

## Scalable High-Availability Libraries

In the world of business, information has become the metaphorical bullion with which organizations have built their empires. Whether their riches are locked in a safe at the downtown bank or within the protective walls of today's tape libraries, the objective is still the same: everything needs to be there when its needed most.

Today, information provides larger challenges than gold prospecting did over a century ago, however, key parallels still exist. First, as riches grow, the rainy day repository in which they are locked must accommodate the expansion. Second, to minimize the exposure to loss, the bank must always be open when a large withdrawal is needed quickly.

For the enterprise of the new millenium, a reliable, scalable and highly available secondary storage solution is paramount to protect the information assets of the organization.

### Benefits of a Scalable, High-Availability Tape Library

Beyond obvious business benefits, including improved availability of information and faster recovery in emergency situations, a scalable, high-availability library offers the following functional advantages:

*Scalable, high-availability libraries provide a number of important benefits to enterprises including decreased downtime and a wide range of expansion capabilities*

- **Reliable Unattended Backups:** Automated libraries were introduced to reduce the cost of an administrator swapping tapes in and out of tape drives. Today, a highly reliable library is able to increase the life of tape drives and process massive amounts of data without any human interaction, freeing IS staff to focus on more pressing matters.
- **Reduced Library Downtime:** With the increasing importance of data, a library must remain on-line, all the time. Today's advanced libraries include fault-tolerant components. When mechanical items - such as fans - fail, another device is prepared to take over without interruption to on-going library operations.
- **Non-Disruptive Service:** As organizations move to around-the-clock business models, planned down-time is costly to the business and its customers. A high availability library enables on-line serving, keeping data available continuously.
- **Investment Protection:** Secondary storage can be a costly investment. As new technology evolves, a library that is able to expand and accommodate new features and functions protects the capital investment over a longer period of time.
- **Scalable Storage Capacity:** As storage volumes increase at exponential and unpredictable rates, the ability to add additional storage bins or link multiple units together for integrated operation are crucial.
- **Increased, On-Demand Throughput Performance:** The addition of new servers and users may require multiple parallel data streams. The ability to add tape drives and upgrade tape technology helps accommodate these increased demands.
- **Ensured Availability When Required:** The combination of redundant components and a robust design increases the chance that critical data can be stored or retrieved when it is needed most.

## Attributes of an Enterprise-Class Library

Tape libraries deployed in the enterprise demand scalability and high-availability features found only in the best libraries. Features required in an enterprise-class library include:

*An advanced robotics mechanism and a solid, fully welded frame will increase the inherent reliability of a tape library*

**Inherent Reliability** — the innate ability to continue to operate for long periods of time. In most libraries, this is a function of quality robotics and a stabilized internal environment. The robotics mechanism should pick and place cartridges with extreme precision and should avoid dragging cartridges along any surfaces — a situation that can contaminate the unit. In addition, the frame of the library should be fully welded to minimize vibrations that could negatively impact the robotics mechanism.

**Redundant Components** — the presence of dual primary components capable of immediately taking over the tasks of their counterparts when a failure occurs. Common redundant components include fans, power supplies and power cords.

**Hot-Swap Components** — the ability for a user to remove a component and insert a replacement while the library remains on-line. True hot-swap components do not interrupt any library operations, including data throughput on a single common SCSI bus. Important hot-swap components include tape drives, fans and power supplies.

**Scalable Performance** — the ability to add additional tape drives to increase data throughput.

**Scalable Capacity** — an on-demand increase of a library's total storage capability. This is traditionally accomplished by installing additional tape bins into a large library and/or coupling several libraries together to function as one large unit.

*An enterprise-class library should be able to communicate with a network framework application and be monitored anywhere in the world via the web*

**System Monitoring** — the process of checking a library to make sure all critical components are fully operational. Some monitoring software extends to the Internet or an SNMP-based application. In addition, some monitoring can predict failures before they occur, empowering administrators to act proactively.

**User Selectable Levels of Tolerance** — the ability to select a customized level of redundancy or fault tolerance. This is commonly manifested in power supply options. There tends to be a direct correlation between low-tolerance solutions and high costs.

**On-Line Serviceability** — the ability to service a library without taking it off-line. In around-the-clock operations, this feature is important.

**Fault Notification** — the ability to notify users of errors as they occur, in real time. Notification methods often include e-mails, pages and user terminal reporting.

**Library Partitioning** — the ability to divide a library into multiple virtual libraries that can be utilized by unique hosts.

**Dynamic Drive Allocation** — the use of a library controller to route data to an alternative path when required.

**Drive Fail-Over** — the combination of dynamic drive allocation and restarting a backup procedure.

## Scalability and High-Availability Library Attributes

Advanced tape libraries are evolving to include fault-tolerant features of their direct cousins: disk arrays. Similar to RAID, when key components within a high availability library fail, the library continues its backup operations unabated. The virtual elimination of single points of failure means the library will continue to be available for backup and restore operations. In addition, the need to bring a library off-line for servicing is eliminated.

An advanced library includes multiple levels of redundancy. In a high-availability library, mechanisms most likely to fail (e.g., tape drives, power supplies and fans) can be configured redundantly.

### Redundant Tape Drives

*Unique to only an elite group of tape libraries, tape drive failover is a powerful feature that assures a library remains on-line and backups are completed during a tape failure*

Redundant tape drives are unique to only the most advanced tape libraries. The best robotics mechanism may still be at the mercy of a tape drive. Research has revealed a tape drive fails primarily during the insertion or ejection of a cartridge. To work around this limitation, a valuable process known as drive failover has been introduced. Drive failover is the process of rerouting data to an extra tape drive (a hot spare) that will complete the backup process. This is extremely valuable in environments that have a fixed backup window because the backup operation will continue on a new device.

### Redundant Power and Fans

Traditional tape libraries include just enough power and fans to satisfy the minimum requirements for operation. This can often expose a library to failure if, for example, a single power supply fails. If the library ceased functioning, a service call would be required before it could return to normal operations. In environments that are sensitive to equipment uptime, there are multiple levels of redundancy that can keep the library on-line in case of a power supply or fan failure.

N+1 Redundancy: The "N" stands for the minimum number of power supplies and fans required to operate a tape library. N+1 is the addition of one extra power supply or fan that is connected together with its similar components. Should a power supply or fan fail, the extra component shares the load, eliminating the downtime that would have otherwise been caused by the failure. The downed component can be replaced at a convenient time.

2N Redundancy: Once again, "N" stands for the minimum component configuration required to operate the library. 2N is the doubling of the base components. In some environments, it is necessary to eliminate as many points of failure as possible. Fully redundant, a 2N system includes a complete backup power system and/or cooling system. In 2N configurations, fans, power supplies, power distribution boards and power cords can be duplicated for complete redundancy.

## Scalability Options

*An enterprise library must be able to scale its capacity, performance and features to the changing demands of its owner*

A scalable library is one that can grow as customer's needs expand. If the library is too small, it could reach its upper limits too quickly, making the investment obsolete. At the opposite extreme, a library that is too large is one that may take years to reach its full potential—an unwise investment in the rapidly changing secondary storage world.

There are many ways a library can be scaled, including its performance, capacity and features/drive technology. Whenever possible, a library should be able to scale within a common frame, significantly reducing the cost of expansion.

Scalable features that are common in today's most advanced libraries include:

- Increasing the number of tape drives for greater performance and throughput
- Increasing the total capacity of the library by increasing the amount of internal storage (typically involves adding more storage bins)
- Upgrading tape technology to newer, faster tape devices, thus directly impacting performance and total capacity
- Connecting multiple libraries together for significant increases in capacity and performance
- Adding enhanced functionality internal to the library, such as Fibre Channel and partitioning

The more scalable the library, the longer it will last and the more value, dollar for dollar, it will bring to its owner.

## **P2000 and P3000 - The World's First Scalable High Availability Libraries**

*The P2000 and P3000 libraries have unparalleled availability, scalability and enterprise functionality*

The P2000 and P3000 libraries from Quantum|ATL are the only solutions that meet the extremely demanding availability and scalability needs of the enterprise.

Their ingenious design has set a new standard for tape libraries. Unique features designed specifically for enterprise users include:

**Inherent Reliability:** With a proven one million swaps before failure, the P2000 and P3000 have the industry's highest reliability. This is a direct result of ATL's patented IntelliGrip™ robotics mechanism and a unibody fully welded frame. This reliability translates into increased data availability—an essential benefit of an enterprise-class library.

**Redundant Components:** Setting a new standard for tape library packaging, the P2000 and P3000 deliver a wide range of redundant components including AC power cords, power supplies, power distribution boards and fans to ensure continuous uptime. In addition, the libraries can be configured with redundant tape drives and internal components, such as Fibre Channel bridges.



P3000

**Hot-Swap Components:** The P2000 and P3000 include cable less, hot-pluggable, user replaceable, hot-swap drives that do not interrupt any library operations, including data throughput on a single common SCSI bus. Additional hot-swap components include fans and power supplies.

**Scalable Performance:** The P2000 and P3000 can be configured with up to 10 and up to 16 DLTtape™ drives, respectively. Native performance throughput rates of up to 180 GB/hour and 288 GB/hour, respectively, can be achieved with DLT7000 tape drives. Plug compatible support for future high-performance technologies such as DLT8000, SuperDLT™ and LTO Ultrium™ tape drives ensure extreme performance scalability. Fibre Channel connectivity is also internally supported for increased performance over a storage area network.



P2000

**Scalable Capacity:** The P2000 and P3000 can house up to 198 and up to 326 DLTtape cartridges for 6.9 TB and 11.4 TB of native storage capacity respectively. Support for future tape drive technologies such as DLT8000, SDLT and LTO Ultrium will further increase the storage capacity within a single library. Additionally, multiple P2000 and P3000 libraries can be linked together for massive storage requirements.

*ATL's advanced system management tools enable easy library configuration, error condition monitoring and firmware upgrades via a network framework application or over the web*

**System Monitoring:** The P2000 and P3000 offer the most comprehensive monitoring tools of any enterprise-class library. ATL software options include WebAdmin™, the Prism ALERT (Active Library Error Reporting Tool) option<sup>1</sup> and ATL's Management Console. These monitoring and control features include the ability to manage the P2000/P3000 from a major network management framework or from anywhere in the world via the web. ATL's library management tools are capable of a wide variety of functionality including library configuration, error condition monitoring and firmware upgrades.



**User-Selectable Levels of Tolerance:** The P2000 and P3000 are the first tape libraries to allow user selectable levels of tolerance. A N+1 load-sharing power supply configuration is available, enabling a single power supply to fail without taking the library off-line. For complete redundancy, power supplies, power cords, power distribution and fans can be ordered in a 2N configuration. This means that a power supply, fan, power distribution board and a power cord could all fail simultaneously and not affect the operation of the P2000 or P3000.

**On-Line Serviceability:** Both the P2000 and P3000 include hot-swap, plug-and-play components without any cabling, enabling hot-plug service replacements without bringing the library off-line.



**Fault Notification:** A P2000 or P3000 with the Prism ALERT option<sup>1</sup> is able to report faults in real time via an SNMP agent to a remote console, via e-mail or over the airwaves to a pager.

<sup>1</sup> Check with an ATL Representative for availability information.

**Library Partitioning:** The P2000 and P3000 are the only libraries in the world that include an optional embedded management card and partitioning software that can divide the library into multiple virtual libraries that can be utilized by unique hosts and/or multiple different applications. This allows the cost of a larger, more cost-effective library to be amortized across multiple servers.

**Dynamic Drive Allocation and Drive Fail-Over:** The P2000 and P3000 are the world's only libraries capable of dynamic drive allocation and drive fail-over<sup>2</sup>. An optional embedded management card brings these features inside the P2000 and P3000. Data can be routed to an alternative drive when required and the backup operation can be restarted, assuring a complete backup operation.



### **Summary: The P2000 and P3000 are Unparalleled Enterprise-Class Library Solutions**

The P2000 and P3000 libraries are the world's most resilient, scalable, high-availability libraries. For enterprises that treat their data like gold, only the P2000 and P3000 provide the availability and reliability to keep data on-line all the time. The libraries have the ability to incur multiple, simultaneous failures yet still operate at full performance. For example, the P2000/P3000 could have a failed power supply; a failed fan and a drive failure while backups are still in process. The library would notify the system administrator that service was required and the library could be fully repaired within minutes, all without interrupting the backup already in progress. The result is reduced library downtime due to common failures and increased data center productivity and efficiency.

The P2000 and P3000 are the industry's most scalable libraries, enabling cartridge, drive, technology and feature enhancements all within a single box. Open your enterprise to its fullest potential with the advanced storage capabilities of a P2000 or P3000 library—the scalable, high-availability leaders.



P3000

<sup>2</sup> Check with an ATL Representative for availability information.