



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.

April 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/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
XS1224D	Dual Controller	LFF 24-disk 4U Chassis
XS1224S	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
XS1216D	Dual Controller	LFF 16-disk 3U Chassis
XS1216S	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
XS1212D	Dual Controller	LFF 12-disk 2U Chassis
XS1212S	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
XS1226D	Dual Controller	SFF 26-disk 2U Chassis

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

Information contained in 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
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.....	15
Test Results	18
Test Case 1: RAID 5 vs. RAID 5EE.....	18
Test Case 2: RAID 60 vs. RAID 60EE	21
Conclusion	24
Apply To	24
Reference.....	24
Appendix.....	25
Related Documents.....	25
Technical Support	25

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 1.3.0.

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. It assumes 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 user 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 system as a mature and reliable data protection standard. However, with the rapid growth in demand of disk drive 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 troubles for enterprise storage

management today. In the past days when the hard disk capacity was only 10GB to 100GB, RAID built was a job that could be completed in 10 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 5 with 8 and 1 parity on 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 priority, the rebuild time will be even longer. The important of all, during time-consuming rebuilding, a large number of access operations could cause the failure of other disk drives in the pool, greatly increasing the probability of disk failure and the risk of data loss.

Limitations of Traditional RAID Architecture

The traditional RAID architecture is composed of a certain number of disk drives which are 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 in such 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 participates in the rebuild job, and the data writing loading 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 a 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 disks 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 healthy disk

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

- **Rebuild Job Priority:** During RAID rebuild, the system still has to assume I/O access to the front-end host. The higher the priority assigned to the RAID rebuild job, the faster the rebuild, but the less the front-end host gains I/O performance.
- **Fast Rebuild:** Enabling fast rebuild function only need to rebuild the actual capacity of the volume, unused disk group space has not to rebuild. If 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 rebuild faster than RAID 5 and RAID 6 with parity calculations.

Given the potential for failure on each disk drive, the more disk drives contain in a disk group, the more possibility of cumulative failure increase, so there is an upper limit on the 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 any 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 spares. 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 rebuilt data, the bottleneck of traditional RAID rebuild is gone, rebuild performance dramatically improved. If new disks are added in, data in spare areas are copied back to new 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 is stands for parity, S is 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

Figure 1 Data Block Distributed of RAID 5EE

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

Figure 2 Disk 2 has Failed

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

Figure 3 Empty Blocks are Rebuilt from the Failed Disk Drive

When a new disk drive is joined into the RAID EE disk group, the data rebuilt in the spare area 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

Figure 4 Data is copied back

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

Example of RAID 60EE with 2 RAID EE spares

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




Figure 5 Data Block Distributed of RAID 60EE

Rebuild and copy back of RAID 60EE is similar as the above; it will not be repeated here.

RAID EE Level Summary

The following is the summary of the RAID EE levels.

Table 1 RAID EE Level Summary

	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 (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)	$(N-2xG-S) \times M$ (e.g., (10 drives - 2 x 2 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 EE pool.

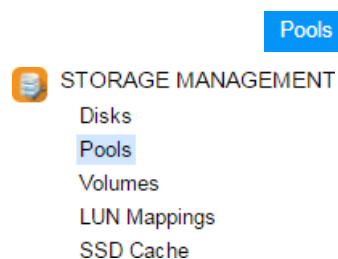


Figure 6 Pools Function Submenu

Create a RAID EE Pool

Here is an example of creating a RAID EE pool with 4 disks configured in RAID 5EE. At the first time of creating a pool, it contains a disk group and the maximum quantity of disk in a disk group is 64.

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



TIP:

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

Create Pool

General

- Disk Selection
- RAID Configuration
- Disk Properties
- Summary

Pool Type

Please select a pool type.

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

Pool Properties

Please enter a pool name and select preferred controller setting.

Pool Name : ⓘ

Preferred Controller : ▾

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

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

Figure 7 Create a RAID EE Pool Step 1

2. Select a **Pool Type**.

3. Enter a **Pool Name** for the pool. 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 drop-down 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. 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.

General
Select Disks

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

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

Disk Properties

<input type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input type="checkbox"/>	0	1	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	2	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	3	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input type="checkbox"/>	0	4	Good	372.36 GB	SAS SSD 12.0Gb/s	SEAGATE	ST400FM0053
<input 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
<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	SEAGATE	ST6000NM0014

Summary

Back
Next
Cancel

Figure 8 Create a RAID EE Step 2

7. Please select disks for pool. The maximum quantity of disk in 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.

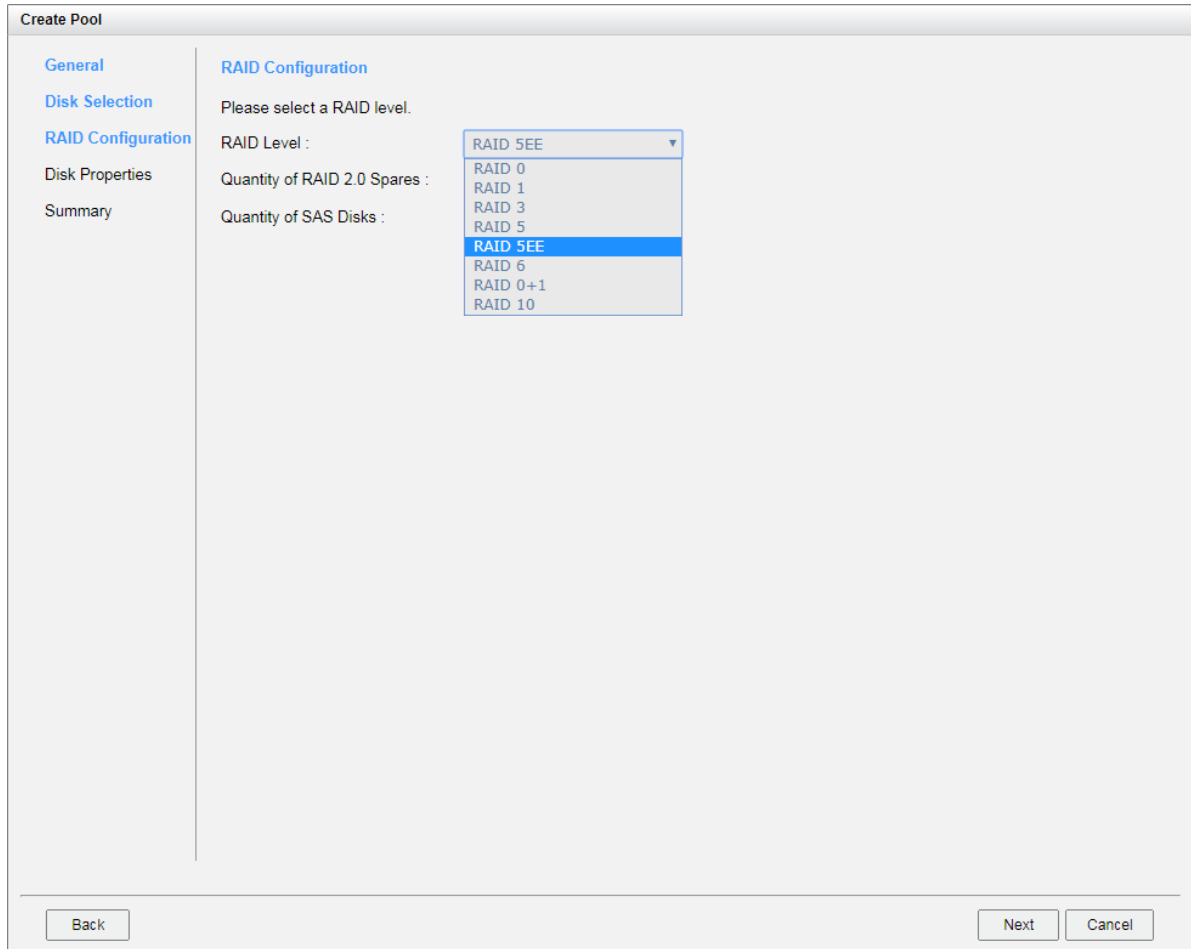


Figure 9 Create a RAID EE Pool Step 3

9. Select a **RAID Level** from the drop-down 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 a **Quantity of Subgroups** if the combination RAID level is selected.
10. Click the **Next** button to continue.

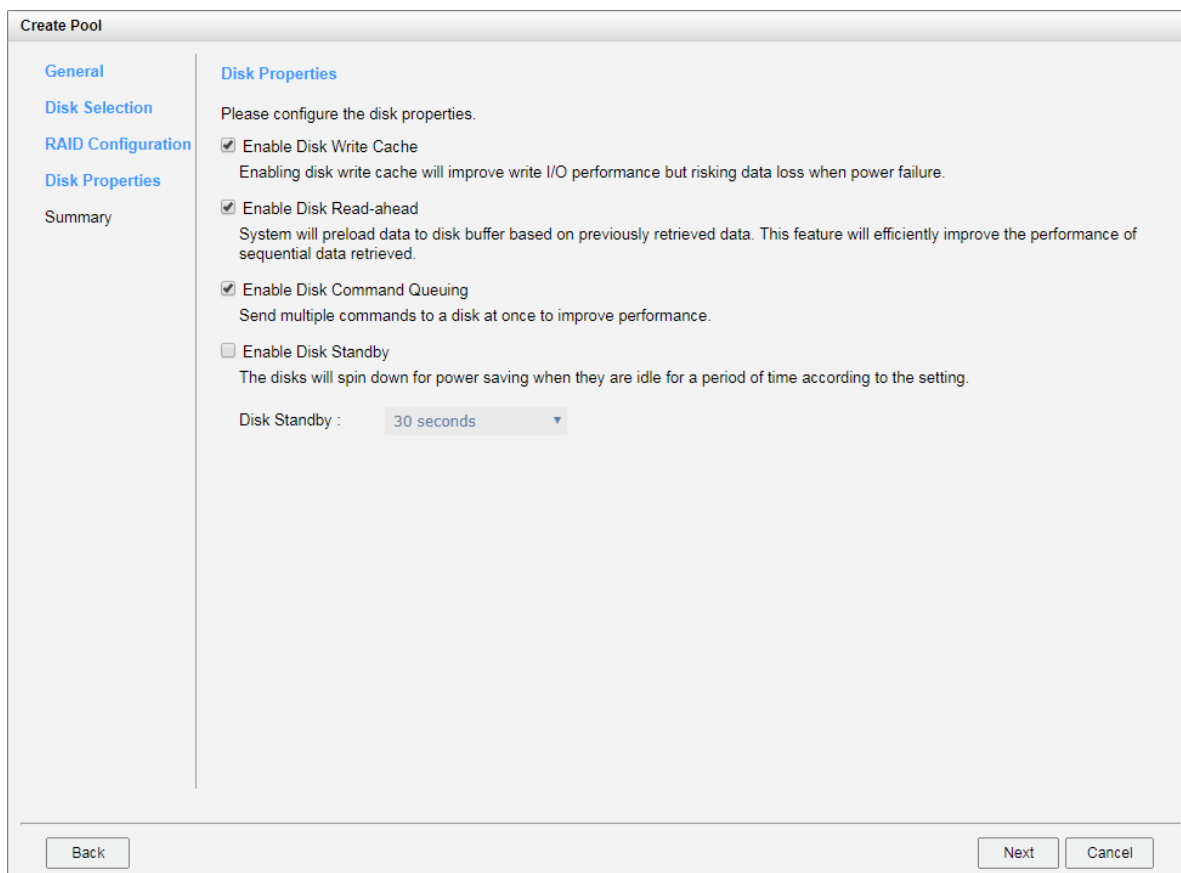


Figure 10 Create a RAID EE Pool Step 4

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.

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

Figure 11 Create a RAID EE Pool Step 5

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

Figure 12 A RAID EE Pool is Created

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

Figure 13 List RAID EE Pools

This table shows the column descriptions.

Table 2 Pool Column Descriptions

Column Name	Description
Pool Name	The pool name.
Status	The status of the pool: <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	The health of the pool: <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 the 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 <i>(This option is only visible when dual controllers are installed.)</i>	The current running controller of the pool.

Table 3 Disk Group Column Descriptions

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. The 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. Display N/A is the RAID level is traditional.

Table 4 Disk Column Descriptions

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 alerts. • 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.

Operations on RAID EE Pools

Most operations are described in the Configuring Storage Pools section in the [XCubeSAN SANOS 4.0 User's 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 regenerate parity for the pool. It supports RAID level 3, 5, 6, 30, 50, 60 and RAID EE level 5EE, 6EE, 50EE, 60EE.

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 8.4.5, Add a Disk Group in a Thick Provisioning Pool section in the [XCubeSAN SANOS 4.0 User's Manual](#).

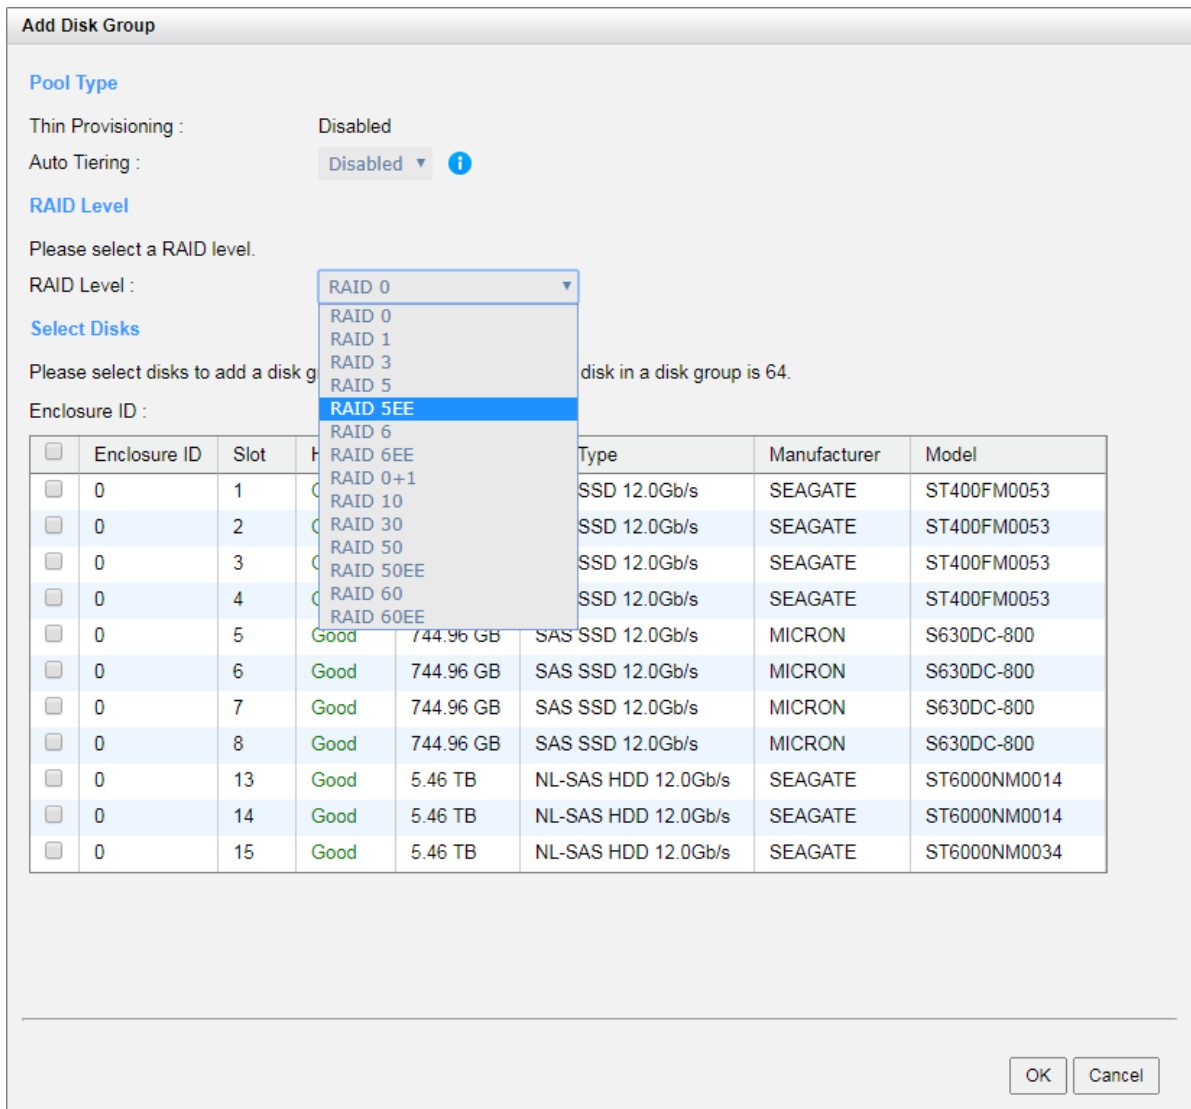


Figure 14 Add a Disk Group into the Pool

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 EE level only.

Migrate Disk Group

RAID Level

Please select a RAID level.

RAID Level : RAID 5EE ▾

Quantity of RAID 2.0 Spares : RAID 5EE

Select Disks

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

Enclosure ID : 0 (Head Unit: XS5216) ▾

<input type="checkbox"/>	Enclosure ID	Slot	Health	Capacity	Disk Type	Manufacturer	Model
<input type="checkbox"/>	0	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
<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	SEAGATE	ST6000NM0034

Figure 15 Migrate a RAID EE Disk Group in a Thick Provisioning Pool



TIP:

A traditional RAID level can be migrated to RAID EE level, but RAID EE level can be migrated to RAID EE level only.

Test Results

Test Case 1: RAID 5 vs. RAID 5EE

This test provides the comparison of rebuild time and copyback time between RAID 5 and RAID 5EE. We assume that the more RAID EE spare disks will have less rebuild time.

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 1.3.0
HDD: 24 x Seagate Constellation ES, ST500NM0001, 500GB, SAS 6Gb/s
 - HDD Pool:
 - RAID 5** Pool with 16 x NL-SAS HDDs in Controller 1
 - RAID 5EE** Pool with 17 (16+1 x RAID EE spare) x NL-SAS HDDs in Controller 1
 - RAID 5EE** Pool with 18 (16+2 x RAID EE spares) x NL-SAS HDDs in Controller 1
 - RAID 5EE** Pool with 20 (16+4 x RAID EE spares) x NL-SAS HDDs in Controller 1
 - RAID 5EE** Pool with 24 (16+8 x RAID EE spares) x NL-SAS HDDs in Controller 1
 - HDD Volume: 100GB in Pool
- I/O Pattern
 - Tool: IOmeter V1.1.0
 - Workers: 1
 - Outstanding (Queue Depth): 128
 - Access Specifications:
 - Backup Pattern** (Sequential Read / Write, 256KB (MB/s))
 - Database Access Pattern** (as defined by Intel/StorageReview.com, 8KB, 67% Read, 100% Random)
 - File Server Access Pattern** (as defined by Intel, refer to the Figure 16, http://www.storagereview.com/articles/200003/20000313OSandBM_5.html)
 - Idle**

Access Patterns			
% of Access Specification	Transfer Size Request	% Reads	% Random
File Server Access Pattern (as defined by Intel)			
10%	0.5 KB	80%	100%
5%	1 KB	80%	100%
5%	2 KB	80%	100%
60%	4 KB	80%	100%
2%	8 KB	80%	100%
4%	16 KB	80%	100%
4%	32 KB	80%	100%
10%	64 KB	80%	100%
Workstation Access Pattern (as defined by StorageReview.com)			
100%	8 KB	80%	80%
Database Access Pattern (as defined by Intel/StorageReview.com)			
100%	8 KB	67%	100%

Figure 16 Access Patterns by StorageReview.com

- Test Scenario
 - First we create a RAID 5 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 5EE with 1 / 2 / 4 /8 x RAID EE spare disks in sequence. After initialization, plug out one disk drive. The 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. Last, count the copyback time.

Test Result

Table 5 The Test Result of RAID 5 and RAID 5EE

	RAID 5 (x16)	RAID 5EE (x16+1)	RAID 5EE (x16+2)	RAID 5EE (x16+4)	RAID 5EE (x16+8)
Sequential Read, 256KB	3'23"	3'14"	2'35"	2'13"	1'46"
Improved		4%	24%	34%	48%
Copyback		4'19"	7'06"	3'31"	4'07"
Sequential Write, 256KB	2'12"	2'57"	2'30"	2'25"	1'51"
Improved		-34%	-14%	-0.1%	16%
Copyback		5'36"	6'39"	4'51"	5'46"
Database Access Pattern	47'17"	52'25"	44'01"	37'03"	32'52"
Improved		-11%	7%	22%	30%
Copyback		136'58"	126'58"	107'49"	81'20"
File Server Access Pattern	35'13"	44'12"	30'13"	26'27"	20'09"
Improved		-26%	14%	25%	43%
Copyback		85'41"	73'50"	69'25"	45'11"
Idle	1'44"	1'43"	1'46"	1'41"	1'43"
Improved		1%	-2%	3%	1%
Copyback		3'32"	3'24"	3'06"	3'53"

Take an example, the rebuild time of RAID 5 with sequential read 256KB is 3 minutes and 23 seconds. Compare to the RAID 5EE with 8 x RAID EE spare disks, the rebuild time is 1 minute and 46 seconds. It improves $((3 \times 60 + 23) - (1 \times 60 + 46)) / (3 \times 60 + 23) = (203 - 106) / 203 = 0.4778 = 48\%$.




Figure 17 The Chart of RAID 5 and RAID 5EE

Summary

- RAID EE can improve rebuild time by up to 48%.
- The more RAID EE spare disks are used, the less rebuild time is.
- Rebuild time is more effective when there are reading accesses.

Test Case 2: RAID 60 vs. RAID 60EE

This test provides the comparison of rebuild time and copyback time 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 1.3.0
 - 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 x 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: 100GB in Pool
- I/O Pattern
 - Tool: IOmeter V1.1.0
 - Workers: 1
 - Outstanding (Queue Depth): 128
 - Access Specifications:
 - Backup Pattern** (Sequential Read / Write, 256KB (MB/s))
 - Database Access Pattern** (as defined by Intel/StorageReview.com, 8KB, 67% Read, 100% Random)
 - File Server Access Pattern** (as defined by Intel, refer to the Figure 16, http://www.storagereview.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. The 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. Last, count the copyback time.

Test Result

Table 6 The Test Result of RAID 60 and RAID 60EE

	RAID 60 (x16)	RAID 60EE (x16+1)	RAID 60EE (x16+2)	RAID 60EE (x16+4)	RAID 60EE (x16+8)
Sequential Read, 256KB	6'08"	3'33"	2'49"	2'37"	2'33"
Improved		42%	54%	57%	58%
Copyback		6'38"	5'40"	4'38"	5'06"
Sequential Write, 256KB	2'59"	4'18"	4'34"	3'49"	3'18"
Improved		-44%	-53%	-28%	-11%
Copyback		8'20"	7'39"	7'21"	6'48"
Database Access Pattern	56'01"	57'01"	50'35"	40'18"	37'21"
Improved		-2%	10%	28%	33%
Copyback		175'56"	168'57"	151'29"	114'00"
File Server Access Pattern	40'35"	43'08"	30'11"	27'26"	21'15"
Improved		-6%	26%	32%	48%
Copyback		125'40"	116'38"	100'59"	75'06"
Idle	1'17"	2'01"	1'52"	1'44"	1'22"
Improved		-57%	-45%	-35%	-6%
Copyback		4'35"	4'29"	4'05"	4'43"

Take an example, the rebuild time of RAID 60 with sequential read 256KB is 6 minutes and 8 seconds. Compare to the RAID 60EE with 8 x RAID EE spare disks, the rebuild time is 2 minute and 33 seconds. It improves $((6 \times 60 + 8) - (2 \times 60 + 33)) / (6 \times 60 + 8) = (368 - 153) / 368 = 0.5842 == 58\%$.

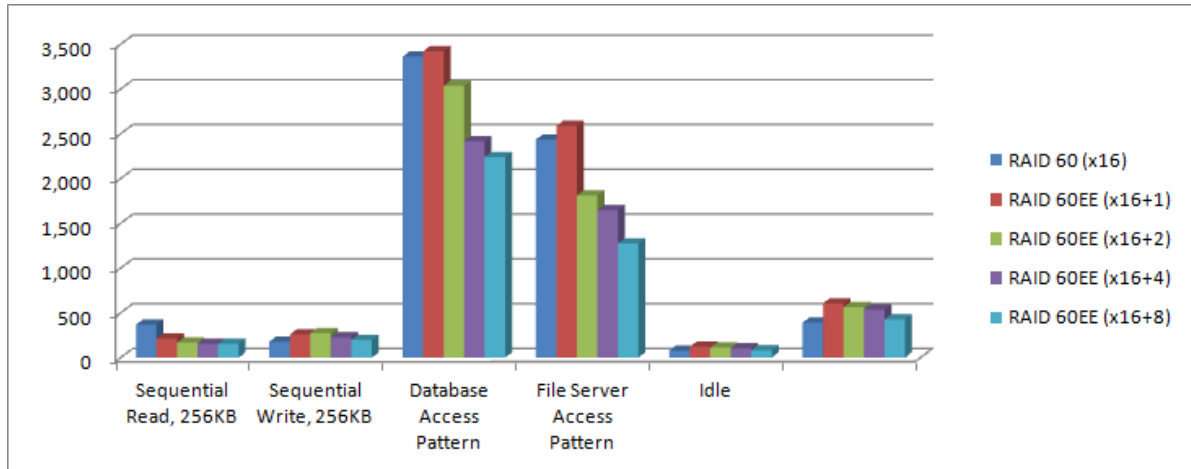


Figure 18 The Chart of RAID 60 and RAID 60EE

Summary

- RAID EE can improve rebuild time by up to **58%**.
- The more RAID EE spare disks are used, the less rebuild time is.
- Rebuild time is more effective when there are reading accesses.

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 increase, so does the increasing impact of the disk drive capacity on the rebuild speed. Using RAID EE technology will greatly reduce these risks.

Apply To

- XCubeSAN XS5200 / XS3200 / XS1200 FW 1.3.0 and later

Reference

XCubeSAN SANOS 4.0 User's Manual

- [XCubeSAN SANOS 4.0 User's Manual](#)

Appendix

Related Documents

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

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

Technical Support

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

- Via the Web: <https://qsan.com/support>
- 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, Summer time: 09:30 - 01:00)
- Via Email: support@qsan.com