

Benefits of Using RAID 50 or 60 in Single High Capacity RAID Array Volumes Greater than 16 Disk Drives

Document Version 1.0

Promise Technology Inc www.promise.com



Copyright

©2007 Promise Technology, Inc. All Rights Reserved.

No part of this document may be reproduced or transmitted in any form without the expressed, written permission of Promise Technology, Inc.

Trademarks

Promise, WebPAM, PerfectRAID and the Promise logo are registered trademarks of Promise Technology, Inc. All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Promise U.S.A. / Headquarters

Promise Technology, Inc. 580 Cottonwood Drive Milpitas, CA 95035 U.S.A.

Phone: 1-408-228-1400 Fax: 1-408-228-1100

Promise EMEA

Science Park Eindhoven 5542 5692 EL Son, The Netherlands

Phone: +31-40-235-2600 Fax: +31-40-256-9463

Promise Taiwan

Promise Technology, Inc. 2F, No. 30, Industry E. Rd. IX, Science-Based Industrial Park, Hsinchu, Taiwan

R.O.C.

Phone: +886-3-578-2395 Fax: +886-3-578-2390

Promise EMEA

Promise Technology Germany Europaplatz 9 44269 Dortmund Germany

Phone: +49 (0) 2 31 56 76 48 - 0 Fax: +49 (0) 2 31 56 76 48 - 29

Promise China

Room 1205, Tower 3 Webok Time-center, No.17 South Zhong Guan Cun Road Hai Dian District Beijing, China

Phone: +86-10-6872-3941 Fax: +86-10-6872-3940



Table of Contents

1	Abstract	4
2	Discussion	5
3	E-Class RAID Level Support of Interest	6
4	Benefits	11
5	Conclusion	13



1 Abstract

This document will explain and demystify the misconception that there is a benefit to creating single RAID 5 and RAID 6 array volumes greater than 16 SAS/SATA disk drives. It will explain the logic behind using less than 16 disk drive spindles in a RAID 5 or RAID 6 array as well as explain the benefit of using RAID 50 or RAID 60 for RAID arrays with more than 16 disk drive spindles. All of these RAID levels accomplish the same objective, data protection.



2 Discussion

The reason for the 16 drive design maximum in the Promise V3 and V4 RAID Engine for RAID 5 and 6 is due to the increased probability of disk drive spindle or bad block failures when multiple drives are configured in the same RAID 5 or RAID 6 array unit. The V4 and V3 RAID Engine are currently used in our VTrak E-Class, VTrak M-Class and SuperTrak EX line of products.

The type of failures can range from simple bad blocks to catastrophic failures, where the disk drive is no longer operational, such as motor failure or head crash. Although these types of issues are less likely to happen with the current, constantly improving SAS/SATA drive, bad blocks are not uncommon, especially on the larger capacity disk drives. The larger the capacity of a disk drive the greater the probability there is of hitting a bad block. Factors can include the platter size, density or both. Since RAID 5 allows for a one drive failure and remain operational running in Critical mode, losing an additional drive can result in data loss, i.e. result in an Offline RAID 5 array. RAID 6 allows for a two drive failure before resulting in a Critical mode thus reducing the chances of the array being in an Offline state. Although RAID 6 has its benefits, RAID 50 and RAID 60 are much more redundant in that the more axles (more explanation on axles in the next section) in an array the more drives can fail before an array will go into an Offline state. Using RAID 50 and 60 also allows the user to surpass the 16 disk drive limitations of RAID 5 and RAID 6. If you are creating a single drive volume of up to 60 disk drives, RAID 50 or 60 is a better alternative.

Below Figure 1 displays the RAID Level and number of drives that are configurable. Expansion units are only applicable on E-Class product configurations with use of J-Class.

Level	Number of Drives		Level	Number of Drives
RAID 0	1 or more		RAID 6	4 to 16*
RAID 1	2 only		RAID 10	4 or more**
RAID 1E	2 or more		RAID 50	6 or more
RAID 5	3 to 16*		RAID 60	8 or more
* Drive counts above 12 require an expansion unit				

Drive counts above 12 require an expansion unit

Figure 1. Depicts possible RAID Level Combinations with minimum drive requirements

^{**} Must be an even number of drives.



3 E-Class RAID Level Support of Interest

RAID 5 – Block and Parity Stripe

RAID 5 organizes block data and parity data across the physical drives. Generally, RAID Level 5 tends to exhibit lower random write performance due to the heavy workload of parity recalculation for each I/O. RAID 5 is generally considered to be the most versatile RAID level.

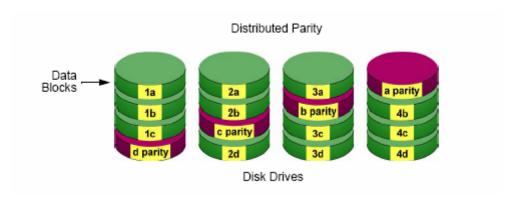


Figure 2. RAID 5 Stripes all drives with data and parity information

The capacity of a RAID 5 disk array is the smallest drive size multiplied by the number of drives less one. Hence, a RAID 5 disk array with (4) 100 GB hard drives will have a capacity of 300GB. A disk array with (8) 120GB hard drives and (1) 100GB hard drive will have a capacity of 800GB. RAID 5 requires a minimum of three physical drives and a maximum of 16.

Recommended applications: File and Application Servers; WWW, E-mail, News servers, Intranet Servers

RAID 6 – Block and Double Parity Stripe

RAID level 6 stores dual parity data is rotated across the physical drives along with the block data. A RAID 6 disk logical drive can continue to accept I/O requests when any two physical drives fail.



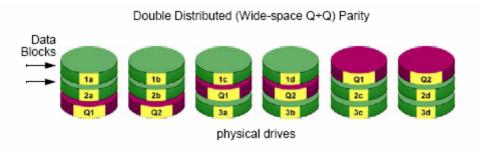


Figure 3. RAID 6 Block and Double Parity Stripe

The total capacity of a RAID 6 disk logical drive is the smallest physical drive times the number of physical drives, minus two. Hence, a RAID 6 disk logical drive with (6) 100 GB hard drives will have a capacity of 400 GB. A disk logical drive with (4) 100 GB hard drives will have a capacity of 200GB. RAID 6 becomes more capacity efficient in terms of physical drives as the number of physical drives increases. RAID 6 offers double fault tolerance. Your logical drive remains available when up to two physical drives fail. RAID 6 is generally considered to be the safest RAID level. RAID 6 requires a minimum of four physical drives and a maximum of 16.

Recommended applications: Accounting, financial, and database servers; any application requiring very high availability.

RAID 50 – Striping of Distributed Parity

RAID 50 combines both RAID 5 and RAID 0 features. Data is striped across disks as in RAID 0, and it uses distributed parity as in RAID 5. RAID 50 provides data reliability, good overall performance and supports larger volume sizes.



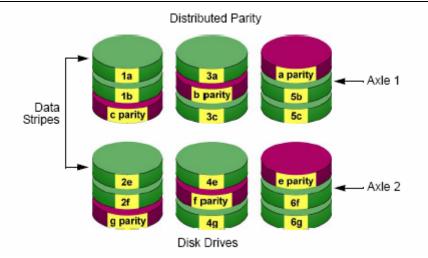


Figure 4. RAID 50 Striping of Distributed Parity disk arrays

RAID 50 also provides high reliability because data is still available even if multiple disk drives fail (one in each axle). The greater the number of axles, the greater the number of disk drives that can fail without the RAID 50 array going offline. RAID 50 arrays consist of six or more physical drives.

Recommended applications: File and Application Servers, Transaction Processing, Office applications with many users accessing small files.

RAID 50 Axles

When you create a RAID 50, you must specify the number of axles. An axle refers to a single RAID 5 array that is striped with other RAID 5 arrays to make RAID 50. An axle can have from three or greater physical drives, depending on the number of physical drives in the array (E-Class has been tested up to 60 physical drives, 1 E310f + 4 J300s). Example: Although not depicted in the table below in a 60 disk drive configuration you could set up to 16 drives per axle for a total of 4 axles. 3 out of the four axles would have 16 disk drives and 1 axle would have 12 disk drives thus resulting in an unbalanced RAID 50 configuration.

The chart below shows RAID 50 arrays with 6 to 16 disk drives, the available number of axles, and the resulting distribution of disk drives on each axle. The VTrak attempts to distribute the number of disk drives equally among the axles but in some cases, one axle will have more disk drives than another.



No. of Drives in RAID 50 Array	No. of Axles in RAID 50 Array	No. of Drives per Axle
6	2	3
8	2	4
9	3	3
10	2	5
12	2	6
	3	4
	4	3
14	2	7
15	3	5
	5	3
16	2	8
	4	4

RAID 60 – Striping of Double Parity

RAID 60 combines both RAID 6 and RAID 0 features. Data is striped across disks as in RAID 0, and it uses double distributed parity as in RAID 6. RAID 60 provides data reliability, good overall performance and supports larger volume sizes.

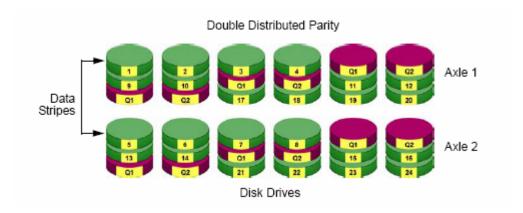


Figure 5. RAID 60 Striping of Double Distributed Parity disk arrays

RAID 60 also provides very high reliability because data is still available even if multiple disk drives fail (two in each axle). The greater the number of axles, the greater the number of disk drives that can fail without the RAID 60 array going offline. RAID 60 arrays consist of eight or more physical drives.



Recommended applications: Accounting, financial, and database servers; any application requiring very high availability.

RAID 60 Axles

When you create a RAID 60, you must specify the number of axles. An axle refers to a single RAID 6 array that is striped with other RAID 6 arrays to make RAID 60. An axle can have from four or more physical drives depending on the number of physical drives in the array (E-Class has been tested up to 60 physical drives, 1 E310f + 4 J300s). Example: Although not depicted in the table below in a 60 disk drive configuration you could set up to 16 drives per axle for a total of 4 axles. 3 out of the four axles would have 16 disk drives and 1 axle would have 12 disk drives thus resulting in an unbalanced RAID 60 configuration.

The chart below shows RAID 60 arrays with 8 to 16 physical drives, the available number of axles, and the resulting distribution of disk drives on each axle. VTrak attempts to distribute the number of disk drives equally among the axles but in some cases, one axle will have more disk drives than another.

No. of Drives in RAID 60 Array	No. of Axles in RAID 60 Array	No. of Drives per Axle
8	2	4
10	2	5
12	2	6
	3	4
15	3	5
16	2	8
	4	4



4 Benefits

RAID 5

Advantages	Disadvantages
High Read data transaction rate	Disk failure has a medium impact on
Medium Write data transaction rate	throughput
Good aggregate transfer rate	
16 Array Limitation per Single RAID 5 volume	

Recommended Applications for RAID 5

- File and Application servers
- WWW, E-mail, and News servers
- Intranet servers
- Most versatile RAID level

RAID 6

Advantages	Disadvantages
 High Read data transaction rate Medium Write data transaction rate Good aggregate transfer rate Safest RAID Level 16 Array Limitation per Single RAID 6 volume 	 High disk overhead – equivalent of two drives used for parity Slightly lower performance than RAID 5

Recommended Applications for RAID 6

- Accounting and Financial
- Any Application requiring very high availability
- Database Servers



RAID 50

	Advantages		Disadvantages
•	Greater than 16 disk drive single volume for high capacity (up to 60 disk drives using J300s)	• I	Higher disk overhead than RAID 5
•	High Read data transaction rate		
•	Medium Write data transaction rate		
•	Good aggregate transfer rate		
•	High reliability		

Recommended Applications for RAID 50

- File and Application servers
- Transaction Processing
- Office Applications with many users accessing small files
- High Capacity Data Storage

RAID 60

Advantages		Disadvantages	
•	Greater than 16 disk drive single volume for high capacity (up to 60 disk drives using J300s)	•	High disk overhead – equivalent of two drives used for parity
•	High Read data transaction rate	•	Slightly lower performance than RAID 50
•	Medium Write data transaction rate Good aggregate transfer rate		
•	Safest RAID level		

Recommended Applications for RAID 60

- Accounting and Financial
- Database Servers
- Any Application Requiring high availability
- High Capacity Data Storage



5 Conclusion

RAID 50 or RAID 60 is best suited for high capacity volume array with greater than 16 disk drive single RAID configurations as opposed to RAID 5 or RAID 6 due to the increased probability of disk drive spindle or bad block failures. Using RAID 50 in these types of configurations allows for one drive failure per one disk drive per axle. RAID 60 allows you to lose up to two disks per axle thus making RAID 50 and RAID 60 the array of choice for safer, more redundant, high capacity Single RAID volumes with greater than 16 disk drives.