

Position Summary: A Backup Appliance Composed of High-capacity Disk Drives

Kimberly Keeton and Eric Anderson
Hewlett-Packard Labs Storage Systems Program
{kkeeton,anderse}@hpl.hp.com

Disk drives are now available with capacity and price per capacity comparable to nearline tape systems. Because disks have superior performance, density and maintainability characteristics, it seems likely that they will soon overtake tapes as the backup medium of choice. In this position summary, we outline the potential advantages of a backup system composed of high-capacity disk drives and describe what implications such a system would have for backup software.

In the past, magnetic tape had higher capacity and lower price/GB than magnetic disk; however, technology trends are reversing this relationship. Tape's capacity lead over magnetic disks has shrunk over the last 15 years, and disk capacity is now on par with tape capacity. Furthermore, disk media price is within a factor of 3X of tape media price, and disks cost less per GB than tapes, once the tape drive and the tapes supported by the enclosure are taken into account. Given these trends, it is time to rethink tape's role as the backup medium of choice.

Disks confer tremendous hardware-related benefits for a backup system:

Performance. Disks' 5X faster sequential performance suggests that disks are better for creating and fully restoring backup volumes, allowing easier verification and more efficient data scrubbing. Furthermore, disk bandwidth scales more cheaply, since each disk adds bandwidth, whereas only expensive tape drives add tape bandwidth. Disks' superior random access performance implies that disk-based backup will be better at partial restorations and at satisfying simultaneous restore requests.

Density. Designing the appliance so that only a fraction of the drives are simultaneously powered on reduces power and cooling requirements, allowing denser packing. Back-of-the-envelope calculations indicate that a disk-based backup appliance could provide roughly 2X more capacity per unit volume than a tape-based system [1].

Support for legacy devices. Restoring data from tape requires finding a matching tape drive, which can be difficult since tapes come in many formats. Disks include their own read/write heads, eliminating the need to search for a separate drive to retrieve data.

Maintainability. Tape drives need to be periodically

cleaned with special cartridges and periodically serviced to ensure that head drift doesn't render a tape unreadable. In contrast, disk drives are enclosed media, which don't require cleaning and don't suffer head drift problems.

Lifetime. Empirical evidence suggests that disks could have a longer shelf life than tapes, implying that disks may ultimately be better archival media. System administration experts advise re-recording tape data every three years. Disks come with warranties for three to five years, and disk experts believe that lifetimes over ten years are possible for backup-optimized disks.

The characteristics of disk-based backup have implications for the creation of backup software:

Design for reliability. Backup software protects data by maintaining a read-only copy that cannot be inadvertently corrupted, and by providing an alternate, simpler software path than a snapshotting file system. Furthermore, we can design the backup system to trade off reliability for performance, by using self- and peer-checking code, storing checksums with each data block and verifying those checksums periodically and when the data is accessed, and pro-actively testing the system.

Design for sharability. A backup system that keeps a fraction of its disks online may be able to approximate the performance of an online snapshot using hierarchical storage management techniques, allowing greater simultaneous sharing, while still maintaining the data protection properties of a backup.

Design for longevity. A final opportunity for backup software is to automatically convert data formats commonly used today into formats that will be easy to read many years in the future, either automatically, or through user control.

Key challenges lie in designing backup software for optimizing reliability and data integrity, scheduling the resources of the backup appliance, and developing APIs for giving users and applications more control over how backups are performed.

[1] K. Keeton and E. Anderson. "A Backup Appliance Composed of High-capacity Disk Drives," HP Laboratories SSP Technical Memo HPL-SSP-2001-3, available from <http://www.hpl.hp.com/research/itc/csl/ssp/papers/>.