



White Paper

Enhancing Storage Performance and Investment Protection
Through RAID Controller Spanning

May 2005



Introduction

It is no surprise that the rapid growth of data within enterprise networks is creating very significant challenges for IT departments. One such challenge is supporting large amounts of data on IT budgets that are not able to keep pace with the growth of this data. When focusing on storage equipment, any technology that can reduce costs while simultaneously improving data I/O access performance and enhancing the protection of investments in storage and staff would contribute greatly to efforts to streamline and optimize. Broadcom's controller spanning technology, which is incorporated in its popular XelCore™ RAID software, does all of this and, as a result, stands out as an ideal candidate in support of storage optimization efforts.

Spanning Defined

Simply stated, controller spanning (or spanning) is a RAID software feature that allows a RAID data array to span hard disk drives attached to multiple RAID controllers. With spanning, a single data array, which appears as one virtual volume or disk to the user, actually resides on disks attached to two or more different controllers. Without spanning, a data array can only include those drives attached to a single RAID controller.

To illustrate, Figure 1 shows a server equipped with four 8-channel controllers and RAID software that does not support spanning. In this case, an array can only incorporate the disks attached to an individual controller creating four RAID arrays apparent to the user, each comprised of eight disks.

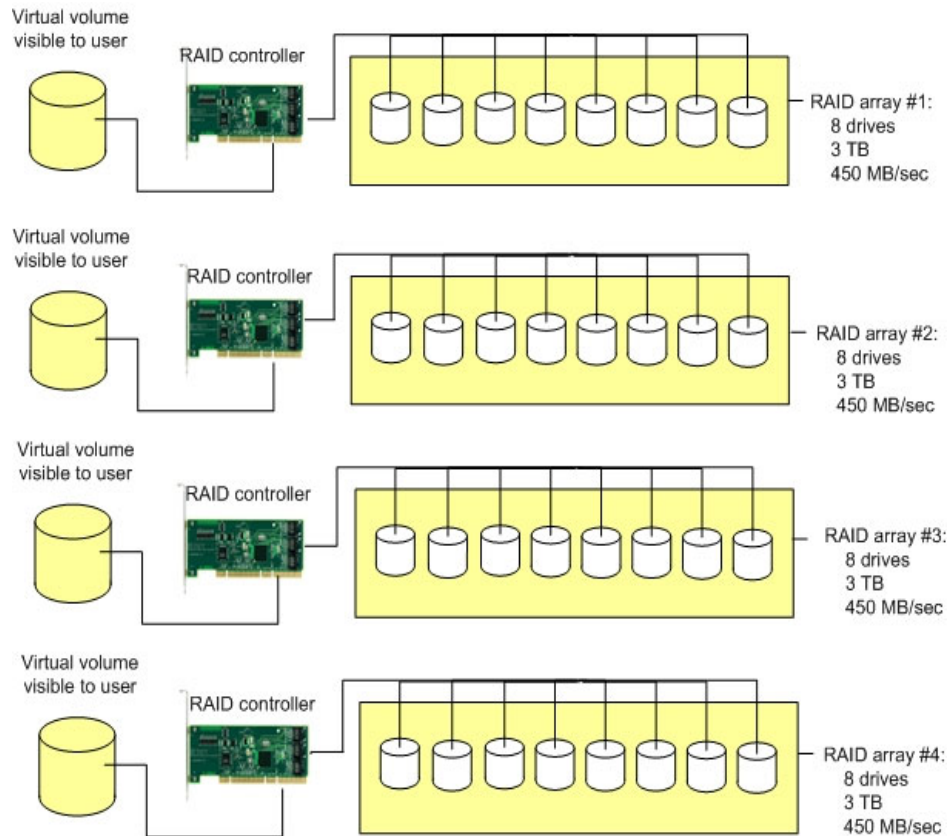


Figure 1: RAID Arrays without Spanning

In contrast, a server featuring RAID controllers and RAID software that do support controller spanning is shown in Figure 2. In this case, an array can span any or all of the controllers so that a single array can include any number of disks up the maximum number of disks supported.




Figure 2: RAID Arrays with Spanning

Not all RAID products in the market support spanning. In fact, the Broadcom® XelCore RAID software is the only product available from a major storage vendor that does support spanning. One unique architectural feature of XelCore RAID software that complements and empowers spanning is its operating system and hardware independence. XelCore is a software RAID stack that is state-driven and, as such, uses very few operating system-specific or hardware-specific commands. This relative independence results in a high degree of flexibility that provides easier porting to any hardware implementation (CPU, RAID-on-Chip, or controller card) or operating system such as Windows® or various Linux® flavors. It also means XelCore can be used in either a host-based or a controller-based RAID implementation. With host-based RAID, the RAID stack (the array configuration data and RAID algorithms) and data cache run on the host CPU and utilize CPU memory. With controller-based RAID, the RAID stack and data cache reside on the controller.)

Technical Benefits of Spanning: Large Array Support and Fast I/O

From a technical perspective, controller spanning provides two obvious benefits: the ability to create very large data arrays and a dramatic boost in I/O performance.

Large Array Support

Spanning allows very large data arrays to be created since an array can span all the disks attached to that server's RAID controllers. Figure 2 shows a configuration using Broadcom's RAIDCore™ BC4852 Serial ATA (SATA) RAID controller. As shown, up to four BC4852 (8-channel) controllers can be configured in a server to support a maximum of 32 SATA disks. With spanning, a single data array can be created to incorporate all 32 drives. If the disks in this configuration were all large 400 GB SATA disks, then an array of over twelve terabytes could be created (400 GB drives x 32 drives). For these configurations, Windows 2003 Service Pack 1 (Build 1222 or greater) or a version of Linux that is based on the 2.6 kernel or higher will support arrays greater than two terabytes.

Improved I/O Performance

There are two ways that controller spanning boosts I/O performance. One is fairly basic — the more disk spindles there are in an array, the higher performance. Since spanning allows more disks to be included in an array, it boosts I/O performance.

The second way spanning boosts performance is less obvious. If two RAID controllers are placed in motherboard slots that are on different PCI-X® bus segments, each controller has its own direct path into memory. As a result, a data array spanned across two controllers that are plugged into slots on different bus segments will have two paths into memory creating higher I/O performance. (Motherboard slot bus assignments are available from motherboard documentation.)

As proof of these performance boosts, independent reviewer Tom's Hardware Guide, tested four BC4852s in a single RAID0 array spanning 32 SATA drives and achieved the unprecedented results of over 1 gigabyte per second sequential reads and writes. (See *THG Pushes Broadcom's RAIDCore Storage Controller Envelope*, <http://www20.tomshardware.com/storage/20041006/index.html>.)

This review confirms the results obtained in Broadcom's internal testing shown in Figure 3.

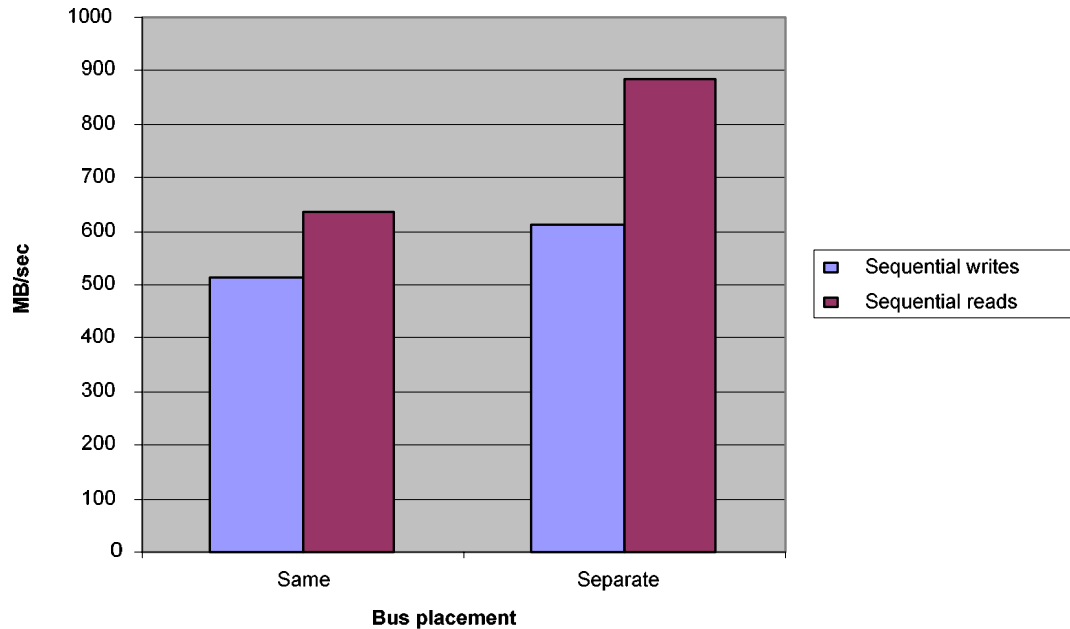


Figure 3: RAID 0 Performance Increases with the Number of Disks in an Array

To demonstrate the performance impact of motherboard slot selection, the I/O performance obtained with two BC4852 controllers on the same bus segment was compared to the performance obtained from the same two BC4852s, but on different bus segments (in all tests, a Gigabit Ethernet adapter shared one of the buses). As seen from the results in Figure 4, by selecting slots for the controllers that are on different bus segments results in a significant performance improvement particularly for RAID0 sequential reads since each controller is afforded its own direct path into main memory.

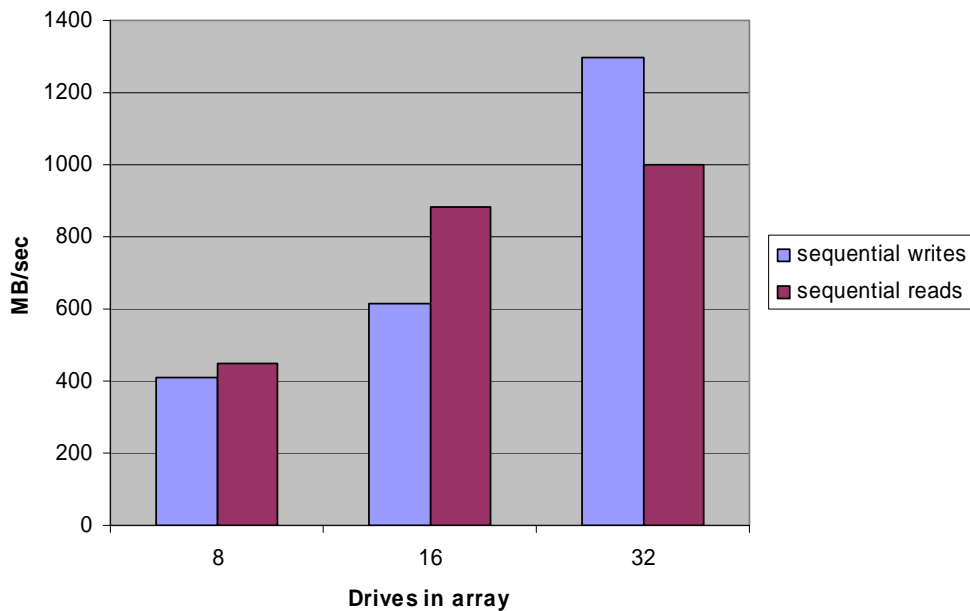


Figure 4: Performance Increases with Spanned Controllers on Different Bus Segments

The performance benefits achieved with controller spanning cannot be realized using alternative approaches to extending arrays across additional drives. For example, some RAID implementations offer port multipliers or expanders as ways of adding disks to a controller beyond those directly supported by the available channels. These are used to expand capacity cheaply; they do not improve maximum I/O throughput because all the disks are still on the same bus segment and the bus speed remains a limiting performance factor. Furthermore, the adaptor port is divided across multiple drives with multipliers and expanders, reducing the I/O bandwidth available to individual drives.

Business Benefits of Spanning: Storage Integration and Investment Protection

From a business perspective, spanning provides impressive benefits, which are also critical to improving storage ROIs strategies. The primary benefits are 1) the easy integration of new storage technologies, which leads directly to 2) the achievement of a previously unattainable level of storage investment protection.

Integrating New Storage Technologies

With controller spanning, new RAID controllers and attached disks can be easily integrated into an existing data array. In addition, when a server is upgraded with a new motherboard that has RAID-on-motherboard (ROMB), the existing data array can be expanded to include both the old disks attached to the RAID controller card and the new disks attached to the ROMB channels. Controller spanning enables this seamless inclusion of a new RAID implementation into an existing RAID array. The current implementations of RAID that can be integrated are ROMB and host bus adapter cards (HBAs).

A typical user scenario begins with an existing RAID data array created and managed using the XelCore RAID software and four channels of ROMB available on some motherboards (MBs), such as MBs with the four-channel HT1000 SATA-II RAID Southbridge chipset from Broadcom. Once the data demands of the server outgrow the four SATA drives, a new implementation of RAID, based on an HBA supporting XelCore and four-to-eight additional SATA drives, could be added to the system and seamlessly integrated into the existing RAID array using online capacity expansion (another XelCore feature that allows new disks to be included in an existing array without taking the array offline). Figure 5 shows these two RAID implementations integrated within a single array using XelCore's spanning feature.




Figure 5: Spanning Integrates Different RAID Implementations

Storage Investment Protection

Easy integration provides technological and staff investment protection. The protection of storage equipment investments is derived from the ability to seamlessly span existing arrays to include new RAID implementations without rendering obsolete the current RAID equipment such as the ROMB RAID controller in our example above. This storage investment protection also applies to HBAs. In this manner, Broadcom customers can purchase a BC4000 series RAID card today and then another card at a later point knowing that both can be easily integrated into the same array using XelCore's controller spanning. Furthermore, this protection benefit holds true even as new serial technologies come on line. That is, a BC4000 controller that supports new Serial Attached SCSI (SAS) disks can be added to an existing SATA array, enabling a smooth and easy SATA to SAS co-existence or migration.

The staff investment protection arises from the common RAID software. With the XelCore RAID stack, server administrators use the same management interface, as well as the same features and functions across the implementations. Since the software is the same, there is no additional training required and the investment protection extends to staff and training investments. In the absence of spanning and XelCore, there may be new management software introduced with its own learning curve and added training expenses, depending upon the product chosen for the new RAID controller.

Summary

Spanning puts an important tool into the hands of IT professionals tasked with optimizing and streamlining their storage. The technical and business benefits that make spanning such a potent weapon are straightforward:

Technical Benefits

- Larger data arrays - with spanning, a single data array can span all the drives connected to all the RAID controllers in a system.
- Better I/O performance - spanning boosts I/O performance by providing an array with more paths into memory and by allowing a greater number of disks to be included in an array, up to 32 with Broadcom.

Business Benefits

- Easy integration - when secure storage capacity growth is required, spanning makes integrating additional RAID controllers and disks into an existing data array easy.
- Investment protection - when data storage growth leverages existing equipment and IT knowledge, the storage ROI is greatly improved.

In closing, spanning supports more data and faster I/O without negating existing equipment investments: goals of any sound storage strategy.



Phone: 949-450-8700
Fax: 949-450-8710
E-mail: info@broadcom.com
Web: www.broadcom.com

Broadcom®, the pulse logo, Connecting everything®, the Connecting everything logo, RAIDCore™, and XelCore™ are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Windows® is a trademark of Microsoft Corporation. Linux® is a trademark of Linus Torvalds. PCI-X® is a trademark of PCI-SIG Corporation. Any other trademarks or trade names mentioned are the property of their respective owners.

BROADCOM CORPORATION
16215 Alton Parkway, P.O. Box 57013
Irvine, California 92619-7013
© 2005 by BROADCOM CORPORATION. All rights reserved.

RAID-CSpan-WP100-R