



WHITE PAPER

Extreme RAID:

**How STORM Xtreme Delivers *True*
“No Single Point of Failure” Data
Availability for Mission-Critical
Applications**

**BY
Dr. Tomlinson Rauscher,
Chief Operating Officer
Digi-Data Corporation**

How Fault Tolerant Do You Need Your RAID System?

RAID storage systems have always been regarded as inherently fault tolerant. In fact, most any RAID system available today offers some level of redundancy and data availability. For instance, typical RAID systems today use parity striping or data mirroring across multiple independent disks to withstand disk failures. They offer multiple, redundant Fibre Channel connections from the RAID controller to the disks and hosts to survive Fibre Channel connection or adapter failures. They also provide hot-swappability of many of the components within the system to maximize system uptime. As a result, this level of fault tolerance has become the accepted norm.

However, for enterprise or mission-critical applications like e-commerce sites, customer relationship management, transaction processing, financial reporting and operational applications, these environments need a higher level of data availability. More specifically, they need "No Single Point of Failure" RAID storage.

Historically, the difference between standard RAID systems and "No Single Point of Failure" RAID systems has been quite significant in terms of hardware complexity, management software, configuration and deployment, management and maintenance, and above all, cost. As a result "No Single Point of Failure" RAID storage systems have been unaffordable and too cumbersome to deploy and manage for most organizations. In turn, this has created a large gap in the market – a market demanding for the protection and peace-of-mind of "No Single Point of Failure" RAID storage, but without all the cost or complexity.

Introducing STORM Xtreme with *truly affordable* "No Single Point of Failure"

With STORM Xtreme, Digi-Data has dramatically narrowed the gap between simple fault tolerance and "No Single Point of Failure". Digi-Data STORM Xtreme is an advanced Fibre-Channel based RAID storage system that uses a unique modular architecture that allows it to scale data availability, capacity or performance by simply adding additional modules to the system such as Disk Controller Modules (DCM) and Disk Enclosure Modules (DEM). For instance, to create a "No Single Point of Failure" STORM Xtreme system, you simply connect a second DCM to a system with multiple DEMs, and configure the system as either Active-Passive or Full-Speed Active-Active™. With either of these configurations, STORM Xtreme systems can survive failure of any or multiple of the following components:

- Disk Controller Module
- Disk Enclosure Module
- Disks
- Fans
- Power Supplies
- Power Cords
- Fibre Channel

STORM Xtreme's "No Single Point of Failure" configurations use dual interconnected Disk Controller Modules with redundant paths to all system components including hosts, disks, disk enclosures, power supplies and the other controller. This is a much different approach than traditional configurations that use dual controller boards within a single

controller unit. In these configurations, the controller boards share a common backplane within the same disk controller unit, making these units more susceptible to failure than STORM Xtreme.

There are two basic 'modes' for configuring STORM Xtreme for "No Single Point of Failure" – Active-Passive and our unique and patented Full Speed Active-Active™. Each configuration offers unique advantages. To help you understand how these configuration modes compare, this paper will go into detail on each one. However, before we can do this, it is important to understand the STORM Xtreme architecture, its physical characteristics and its operational capabilities.

STORM Xtreme Primer

STORM Xtreme is an enterprise-class RAID storage system engineered to deliver uncompromising performance, scalability, data availability, and ease of management. It is an open, Fibre Channel based storage system, it combines advanced, modular technology with simplified management to deliver increased investment protection while lowering total cost of ownership.

STORM Xtreme has 8 fully independent 2 Gb Fibre Channel connections, providing it inherent flexibility for connecting to hosts and disks. It can be directly connected to as many as 7 application servers, NAS servers or SAN switches, or it can scale up to 840 drives for over 120 TB of raw storage – or any configuration in between. It also uses a dual special purpose processor architecture and multi-bus design to accelerate RAID operations and improve read and write performance.

Even without its "No Single Point of Failure" configurations, STORM Xtreme offers very high data availability and worry-free maintenance. It uses a fully redundant, hot-swap design and continuously monitors all subsystem components. It can withstand failures to disks, power supplies, fans, even entire Disk Enclosure Modules without interrupting system operation. STORM Xtreme also offers advanced RAID levels such as RAID 30 and 50 to concurrently boost performance and availability with two dimensional disk striping – both across and along channels. In addition, STORM provides graceful recoveries from soft and hard disk failures with features like automatic drive failure detection and reconstruction.

Equally important to its technical capabilities is its ease of management. STORM Xtreme simplifies configuring, deploying, and managing sophisticated RAID systems using an innovative embedded RAID manager called RECON. This eliminates the need to buy or install any software, and uses an intuitive browser user interface to provide storage managers with unparalleled control, flexibility and visibility.

Comparing STORM Xtreme's Dual Controller Configurations

With that overview of STORM Xtreme, now let's compare the dual controller configuration modes.

Active-Passive Configuration

Active-Passive is a dual controller configuration where two controllers provide full redundancy to all disks, disk enclosures, and Fibre Channel host connections. In an Active-Passive configuration, the primary (active) controller services all host I/O requests

and performs all data transfers, while the passive controllers remains alert to the active controller's status using bi-directional heartbeat communications. Typically, the available space in the RAID array is divided up into an arbitrary number of logical units (LUNs). The capacity of each LUN can be spread across multiple controller Fibre Channels and disk drives. In this configuration, both the active and passive controller know the logical volume configuration.

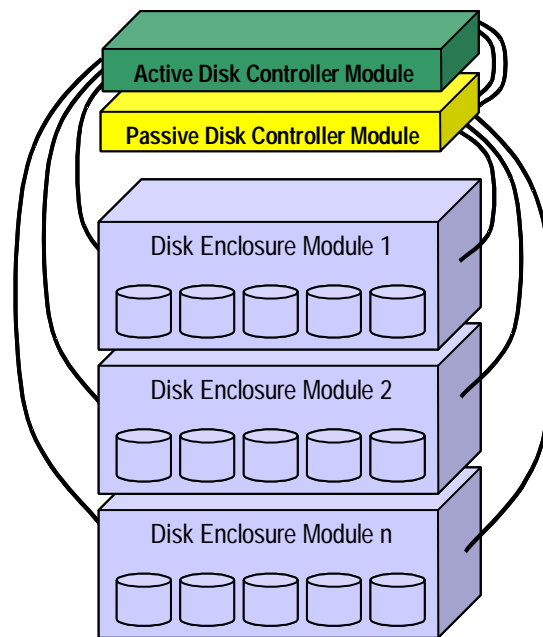
With the STORM Xtreme Active-Passive configuration, while the active controller is performing I/O and data transfer operations, the passive controller continually monitors the "heartbeat" of the active controller. It also runs internal diagnostics to assure its readiness and responds to the active controller with its own heartbeat. This continuous monitoring ensures uninterrupted operation in the event of a controller failure.

In the event of a primary controller failure, the passive controller automatically and seamlessly assumes I/O and data transfer activities without interrupting system performance or operation. It is important to note that one of the advantages to Active-Passive is that there is no degradation of performance when one controller fails or is taken off-line for maintenance.

As you can see in Diagram 1, the Active-Passive system provides full redundancy to all disks and host connections, as well as dual redundant connections between controllers on a dedicated private interface for intercontroller communications, called the DIGINET ports, which are separate from the controller's Fibre Channel connections.

In the STORM Xtreme Active-Passive configuration, the active controller uses write-through caching, which performs a write to the cache on the Fibre Channel disk (which passes the data directly to the Fibre Channel disk) before sending notification back to the originating application or host.

Diagram 1: Active-Passive



Some RAID systems use write-back caching in their Active-Passive dual controller configurations in an effort to improve data transfer rates (which requires the primary controller to write data to the second controller's cache to get authorization). However, having data stored in controller cache causes other concerns and adds complexity in both controller design and the interconnection between the two controllers. With write-back caching, the two controllers need to continually synchronize the information in cache, so no data is lost during the transition of control from one controller to the other. In effect, the primary controller needs to do a write to the cache of the second controller before the primary controller can send notification to the originating application or host that a write is complete. Hence, the overhead associated with this additional

synchronization reduces the performance of the system. In addition, since many of the systems that use write-back cache also have dual controller cards running on a common backplane, there is much higher risk of data loss in the event of a backplane or power failure. To compensate for this, systems that use write-back cache have needed to add the complexity of battery-backed cache as a safeguard. This only serves to introduce additional points of failure in the system since batteries don't last forever, and the cache using the battery backup is not the only cache on the system containing volatile data. For instance, Fibre Channel chips on the controller have memory that if lost will cause an "unknown state" from which the controller cannot recover.

Some other redundant controller solutions use write-back with additional cache writing synchronization with the other controller. However, in dual controller configurations, it takes just as long to perform a write to the cache on the passive controller (for write-back configurations) as it does to perform a write to disk cache (for write-through configurations). In addition, added complexity and points of failure created by write-back caching in terms of battery backup, software and configuration, make write-through caching a simpler, safer and faster solution for most environments. In turn, write-through caching improves data integrity, reduces points of failure, and reduces calculation and data synchronization overhead on both controllers.


Traditional Active-Active Configuration

As you can see from Diagram 2, in the traditional Active-Active configurations the disk channels are "hubbed" together so that both controllers can access all the disks.

In normal operation, traditional Active-Active controller configurations compete for disk channels which cause conflicts that degrade system performance. Typical performance gains of traditional active-active over active-passive configurations peak at about 40% and decline for larger block sizes. Chart 1 at end of this white paper illustrates this point.

In a traditional Active-Active configuration, both controllers are working concurrently to serve host I/O requests and transfer data. In this mode, when both controllers are operating normally, the system is theoretically able to handle twice the workload and traffic, doubling the speed of the system compared to the Active-Passive configuration. However in practice, performance increases are much less significant.

Diagram 2: Traditional Active-Active



In the event of a controller failure in traditional Active-Active configurations, the remaining controller automatically assumes responsibility for handling all I/O requests

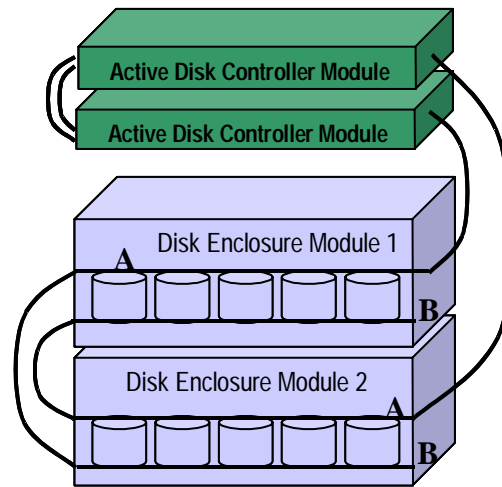
and data transfer. Once the failed controller is replaced, the controllers will automatically read the configuration of drives and LUNs in the system, and return to normal operation.

Full Speed Active-Active™ Configuration

Full-Speed Active-Active™ is an innovative dual controller configuration patented by Digi-Data and is available only on the STORM Xtreme system. It offers all the “No Single Point of Failure” protection of dual controllers, yet surpasses traditional Active-Active configurations.

As you can see in Diagram 3, STORM Xtreme's Full Speed Active-Active uses an innovative wiring design to provide full redundancy to all disks and hosts. From each controller, the Fibre Channel first connects to the “A” ports on each drive in one DEM, then continue on to connect to the “B” ports in the second DEM, and so on. By doing so, each controller uses the dual ports on the disk drives themselves to create redundancy. This configuration eliminates the channel conflicts of traditional Active-Active configurations, enabling the system to run at full operational capacity. With STORM Xtreme, with up to 7

Diagram 3: Full Speed Active-Active™



available disk connections per controller, this equals the bandwidth of 14, 2Gb Fibre Channels or 2,800 Gb per second throughput capacity. Chart 1 at the end of this white paper illustrates the performance advantages of the Full Speed Active-Active vs other dual controller configurations.

In normal operation, the NAS server, application server or SAN switch will use both controllers for accessing and carrying out data requests. If one controller is unavailable, the request is simply handled by the second controller. This is also how the system handles a controller failure. The surviving controller automatically takes on the identity of the failed controller to handle all its operations. Compared to traditional Active-Active configurations that use write-back caching requiring continuous intercontroller cache writes to synchronize controller cache, STORM Xtreme uses write-through caching. As detailed earlier, this cache handling approach is a simpler, safer and faster solution for most environments. With STORM Xtreme's Full Speed Active-Active configuration, the dual controllers use the fully redundant DIGINET dedicated private interface to monitor each controller's health and operation.

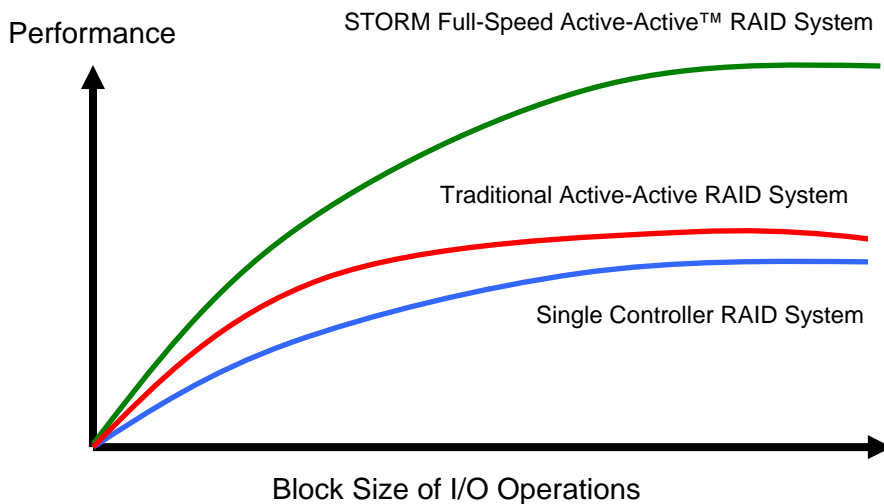
Full-Speed Active-Active is unique to the Digi-Data STORM Xtreme system. It is possible only because of the modularity and innovative capabilities of the system components. For instance, the Disk Enclosure Modules were designed to have multiple ports enabling more flexible Fibre Channel loop configurations and daisy-chain capability. The Digi-Data STORM system takes advantage of the dual-ported Fibre Channel drives in a unique way. Then the Disk Controller Module with its integrated RAID manager has the intelligence needed to understand what data is where within the

RAID system without requiring additional management software or special drivers on the network.

Table 1: Dual Controller Comparison

	STORM Active-Passive	Traditional Active-Active	STORM Full Speed Active-Active™
Data Availability	Very High	High	Very High
Full Operational Performance	Medium: I/O processing and data handling speed of single controller	High: I/O processing and data handling speed of dual controllers, but with shared disk channels	Very High: I/O processing and data handling speed of dual controllers, plus no competition for disk channels
Controller Failure Handling	Passive controller automatically and seamlessly assumes handling of all I/O and data handling requests from host(s)	Surviving controller assumes all I/O and data handling requests from host(s)	Surviving controller assumes all I/O and data handling requests from host(s)
Performance with Failed Controller	No degradation in performance	Performance is reduced, and operates at the full performance level of a single controller	Performance is reduced in half, and operates at the full performance level of a single controller

Chart 1: Performance Comparison of Dual Controller Configurations



Conclusions

For many organizations, a growing concern for data protection, business continuity and data availability has accelerated the need for RAID storage systems that can deliver *true*

"No Single Point of Failure". At Digi-Data Corporation, we have engineered our STORM Xtreme RAID systems to offer the highest levels of data availability at a much lower Total Cost of Ownership (TCO) than other redundant controller systems which have single points of failure.

STORM Xtreme offers multiple "No Single Point of Failure" configurations to meet the unique requirements of each environment. From the full redundancy fail-over protection of the Active-Passive configuration to the extreme performance of Full Speed Active-Active™ configuration, each approach offers distinct advantages.

If you would like more information on these configurations, or need help determining what the right configuration is for your needs, please contact Digi-Data at 410-730-6880, or contact your authorized Digi-Data reseller.

This white paper was provided as a courtesy of:

Digi-Data Corporation
8920-D Route 108
Columbia, MD 21045
410-730-6880
www.digidata.com