



# XCubeNAS Series White Paper

## Auto Tiering

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

© 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.



## **Edition 1.0 (January 2018)**

This edition applies to QSAN XCubeNAS NAS storage systems. 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, XCubeNAS, and QSAN.com are trademarks or registered trademarks of QSAN Technology, Inc.

Intel, Xeon, Pentium, Core, Celeron and Intel Inside logo are trademarks of Intel Corporation in the United States and 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

---

The 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 products. 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
Preface .....	3
About This Manual .....	3
Related Documents .....	3
Technical Support .....	4
Information, Tip and Caution .....	4
<b>1. Introduction .....</b>	<b>5</b>
1.1. Exclusive Summary .....	5
1.2. Overview .....	5
1.3. Tier Categories .....	6
1.4. Flexible RAID and Disk Configurations .....	6
1.5. Intelligent Auto Tiering Mechanism .....	7
1.6. Tiering Policies .....	8
<b>2. Configuration and Operation .....</b>	<b>11</b>
2.1. Create an Auto Tiering Pool .....	11
2.2. Operations on Auto Tiering Pool .....	14
2.3. Create Volumes and iSCSI LUNs .....	16
<b>3. Best Practice .....</b>	<b>20</b>
3.1. Overview .....	20
3.2. Configuration Planning Advice .....	20
3.3. Case Study .....	23
<b>4. Conclusion .....</b>	<b>31</b>
4.1. Auto Tiering Notices .....	31
4.2. Apply To .....	31
<b>5. Support and Other Resources .....</b>	<b>32</b>
5.1. Getting Technical Support .....	32
5.2. Documentation Feedback .....	33
<b>Appendix .....</b>	<b>34</b>
End-User License Agreement (EULA) .....	34

# Preface

---

## About This Manual

QSAN Auto-Tiering technology helps the company to get more benefit from their IT investment. It places the most frequent access data (Hot data) on higher speed disks; moreover, the less frequent data is placed on high-capacity drives, which helps the company optimize application performance without sacrificing capacity or increasing the budget. This document is intended for use by system administrators, NAS designers, storage consultants, or anyone who has purchased these products and is familiar with server and computer network, network administration, storage system installation and configuration, network attached storage management, and relevant protocols.

**CAUTION:**

Do NOT attempt to service, change, disassemble or upgrade the equipment's components by yourself. Doing so may violate your warranty and expose you to electric shock. Refer all servicing to authorized service personnel. Please always follow the instructions in this owner's manual.

---

## Related Documents

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

- XCubeNAS QIG (Quick Installation Guide): <https://www.qsan.com/en/download.php>
- Compatibility Matrix: <https://www.qsan.com/en/download.php>
- XCubeNAS QSM 3.0 User Manual: <https://www.qsan.com/en/download.php>
- XCubeNAS Hardware Manual: <https://www.qsan.com/en/download.php>
- Online FAQ: <https://www.qsan.com/en/faq.php>
- White Papers: <https://www.qsan.com/en/download.php>
- Application Notes: <https://www.qsan.com/en/download.php>

## 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: [http://www.qsan.com/en/contact\\_support.php](http://www.qsan.com/en/contact_support.php)
- Via Telephone: +886-2-7720-2118 extension 136  
(Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support  
(Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summertime: 09:30 - 01:00)
- Via Email: [support@qsan.com](mailto:support@qsan.com)

## Information, Tip and Caution

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



**INFORMATION:**

INFORMATION provides useful knowledge, definition, or terminology for reference.



**TIP:**

TIP provides helpful suggestions for performing tasks more effectively.



**CAUTION:**

CAUTION indicates that failure to take a specified action could result in damage to the system.

---

# 1. Introduction

---

## 1.1. Exclusive Summary

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

All of this is managed in the background without user intervention. This tiered pool will also function the same as any standard QSAN pool, and access to our enterprise features such as 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 the XCubeNAS series. This document discusses the Auto tiering technology and describes its features, functions, management, and best practice.

## 1.2. Overview

From the perspective of storage features, the performance of SSD is high, but the cost is also high per GB. Relatively speaking, the cost of traditional hard drive is low, so as performance is relatively poor. If we follow the 80/20 rule to configure storage systems, all-SSD configuration is unreasonable 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 being 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 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 pool is a simple and elegant solution for dynamically matching storage requirements with changes in the frequency of data access.

## 1.3. Tier Categories

As the name suggestion, auto tiering must have two tiers at least. Automated tiering pool segregated disk drives into four categories.

- Tier 1: PCIe SSD drives for extreme performance tier
- Tier 2: SATA SSD drives for high performance tier
- Tier 3 & 4: SAS & SATA drives for capacity tier



#### INFORMATION:

SAS drives are supported when connecting with XCubeDAS JBOD with add-on card installed on XCubeNAS, SAS drives can then be installed in the JBOD enclosure.

---

## 1.4. Flexible RAID and Disk Configurations

Auto Tiering supports flexible RAID and disk configurations. You can create each tier with different RAID level and different quantity of disk. For example, PCIe SSD tier uses 2 disks with RAID 1 for extreme performance, SATA SSD tier uses 4 disks with RAID 5, and SAS or SATA HDD tier uses 12 disks with RAID 6 for capacity. This feature is very important for IT administrator to arrange storage plan flexibly.



## 1.5. Intelligent Auto Tiering Mechanism

Auto tiering storage management system manages the data relocation and monitors the data hotness ratio using half-life 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. XCubeNAS system maintains a cumulative I/O count and weights each I/O by how recently it arrived. The new coming I/O is given a full weight. After approximately 24 hours, the weight of this IO is nearly cut in half and continues to decrease. The reduction weight is processing per hour by our precision 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

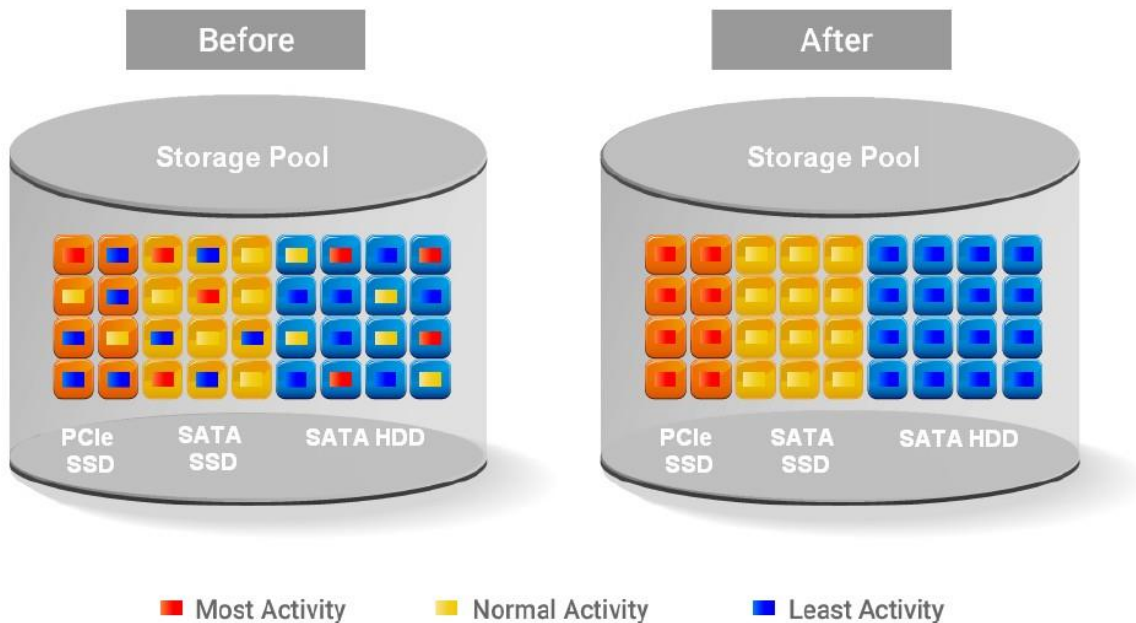
### Relocation

According to the hotness analysis, relocation is processed during the user-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. Valid value of relocation rate is High, Normal, and Low.

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 the 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. Lower tiers are used for capacity when needed. The entire relocation process is completed automatically based on the user-defined relocation schedule, or manually if user triggers by himself. The following figure provides an illustration of how auto tiering can improve sub LUN placement in a pool.



## 1.6. Tiering Policies

For a best performance in various environments, auto tiering has 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 per volume to create a candidate list for data movement. The following volume policies are available:

**Auto-tier (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 data. Sub LUNs are relocated based on the highest performance disk drives available and its hotness. Although this setting relocates data based on the performance statistics of the volume, the volume sets with “Always 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.

**Begin high after auto-tier**

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

**Always highest available tier**

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

**Lowest available 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 available tier” policy would always reside in the lowest tier.

Changing policy of a volume with data in higher tiers to “Lowest available tier” would cause all its data in higher tier to be relocated down to the lowest tier.

### **No data movement**

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

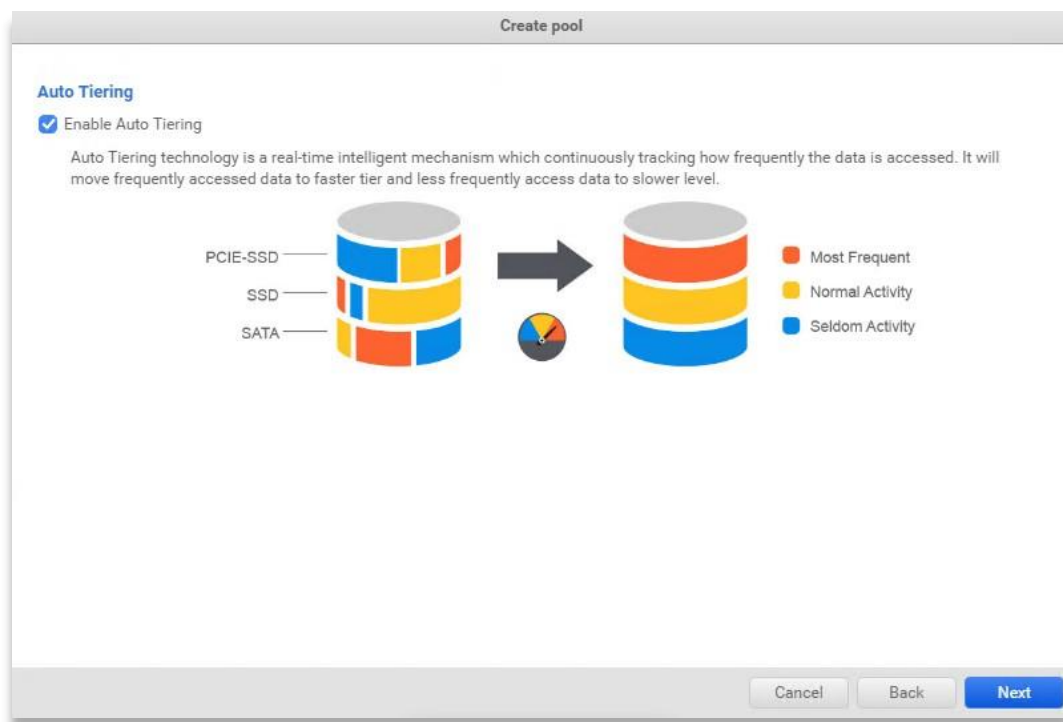
## 2. Configuration and Operation

---

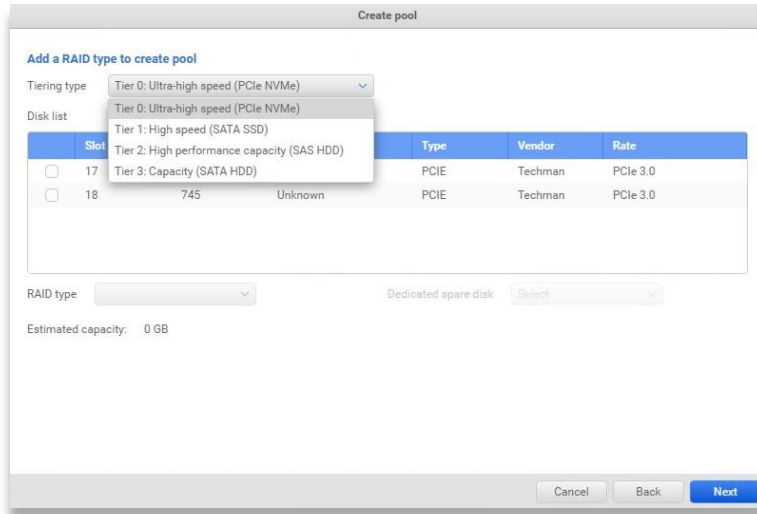
### 2.1. Create an Auto Tiering Pool

We installed PCIe SSD \* 1 to be created as RAID0 (In real case this is recommended to be RAID1 with two PCIe SSD), SATA SSD \* 2 to be created as RAID1, SAS HDD \* 4 to be created as RAID6, and SATA HDD \* 4 to be created as RAID6 in this example.

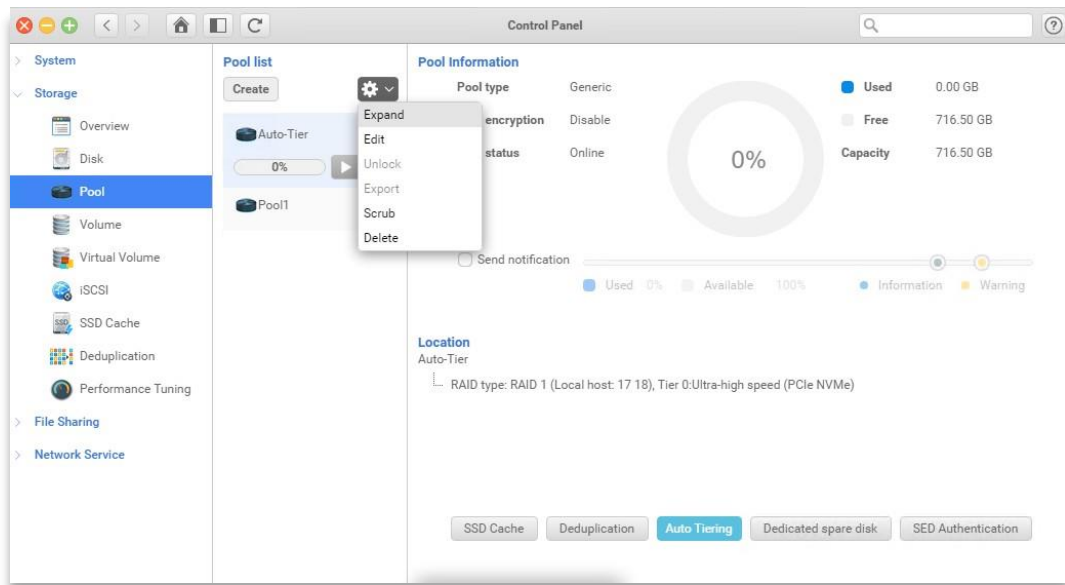
1. Create a pool with Auto Tiering function enabled



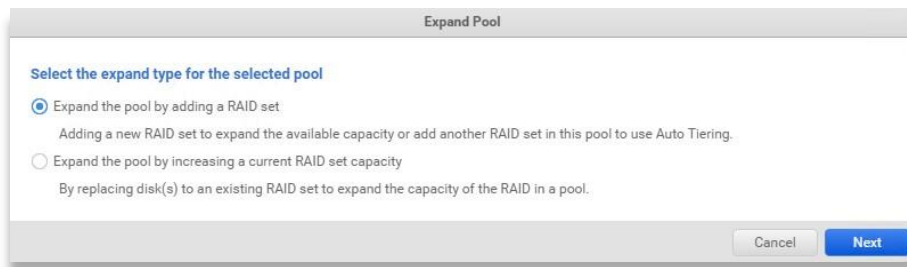
2. Select the preferred tier you would like to create, and select the corresponding disks, then specify the RAID level you can use according to the disk quantity you selected



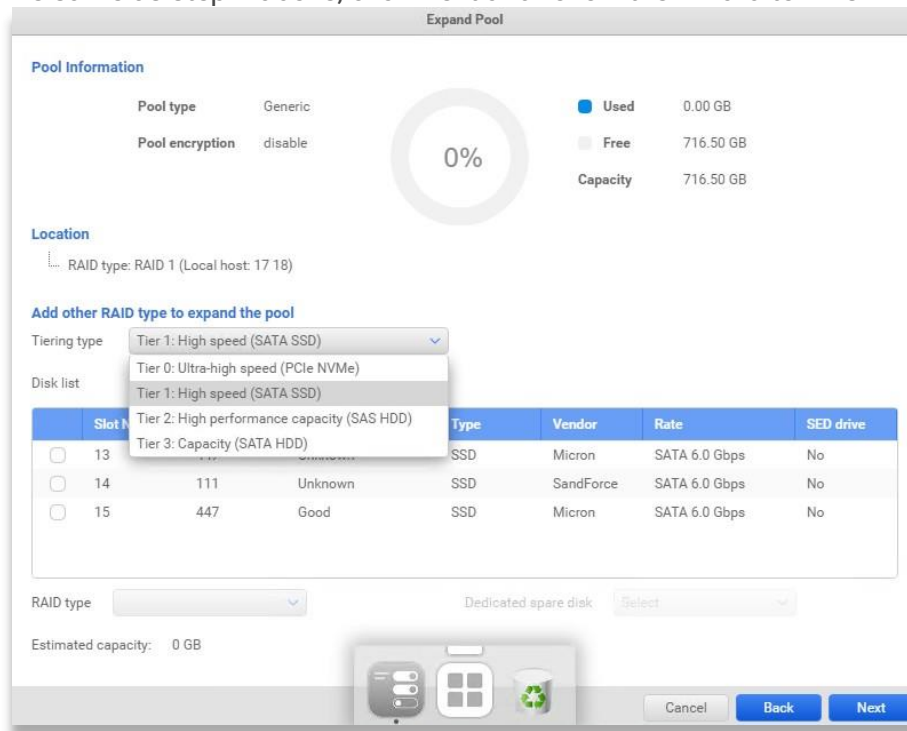
- After the pool is created, select the pool and click the setting button, choose Expand to add another tier



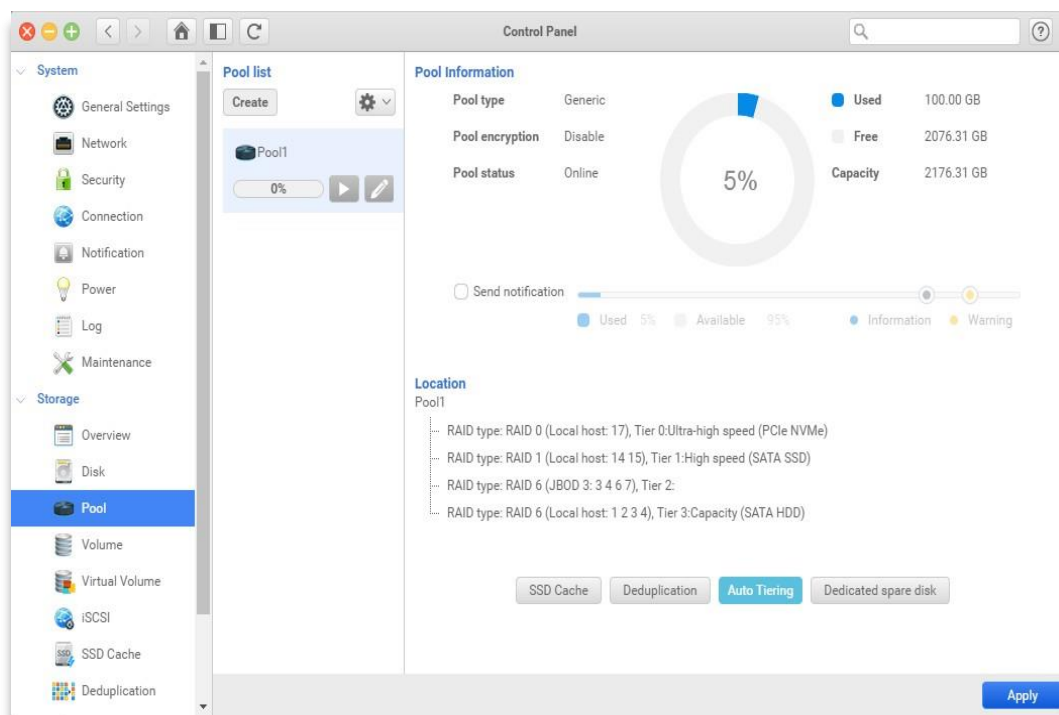
- Select the first option to add another RAID set (tier)



5. The same as step.2 above, click Next and follow the wizard to finish

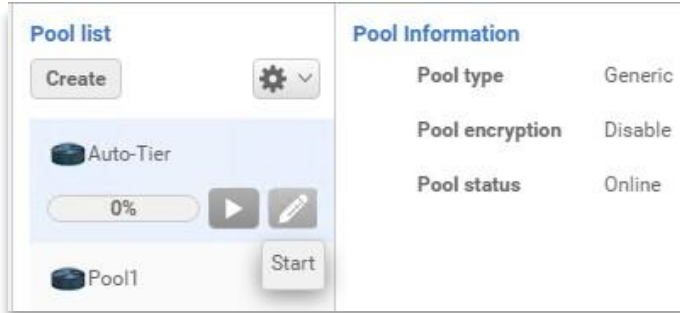


6. You shall be able to see the result like below when you finish all the expanding process for each tier you need



## 2.2. Operations on Auto Tiering Pool

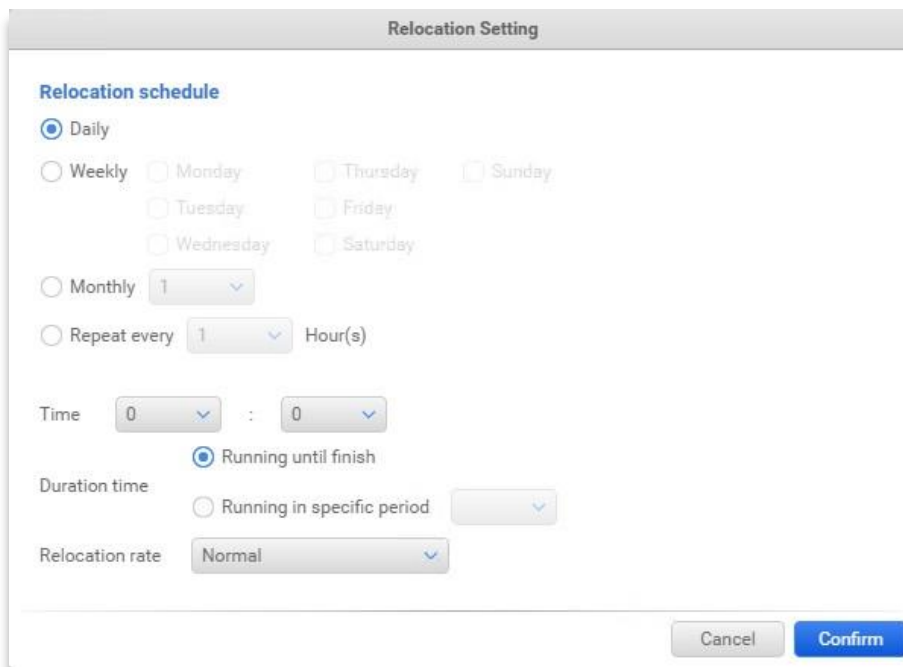
You shall see two buttons shown on the created auto tiering pool, the buttons are used for triggering relocation process manually or automatically by schedule.



### INFORMATION:

Relocation schedule is configured as Daily and execute at midnight by default.

The written data won't be relocated to other tiers after a snapshot is taken on the volume/LUN no matter what tiering policy is configured on it, but data written after the snapshot is taken is still capable of relocating.





For the other operations/options that an auto tiering pool can do are just like the normal pool, and you may click on the Auto Tiering button at the right-hand side to have an overview of the whole pool and the historical operations that have been done on this pool.

### Pool list

Create ⚙️

- Auto-Tier
  - 0% ▶️ ✎
- Pool1

### Pool Information

Pool type	Generic	<span style="color: blue;">■</span> Used	0.00 GB
Pool encryption	Disable	<span style="color: gray;">■</span> Free	904.00 GB
Pool status	Online	<b>Capacity</b>	904.00 GB

0%

Send notification

Used 0%
Available 100%
Information
Warning

### Location

Auto-Tier

- RAID type: RAID 1 (Local host: 17 18), Tier 0:Ultra-high speed (PCIe NVMe)
- RAID type: RAID 5 (Local host: 13 14 15), Tier 1:High speed (SATA SSD)

SSD Cache
Deduplication
Auto Tiering
Dedicated spare disk

### Tiering information

#### Tiering information

[Status](#) [History](#)

Relocation status: 0%

<p>Tier 0: Ultra-high speed (PCIe NVMe)</p> <div style="display: flex; align-items: center; margin-bottom: 10px;"> <div style="margin-left: 10px;"> <p><span style="color: teal;">■</span> Used 0.00GB</p> <p><span style="color: blue;">■</span> Free 728.00GB</p> <p><b>Capacity</b> 728.00GB</p> </div> </div>	<p>Auto-Tier</p> <ul style="list-style-type: none"> <li>RAID 1(Local host:17 18)</li> </ul> <div style="display: flex; justify-content: space-between;"> <span>Move in 0.00GB</span> <span>Move down 0.00GB</span> </div>
<p>Tier 1: High speed (SATA SSD)</p> <div style="display: flex; align-items: center; margin-bottom: 10px;"> <div style="margin-left: 10px;"> <p><span style="color: teal;">■</span> Used 0.00GB</p> <p><span style="color: blue;">■</span> Free 286.00GB</p> <p><b>Capacity</b> 286.00GB</p> </div> </div>	<p>Auto-Tier</p> <ul style="list-style-type: none"> <li>RAID 5(Local host:13 14 15)</li> </ul> <div style="display: flex; justify-content: space-between;"> <span>Move up 0.00GB</span> <span>Move in 0.00GB</span> <span>Move down 0.00GB</span> </div>
<p>Tier 2: High performance capacity (SAS HDD)</p> <div style="display: flex; align-items: center; margin-bottom: 10px;"> <div style="margin-left: 10px;"> <p><span style="color: teal;">■</span> Used 0.00GB</p> <p><span style="color: blue;">■</span> Free 0.00GB</p> <p><b>Capacity</b> 0.00GB</p> </div> </div>	<p>Auto-Tier</p> <div style="display: flex; justify-content: space-between;"> <span>Move up 0.00GB</span> <span>Move in 0.00GB</span> <span>Move down 0.00GB</span> </div>

Tiering information

**Tiering information**

Status [History](#)

Task	Status	Data moved	Total data ...	Start time	End time	Duration
1	Finish	▼ 0.00GB ▲ 0.00GB	0.00GB	2017-05-01 22:29:...	2017-05-01 22:29:...	00:00:00

## 2.3. Create Volumes and iSCSI LUNs

Please visit “Storage\Volume” page on web UI in Control Panel to create a volume with auto-tiering policy specified.

Control Panel

- > System
- > Storage
  - Overview
  - Disk
  - Pool
  - Volume
  - Virtual Volume
  - iSCSI
  - SSD Cache
  - Deduplication
  - Performance Tuning
- > File Sharing

**Volume list**

Create ⚙️

- > Auto-Tier
- > Pool1

**Information**

Volume name: Volume1

Tiering policy: Auto-tier

Compression: Disable

Send notification

---

Used 0% Available 100%

The image shows two parts of the QSAN management interface. The top part is a 'Create Volume' dialog box with the following fields: 'Volume name' (Volume1), 'Location' (Auto-Tier), and 'Capacity' (0 GB). Below these is a progress bar and a status indicator showing 'Pool capacity 904 GB', 'Used 0 MB', and 'Available 904 GB'. The 'Tiering policy' dropdown menu is open, showing options: 'Auto-tier', 'Always highest available tier', 'Begin high after auto-tier', 'Lowest available tier', and 'No data movement'. The 'Next' button is visible.

The bottom part is the 'Volume list' section, which includes a 'Create' button and a settings icon. The list shows three volumes: 'Auto-Tier', 'Volume1', and 'Pool1'. The 'Volume1' entry is selected. To the right, the 'Information' section for 'Volume1' is displayed in a table:

Volume name	Volume1
Tiering policy	Auto-tier
Compression	Disable

Below the table is a 'Send notification' checkbox. To the right of the table is a circular progress indicator showing '0%' and a summary of storage usage: 'Used 0 MB', 'Free 99.99 GB', and 'Capacity 99.99 GB'.

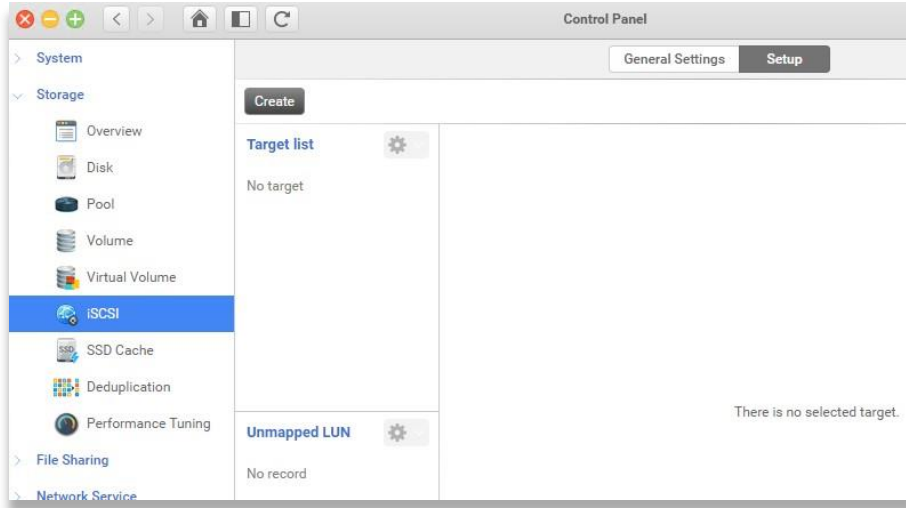


**INFORMATION:**

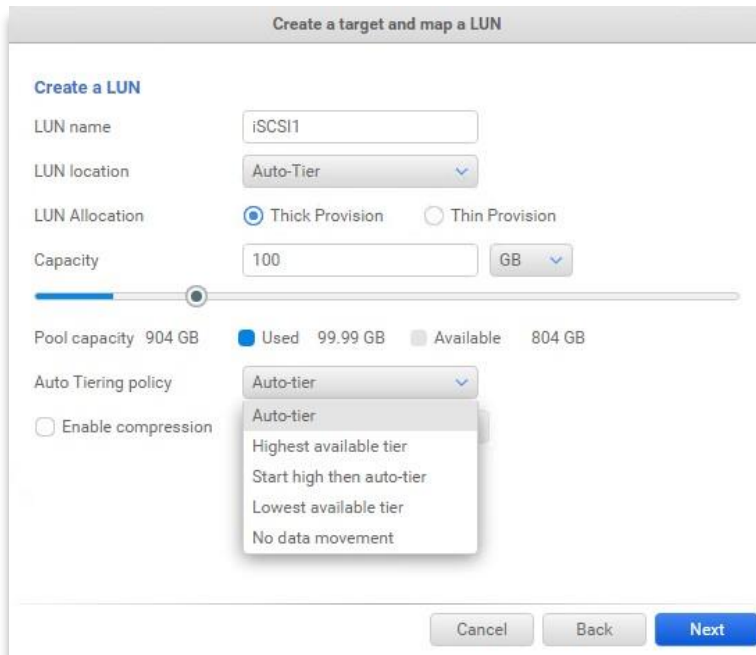
Auto Tiering policy can be modified on the fly, but the collected statistics will be cleared for re-collecting with new policy if the policy is changed afterwards.

The processing Relocating process will be aborted if the tiering policy is changed.

To create an iSCSI LUN, please visit **Storage > iSCSI > Setup** on web UI,



Select the auto-tiering pool while creation, and specify the Auto Tiering policy.



General Settings   Setup

Create

**Target list** ⚙️

- ⌵ iSCSI1
- iSCSI1

**LUN information**

LUN type      Thick Provision

Target name    iSCSI1

Tiering Policy   Auto-tier

Capacity        99.99 GB

**Unmapped LUN** ⚙️

No record

**LUN Masking**

Policy name	Initiator IQN	Rea
Default_1493651720	default_policy_initiatorname	

## 3. Best Practice

---

### 3.1. Overview

Auto tiering technology provides a solution to achieve optimal storage efficiency and improves 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 (PCIe SSD, SATA SSD, SAS HDD, or SATA HDD) 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 lifecycle.

---

### 3.2. Configuration Planning Advice

**PCIe SSD / SATA SSD / HDD 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.

- **PCIe SSD Tier (\$\$\$)**

Suggest SSD tier using at least 2 disks with RAID 1 (better) 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.

- **SATA SSD Tier (\$\$)**

Suggest SATA SSD tier configuring with RAID 6 (better) or RAID 5. Prepare about 30% of the total storage capacity.

- **SAS/SATA HDD Tier (\$)**

For capacity tier, suggest SAS/SATA HDD using RAID 6 (better) or 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 every day. Assuming 666GB per day, the recommended PCIe SSD tier capacity is at least 1.5 times,  $1.5 \times 666\text{GB} = 1\text{TB}$ , as a conservative estimation. And then calculate the SATA SSD tier capacity about 3 times of the PCIe SSD tier capacity,  $3 \times 1\text{TB} = 3\text{TB}$ , if the PCIe SSD tier full of the buffer, the performance does not drop too much. The remaining space left for SAS/SATA HDD tier. The following table is the summary for reference.

Table Tier RAID Level and Capacity Ratio

Tier	Capacity per Drive	Quantity	RAID Level	Capacity per Tier	Capacity Ratio
PCIe SSD Tier	500GB	2	RAID 1	$(2) \times 500\text{GB} = 1\text{TB}$	10%
SATA SSD Tier	1TB	5	RAID 6	$(5-2) \times 1\text{TB} = 3\text{TB}$	30%
SATA/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 user to calculate the performance and necessary capacity. Of course, if more capacity is needed, you can also extend capacity for each tier.

### Relocation and Its Effect

There are three major functions in auto tiering technology, Statistics Collection and Ranking operate automatically, but Relocation can be configured manually. We would like to suggest that users can set the relocation schedule at midnight every day (**Daily 00:00**), the relocation

period sets to 8 hours (**08:00**), and the relocation rate sets to **Fast**. So you can ensure that the performance at working hours will not be affected.



**TIP:**

If the storage needs to provide 7 x 24 hours of data access services, may or may not find a long period without data access, please try to find a time frame with a slight I/O of inbound and outbound data flow, execute the relocation rate with **Normal** or **Slow** by either schedule or manual for eliminating the possible performance impact.

Also note that performance improvements may not be obvious when using a relocation rate with **Normal** or **Slow** in comparison with **Fast**.

## Auto Tiering Policies and Their Effect

There are five policies described in the above section, each policy has a suitable situation.

- **Auto-tier (Default)**

This can be used in a large volume of storage structure. Usually the user does not know how to put the data to the right tier; it is entirely handled by the storage system. By default, the data will be relocated at midnight. At this case, hot data calculations take a long time to accumulate and move up, and a few fixed blocks require extreme high performance (but usually the user does not understand the situation). Using this policy will have a significant effect.

- **Begin high after auto-tier**

This can be used for hot data in a short time, such as video editing. The new coming films are often edited at the beginning. After the editing is complete, the files are not always used and eventually moved to the archive. In this scenario, you need to understand the capacity of the hot data and prepare the capacity of the SSD tier. Then this policy can maximize the efficiency.

- **Always highest available tier**

This allows users to allocate resources in a timely manner. Assuming that some volumes will be frequently accessed tomorrow, the IT administrator can manually adjust to this policy. As a result, the data will be relocated to the highest available tier at midnight. In this case, you can get better efficiency under the same resources. Of course, the premise is that the capacity of the volume needs to be controlled.



- **Lowest available tier**

It is for the purpose of data backup, for those volumes which do not need the performance, and the need for large capacity storage of data. It can be set to this policy.

- **No data movement**

This should be least used. The data in the volume using this policy will not operate any hotness analysis. It is suitable for infrequently used data.

As mentioned above, you can choose the right policy based on your application. Or you are unsure, it is recommended to use Auto Tiering policy when creating a volume, and the relocation schedule remains in daily. Then observe the usage of every volume via the performance monitor for a while. And then set the required policy for each volume.

### 3.3. Case Study

- **Case 1: Database, and Archive the historical data**

In the environment of database, people may require high random IOPS but the data which has been written won't be accessed frequently, with the policy – **Begin high after auto-tier** – users in this environment are capable of getting the high random IOPS delivered by PCIe SSD tier, and have large storage capacity to keep the historical data as the cold data will be moved to the high capacity, SATA/SAS tier. So users may grant high random IOPS as what All Flash Array can do while using a much lower budgets to do so by using QSAN XCubeNAS Auto Tiering feature.

#### **Test Equipment and Configurations**

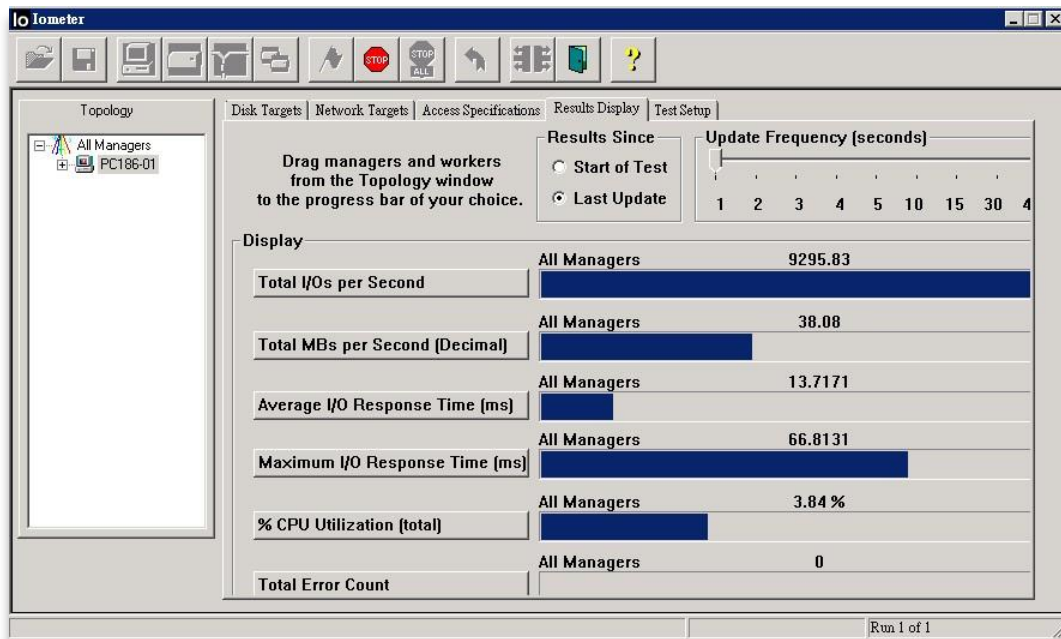
- Server
  - Model: ASUS RS700-E6/ERS4 (CPU: Intel Xeon E5620 2.4GHz / RAM: 24GB)
  - Network: GbE link
  - OS: Windows Server 2008 R2
- Storage
  - Model: QSAN XCubeNAS XN8012R
  - Firmware 3.0.1

- PCIe SSD: 1
- SATA SSD: 2
- SAS HDD: 4
- SATA HDD: 4
- Auto Tiering Pool: 2176.30GB
- PCIe SSD Tier: RAID 0 with 1 x PCIe SSD (Suggest to be RAID1 in real world)
- SATA SSD Tier: RAID 1 with 2 x SATA SSD
- SAS HDD Tier: RAID 6 with 4 x SATA HDD
- SATA HDD Tier: RAID 6 with 4 x SATA HDD
- Volume: 1 x 100GB in Auto Tiering Pool
- Auto Tiering Policy: **Begin high after auto-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 a volume with 100GB capacity, and create a Shared Folder based on this volume
2. Login CIFS session on the server and map as a network drive
3. Test the performance (100% random write) through IOmeter with 128 Outstanding

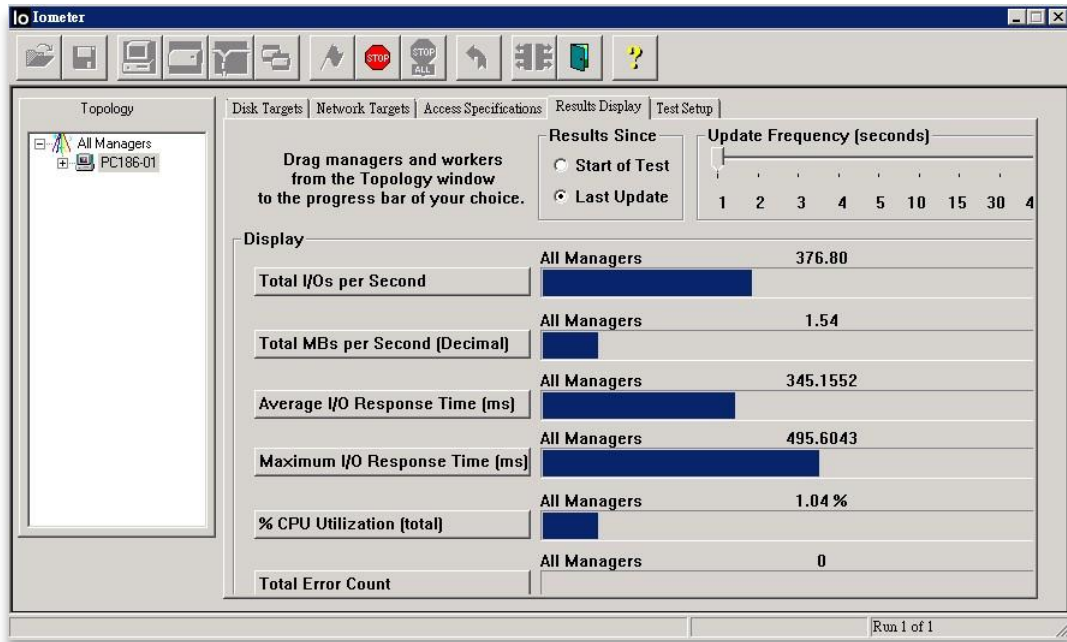
4. Check the result (1)



As the data is now written in the SSD tier, the random IOPS value is delivered from this member SSD(s), we then realize that the newly written data will keep obtaining this value as they are all stored in the SSD tier.

Gradually, the SSD tier is full of the written data, the coldest data (the first written one) will be relocated to the lower tier (SAS or SATA) by either schedule or manual, so if we test the performance on the coldest data again, it would become much lower.

5. Check the result (2)



- **Case 2: Big volume, no idea what tier should be used**

If the environment where you (or users) used to use storage is complex and full of various data access types, usually it is hard to determine what kind of tiering policy should be used for each volume, if you don't want to manage this hassle, **Auto-tier** policy will be the right choice for this scenario.

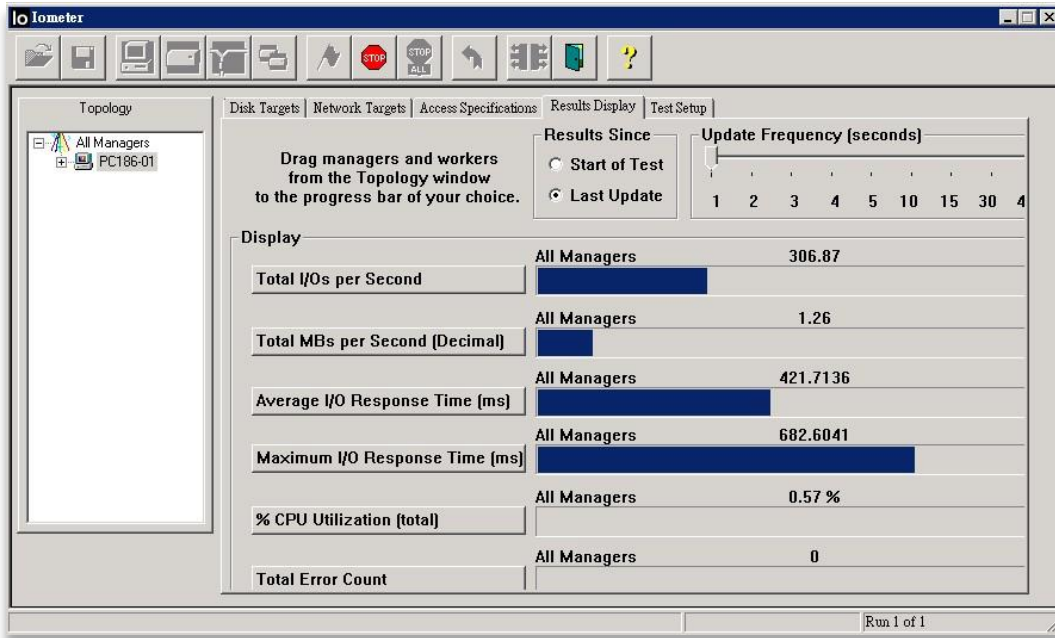
### Test Equipment and Configurations

- Server
  - Model: ASUS RS700-E6/ERS4 (CPU: Intel Xeon E5620 2.4GHz / RAM: 24GB)
  - Network: GbE link
  - OS: Windows Server 2008 R2
- Storage
  - Model: QSAN XCubeNAS XN8012R
  - Firmware 3.0.1
  - PCIe SSD: 1
  - SATA SSD: 2
  - SAS HDD: 4
  - SATA HDD: 4

- Auto Tiering Pool: 2176.30GB
- PCIe SSD Tier: RAID 0 with 1 x PCIe SSD (Suggest to be RAID1 in real world)
- SATA SSD Tier: RAID 1 with 2 x SATA SSD
- SAS HDD Tier: RAID 6 with 4 x SATA HDD
- SATA HDD Tier: RAID 6 with 4 x SATA HDD
- Volume: 1 x 2176.30GB in Auto Tiering Pool
- Auto Tiering Policy: **Auto-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

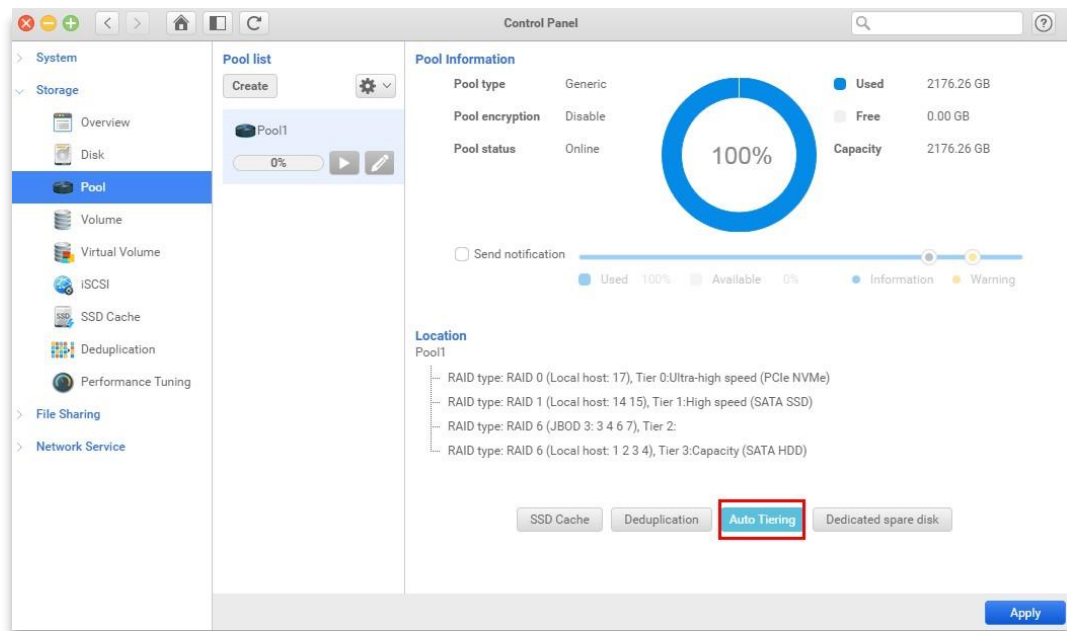
### **Test Scenario and Result**

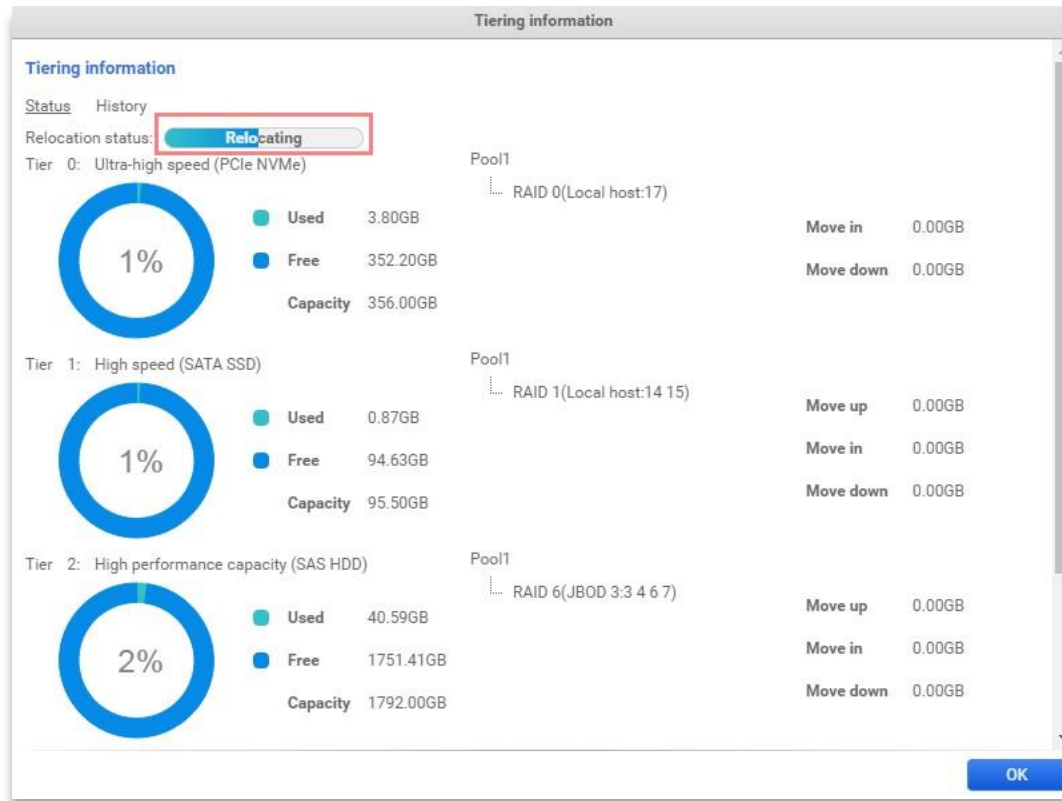
1. Create a volume with full pool capacity, and create a Shared Folder based on this volume
2. Login CIFS session on the server and map as a network drive
3. Test the performance (100% random write) through IOmeter with 128 Outstanding
4. Check the result (1)



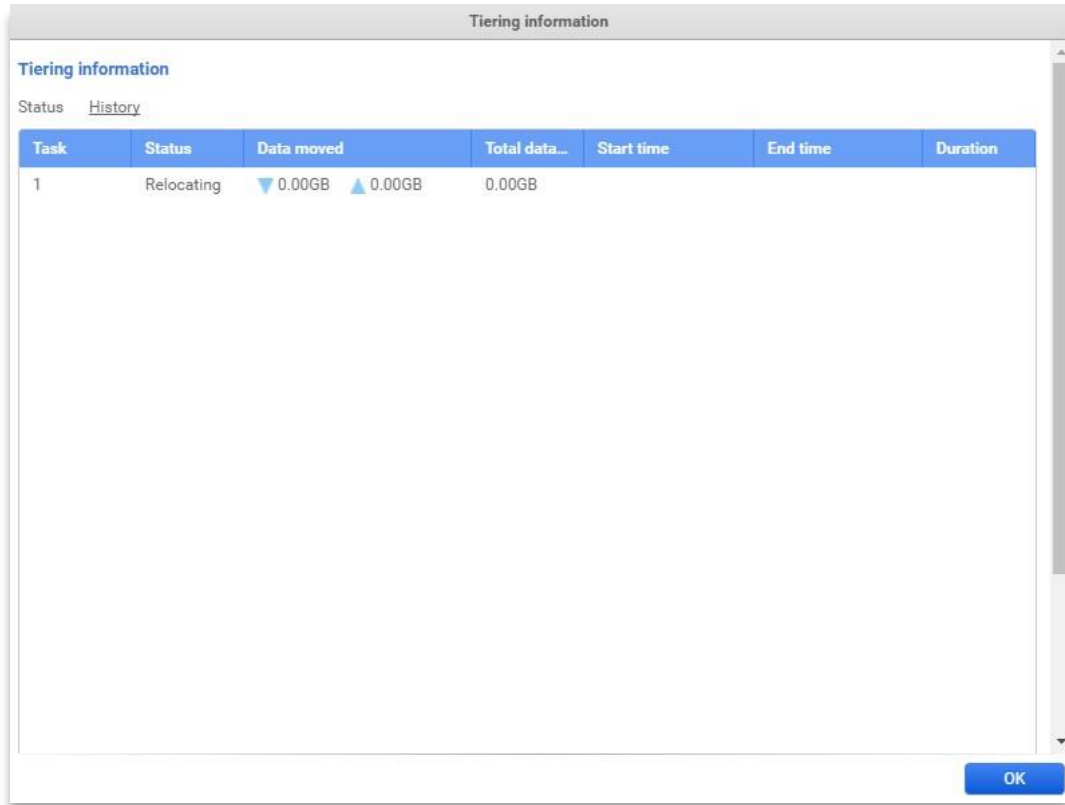
The random IOPS value is quite low as the initial data is written in the healthy and big capacity tier, which is SATA HDD tier at the moment.

5. Copy and paste the data again and again, leave the first written data being accessed continuously
6. As the time passes, there will be a progress in the Auto Tiering page shown on the QSM UI

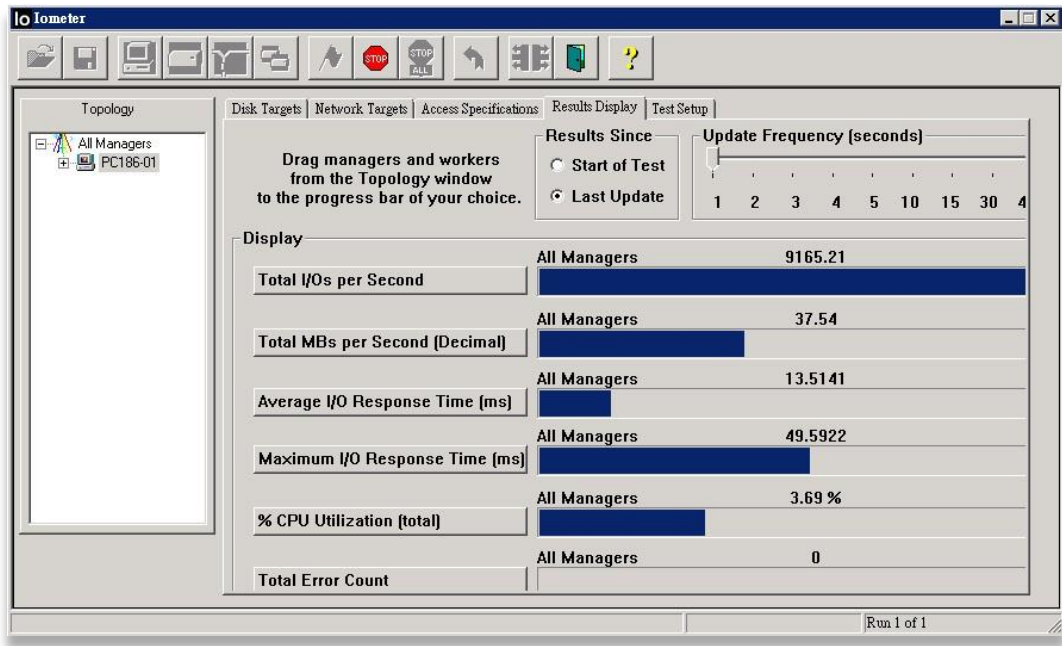




- If you have configured Relocation schedule, hot data (frequently accessed) will be moved into the SSD tier, the data that didn't be read/written anymore since they were inputted will stay in the original tier as what they were, you may check the History page afterwards to check what movements have been done at what time point



8. Check the result again after the relocation is finished





## 4. Conclusion

---

### 4.1. Auto Tiering Notices

The above case studies might not meet the reality due to the variety of different environment, please read thoroughly this technical white paper to understand what tiering policy should be applied to best-fit each scenario.

### 4.2. Apply To

All XCubeNAS series products.

## 5. Support and Other Resources

---

### 5.1. Getting Technical Support

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

1. Via the Web: [http://www.qsan.com/en/contact\\_support.php](http://www.qsan.com/en/contact_support.php)
2. Via Telephone: +886-2-7720-2118 extension 136  
(Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
3. Via Skype Chat, Skype ID: qsan.support  
(Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summertime: 09:30 - 01:00)
4. Via Email: [support@qsan.com](mailto:support@qsan.com)

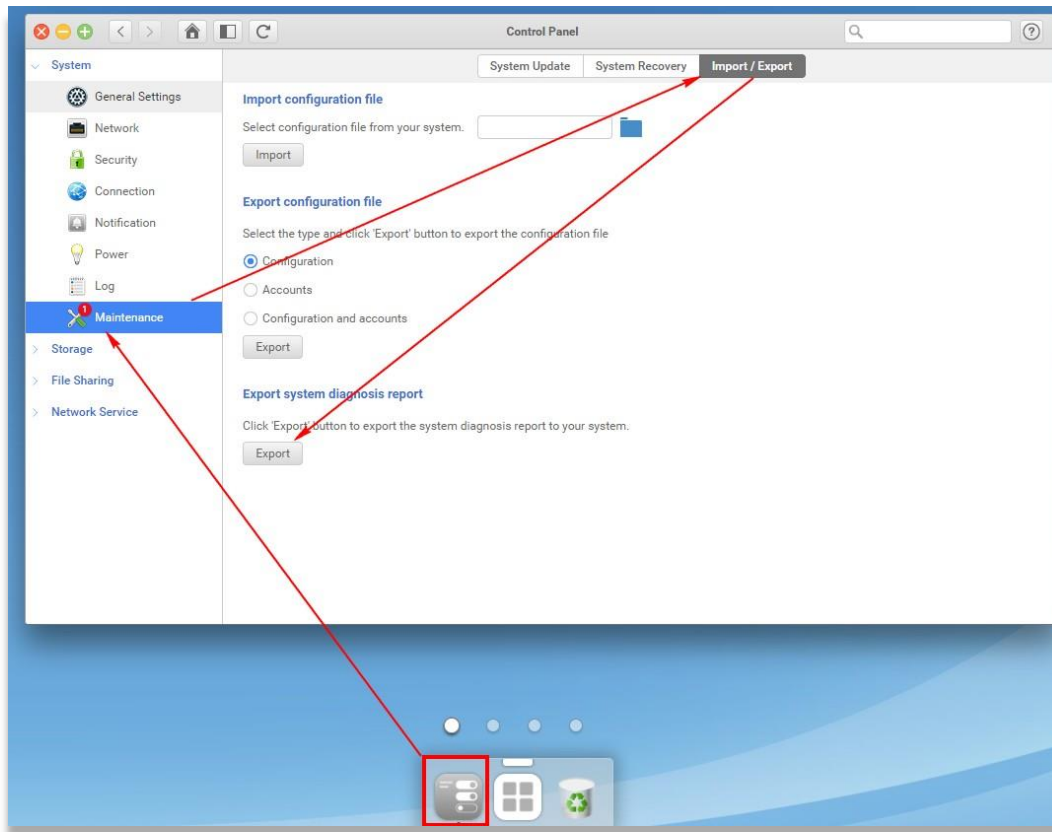
#### Collect Information for Analysis

1. Product name, model or version, and serial number
2. Firmware version
3. Error messages or screenshot images
4. Product-specific reports and logs
5. Add-on products or components installed
6. Third-party products or components installed

#### Information for Technical Support

The following system information is necessary for technical support, please refer to following for what and where to get the information of your XN3002T model.

If the technical support requests you to download the service log, please navigate to the QSM UI → **Control Panel** → **System** → **Maintenance** → **Import/Export** → **Export system diagnosis report**, and then click the **Export** button.



## 5.2. Documentation Feedback

QSAN is committed to providing documentation that meets and exceeds your expectations. To help us improve the documentation, email any errors, suggestions, or comments to [docsfeedback@qsan.com](mailto:docsfeedback@qsan.com).

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

## Appendix

---

### End-User License Agreement (EULA)

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

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

#### **General**

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

#### **License Grant**

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

#### **Intellectual Property Right**

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

#### **License Limitations**

The user may not, and may not authorize or permit any third party to (a) use the Product for any purpose other than in connection with the Product or in a manner inconsistent with the design or documentations of the Product; (b) license, distribute, lease, rent, lend, transfer, assign or otherwise dispose of the Product or use the Product in any commercial hosted or

service bureau environment; (c) reverse engineer, decompile, disassemble or attempt to discover the source code for or any trade secrets related to the Product, except and only to the extent that such activity is expressly permitted by applicable law notwithstanding this limitation; (d) adapt, modify, alter, translate or create any derivative works of the Licensed Software; (e) remove, alter or obscure any copyright notice or other proprietary rights notice on the Product; or (f) circumvent or attempt to circumvent any methods employed by QSAN to control access to the components, features or functions of the Product.

### **Disclaimer**

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

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

### **Limitation of Liability**

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

### **Termination**

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

## **Miscellaneous**

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