



## ► What If Disk Space Were Free?

Richard J. Cichelli, SCS President

Abstract: An active archiving algorithm is presented for managing hierarchical storage and taking advantage of new low-cost, high-capacity Serial ATA (SATA) disk technology. A newspaper publishing application provides a concrete example.

As of February 2003, retail prices published on the web showed costs per terabyte of disk space ranging from \$4,360 for an Apple Xserve RAID to over \$75,000 for a Sun StorEdge T3 system. Dell's EMC CX200 unit costs more than \$13,600. Our suppliers offer solutions for less than \$2,000 per terabyte. Now that the rules for configuring systems have changed, how can you take advantage of this new situation?

Software Consulting Services, LLC (SCS) provides IT solutions to newspapers. Managing the production of display advertisements and archiving them are mission-critical tasks for newspapers. Recently, we have been able to quadruple the storage space in our ad production systems while significantly reducing both platform costs and IT support costs. What we have done could be applied to other application areas and, I would expect, yield similar advantages.

Here is what's special:

- What you are making involves big files. In our case, it's ads and their contents. Art work, scanned images and graphics all fill lots of disk space.
- You have some mission-critical data. For our application, about 100GB is used for the storage of ads that are being made up for future publication. If these weren't available, you couldn't put out the paper.
- You have an archive of less critical data. In SCS/Track, the archive area is from 900GB to 1,900GB. It holds years' worth of published ads and is used to hold old ads for selling new insertions and rebuilding modified versions of them.

How do you convince a publisher to invest in digital ad asset management when, for 20 years or more, ads

Enterprise-wide digital asset management (DAM) is a myth for media companies. Publishers and networks should attack specific workflows with open, cheap, modular tools--not with galactic DAM deployments.

**Forrester Research,  
TechStrategy Report  
January 2003**

were discarded two or three weeks after expiration? You make it very cost-effective to save old ads.

We configure terabyte-capacity servers for our applications. Each costs us between \$8,000 and \$12,000. The low-end servers have a single 2.4GHz Intel Pentium 4 Xeon processor, 2GB of RAM, ten 250GB Maxtor SATA drives with two 3Ware SATA RAID Controllers, DVD/CD reader/writer, 50/100GB QIC tape drive, floppy, dual ethernet, keyboard, monitor, mouse, UPS and Linux. The high-end servers have two 2.8GHz Pentium 4 processors with 4GB of RAM, eight 250GB SATA drives on a 3Ware SATA RAID 5 Controller and two Seagate 146GB SCSI drives on an LSI RAID 1 Controller. (RAID 1 mirrors; RAID 5 has parity. Both offer extra redundancy and improve performance.) We put up to 35 users on the low-end servers while the high-end servers support 90 or more.

Having two controllers is a big advantage. Accessing mission-critical data on a RAID 1 SATA drive pair is faster than doing so on six SCSI drives in RAID 5. A pair of SCSI drives in RAID 1 with a 128MB caching controller benchmarks at three times the speed of six SCSI drives in RAID 5.

From 1998 to 2001, we configured systems with tape auto-changers and virtual CD towers. Full backups at 50GB/hr took far too long, and the proprietary

Software Consulting Services, LLC  
630 Selvaggio Drive, Suite 420  
Nazareth, PA 18064  
Sales: 1-800-568-8006  
Fax: 610-746-7900  
E-mail: sales@newspapersystems.com  
www.newspapersystems.com

**SCS BUILDS TRUSTED NEWSPAPER SYSTEMS.**



virtual CD tower required high IT maintenance. The new configuration eliminates these devices.

Permit me to usurp a catch phrase from Carnegie Mellon University, "Lots of copies keep stuff safe" (LOCKSS). Here's how our version of LOCKSS works:

- Define three data areas: current, expired and archived. For us, current stores ads being built, as well as those already built and ready for publication. Expired holds those ads which have run all their insertions, and archived holds those ads that have reached their "purge to archive" date.
- Implement active archiving by moving current ads to the expired area when their publishing insertion sequence is completed and from expired to archived when their purge date is reached. For this to work, it's best that the database hold pointers to the objects (ads, in our case) and not hold the objects themselves. (Besides facilitating the movement of the ads within the storage hierarchy, this assures that, in the event of database corruption, ads can still be found, built and forwarded to pagination.) A RAID 1 mirroring configuration assures both high availability and quick access. Current ads, the ad database indexes, thumbnails, JPEG images and plain text for all ads are kept in high-performance storage. (Plain text is kept for full-text searches of ads.)
- Move ads from the expired area to the archived area in stages. Change the organizational structure from being task-oriented (i.e., a collection of advertiser folders, each containing zero or more ad folders) to being time-oriented (i.e., ads grouped by expiration date). (Each ad folder contains one or more files. These include input graphic and text files, QuarkXPress®, MultiAd®, InDesign® and other documents and EPS (Encapsulated PostScript) and PDF (Portable Document Format) files.)
- Build archivable collections of ads designed to fill a DVD and write them to DVD, verify and then move them to the archive. To provide off-site storage, burn a second DVD and write two tape backups of current ads. By having the archive grow by DVD-sized increments, no all-at-once backup of it is needed. When pickups are needed, they are identified by searching the ad database. Once found, they are copied into new ad folders, and the ad building process starts over again from there.

We provide another level of redundancy by having dual replicated servers. Replication supports our preferred, simple "shared-nothing" architecture and allows the streaming log files to be filtered for errors as they are sent to update the replication server. We also delay deletions and purges for later processing on the second server. (Another safety check.)

Expanding system storage is easy because we allow the storage areas to be spread over multiple volumes and servers.

Low-cost SATA drives are a "fill the chassis" technology. They cost so little that you just install as many as you have room for. They can replace SCSI technology in many applications. Thanks to Intel for its motherboards and processors, Antec for building a 14-drive bay chassis, 3Ware for its SATA RAID controllers, LSI for its SCSI controller, Tanberg for its tape technology, the open-source community for Linux and Cobalt Computers for assembling servers to our specifications. Eight SCS terabyte-capacity servers are already being used by daily newspapers. Our thanks to *Kingsport (TN) Times-News*, *Danbury (CT) News-Times*, *High Point (NC) Enterprise* and *Hyannis (MA) Cape Cod Times* for willing to be first.

In summary, we showed how \$3,000 SATA terabyte disk subsystems can be configured for certain applications to out-perform SCSI disk subsystems costing many times as much.

## Your Questions Answered

### Why did we want terabyte-capacity servers?

Our customers are newspapers. We help newspapers sell, create, publish and archive display advertising. The Cobalt servers are a very cost-effective platform for display ad production. Terabyte capacity servers make these tasks easier. We use about 100GB for ad production, and the remaining 900GB is used for ad archiving (i.e., digital ad asset management).

Our customers usually buy two identically configured servers. Each has a 3Ware 8500 series, eight port, Serial ATA controller and eight Serial ATA 120GB Seagate drives. (We can now get 160GB drives from Seagate and 250GB from Maxtor.) The eight drives are set up in a RAID 5 configuration. The Antec chassis holds 10 drives. The remaining two drives are either set up as RAID 1 Serial ATA with 120GB drives, or RAID 1 SCSI with 146GB SCSI drives. The RAID 1 configurations give very high performance for the production ad building environment.



### **What problem did the terabyte-capacity servers solve?**

With more disk space, an ad archiving system can be easier to manage. Unfortunately, as we increased disk capacity, tape backups took too long. The new server architecture allows us to do disk to disk backup across the network. This allowed us to eliminate the tape auto-changer. An auto-changer costs nearly as much as a server! Having lots of inexpensive disk space ultimately solved all of the capacity and support problems.

### **How are the systems being used?**

We now have Cobalt servers in production at eight newspaper sites. The three described below show the different ways the systems are used. We'll start with the simplest and work toward the more sophisticated.

At the *Cape Cod Times* in Hyannis, MA (an Ottaway/Dow Jones newspaper), SCS installed an SCS/Track system early in 2002. They recently purchased a single Cobalt server from us set up as a CD/DVD virtual jukebox. We used open-source software to build the virtual jukebox. It runs Linux and costs about the same as the proprietary virtual jukebox we previously sold, yet it has three times the capacity. An added benefit is that maintenance is easier on Linux.

At the *Danbury (CT) News-Times*, servers are configured to do data replication for SCS/Track. *The*

*News-Times* not only went live with SCS/Track on April 25, 2003, they also achieved full digital pagination of the newspaper for the first time. At the *News-Times*, having SCS/Track to manage workflow and Cobalt servers to store and move large amounts of digital data provided what was needed. All those involved regarded this as an extraordinary success.

### **How do we measure success?**

Our greatest concern is that our systems be reliable. To be reliable, the systems must have fail-soft redundancy and be easy to manage. Having lots of inexpensive, reliable disk space allows us to eliminate most operator-performed system management tasks. No more tape autochanger tasks, no more proprietary CD/DVD jukebox, etc. Having fewer devices means the systems not only cost less but are more reliable and easier to manage. Using SATA technology and Linux has allowed us to dramatically reduce platform costs. All ownership costs are lower for our customers with the Cobalt Computer servers. SATA technology works great and costs less.

### **Customer feedback?**

Our customers love the new servers. There is no end to the praise they have for the speed, capacity and manageability of the new systems. On Friday, May 30, 2003, we received our fifth contract for SCS/Track running on Cobalt servers. Five sites in less than 50 days! It's a rapidly growing success.