

The Benefits of RAID 6 in Storage Builder[®] Appliances

RAID is an acronym for Redundant Array of Independent (or Inexpensive) Disks. It is a technology that allows for a combination of two or more hard disks to deliver increased fault tolerance and performance compared to what could be achieved with one drive or simply aggregating multiple drives in a JBOD (Just a Bunch of Disks) configuration. Fault tolerance, for the purposes of this discussion, is the ability of the storage system to withstand a hard disk failure without losing data. A relative newcomer to this family of technologies is RAID 6, also known as double-parity RAID.

For large capacity storage devices, RAID 5 is the most common implementation. In RAID 5, data is striped across all the drives, with one data block containing the parity data, as shown in **Figure 1**. The parity data is what allows the data to be reconstructed and accessed in the case of a drive failure. A RAID 5 array has the capacity of n-1 hard drives, where n equals the total number of drives in the array.




Figure 1: RAID 5 Configuration

Figure 1 shows how data is striped across the disks in a RAID 5 array. Stripe 1 consists of data blocks 1A, 1B, 1C, 1D, and parity block P1. Block P1 contains the parity information that is calculated for the data blocks in stripe 1; block P2 contains the parity information for stripe 2, and so on. If one of the hard disks (1, 2, 3, or 4) were to fail, the parity information would be used to reconstruct the missing data during a read operation and to rebuild the array once a replacement disk has been provided. For example, if disk 3 were to fail, the parity information in P1, P2 and P4 would be used to recreate the data that was in blocks 1C, 2C and 4B respectively. Likewise, the data in blocks 3A, 3B, 3C, and 3D would be used to recreate parity block P3. For subsequent stripes, the parity block is rotated between the disks in the array. Rotating the parity blocks keeps all parity from being written to the same disk, which would result in that disk becoming a performance bottleneck.

If a hard disk in a RAID 5 array fails, the array becomes degraded. The data on the array is still accessible, but the array is no longer fault tolerant until it is completely rebuilt. Until the array is completely rebuilt there exists a window of vulnerability, and if a second hard disk were to fail all data that existed on that array would be lost. Even without a second disk failing, if any of the remaining array members have developed bad sectors, (i.e., the data on those sectors is unreadable), the stripes which those sectors were a part of will not be able to be rebuilt. Periodic array scans/verifications can minimize the risk of this occurring by proactively detecting a bad sector, using the parity data to recreate the data that existed on that sector, and writing the data to another location on the drive.

RAID 6

RAID 6 greatly reduces the window of vulnerability by utilizing a second set of parity information striped across the drive members of the array. A RAID 6 array can withstand both of the scenarios that would cause a RAID 5 array to fail — the failure of two hard disks, or bad sectors encountered during a rebuild.




Figure 2: RAID 6 Configuration

As with RAID 5, data in a RAID 6 array is striped across the disks. In addition to the parity block P, another block of parity information, parity block Q, is calculated as shown in **Figure 2**. A different method of calculating parity is used for Q, which enables a RAID 6 array to survive the failure of two disks.

RAID 5 or RAID 6?

There are trade-offs to consider when deciding whether to use RAID 6; the first is usable storage capacity. Whereas a RAID 5 array will provide usable capacity of $n-1$ hard disks in an array, RAID 6 is less efficient and provides usable capacity of $n-2$ hard disks. As an example, if there are six (6) 500 GB hard disks in an array, it has a raw capacity of 3000 GB (3 TB). A RAID 5 array will provide 2500 GB (2.5 TB) of usable capacity. A RAID 6 array with the same drives would provide 2000 GB (2 TB) of usable capacity.

Another consideration is write performance. When data is written to an array that uses parity, there is some write latency due to the parity calculations. The data is first written into the RAID controller's cache memory, the parity is calculated, and finally the data and parity blocks are written to the array. Because a RAID 6 array has additional parity information that is calculated, it incurs some additional latency for the write operation. The use of optimized RAID algorithms minimize the write penalty inherent in RAID 5 and 6.

Additional Configuration Possibilities With WasabiRAID™

Wasabi Storage Builder® features an integrated RAID capability powered by WasabiRAID™. WasabiRAID can be used in conjunction with two RAID controllers to create RAID 60 or RAID 61 volumes, as shown in **Figure 3**.

WasabiRAID can create RAID 60 or RAID 61 Volumes by striping across or mirroring the RAID 6 arrays

RAID 6 arrays created with the RAID controller

RAID controllers

Physical disks attached to the RAID controllers




Figure 3: RAID 60 or RAID 61 Configuration

RAID 60 leverages the double-disk protection of RAID 6 and combines that with RAID 0, which enhances performance by splitting data evenly across each of the RAID 6 arrays so that data can be written to and read from each array in parallel. This configuration can withstand failure of up to two drives behind each RAID controller without losing data. However, if either of the controllers were to fail, then all of the data on the RAID 60 volume would be lost. RAID 61 mirrors the RAID 6 arrays so that even if one of the controllers fails no data loss will occur. Assuming the disks in the illustrations above were 500 GB, the usable capacity with the RAID 60 array would be 2TB (a total of 4 disks worth of capacity), and the RAID 61 usable capacity would be 1TB. As with the RAID 5 vs. RAID 6 decision, the optimal balance between performance and protection needs to be taken into consideration when deciding between RAID 60 and RAID 61. The decision will be based on how critical the application is, available budget, and comfort level as well as other factors as defined by the user.

Summary

RAID 6 protects against the failure of two hard disks and adds an additional level of data protection to Storage Builder appliances. WasabiRAID enables additional configuration options when used in conjunction with RAID controllers, providing the flexibility to choose the right balance of performance, protection, and capacity for any application.

Copyright © 2007 Wasabi Systems Inc., 500 East Main Street, Suite 1520, Norfolk, VA, 23510. All rights reserved. This document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Wasabi Systems Inc. and its licensors, if any. Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California; full license and copyright information is available at: <http://www.netbsd.org/Goals/redistribution.html>. NetBSD® is a registered trademark of The NetBSD Foundation. Wasabi®, the Wasabi logo, Flashware®, and Storage Builder® are registered trademarks of Wasabi Systems Inc. All other brand and product names are trademarks of their respective owners. REV. WSB_RAID6_20070205