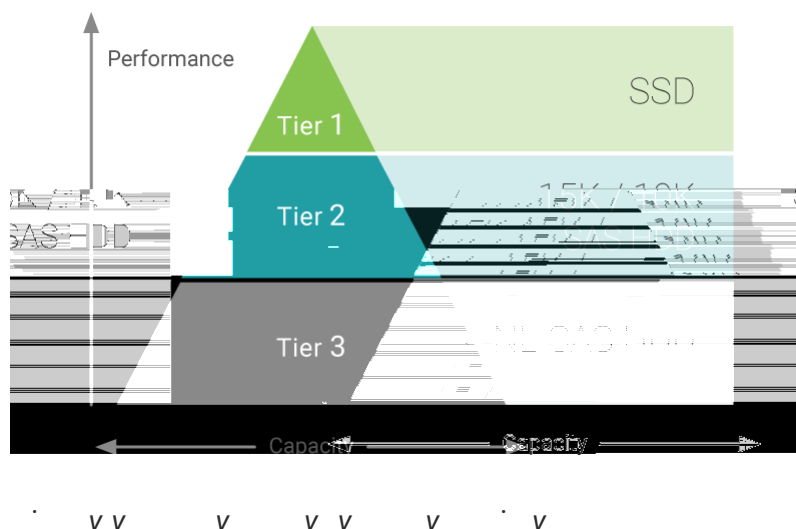
Therefore, to balance performance and cost factors, adapting hybrid storage architecture with a mixture of SSD and traditional HDD seems to be the most reasonable approach for modern IT environments. Generally, SSD-based storage capacity is 10 to 15% of the total storage capacity should be enough to fulfill the requirements of critical high I/O applications. An automated tiering pool 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. In a segmented 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 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 performance resources because it automatically relocates "hot" data to the SSD tier.

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 of performance with consistent response time, high throughput and good bandwidth at moderate price.

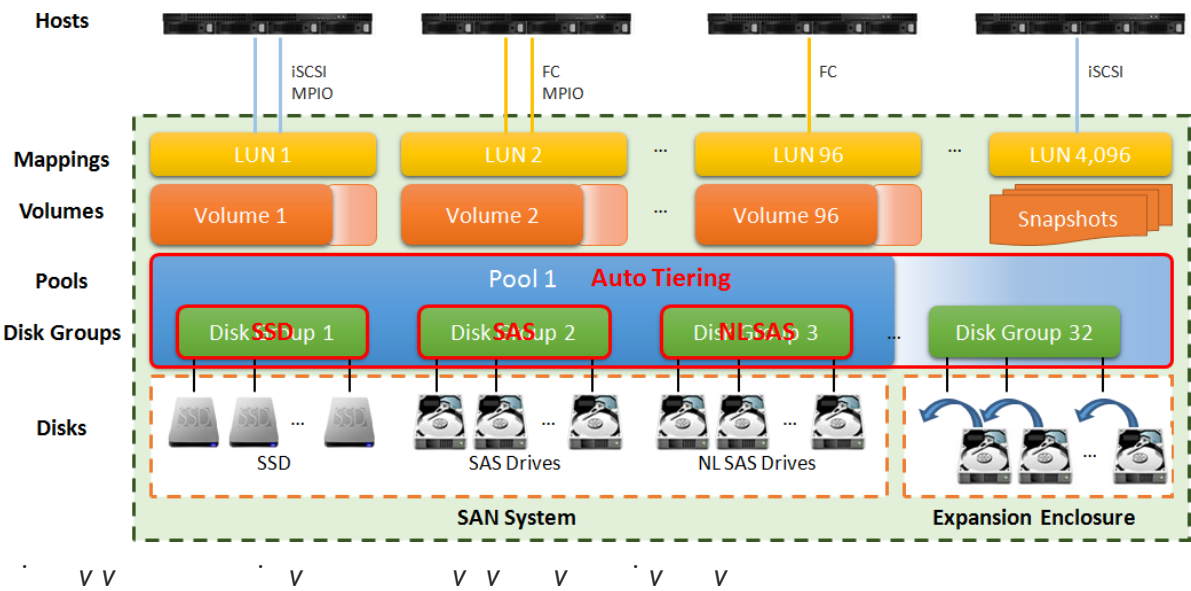
Tier 3 / NL-SAS HDD Tier / Capacity Tier

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

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 256TB. 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 disk types. 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 relocating 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 I/O access is collected and analyzed per hour. This is also called subLUN. Activity level of a subLUN is determined by counting the quantity of read and write access on the subLUN. 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 subLUN within the pool. Note that the policies of volumes would affect how subLUNs 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 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 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

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 **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.

Auto Tiering License

Download Request License file and send to your local sales to get a License Key.

Select the license file to enable Auto Tiering:

Choose File

No file chosen

Apply

Request License

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 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 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.

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)

| | 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

RAID Configuration

Disk Selection

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

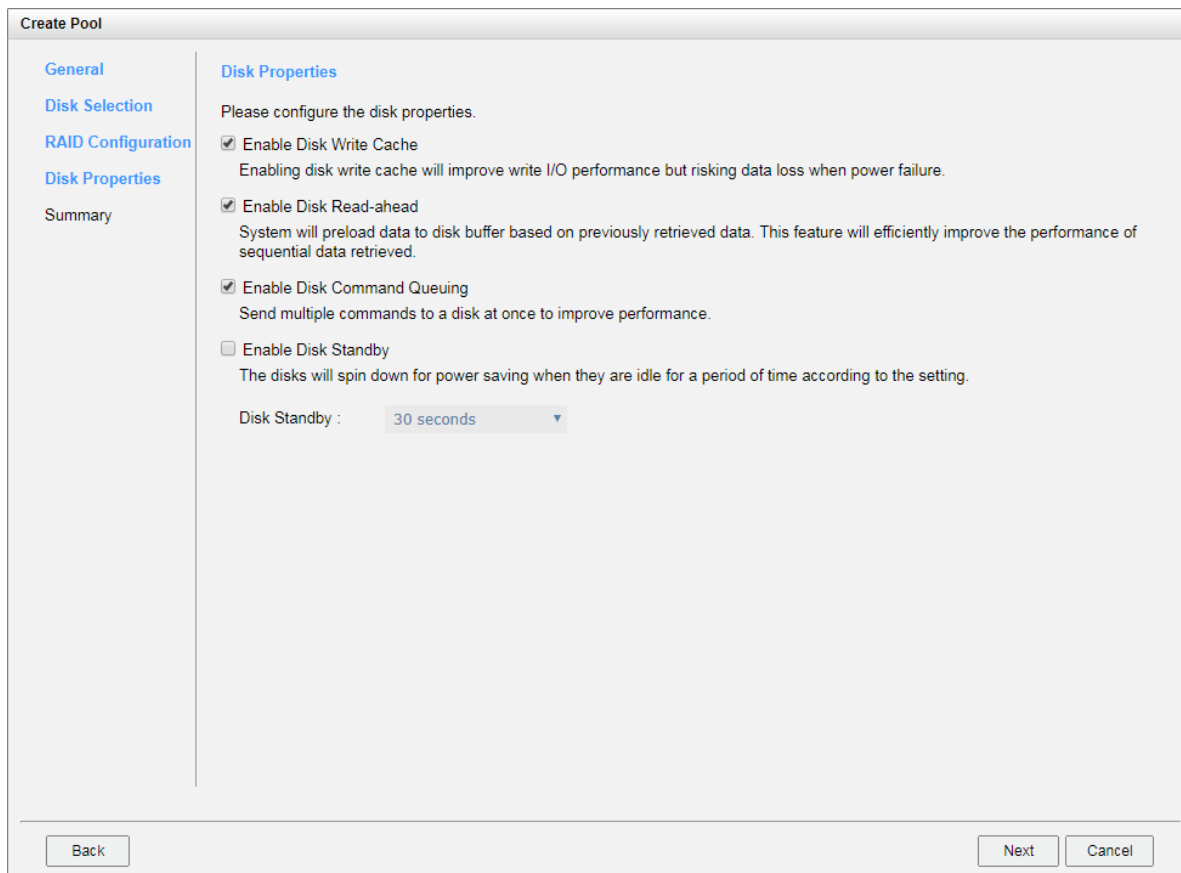
Quantity of NL-SAS Disks :

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



· V V V V V V V

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

· V V V V V V V V

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 a 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

· V V V V V V V

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 drive. 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

· V V V V · V V

This table shows the column descriptions.

| Column Name | Description |
|-----------------------------------|---|
| Pool Name | The poolname. |
| Status | <p>The status of the pool</p> <ul style="list-style-type: none"> Online: The pool is online. Offline: The pool is offline. Rebuilding: The pool is being rebuilt. Migrating: The pool is being migrated. Relocating: The pool is being relocated. |
| Health | <p>The health of the pool</p> <ul style="list-style-type: none"> Good: The pool is good. Failed: The pool is failed. Degraded: The pool is not healthy and not complete. The reason could be missing or failed disks. |
| Total | Total capacity of the pool. |
| Free | Free capacity of the pool. |
| Available | Available capacity of the pool. |
| Thin Provisioning | <p>The status of Thin provisioning:</p> <ul style="list-style-type: none"> Disabled. Enabled. |
| Auto Tiering | <p>The status of Auto Tiering:</p> <ul style="list-style-type: none"> Disabled. Enabled. Not Supported: The pool contains disk groups with mixed disk type. |
| Volumes | The quantity of volumes in the pool. |
| Current Controller (v v v v) | 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"> • 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. • Relocating The disk group is being relocated. |
| Health | The health of the disk group <ul style="list-style-type: none"> • 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 | 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"> • 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. • Check Done The disk drive has been checked the disk health. |
| Health | The health of the disk drive <ul style="list-style-type: none"> • Good: The disk drive is good. • Failed: The disk drive is failed. • Error Alert S.M.A.R.T. error alert. • Read Errors The disk drive has unrecoverable read errors. |
| Capacity | The capacity of the disk drive |

| | |
|--------------|---|
| Disk Type | The type of the disk drive <ul style="list-style-type: none"> [SAS HDD NL-SAS HDD SAS SSD SATA SSD] [12.0Gb/s 6.0Gb/s 3.0Gb/s 1.5Gb/s] |
| 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 net 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"> Light Blue Used capacity. |

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

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 for 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

· V V I V V

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 :

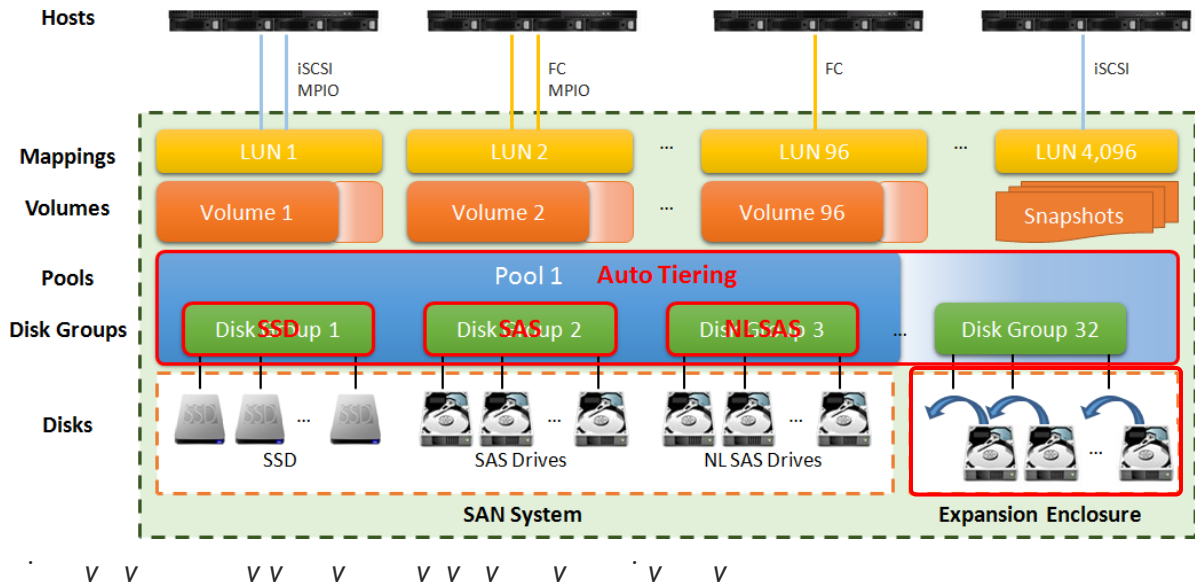
Fast ▼

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 and click ▼ -> Add Disk Group to add a disk group in the auto tiering pool

Add Disk Group

Pool Type

Revolving

Enabled

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)

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

OK

Cancel

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 spare 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 |
| ▼ | 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 | | MICRON | S630DC-800 |
| ▼ | 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 | ST1200MM0088 |
| ▼ | 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 | Pool-3 | SEAGATE | ST6000NM0014 |
| ▼ | 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 |

· V V V V V · V V

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 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 **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.

Create Volume

General
Advanced
Summary

Volume Advanced Settings

Please configure the volume advanced settings.

Block Size : 512 Byte

Priority : High

The priority is the comparison with the other volumes.

Background I/O Priority : High

Background I/O priority will influence volume initialization, rebuild, and migration.

Tiering Policy : Auto Tiering

☒ Enable Cache Mode (Write-back Cache)
Write back optimizes the system speed but comes with the risk which the data may be inconsistent between the cache and disks in a short time interval.

☐ Enable Video Editing Mode
Please enable it when the application is in the video editing environment. It sacrifices a bit of performance but is stable.

☒ Enable Read-ahead
The system will identify what is needed next, buffer. When the data to be transmitted is contiguous, this feature will improve performance.

☒ Enable Space Reclamation

Next Cancel Back

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 [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

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's 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

Video Editing Mode:

Disabled

Read-ahead:

Enabled

Space Reclamation:

Enabled

Volume Type:

RAID Volume

OK

Cancel

Reclaim Space with Thin Provisioning Pool

Click ▼ -> Space Reclamation to reclaim space from the volume when the volume is in an auto tiering pool. For more information about space reclamation, please refer to the chapter 9.2.2, Space Reclamation section in the [XCubeSAN SANOS 4.0 User's Manual](#).

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 the chapter 8.6, Configure LUN Mappings section in the [XCubeSAN SANOS 4.0 User's Manual](#) for detail. For more information about host initiator, please refer to the chapter 8.7, Connect by Host Initiator section in the [XCubeSAN SANOS 4.0 User's Manual](#) for detail.

Transfer to Auto Tiering Pool

This section describes thick provisioning pool or thin provisioning pool transfer to auto tiering. If an auto tiering license is enabled, the thick or thin provisioning pool without disk group of mixed disk type can be transferred to the auto tiering **Add Disk Group** option.

Also note that the thin provisioning pools preconfigured the space, after transferring to the auto tiering, the original disk group in the thick provisioning pool will be the lowest tier. When auto tiering mechanism is running, the hot data are copied to higher tier, but still occupy the space of the original block. If the data is cold, return to the original block space. So the total capacity of the pool does not change even adding the capacity of higher tiers.

· v v v v v v · v v · v v v · v

- 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

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 disks in a disk group is 64.

Enclosure ID :

0 (Head Unit: XS5216)

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

- 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

· V V V V · V V V V · V V V

- 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 thin provisioning pool to auto tiering one.

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

| | Pool Name | Status | Health | Total | Free | Available | Thin Provisioning | Auto Tiering | Volumes | Current Controller | |
|--------------|--------------|--------|--------|---------|----------|-----------|-------------------|--------------|----------|--------------------|--------|
| <div>▼</div> | Pool-1 | Online | Good | 2.18 TB | 2.18 TB | 2.18 TB | Disabled | Enabled | 0 | Controller 1 | |
| ler 1 | <div>▼</div> | Pool-2 | Online | Good | 10.92 TB | 10.92 TB | 10.92 TB | Enabled | Disabled | 0 | Contro |

Disk Groups

| | No. | Status | Health | Total | Free | Disks Used | RAID |
|--------------|-----|--------|--------|----------|----------|------------|--------|
| <div>▼</div> | 1 | Online | Good | 10.92 TB | 10.92 TB | 3 | RAID 5 |

Disks

| Enclosure ID | Slot | Status | Health | Capacity | Disk Type | Manufacturer | Model |
|--------------|------|--------|--------|----------|---------------------|--------------|--------------|
| 0 | 13 | Online | Good | 5.46 TB | NL-SAS HDD 12.0Gb/s | SEAGATE | ST6000NM0014 |
| 0 | 14 | Online | Good | 5.46 TB | NL-SAS HDD 12.0Gb/s | SEAGATE | ST6000NM0014 |
| 0 | 15 | Online | Good | 5.46 TB | NL-SAS HDD 12.0Gb/s | SEAGATE | ST6000NM0014 |

Create Pool

- 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 |
| ▼ | Pool-2 | Online | Good | 11.64 TB | 11.64 TB | 11.64 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 | 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 |

Create Pool

· V V V V · V V V V · V V V

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

SSD cache is used for highly frequent data access environments and is effective short term such as virtualization or video editing applications. However, auto tiering is used for predictable I/O workloads and is effective in long term. It's suitable for web, file, or email server applications.

| | SSD Cache | Auto Tiering |
|---------------------|-------------------------|------------------------|
| Total Capacity | HDD | HDD + SSD |
| When SSD is Damaged | Pool Works Fine | Pool Fails |
| Performance | Effective in Short Term | Effective in Long Term |

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 or RAID 6 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 data or high I/O your application uses everyday. Assuming 666GB per day, the recommended SSD tier capacity is at least 1.5 times $5 \times 666\text{GB} = 1\text{TB}$ 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 spaces 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

- 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.



· V V V V W · V

- 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.

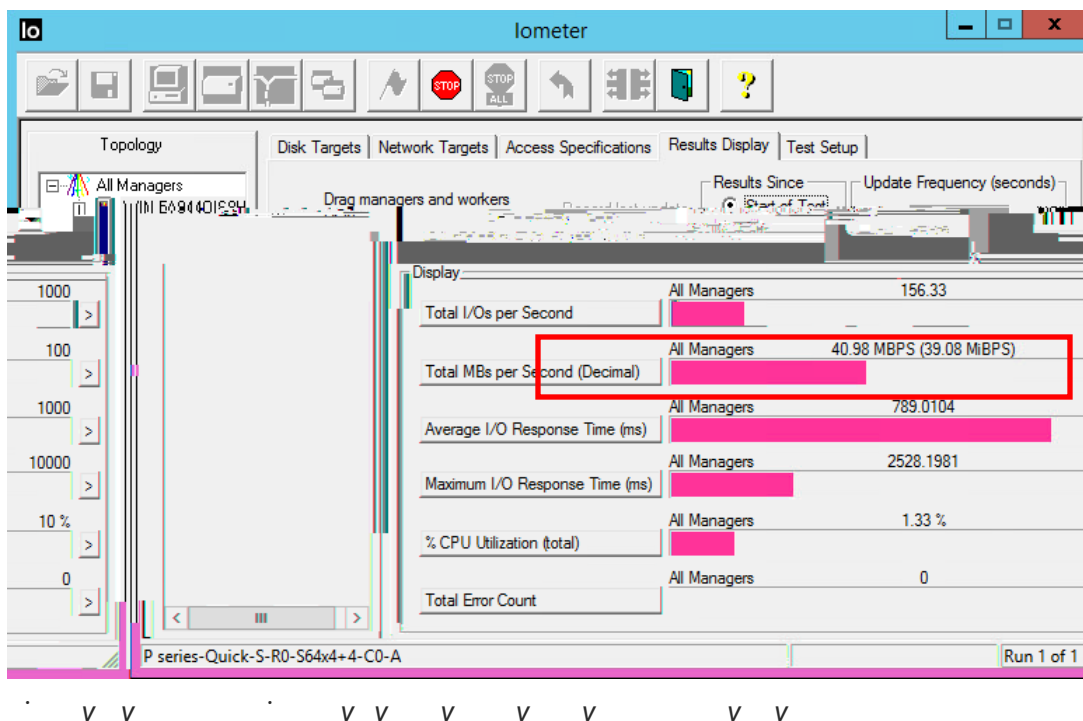


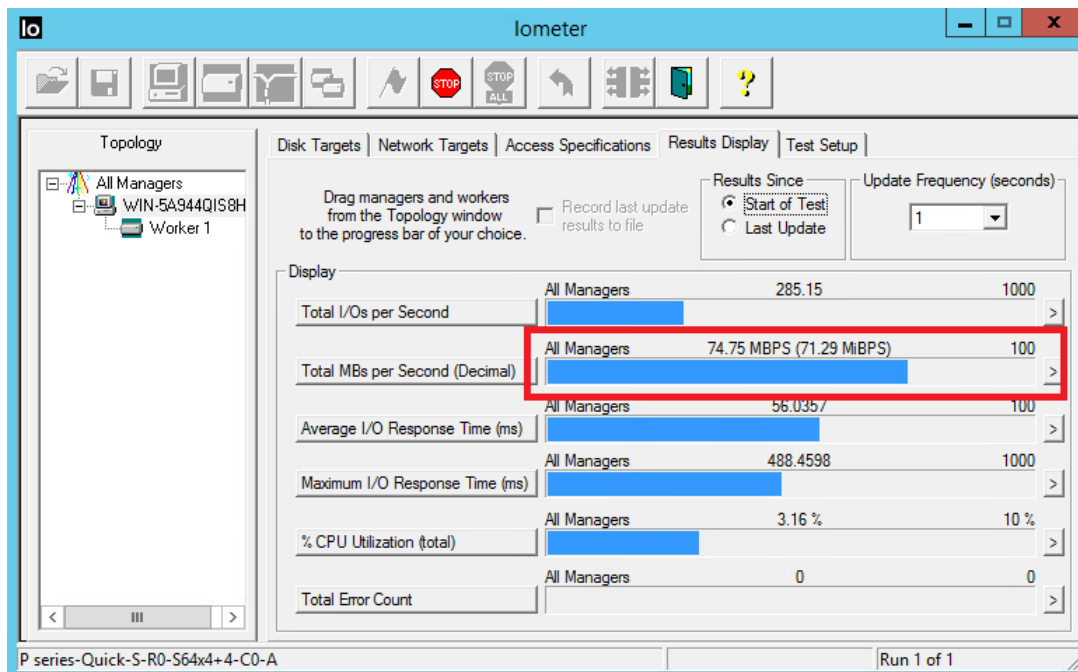
· V V W V

- 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 the 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 which tier the data is located.

- 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.





· V V · V V V

- VM3: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 68.78 MB/s
- VM4: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 63.59 MB/s
- VM5: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 60.03 MB/s
- VM6: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 57.12 MB/s
- VM7: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 10GB the throughput is 54.90 MB/s
- VM8: 256KB, 100% Write, 100% Random, Outstanding Maximum Disk Size 20GB the throughput is 54.18 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 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 and 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, an administrator can move the required data to the SSD tier in advance. In general, we recommend setting the auto tiering policy to **Lowest Tier**. The day before the activity, IT administrators manually set the volume containing the required data to **Highest Available Tier** and then perform **Relocation Now** manually to force relocating data.

Test Equipments and Configurations

- Server
 - Model: ASUS RS700E6/ERS4 (CPU: Intel Xeon E5620 2GHz / RAM: 24GB)
 - 10GbE HBA: Intel Ethernet CNA X710-DA4 FH
 - OS: Windows Server 2012 R2
- Storage
 - Model: QSAN XCubeSAN XS5216
 - Memory: 16GB (2 x 8GB in bank & 3 per controller)
 - Firmware 1.2.1
 - SAS SSD: 4 x HGST Ultrastar SSD800MH.B, HUSMH8010BSS200, 100GB, SAS 12Gb/s
 - SAS HDD: 4 x HGST Ultrastar C15K60Q HUC156030CS42Q000GB, SAS 12Gb/s
 - NL-SAS HDD: 4 x Seagate Constellation ES ST500NM000, 1500GB, SAS 6Gb/s
 - Auto Tiering Pool: 2.09TB
 - SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
 - SAS Tier: RAID 6 with 4 x SAS HDD, 558GB

- NL-SAS Tier: RAID 5 with 4 x SAS SSD, 1.36TB
 - Volume: 1 x 2.09TB in Auto Tiering Pool
 - Auto Tiering Policy: Lowest Tier then Highest Available Tier
- I/O Pattern
 - Tool: IOMeter V1.1.0
 - Workers: 1
 - Outstanding (Queue Depth): 128
 - Maximum Disk Size: 50GB
 - Access Specifications: 4KB, 100% Write, 100% Random

Test Scenario and Result

1. Create an auto tiering pool with the following configurations.

- Auto Tiering Pool: 2.09TB
- SSD Tier: RAID 10 with 4 x SAS SSD, 185GB
- SAS Tier: RAID 6 with 4 x SAS HDD, 558GB
- NL-SAS Tier: RAID 5 with 4 x SAS SSD, 1.36TB

| Pools | | Auto Tiering | | | | | | | |
|-------|-----------|--------------|--------|---------|---------|-----------|-------------------|--------------|--|
| | Pool Name | Status | Health | Total | Free | Available | Thin Provisioning | Auto Tiering | |
| ▼ | Test | Online | Good | 2.09 TB | 2.09 TB | 2.09 TB | Enabled | Enabled | |

| Disk Groups | | | | | | | | |
|-------------|-----|--------|--------|-----------|-----------|------------|------------|---------|
| | No. | Status | Health | Total | Free | Tier Level | Disks Used | RAID |
| ▼ | 1 | Online | Good | 1.36 TB | 1.36 TB | NL-SAS | 4 | RAID 5 |
| ▼ | 2 | Online | Good | 558.00 GB | 558.00 GB | SAS | 4 | RAID 6 |
| ▼ | 3 | Online | Good | 185.00 GB | 185.00 GB | SSD | 4 | RAID 10 |

Create Pool

2. Create a volume of the capacity 2.09TB, and the tiering policy sets as Lowest Tier.

Create Volume

General
Advanced
Summary

Volume Advanced Settings

Please configure the volume advanced settings.

Block Size : 512 Byte

Priority : High

The priority is the comparison with the other volumes.

Background I/O Priority : High

Background I/O priority will influence volume initialization, rebuild, and migration.

Tiering Policy : **Lowest Tier**

☒ Enable Cache Mode (V)
 Write back optimizes the disk write performance, but it has a risk which the data may be inconsistent between the cache and disks in a short time interval.

☐ Enable Video Editing Mode
 Please enable it when the application is in the video editing environment. It sacrifices a bit of performance but is stable.

☒ Enable Read-ahead
 The system will identify what is needed next, based on the content just retrieved from the disk, and then preload the data into the disk's buffer. When the data to be transmitted is continuous, this feature will improve performance.

☒ Enable Space Reclamation

Back Next Cancel

3. Run IOMeter to observe the performance. IOMeter parameters are on the following.

- Tool: IOMeter V1.1.0
- Workers: 1
- Outstanding (Queue Depth): 128
- Maximum Disk Size: 50GB
- Access Specifications: 4KB, 100% Write, 100% Random

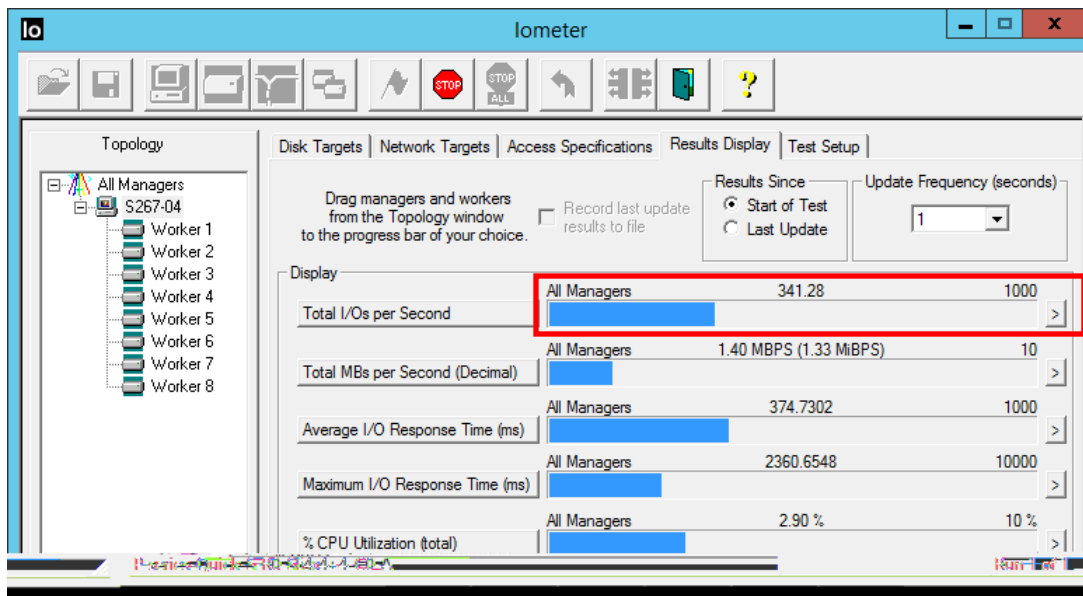
Because the tiering policy sets **Lowest Tier**, the I/O file is located in the NL-SAS tier, and the IOPS is 3428.

Pools **Auto Tiering**

| | Pool Name | Status | Health | Total | Free | Available | Volumes | Disks | Current Controller |
|---|-----------|--------|--------|---------|---------|-----------|---------|-------|--------------------|
| ▼ | Test | 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 | 0 | |
| SAS | 557 | 0 | 0 | 0 | 0 | |
| NL-SAS | 1395 | 52 | 0 | 0 | 0 | |



- Assume that the data in this volume will be used frequently and 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 or 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

· V V M V V V

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

Pools **Auto Tiering**

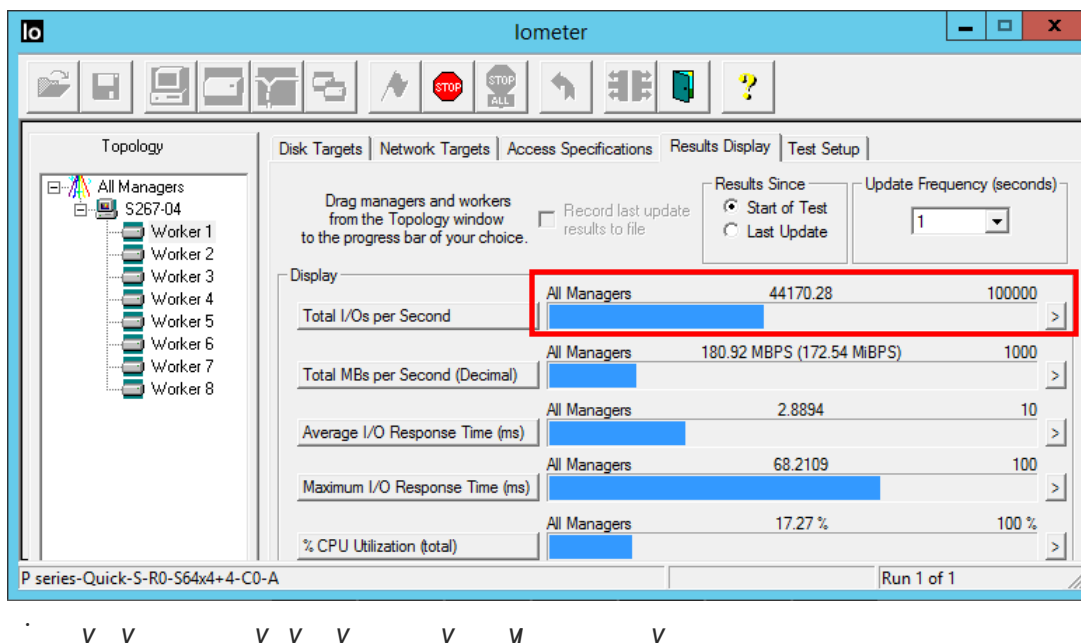
| | Pool Name | Status | Health | Total | Free | Available | Volumes | Disks | Current Controller |
|---|-----------|--------|--------|---------|---------|-----------|---------|-------|--------------------|
| ▼ | Teat | Online | Good | 2.09 TB | 2.04 TB | 2.04 TB | 1 | 12 | Controller 1 |

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 | |

· V V M V

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 system, 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 need, and 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.

The 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 \(QuickInstallation Guide\)](#)
- [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: 09:30-20:00, Monday-Friday, UTC+8, Summer time: 09:30-21:00)
- Via Email: support@qsan.com