Seagate
We turn on ideas

TP-544 • From: Global Product Marketing • September 2005

Technology Paper

# Nearline Goes Native: Tiered Storage With FC Disc Drives

### Introduction

The advent of Fibre Channel (FC) prompted a new network paradigm: the storage area network (SAN). With every server able to see every storage device in an FC SAN, data accessibility skyrocketed-and network efficiency with it. Now tiered storage has rewritten the rules for FC storage efficiency, complementing established online and offline storage platforms with a third tier of nearline storage to cost-effectively house the rarely accessed but fundamentally essential data that's flooding the enterprise.

The ideal nearline disc drive combines economical capacity, seamless compatibility with existing network infrastructures, and nearline-class reliability. A growing number of businesses have turned to high-capacity, low-cost Serial ATA (SATA) drives for nearline duty in their FC SANs. However, such devices require a separate SATA infrastructure or FC-to-SATA bridge (eliminating advanced capabilities of the FC protocol), and are not always designed for rigorous nearline workloads.

The solution is high-capacity, low-cost disc drives purpose-built for Fibre Channel networks. As one would expect, these nearline-ready drives deliver the exceptional reliability (1.0 million hours MTBF) that nearline duty demands. But more importantly, native FC nearline drives make full use of the FC protocol, enabling tighter integration of online, nearline and offline storage tiers. Such interface rationalization streamlines FC SAN infrastructures, reducing costs, easing management and improving performance.

### Tiered Storage Revolution

The traditional storage model acknowledges only two types of data storage: online and offline. Online describes high-availability, transactional data that demands the performance and reliability of enterprise-class disc drives, while offline refers to archival data that is infrequently accessed and stored in libraries of high-capacity tapes.

But what of the vast quantity of data that falls between these two categories? The following applications all involve data that is below the threshold of high-availability, mission-critical storage, but must still be readily accessible to multiple users:

- File serving
- Fixed-content data
- Disc-to-disc backup
- Bulk storage
- Short-term archiving

In addition, regulatory compliance (Sarbanes-Oxley, HIPAA) stipulates rapid retrieval of enormous quantities of financial and medical data. An Enterprise Strategy Group study found this "in-between" information will soon comprise the majority of enterprise data.

In the past, there were only two options for this plethora of data: online storage (optimized for performance, not capacity) with its relatively high cost-per-Gbyte; or offline storage (optimized for capacity, not accessibility) with its slow, labor-intensive data retrievals. In both cases, the net result was gross inefficiency, wasting either costly hardware or overtaxed IT staff resources.

Tiered storage eliminates such inefficiency by introducing a new class of nearline disc drives. Melding high capacity, low cost-per-Gbyte and easy integration into existing enterprise infrastructures, nearline-ready drives fill the gap between performance (15K disc drive) and archival (tape) solutions, enabling comprehensive tiered storage strategies that ensure optimal cost/performance for every type of data.

### Nearline Drives: Workload

While nearline applications don't require the high level of data availability and input/output per second (IOPS) demanded by online applications, they do share the need for around-the-clock data accessibility. And though nearline data activity is far less frequent than online activity, both are highly random in nature. These random reads/writes force drive heads to rapidly and repeatedly traverse a drive's discs.

To deliver the nearline-class reliability standard of 1.0 million hours MTBF, Fibre Channel nearline-ready drives are specifically designed to withstand the rigors of random reads/writes and 24x7, always-on operation. In contrast, the typical 600,000 hours MTBF rating of desktop-class SATA drives is obtained in the mild environment of sequential reads/writes and 8x5 power-on hours, and thus has no relevance when considering the use of such drives in nearline applications.

To be sure, SATA nearline-ready drives are available that achieve nearline-class reliability of 1.0 million hours MTBF, albeit in less demanding SAS environments. SAS-based tiered storage typically supports point applications, characterized by more intermittent workloads and fewer users than FC networks. But even if reliability isn't an issue, SATA nearline-ready drives still fall short due to poor integration in FC SANs.

### Nearline Drives: Integration

There are three fundamental ways to add SATA-based nearline storage to an FC SAN, all with significant drawbacks:

- Deploy a SATA infrastructure, separate and discrete from the FC network. This entails the redundancy (and corresponding expense) of qualifying, purchasing, inventorying and maintaining a range of SATA parts/components.
- Install FC-to-SATA conversion bridges to accommodate SATA drives within an FC SAN. This obviates the need for a separate SATA infrastructure (and its associated costs), and for that reason this configuration is preferable.
- Utilize the upcoming FC-SATA tunneling protocol, which will eliminate the need for FC-to-SATA conversion bridges altogether. In all other respects this tunneling method will share the disadvantages of bridged FC-to-SATA systems (below).

Not surprisingly, the separate SATA infrastructure in the first approach cannot avail itself of the advanced capabilities the FC protocol provides. Unfortunately, the very same limitation holds true for the bridged FC-SATA system in the second approach.

For example, the bridged system loses Fibre Channel's error recovery capabilities, advanced multi-host command queuing features and the ability for a single host to simultaneously communicate with multiple drives. (Also lost is the ability for a single drive to simultaneously communicate with multiple HBAs). All of the commands in a SATA drive's queue must originate from a single controller-a SATA drive recognizes only one host at a time.

To enable a single host to concurrently talk with multiple SATA drives, the bridged system would require a multichannel HBA. This device allows the host to establish individual channels with multiple drives. Adding more electronics and another piece of hardware to the system, this approach brings capital and maintenance expenses that undercut the cost savings associated with lower-price SATA drives.

Finally, the SCSI-based storage management, diagnostic and application software that so effectively drives the day-to-day functionality of FC SANs is effectively wasted in a SATA-based nearline environment.

## Native FC Nearline Advantage

Unlike SATA drives, native Fibre Channel nearline drives extract maximum value from a Fibre Channel SAN by taking full advantage of the FC protocol. This in turn enables FC tiered storage solutions to deliver optimal efficiency and cost-effectiveness, with none of the compromises inherent in bridged SATA solutions. As shown below, the advantages of deploying native FC nearline drives are comprehensive:

- Integration

  Native FC nearline drives integrate more easily with FC SAN infrastructures, without the added expense and management complexity of a separate SATA infrastructure, protocol bridges and an extra interface (beyond SCSI and Fibre Channel). No hardware (cabinets, carriers, power supply) or software (drivers, diagnostics) changes are needed.
- Scalability

  An FC HBA proactively establishes communications with one FC drive after another, issuing read commands to them and then feeding data to their cache memories. When one cache is full, the HBA retrieves data from another drive in the Fibre Channel network. Hence the HBA keeps multiple FC drives working concurrently, maintaining throughput as drive counts increase.
- Performance

  In pure Fibre Channel deployments, an HBA's ability to manage multi-drive queues greatly increases throughput for data-intensive archiving applications. When SATA drives are bridged to an FC infrastructure, overall performance is limited to the speed of a single drive; the host can only talk to one drive at a time.
- Availability

  Fibre Channel nearline drives employ dual ports to enable failover capability in high-availability network environments. SATA nearline drives only offer single ports; should the drive's HBA fail, no backup HBA connection is possible.

  Note that a SATA Port Multiplier can enable active/active connections between a single drive and two hosts, thus providing failover capability. But this approach requires additional electronics (and space) on either the backplane or an interposer board, driving up costs accordingly.
- Simplified Management

  Native FC nearline drives share common management systems with the rest of the infrastructure. When SATA drives are deployed in a Fibre Channel infrastructure, subsystem management capabilities are reduced to the level supported by SATA technology.

| Summary: The Native FC Nearline Advantage | | |
|---|---|---|
| | **Native FC Nearline** | **Bridged or Discrete SATA Nearline** |
| **Integration** | No hardware or software changes needed | Requires costly bridge or separate SATA infrastructure |
| **Scalability** | Throughput increases as drive count increases | Throughput unchanged as drive count increases |
| **Performance** | Simultaneous communication w/multiple drives maximizes throughput | One-at-a-time communication w/single drives limits throughput |
| **Availability** | Dual ports enable seamless failover protection | Single port lacks failover protection (unless PM added) |
| **Simplified Management** | Leverages full capability of FC management systems | Limited to management capabilities supported by SATA |

## Conclusion

Tiered storage is built on a fundamental concept: Use the right tool for the job. By employing storage devices optimized for their specific applications, tiered storage delivers greater efficiency, performance and reliability. In that same manner, Fibre Channel tiered storage is most effective when it utilizes nearline storage devices optimized for FC networks-native Fibre Channel disc drives. Complementing rather than complicating FC SANs, FC nearline drives help maximize the value of tiered storage solutions.

**AMERICAS** **Seagate Technology LLC** 920 Disc Drive, Scotts Valley, California 95066, United States, 831-438-6550
**ASIA/PACIFIC** **Seagate Technology International Ltd.** 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
**EUROPE, MIDDLE EAST AND AFRICA** **Seagate Technology SA** 130-136, rue de Silly, 92773 Boulogne-Billancourt Cedex, France, 33 1-41 86 10 00


Publication Number: TP-544, September 2005