Software Consulting Services, LLC

Comparing Storage Alternatives for Digital Asset Management

By Richard J. Cichelli, SCS President

We have found that a combination of large capacity serial and parallel ATA disk drives can be configured to yield a very large storage subsystem that is reliable, fast and low-cost. Ample, dependable, affordable storage is a critical resource for the digital ad asset management component of an ad tracking production system.

We tested many configurations and now have two new lower-cost platforms that meet the needs of small (5 to 15 ad builders) and mid-market (up to 60 ad builders) newspapers. The first is built to our specifications by Dell Computer Corporation and utilizes the Dell 600SC departmental server with four 120GB ATA parallel drives in a RAID 5 configuration using the Dell/LSI four-channel controllers. It offers a capacity of 360GB and the comfort of the Dell name.

The second is built by Cobalt Computers (a Coulbourn Company) which makes white box systems for hospitals, universities and pharmaceutical companies from here in the Lehigh Valley. It features a 12-bay chassis with up to two Pentium 4 Xeon processors, an Intel motherboard and two 120GB serial ATA drives in a RAID 1 configuration on a four-port 3Ware controller and five to eight 200GB ATA parallel drives in RAID 5 configuration on an eight-port 3Ware controller, yielding over one terabyte of storage.

Here are some of the technical details behind our recommended storage architecture.

Overview

We have been installing our ad tracking systems in newspapers since 1986. As disk space became less expensive, we recommended increased online storage. This allows SCS/Track to more easily manage ads in production. SCS/Track uses the term "current ads" for the collection of ads being built, ads still to be published and unexpired ads. Pick-ups and sales catalogs will likely need expired (i.e., archived) ads. The bigger the capacity of the disk subsystem, the bigger the possible ad archive and thus the longer time ads can be retained online.

Costs

Each tracking system we sell has dual, replicated, redundant, shared-nothing servers. The new disk storage subsystems cost $1,440 for the 360GB Dell disk subsystem or $4 per gigabyte and $3,050 for the 1,120GB Cobalt disk subsystem or $2.75 per gigabyte. The total dual-server platform costs are $19,342 and $19,288 respectively. The Dell solution uses a tape vault for backup while the Cobalt configuration's greater disk capacity allows a more easily maintained "backup to disk" architecture.

For comparison, current pricing on a Dell 4600 configured with eight 73GB SCSI drives and a RAID 5 controller has a disk subsystem costing $6,431 and a total dual-server cost with tape vault, virtual CD tower and rack of $49,600. The per gigabyte cost of the 511GB (net) SCSI disk subsystem is $12.60, over 4.5 times as expensive as our preferred ATA solution. The new configurations cost $20,000 to $30,000 less than our legacy solution.

Performance

In our tests, the two new configurations we recommend performed faster and offered greater capacity than other options. To differentiate among possible solutions, we needed a way to load the systems with activities similar to what SCS/Track might do under the most demanding circumstances. We created a Spice database update task to compare the disk subsystems. The database build involved nearly one million rows with many foreign keys. (Each row is a record describing an ad.) The resulting database fills over 750 megabytes. With this load test, we could compare the performances of prior working SCS/Track architectures with candidate solutions. Doing better with a new solution would indicate acceptable performance.

The new subsystems use ATA drives instead of SCSI drives. The first workstation PCs came with ATA drives while their Intel-based server cousins were

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

SCSI-based. Like most computer technology, SCSI system performance has improved greatly over the years. ATA technology has improved even more so. While typical SCSI drives lead in rotational speed (10,000 RPM) vs. ATA (7,200 RPM), many other performance-related factors are nearly equal, and ATA capacities and cost-effectiveness beat SCSI by almost three to one.

Parallel ATA drives (they have short, wide, flat cables) are being replaced by Serial ATA subsystems (round, thin, possibly longer cables). Serial ATA transfer rates will surpass SCSI in the next few years. Because thin, round cables allow better chassis air flow, bigger chassis and denser drive configurations are possible. Most storage technologists feel Serial ATA technology could overtake SCSI for nearly all applications.

Ad builders need quick access to current ads and less quick access to archived ads. For many newspapers, doing a pickup of a six-month-old expired ad means sifting through a pile of CDs in another room. Hard disk subsystems are not only better than that, they are 10 times as fast as a CD jukebox. They even beat virtual CD jukeboxes. (These are dedicated servers with lots of hard disks interfaced in the same manner as a CD jukebox.)

Our benchmark traditional configuration is typified by a Dell 4600 with a RAID 5 controller with onboard cache of 128 megabytes and eight 73GB SCSI drives. The rest of the configuration included a 2.0GHz Pentium 4 Xeon and a gigabyte of RAM. Task completion time for this was 61 minutes.

The high-end Cobalt servers had a 2.4GHz Pentium 4 Xeon processor, 2GB of RAM and two disk subsystems: a pair of 120GB Serial ATA150 drives in a RAID 1 (mirror) configuration and six 200GB parallel ATA100 drives in a RAID 5 configuration. The former is for current ad storage, and the latter is for archived ads.

The load test on the Serial ATA subsystem of the Cobalt server ran in 55 minutes. This is a 10% improvement over the SCSI subsystem. 120 gigabytes is a generous amount of space for current ads. A terabyte will store many years of archived ads.

A number of factors help explain this happy improvement. Each Serial ATA drive has its own 8MB of cache, each ATA drive gets its own channel on a controller, etc.

Running the load test on an uncomplicated Intel Pentium 4 server with four Serial ATA drives in a RAID 0 configuration yielded a surprising time of just 18 minutes. Using Serial ATA drives in a RAID 5 configuration alone yielded a 110-minute run time. The corresponding time for parallel ATA drives RAID 5 is 186 minutes. All these systems were running Linux 7.3 with the Ext3 journaling file system. Ext3 was set at the default "ordered" mode. We have and are testing other configurations. For example, a Sun Sunfire V120 server running Solaris 9 yielded a time of 63 minutes.

Reliability

The configurations we offer demonstrate a high level of reliability. SCSI drives are rated as having a mean time between failures (MTBF) of 1.2 million hours (i.e., 136 years)! Current ATA drives have a MTBF of 68 years. Their greater capacity means that on a per gigabyte basis, they are more reliable than SCSI by a ratio of 3 to 2.

When you consider the squaring effect due to the redundancy of RAID 1 or RAID 5, MTBF calculations rise as much as 50,000 times. Dual servers square that again.

Perhaps most importantly, these MTBF values indicate that the drive subsystem is unlikely to be the weak link in system reliability. Application, OS and human errors, natural disasters and equipment obsolescence are more likely sources of difficulties. We seek to minimize the impact of all these issues as well.

Conclusion

We have three newspapers enthusiastically using SCS/Track on Dell 500SC/600SC servers with parallel ATA drives. They are the *Norwich (CT) Bulletin* with 10 ad builders, The *Leesburg (VA) Times Community Newspapers* with 12 ad builders and the *Plymouth (MA) Memorial Press Group* with 8 ad builders.

The components used in Cobalt servers are the same types of components from the same manufacturers that Dell uses. Intel processors and motherboards, Seagate disk drives, etc. are integrated by both suppliers.

Why not just have Dell build a Cobalt-like system? We tried but were told getting exactly what we wanted wasn't yet possible.

The SCS-configured Cobalt solution is bigger, faster, more economical and arguably better than other alternatives we've tested.

Cobalt Server Pentium 4 2530MHz 1024MB RAM (DDR266) 3ware Escalade 8500-4 Serial ATA RAID controller four Seagate ST31200 120GB Serial ATA hard drives	All of these timing tests used ext3 file system in default mode. Single drive, no RAID 56m26s RAID 0 four drives 18m47s RAID 1 two drives 55m31s RAID 5 three drives 110m57s
Dell PowerEdge 4600 Server XEON 2000MHz 1024 MB RAM (ECC DDR) DELL PERC 3/DC SCSI RAID controller (128MB cache) six 73GB 10K RPM SCSI hard drives in a RAID 5	61m32s
Dell PowerEdge 600SC Pentium 4 2400Mhz 1024 MB RAM (ECC DDR) DELL/LSI ATA 100 four channel parallel ATA controller (16MB cache) three 120GB (IBM Deskstar) 7200 RPM parallel ATA hard drives in a RAID 5 array	100m42s
Sun Sunfire V120 650 MHz with 512MB RAM one 36GB SCSI Drive	63m