

# Whitepaper



## **Selecting the Right SATA Storage Array for Business Critical Applications**

A checklist on profiles and benefits of Serial ATA-based storage arrays that meet the availability requirements of business critical applications

*May 2004*

## **IMPORTANT INFORMATION**

© 2004 iQstor Networks, Inc. The information contained in this white paper does not constitute a contractual agreement with iQstor Networks and may be modified at any time without advance notice. This white paper is supplied on an "as is" basis with no warranty and no support. iQstor Networks makes no express warranty, whether written or oral with respect to this white paper.

No trademark, copyright, or patent licenses are expressly or implicitly granted (herein) with this white paper.

## Overview

As the volume of business information explodes, companies are struggling to find cost-effective solutions to store their valuable data. The combination of exponential data growth, tight storage budgets and the complexity of data management has, in the past, required companies to choose between expensive, enterprise-level solutions that offer redundant, and complex storage networks (with associated high management costs) and simple RAID storage solutions that provide lower hardware costs but offer no real paths for network growth or consolidated storage network management.

That is, until now. SATA storage solutions are maturing to the point where they can deliver enterprise-level functionality and management advantages at a cost-effective price. Combining SATA and Fibre Channel solutions, the triple punch of data growth, budgetary restrictions and data management complexity can be effectively managed if storage administrators employ two concepts in planning for growth.

## Not all Information is Created Equal

---

First, it is important to realize that not all information is created equal – some data is simply more valuable than other data – and the ability to understand the change in value of its lifecycle, protecting when needed, allows the smart application of budgetary resources. Secondly, the ability to deploy multiple classes of storage devices (or storage tiers with varied costs, access patterns and response times) can greatly increase the economies of storage utilization, allowing companies to implement more inexpensive hardware for secondary data.

Let's look at how this can be applied to support business. Primary storage is used to store mission-critical information that has the highest value and importance. In today's SAN environments, this data is usually stored on Fibre Channel disks that provide the highest availability and performance. Secondary storage is used to store business-important information that does not require the level of protection, availability and performance of mission-critical data. Generally, this would include operational data such as email, departmental file data, reference information, certain CRM/ERP data and multimedia content.

In the past, there was not an ideal solution for secondary storage. Storage solutions with enterprise-level robustness and performance were expensive because of the Fibre Channel drives. On the other hand, systems with lower cost disk drives did not offer sufficient functionality or reliability. Companies made tradeoffs and compromised among unfavorable choices.

In the past twelve months, that has changed. SATA storage solutions are now available for environments that need lower cost storage, where administrators don't want to give up the availability, protection and functionality found in Fibre Channel based storage systems. With the wider availability of SATA solution, administrators need to weigh the characteristics of enterprise-class storage solutions to see whether SATA solution functionalities can match their older, and much more expensive, brothers.

---

**Weighing the  
Features and  
Functionality:  
A Checklist**

The following sections describe characteristics of enterprise-class storage solutions that have returned value and cost savings to large and small data centers. When reviewing SATA solutions for your implementation, having these characteristics can mean the difference between a system that provides a mere stopgap solution and one that provides a flexible and robust growth path for continuous return on investment.

---

**Reliability,  
Availability  
and  
Serviceability**

In today's mission-critical computing environments, service level agreements (SLAs) commit IT organizations to measurable service metrics such as application performance, end-user response time and system availability. Since storage can directly impact the performance, availability and functionality of most applications, storage plays a vital role in an IT organization's ability to deliver IT services.

As companies consolidate storage with the latest technology, this has dramatically increased the costs of any storage system failure. It is therefore critical that even in the face of component failures, such as disk drive or power supply, storage systems be able to deliver high levels of performance and availability without compromise. A robust storage solution that delivers these uncompromised levels of service must have reliability, availability and serviceability (RAS) features inherent to its architectural design.

Data centers, large and small, require solutions that provide uncompromised levels of service, backed by quality support and a migration path to new technology that protects current investments. They also require storage solutions with robust architectural designs that withstand potential failures without impacting data availability. Components of these RAS features should include:

- No single point of failure. All critical components (blowers, power supplies, storage controllers, and drives) are redundant

- Redundant I/O paths to SATA drives
- High performance RAID for protection against disk failure
- Proactive error detection and response, e.g. memory scrub and cache vaulting
- Non-disruptive operating environment upgrades
- Remote “call home” capability to enable prompt field service response
- Enclosure environmental monitoring
- Modular design enabling non-disruptive component replacement

### Data Services for Data Asset Protection

---

While RAS features and redundancy techniques are vital and necessary components for a storage solution, they only protect against hardware component failures, and for this reason, are insufficient. Beyond hardware failures, disruptions to data access can also come from human errors, data corruption or natural disasters. And, because companies increasingly rely on information access as the lifeblood of their business, disruptions to data access puts the entire enterprise at risk. This risk can be quantified in a number of ways, including lost revenue and lost business opportunities.

Storage systems -- sometimes seen as a container of corporate data assets -- have emerged as a critical resource that has evolved beyond a traditional capacity model and should be view as more than just a receptacle.

With today’s advancements in network technologies, storage has become a platform for data services to manage and mitigate risk to data access. A tightly integrated combination of storage hardware, embedded intelligence and data services are the means to practical risk management and data asset protection; a storage solution that supports business continuance, backup and restore, and capacity management helps to ensure that a company’s core business activities continue to function with minimal interruption, regardless of catastrophic events, malicious attacks, upgrades or common everyday mistakes.

For companies looking to suitably protect their information, data availability and risk management features of storage systems need to include:

## Common Management Tools

- Mirroring and remote replication to enable data to be replicated within and outside of a data center to protect against various threats and failures. Top disaster recovery and business continuity plans leverage features such as local mirroring and remote replication to copy and migrate data and mitigate the loss from any kind of outage.
- Snapshot services to provide the ability to recover from logical errors (such as corrupted data or deleted files) by preserving data at multiple points in time. In addition, snapshots also serve to increase data availability by maximizing the limited backup windows of today's 24x7 business applications.
- Virtualization of storage to enable provisioning of storage on demand, thus increasing system availability. By also enabling the consolidation of storage, management complexity and costs are reduced while overall capacity usage is enhanced.
- Capacity management services to help avoid costly downtime due to unforeseen changes in storage capacity requirements, such as file system growth. By monitoring and automatically expanding storage space from available capacity, downtime due to physical disk upgrades, rebuilding and reformatting of operating system is eliminated.

---

Storage management software serves as the foundation for networked storage, enabling IT organizations to tie together all their storage resources, to monitor them and to control them. To maximize return on investment (ROI), it is important that the configuration, analysis, and control be provided by a central, easy-to-use graphical user interface (GUI) that can be securely and remotely accessed anywhere.

In a tiered storage environment, maximizing ROI means that SATA-based storage systems must be tightly integrated with primary tier storage to increase capacity utilization and overall efficiency. From device discovery to volume management, capacity planning to data management, lower tiers of storage should be treated as part of the entire storage hierarchy in order to bring the highest possible return. The tight integration of SATA based storage into the storage infrastructure is best leveraged by common management tools which increase interoperability and data movement between different storage tiers.

---

## Investment Protection

As companies invest in new storage systems and technologies to meet the needs of their growing IT infrastructure, it is essential that they deploy storage solutions that will evolve with their business needs. This is important since deploying a storage solution developed with long-term product continuity in mind will return the highest level of investment protection. Companies can leverage storage solutions to maximize ROI, protecting their initial investments in hardware, software and training.

Key areas that such a storage solution should address include:

- Modular interfaces which offer multiple connectivity choices, e.g. Fibre Channel or iSCSI and the ability to migrate between protocols with data intact.
- Modular architectures which enable performance upgrades to newer and higher performance controllers. A seamless upgrade with data intact eliminates the downtime and risk of data migration.
- Hardware and software compatibility across storage platform generations which eliminates the burden of retraining IT staff.
- Compatibility and data movement across hardware generations to enable older systems to be utilized at secondary disaster recovery sites as new systems are used for production.
- Common reusable components (such as storage enclosures, disk drives and power supplies) to preserve majority of hardware investments across storage platforms.

## Modular Storage

---

Modular storage uses building blocks of controller and disk drive shelves to allow very customized and scalable solutions. Choosing a modular architecture offers a "build as you grow" approach, such that as more capacity is needed, additional drive shelves can be added non-disruptively. This approach lowers initial acquisition costs for cost-conscious environments such as SMBs (small to mid-sized businesses). At the same time, modular storage can also allow demanding environments that require greater I/Os, throughput and capacity, to easily scale as requirements grow.

The modular storage concept also applies to the data services components – the intelligence – of an intelligent storage system. The ability to license data services components on a usage and capacity basis makes storage system's acquisition costs more affordable while providing a path for future growth requirements.

For example, companies can initially license virtualization services to enable storage consolidation for a starter SAN. Later, snapshot services can be enabled as part of a centralized SAN backup strategy. Eventually as an organization invests in a secondary recovery site, remote replication can be licensed. A modular storage approach such as this enables affordable, scalable and customized solutions to accommodate business growth in a controlled fashion.

## Product Stability

---

As IT organizations implement new technology, stability of the product is of the highest concern. While no test environment can reproduce all customer uses and all application scenarios, new storage architectures and storage solutions see early field problems decline over time after the product matures in data centers. This product refinement is based on two factors: the complexity of the product and the amount of resources a vendor can apply.

One way vendors can mitigate risk factors with the introduction of their SATA-based storage systems is by using the same controllers as are in their Fibre Channel-based storage systems. Some minor hardware changes are required to convert from a Fibre Channel protocol chip to a SATA backend interface, and additional code is required to address characteristics specific to the SATA interface (e.g. SATA protocol and error recovery). However, after these items are addressed, base functionality for SATA products can be the same as for Fibre Channel products that have been shipping for a much longer time. By utilizing the same controllers as in the Fibre Channel storage systems, new SATA-based products can avoid the stability issues inherent in a brand new architectural implementation that has been mentioned above, bringing about two substantial advantages: giving the customer much greater confidence in the new product, and bringing lower development and support costs for the vendor.

## Native Command Queuing Support

---

Unlike most electrical components in modern systems, hard disk drives are electromechanical devices which are hybrids of electronic and mechanical components, and are hampered by the inertia of their mechanical components.

Accessing information stored on mechanical storage media can have a negative impact on overall system performance due to the overheads from seek latencies and rotational latencies. Seek latencies represent the time it takes to mechanically position and settle the read/write heads from the outer diameter to inner diameter, over the correct

track. Rotational latency is the time that it takes for the correct block of data to rotate under the head after the head has been positioned on the right track.

Because applications rarely request data in the order that it is written, the ability to minimize mechanical overhead in a drive can significantly enhance overall performance.

Native Command Queuing (NCQ) is a widely anticipated feature introduced in the Serial ATA II extension to the Serial ATA 1.0 specification. This command protocol allows multiple commands to be outstanding within a drive at the same time.

Support for NCQ allows drives to dynamically reorder command execution and optimize the workload while minimizing mechanical latencies. It follows that intelligently managing the sequence of mechanical motions not only improves performance, but it can reduce mechanical wear and enhance endurance. A storage solution that supports and leverages the NCQ feature in SATA drives delivers higher overall system performance, as well as greater reliability in the storage infrastructure.

---

The Storage Network Industry Association (SNIA) has developed and released the first set of specifications to provide common models for devices and a common protocol for managing storage networking devices. The SMI-S (v1.0) industry standard defines a model for discovering and managing arrays, switches and HBAs. SMI-S also makes it possible to eliminate the need for vendor-specific APIs and to do away with the myriad of management protocols used today to manage storage environments.

The benefits to end users from the SNIA initiative are significant. End users benefit through an increased flexibility to choose both hardware and software vendors. Vendor lock-in is eliminated because SMI-S makes it easier to implement heterogeneous solutions. Implementation times are reduced as new storage devices seamlessly integrate into the management interface, and end users will be able to realize both management and operational cost savings by managing heterogeneous solutions as though they were from a single vendor. Given all these benefits, support for SMI-S should therefore be important to end users as they make purchasing decisions regarding storage and management solutions.

## Summary

---

SATA technology is an ideal solution for the workload, capacity and cost constraints of secondary storage. When properly implemented, SATA technology can be used as the foundation of a cost-effective solution which brings great value to the data center. A SATA-based storage system which is enterprise-capable would integrate RAS features, embed storage intelligence for application data protection, address investment protection concerns and interoperate with primary storage through common management tools.

In our work developing a new generation of intelligent data storage solutions, such as iQstor's iQ1200 SATA Storage System, we continue to be committed to incorporating intelligent, enterprise-level functionality into our products – in fact, iQstor's SATA and Fibre Channel storage systems deliver identical data services features to storage administrators. Using this approach allows administrators to select an intelligent system based on their individual data center needs, and tailor a primary or secondary data solution that can be implemented to fit needed performance characteristics within specified budgets.

With the exponentially increasing volume of business information, the use of SATA storage solutions with intelligent data services can bring a cost-effective answer to address the enterprise-level needs of a wide range of businesses. With intelligent SATA solutions, business can address the exponential data growth that plagues data centers around the world. Using intelligent SATA, business can deploy storage networks with enterprise-level functionality at a reduced price. And with intelligent SATA, business can protect information in tiered approaches that bring intelligent data protection to address business continuity and disaster recovery requirements with a solution that embraces a strategy of long-term product continuity combined with the highest levels of investment protection.

Compliments of:  **iQstor**



3835R East Thousand Oaks BLVD. #315  
Westlake Village, CA 91365  
Tel 877.230.2837 / Fax 805.435.2500 / [www.ess-direct.com](http://www.ess-direct.com)

# Whitepaper