



XCubeSAN Series White Paper

RAID EE Technology



QSAN Technology, Inc.
www.QSAN.com



Copyright

© Copyright 2019 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

October 2019

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 other countries.

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

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

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

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

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

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

Notices

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

XCubeSAN Storage System 4U & 6U Rack Mount Models

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

XCubeSAN Storage System 3U 19 Rack Mount Models

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

XCubeSAN Storage System 2U 19 Rack Mount Models

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

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

Information contained in this document has been reviewed for accuracy. It could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the product's features, functionality, and product specifications, which are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should 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
RAID EE Technology.....	1
Executive Summary.....	1
Audience.....	1
Overview.....	1
Theory of Operation.....	3
Configure RAID EE Pools.....	7
Create a RAID EE Pool.....	8
List RAID EE Pools.....	13
Operations on RAID EE Pools.....	16
Test Results.....	18
Test Case 1: RAID 5 vs. RAID 5EE.....	18
Test Case 2: RAID 60 vs. RAID 60EE.....	22
Conclusion.....	25
Apply To.....	25
References.....	25
Appendix.....	27
Related Documents.....	27
Technical Support.....	27

RAID EE Technology

Executive Summary

The RAID architecture which has been in existence for over 30 years now is undergoing a wave of transformation. The original RAID technology has failed to address the problem of excessive rebuilding time due to TB-class large capacity hard disk drives. The new generation RAID technology based on the traditional block technology, which we call RAID EE, is seen as a path to solving the traditional RAID flaw.



INFORMATION:

RAID EE technology is available in SANOS firmware 1.0 and the performance is greatly improved in SANOS firmware 1.4.1.

Audience

This document is applicable for QSAN customers and partners who are interested in learning about RAID EE for solving the problem of excessive rebuilding time. The reader is familiar with QSAN products and has general IT experience, including knowledge as a system or network administrator. If there is any question, please refer to the manuals of products, or contact QSAN support for further assistance.

Overview

RAID (Redundant Array of Independent Disks) is to combine multiple independent physical disks based on a certain algorithm to form a virtual logical disk that provides a larger capacity, higher performance, or better data error tolerance. RAID has been the basic technology of storage systems as a mature and reliable data protection standard. However, with the rapid growth in demand for data storage and the advent of high-performance applications in recent years, traditional RAID has gradually revealed its defects. As hard disk capacity increases, the amount of time required to rebuild RAID data has also dramatically increased. This makes one of the most troublesome for enterprise storage.

management today. In the past days when the hard disk capacity was 10GB to 100GB, RAID rebuild was a job that could be completed in minutes or even more than 10 minutes, which was not yet a problem without special concern. However, as disk capacity grows to hundreds of GB and even TB, RAID rebuild times have increased to hours or even days, it becomes a major problem in storage management.

For example, a traditional RAID with 8 and 1 parity 6TB NL-SAS disk drives takes 2.5 days to rebuild data. The rebuild process consumes system resources, reducing the overall performance of the application system. If users restrict the rebuild, the rebuild time will be even longer. The important of adding time-consuming rebuilding, a large number of access operations could cause the failure of other disks in the pool, greatly increasing the probability of disk failure and risk of data loss.

Limitations of Traditional RAID Architecture

The traditional RAID architecture is composed of a certain number of disk drives selected to form a disk group (also known as RAID group). You may also assign some disk drives as idle hot spare disk drives. A storage pool is grouped to provide capacity for volumes, and then finally map the LUN to the host to become the storage space on the host.

There are several limitations such as RAID architecture:

- First of all, when a disk drive of the disk group is damaged and the rebuild is required, only the member disks of the disk group participate in the rebuild job, and the data writing load at the time is concentrated on the spare disk to form a bottleneck.
- Second, volume data access is limited to the member disks belonging to the disk group; this restricts the performance available to the host because the storage is executing both accessing and rebuilding I/O.

Why RAID Rebuild Time-Consuming

As drive capacity grows, RAID rebuild time grows linearly, raising the rebuild time required by traditional RAID architectures to tens of hours when using RAID with more than 4TB HDD capacity.

There are several factors that affect the RAID rebuild time:

- **HDD Capacity:** The HDD capacity makes up the disk group, the larger the HDD capacity, the longer the rebuild time is required.
- **Quantity of Disk Drives:** The quantity of disk drives included in a disk group affects the amount of time it takes for the system to read data from the remaining disk.

drives and write them to the spare disk drives. The more disks, the longer the rebuild time.

- **Rebuild Job Priority:** During RAID rebuild, the system still has to assure access to the front end host. The higher the priority assigned to the RAID job, the faster the rebuild, but the less the front end host gains performance.
- **Fast Rebuild:** Enabling fast rebuild function only need to rebuild the actual capacity of the volume. Unused disk group space has not to be rebuilt. Only part of the space in a disk group is used by the volume, the rebuild time will be shortened.
- **RAID level:** RAID 1 and RAID 10 with direct block-to-block replication will be faster than RAID 5 and RAID 6 with parity calculations.

Given the potential for failure on a disk drive, the more disk drives contain in a disk group, the more possibility of cumulative failure increase, so there is an upper limit quantity of disk drives in a disk group. Compared with the previous factors, the increasing impact of the disk drive capacity on the rebuild speed has become the primary factor. Such a long rebuild time is apparently not acceptable to a user. To solve the problems of traditional RAID, we implement RAID EE technology.

Theory of Operation

RAID EE adds more spare disks in a disk group we call them **RAID EE spares** to separate the original global, local, and dedicated spare areas are preserved in each stripe of the disk group and are distributed in the disk group by means of disk rotation. When disks have failed in the disk group missing data is rebuilt into the preserved spare areas. Since all disks in the set are destination of rebuild data, the bottleneck of traditional RAID rebuild is gone, rebuild performance dramatically improved. If new disks are added, data in spare areas are copied back to joined disks.

Four new RAID levels are provided for RAID EE, there are:

- **RAID 5EE** (E stands for Enhanced), requires a minimum of 4 disk drives with one RAID EE spare disk which can tolerate 2 disk drives failure. Adding more RAID EE spares will tolerate more disk drives failure.
- **RAID 6EE** requires a minimum of 5 disk drives
- **RAID 50EE** requires a minimum of 7 drives
- **RAID 60EE** requires a minimum of 9 drives



INFORMATION:

The RAID EE spare quantity in a disk group is 1 to 8 disk drives.

Example of RAID 5EE with 1 RAID EE spare

Now we take an example to describe how it works. The following example is a RAID 5EE with 5 disks. 4 disks are for RAID disks, and additional one disk is for RAID EE spare. After initialization, data block distribution is as follows: P stands for parity, S stands for RAID EE spare, and it is empty now.

D1	D2	D3	D4	D5
1	2	3	P	S
S	4	5	6	P
P	S	7	8	9
10	P	S	11	12
13	14	P	S	15

Assume that disk 2 has failed. RAID 5EE is under degraded mode.

D1	D2	D3	D4	D5
1	2	3	P	S
S	4	5	6	P
P	S	7	8	9
10	P	S	11	12
13	14	P	S	15

The spare areas are rebuilt with data from the failed disk drive. This action is called **Rebuild**. After rebuild, data distributed is like RAID 5. It can tolerate another failed disk drive. As we can imagine, the more RAID 5E spare disks, the faster it rebuilds.

D1		D3	D4	D5
1		3	P	2
4		5	6	P
P		7	8	9
10		P	11	12
13		P	14	15


When a new disk drive is joined into the RAID 5E disk group, the data is rebuilt in the spare area. It will be copied back to the new disk. This action is called **Copyback**.

D1	D2	D3	D4	D5
1	2	3	P	S
S	4	5	6	P
P	S	7	8	9
10	P	S	11	12
13	14	P	S	15

After copied back, it is back to RAID 5E normal state.

Example of RAID 60E with 2 RAID 5E spares

Take another example of RAID 60E with 10 disks. 8 disks are for RAID disks, and 2 disks are for RAID 5E spares. After initialization, the block distribution is as follows



Rebuild and copyback of RAID 60EEs similar as the above; it will not be repeated here.


RAID EE Level Summary

The following is the summary of the RAID EE levels.

	RAID 5EE	RAID 6EE	RAID 50EE	RAID 60EE
Min. # Drives	4	5	7	9
Fault Tolerance (G = subgroups, S = RAID EE spares = 1 ~ 8)	2 ~ 9 drive failures (e.g., 1 (RAID 5) + S spares)	3 ~ 10 drive failures (e.g., 2 (RAID 6) + S spares)	G+1 ~ G+8 drive failure (e.g., 2 subgroups (RAID 50) + S spares)	2xG+1 ~ 2xG+8 drive failures (e.g., 2x2 subgroups(RAID 60) + S spares)
Read Performance	Very Good	Very Good	Very Good	Very Good
Write Performance	Good	Fair to Good	Good	Fair to Good
Capacity (N = drive quantity, M = drive capacity G = subgroups, S = RAID EE spares)	$(N-1-S) \times M$ (e.g., (10 drives 1 ~ 2 spares) x 1TB = 7TB)	$(N-2-S) \times M$ (e.g., (10 drives - 2 - 2 spares) x 1TB = 6TB)	$(N-G-S) \times M$ (e.g., (10 drives - 2 subgroups - 2 spares) x 1TB = 6TB)	$(N-2 \times G-S) \times M$ (e.g., (10 drives - 2x2 subgroups - 2 spares) x 1TB = 4TB)
Capacity Utilization (Min. ~ 26 drives)	18% ~ 92% (e.g., 7/10 = 70%)	17% ~ 88% (e.g., 6/10 = 60%)	29% ~ 88% (e.g., 6/10 = 60%)	25% ~ 80% (e.g., 4/10 = 40%)
Typical Applications	Data warehouse Web service, Archive	Data archive, High Availability solution, Server with large capacity requirement	Large database, File server, Application server	Data archive, High Availability solution, Server with large capacity requirement

Configure RAID EE Pools

This section will describe the operations of configuring a RAID pool.



Create a RAID EE Pool

Here is an example of creating RAID EE pool with 4 disks configured in RAID 5E. At the first time of creating a pool, it contains a disk group and the maximum quantity of a disk group is 4.

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



TIP:

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

Create Pool

General

- Disk Selection
- RAID Configuration
- Disk Properties
- Summary

Pool Type

Please select a pool type.

- Thick Provisioning
- Thin Provisioning
- Auto Tiering (Thin Provisioning Enabled)

Pool Properties

Please enter a pool name and select preferred controller setting.

Pool Name : ⓘ

Preferred Controller : ▾

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

SED Pool

Enable SED Pool

Enabling SED pool will use the secure SEDs to create a pool. Intermixing SEDs and non-SEDs are not supported in a pool.

Next Cancel

2. Select a **Pool Type**.

3. Enter a **Pool Name** for the pool. Maximum length of the pool name is 15 characters. Valid characters are [A-Z | a-z | 0-9-_|<>].
4. Select a **Preferred Controller** from the drop-down list. The backend I/O resources in this pool will be processed by the preferred controller which you specify. This option is available when dual controllers are installed.
5. Check the **Enable SED Pool** checkbox. Enabling SED pool will use the secure SEDs to create a pool. Intermixing SEDs and non-SEDs are not supported in a pool.
6. Click the **Next** button to continue.

Create Pool

General | **Select Disks**


Disk Selection

RAID Configuration: Enclosure ID: 0 (Head Unit: XS5216)

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

Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model	Quantity	Health	Capacity	Disk Type
	1	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053	<input type="checkbox"/>	0		
	2	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053	<input type="checkbox"/>	0		
	3	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053	<input type="checkbox"/>	0		
					SEAGATE	ST400FM0053	<input checked="" type="checkbox"/>	0	4	Good 372.36 GB SAS SSD 12.0Gb/s
					MICRON	S630DC-800	<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
					SEAGATE	ST1200MM0088	<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	ST6000NM0014	<input type="checkbox"/>	0	13	Good 5.46 TB NL-SAS HDD 12.0Gb/s
					SEAGATE	ST6000NM0014	<input type="checkbox"/>	0	14	Good 5.46 TB NL-SAS HDD 12.0Gb/s
					SEAGATE	ST6000NM0014	<input type="checkbox"/>	0	15	Good 5.46 TB NL-SAS HDD 12.0Gb/s

7. Please select disks for pool. The maximum quantity of disk a disk group is 64. Select an **Enclosure ID** from the drop-down list to select disks from expansion enclosures.
8. Click the **Next** button to continue.



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

11. 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.
- **Enable High Latency Disk Warning:** Check to enable the high latency disk warning function in a disk group. When a disk drive is abnormal with high latency, the system

will send event logs to notify administrators in advance to avoid affecting overall performance.

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

Create Pool

[General](#)

[Disk Selection](#)

[RAID Configuration](#)

[Disk Properties](#)

[Summary](#)

Pool Properties

Pool Type : Thick Provisioning

Pool Name : Pool-5

Preferred Controller : Controller 1

RAID Configuration

RAID Level : RAID 5EE

Quantity of RAID EE Spares : 1

Quantity of SAS Disks : 4

Disk Properties

Disk Write Cache : Enabled

Disk Read-ahead : Enabled

Disk Command Queuing : Enabled

Disk Standby : Disabled

Back
Finish
Cancel

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	Encryption	Volumes	Current Controller
▼	Pool-5	Online	Good	2.18 TB	2.18 TB	2.18 TB	Disabled	Disabled	Disabled	0	Controller 1

Create Pool

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

List RAID EE Pools

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

	Pool Name	Status	Health	Total	Free	Available	Thin Provisioning	Auto Tiering	Encryption	Volumes	Current Controller
▼	Pool-5	Online	Good	2.18 TB	2.18 TB	2.18 TB	Disabled	Disabled	Disabled	0	Controller 1

Disk Groups

	No.	Status	Health	Total	Free	RAID	Disks Used	RAID EE Spares
▼	1	Online	Good	2.18 TB	2.18 TB	RAID 5EE	4	1

Disks

Enclosure ID	Slot	Status	Health	Capacity	Disk Type	Manufacturer	Model
0	9	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
0	10	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
0	11	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088
0	12	Online	Good	1.09 TB	SAS HDD 12.0Gb/s	SEAGATE	ST1200MM0088

Create Pool

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. • EE Rebuilding: The pool is being RAID EE rebuilt.
Health	<p>The health of the pool</p> <ul style="list-style-type: none"> • Good: The pool is good. • Failed: The pool has 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	The status of Thin provisioning: <ul style="list-style-type: none"> • Disabled: The pool is thick provisioned. • Enabled: The pool is thin provisioned.
Auto Tiering	The status of Auto Tiering: <ul style="list-style-type: none"> • Disabled: The pool is auto tiering disabled. • Enabled: The pool is auto tiering enabled. • Not Supported: The pool contains disk groups with mixed disk type.
Encryption	The Data Secure Mode: <ul style="list-style-type: none"> • Disabled: The pool is not encrypted. • Enabled: The pool is encrypted.
Volumes	The quantity of volumes in the pool.
Current Controller ()	The current running controller of the pool.

Column Name	Description
No.	The number of the 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. • EE Rebuilding: The disk group is being RAID EE rebuilt.
Health	The health of the disk group <ul style="list-style-type: none"> • Good: The disk group is good. • Failed: The disk group has failed. • Degraded: The pool is not healthy and not complete. Reason

	could be missing or failed disks.
Total	Total capacity of the disk group
Free	Free capacity of the disk group
RAID	The RAID level of the disk group.
Disks Used	The quantity of disk drives in the disk group.
RAID EE Spare	The quantity of RAID EE spare disk drives in the disk group. Traditional RAID level is display as N/A.

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 • Copying Back The disk drive is being copied back
Health	The health of the disk drive <ul style="list-style-type: none"> • Good: The disk drive is good. • Failed: The disk drive has 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 NSAS HDD SAS SSD SATA SSD] • [12.0Gb/s 60Gb/s 3.0Gb/s 1.5Gb/s]
Manufacturer	The manufacturer of the disk drive
Model	The model name of disk drive

Operations on RAID EE Pools

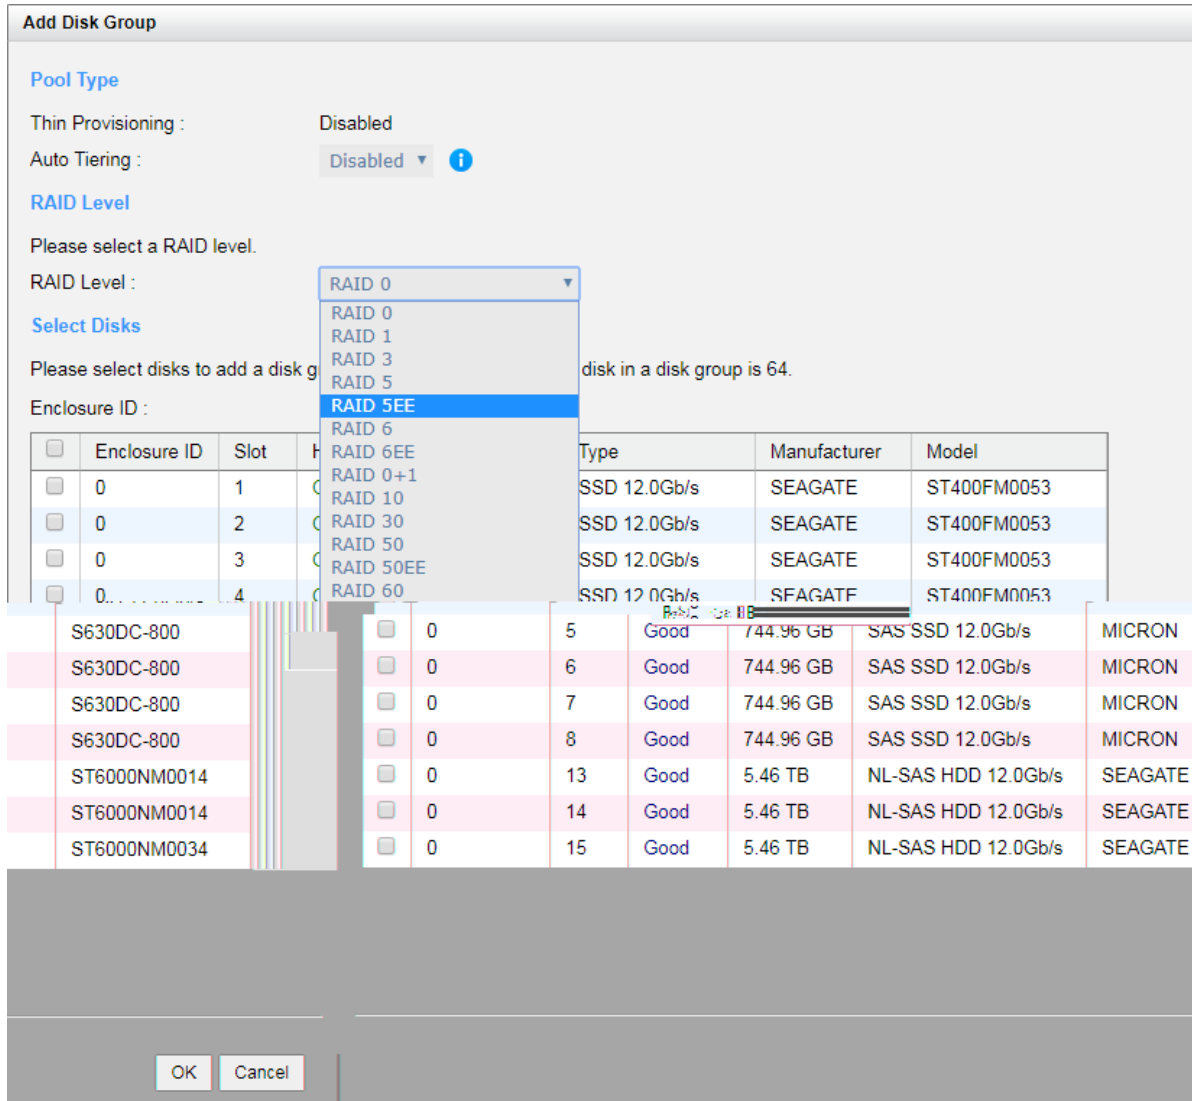
Most operations are described in the Configuring Storage Pools section in the [XCubeSAN SANOS 4.0 Software Manual](#). We describe the restrictions about RAID EE pool in the following.

Verify Parity of the Pool

Click ▼ -> **Verify Parity** in one of pools to generate parity for the pool. It supports RAID level 3, 5, 6, 3Q, 5Q, 6Q and RAID EE level 5EE, 6EE, 5OEE, 6OEE.

Add a Disk Group into the Pool

Click ▼ -> **Add Disk Group** in one of pools to add another disk group into the pool. The new added disk group can be RAID EE level or traditional RAID level. For more information, please refer to the chapter 8.4 Add a Disk Group in a Thick Provisioning Pool in the [XCubeSAN SANOS 4.0 Software Manual](#).




Migrate a Disk Group in a RAID EE Pool

The **Migrate Disk Group** function changes the disk group to a different RAID level or adds the member disks of the disk group to increase the capacity. A traditional RAID level can be migrated to RAID EE level, but RAID EE level can be migrated to RAID level only.

Test Equipments and Configurations

- Server
 - Model:ASUS RS700 X7/PS4CPU: IntelXeon E5-2600 v2/ RAM:8



Summary

- RAID EE can improve rebuild time by up to 96%.
- The more RAID EE spare disks are used, the less rebuild time.
- Rebuild time is more effective when there are reading accesses.
- If the access pattern is random, the copyback time is longer.
- When the data is copying back, the system resources will be reserved preferentially to the front end I/O.

Test Case 2: RAID 60 vs. RAID 60EE

This test provides the comparison of rebuild time and copyback between RAID 60 and RAID 60EE. The same, we assume that the more RAID EE spare disks will have less rebuild time and RAID 60EE will have better efficiency.

Test Equipments and Configurations

- Server
 - Model: ASUS RS700 X7/PS4
 - CPU: Intel Xeon E5-2600 v2/ RAM: 8GB)
 - iSCSI HBA: Intel 82574L Gigabit Network Connection
 - OS: Windows Server 2012 R2
- Storage
 - Model: XCubeSAN XS5224D
 - Memory: 16GB (2 x 8GB in bank 1 & 3) per controller
 - Firmware 14.1

HDD: 24 x Seagate Constellation ES ST500NM0001, 500GB, SAS 6Gb/s

• HDD Pool:

RAID 60 Pool with 16 x NL-SAS HDDs in Controller 1

RAID 60EE Pool with 17 (16+1 RAID EE spare) x NL-SAS HDDs in Controller 1

RAID 60EE Pool with 18 (16+2 x RAID EE spares) x NL-SAS HDDs in Controller 1

RAID 60EE Pool with 20 (16+4 x RAID EE spares) x NL-SAS HDDs in Controller 1

RAID 60EE Pool with 24 (16+8 x RAID EE spares) x NL-SAS HDDs in Controller 1

• HDD Volume: 1TB in Pool

• I/O Pattern

• Tool: IOmeter V1.1.0

• Workers: 1

• Outstanding (Queue Depth): 128

• Access Specifications:

Backup Pattern (Sequential Read / Write, 2KB (MB/s))

Database Access Pattern (as defined by Intel/StorageReview.com, 8KB, 67% Read, 100% Random)

File Server Access Pattern (as defined by Intel/StorageReview.com, refer to the Figure 16,

http://www.storage-review.com/articles/200003/200003130SandBM_5.html)

Idle

• Test Scenario


• First we create a RAID 60 pool. After initialization, plug out and then plug in one disk drive. Count the rebuild time with different I/O access patterns.

• Continue to create RAID 60EE with 1 / 2 / 4 / 8 x RAID EE spare disks in sequence. After initialization, plug out one disk drive. RAID EE starts rebuilding. Count the rebuild time with different I/O access patterns. Then plug in one disk drive and set it as dedicated spare, it starts copying back. Count the copyback time.

Test Result

	RAID 60 (x16)	RAID 60EE (x16+1)	RAID 60EE (x16+2)	RAID 60EE (x16+4)	RAID 60EE (x16+8)
Sequential Read, 256KB	24'58"	13'04"	11'37"	9'23"	4'02"
Improved		48%	53%	62%	84%
Copyback		46'41"	41'38"	34'01"	35'43"
Sequential Write, 256KB	20'16"	13'59"	12'28"	9'40"	4'26"
Improved		31%	38%	52%	78%
Copyback		33'32"	44'33"	31'54"	35'04"
Database Access Pattern	623'22"	93'24"	58'17"	35'34"	16'52"
Improved		85%	91%	94%	97%
Copyback		843'11"	876'59"	492'25"	383'16"
File Server Access Pattern	458'03"	63'36"	55'23"	34'24"	15'51"
Improved		86%	88%	92%	97%
Copyback		1215'34"	1087'47"	673'21"	478'57"
Idle	13'47"	10'49"	10'50"	7'24"	3'30"
Improved		22%	21%	46%	75%
Copyback		26'46"	27'32"	29'30"	33'22"

Take an example, the rebuild time of RAID 60 with sequential read 256KB is 24 minutes and 58 seconds. Compare to the RAID 60EE with RAIDEE spare disks, the rebuild time is 13 minutes and 4 seconds. It improves $((24 \times 60 + 58) - (13 \times 60 + 4)) / (24 \times 60 + 58) = (1498 - 784) / 1498 = 0.4766 = 48\%$.



Summary

- RAID EE can improve rebuild time by up to 97%.
- The more RAIDEE spare disks are used, the less rebuild time is.
- Rebuild time is more effective when there are reading accesses.
- If the access pattern is random, the copyback time is longer.
- When the data is copying back, the system resources will be reserved preferentially to the front end I/O.

Conclusion

As drive capacity grows, RAID rebuild time grows linearly. The more disk drives contain in a disk group, the more possibility of cumulative failure increases. As the increasing impact of the disk drive capacity on the rebuild speed is reduced, using RAIDEE technology will greatly reduce these risks.

Apply To

- XCubeSAN XS5200 / XS3200 / XS1200 1.4.1 and later

References

XCubeSAN SANOS 4.0 Software Manual

- [XCubeSAN SANOS 4.0SoftwareManual](#)

Appendix

Related Documents

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

- [All XCubeSAN Documents](#)
- [XCubeSAN QIG \(Quick Installation Guide\)](#)
- [XCubeSAN Hardware Manual](#)
- [XCubeSAN Configuration Worksheet](#)
- [XCubeSAN SANOS 4.0 Software 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://www.qsan.com/technical_support
- Via Telephone: +862-77206355
(Service hours: 09:30-18:00, Monday-Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support
(Service hours: 09:30-18:00, Monday-Friday, UTC+8, Summer time: 09:30-17:00)
- Via Email: support@qsan.com