Building Reliable Newspaper Systems - The Sequel

By Richard J. Cichelli, SCS President

A prior white paper, "Building Reliable Newspaper Systems - Protecting Data," gave a high-level description for preventing data loss. There are two parts to this follow-up: The first provides computational support for the advocated reliability improvement methodology, and the second discusses issues regarding technological innovation.

Part 1 - The Math

We start by revisiting Dan Gardner's "Top 10 Causes of Data Loss;" however, this time in addition to the weights derived using Zipf's Law, we add a column showing "improved weights." The derivation of the values for the improved weights will be explained shortly.

Rank	Cause	Wt	Improved Weight
1	Mechanical hard-drive failure	24.7	5
2	Data structure corruption	16.5	3
3	Accidental or intentional data deletion	12.4	2
4	RAID server failure	9.9	9.9
5	Backup tape malfunction	8.2	0
6	Physical tape damage	6.3	0
7	Accidental overwrites	6.2	1
8	Software corruption	5.5	2
9	Viruses	4.9	1
10	Natural disasters	4.4	4.4

The weights and other reliability factors are based on a set of reasonable assumptions. Here are some:

1. Reliability factors such as manufacturer's reported mean time between failures (MTBF) are, for practical purposes, valid. Direct valid statistical performance measures for components like disk drives, which have very long MTBF's (125+ years), are impractical. (The drives would be obsolete long before even a single test cycle was completed.) Testing procedures that quickly stress components using excessive heat, vibration and power fluctuation produce acceptable estimates of durability.

2. Gardner's table reflects the reliability of a typical computer. We would expect these to be workstation PCs with an Intel processor and motherboard, an IDE parallel ATA drive, Microsoft system software, etc.

3. The data loss occurs on computers that are in production use and thus have passed their infant mortality stage.

4. The probability for data loss causes conforms to Zipf's Law for ranked items. That is, for the ith ranked item, the frequency of occurrence is proportional to 1/i.

5. The probability of two independent events occurring is the product of the probability of each.

6. The mean time to repair (MTTR) for typical failures is one day.

7. Both Dell and we (through Cobalt Computers) produce equally reliable systems using the same types of parts - Intel processors and motherboards, Seagate or Maxtor drives, etc. The commodity non-proprietary, brand-name components that Dell buys for their systems are available from many sources.

8. Exact prognostication is impossible; however, approximate and back-of-the-envelope calculations can yield valuable insights and useful rules of thumb.

Here's how the improved weights were derived:

Mechanical Drive Failure
From 24.7 to 5. Unlike Gardner's hypothetical typical computer, our configurations are for servers, not workstations. Servers have multiple drives and RAID controllers. They have more drives that can

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

fail and require the extra built-in redundancy that RAID technology provides. Starting with 1,100,000 hour MTBF for Maxtor's 250GB SATA drives, we get a 0.8% annual failure rate. (The annual failure rate is the inverse of the MTBF. 1,100,000 hours is 125 years. 1/125 = 0.8%.)

We configure servers three different ways. Remote site servers have a pair of 100+GB drives in a RAID 1 configuration. Mid-market terabyte-capacity servers have five 250GB SATA drives in a RAID 5 configuration. Enterprise servers with a two terabytes capacity have 10 drives - two SATA or SCSI drives set up in RAID 1 and the remaining eight in RAID 5.

For two RAID paired drives, the failure rate is (2 x 0.8%)2 or 0.026% for both. ((2 x 0.008)2 = 0.00026 = 0.026%) This would reduce the 24.7 weight to 0.8.

More typical is the five-drive RAID 5 configuration. This is (5 x 0.8%)2 or 0.16% ((5 x 0.008)2 = 0.0016 = 0.16%) for two to fail on the same day within a year. This yields an expected improved weight of 4.94 ((24.7 x 0.16)/0.8 = 4.94) or 5 as listed.

The dual RAID configuration for 10 drives would have an improved weight of 0.8 for the essential RAID 1 current ads part and a 12.65 for the ad archive part.

You could further improve the expected reliability by keeping a spare drive available to reduce the MTTR to an hour or two and get a 10-fold drive availability improvement. For support, our supplier keeps one drive plus an additional drive per year of expected server life per 100 drives of each drive model sold.

Data Corruption
From 16.5 to 3. In addition to the journaling file system support provided by Linux, which is not available on typical Windows desktops, there is also added redundancy within the SCS/Track application. While the underlying B-tree-based record manager for SCS/Track's relational database management system (RDBMS) is very stable, were it to fail, ad production and pagination could continue without the database by just accessing the ad files directly in their conveniently named file folders.

Accidental or Intentional Data Deletion
From 12.4 to 2. SCS/Track uses staging to schedule deletions on secondary servers and does not allow users to directly delete data from secondary servers.

RAID Server Failure
Unchanged from 9.9. Note: Our experience with RAID controllers and other non-mechanical electronic parts is that they seldom fail.

Backup Tape Malfunction
From 8.2 to 0. Eliminated as a single point-of-failure for data loss.

Physical Tape Damage
From 6.3 to 0. Eliminated as a single point-of-failure for data loss.

Accidental Overwrites
From 6.2 to 1. The design of SCS/Track prevents accidental file overwrites. It is a controlled environment that protects users from such errors.

Software Corruption
From 5.5 to 2. For the same reason as cause 7 above. In addition, Linux and SCS's tools are much more stable than Windows-based technology.

Viruses
From 4.9 to 1. Viruses seldom occur in any context other than online Windows systems.

Natural Disasters
Unchanged at 4.4.

The sum of the weights is 99, indicating that these factors account for 99% of all data loss causes. The improved weights total about 30 or more than a two-thirds reduction. Given that the failure rate for a single drive in a workstation is 0.75% per year, and this type of failure probably accounts for one-quarter of all failures, the annual workstation failure rate would be 0.75% x 4 or 3% per year. This suggests about one data loss event per 100 days per Windows PC.

Our server configuration is at least three times better. (It needs to be, since enterprises depend on it.) Fortunately, the reliability is even better than one event per year. The configuration has paired, independent servers so, if the failure rate for one is one per 300+ days, then both at the same time would be 1/300 x 1/300 or 1/90,000. This rate is 0.000011 or a 99.999% availability rate. This is five-nines availability.

This computation argues against mirrored or clustered server configurations since replication preserves server independence while these techniques do not. Mirroring can duplicate a problem from the primary server to the secondary. It is usually a more complex, less reliable, less efficient and more costly solution.

Building a reliable platform is the starting place for building forgettably reliable systems.

Part 2 - Innovation

Do I expect everyone who is buying an ad tracking system to want Cobalt servers configured by SCS? Nope.

The value of innovation, no matter how well supported, is always weighed against the cost and fear of change. We usually sell Dell equipment. Unfortunately, Dell hasn't yet found a large enough market for SATA drive technology. (To say nothing of the fact that it might disadvantage their sales of high-margin EMC data storage subsystems.)

So why bother innovating?

One makes choices in life: lead or follow. If you want things to be better, you support the creativity that is a prerequisite for innovation and guide it with the leadership necessary to implement change.

Having long-term perspective helps. A number of years ago, SCS was an approved reseller for Compaq, IBM, Apple and AST. I found a supplier who built systems exactly to my liking and had better prices and faster delivery. Further, their chasses were of heavier metal, their cards were mechanically fastened onto motherboards and their shipping cartons were meant for rough next-day freight handling. I bought bunches of PCs Limited computers and eventually began reselling them to customers.

Believe me, there were those within SCS who thought not using a top name-brand supplier was very poor judgment. It didn't matter to me that the company was only 18 months old and started by a college drop-out in his dorm room. The computers were good and cheap. I was happy. The customers were happy, and, to my surprise, PCs Limited's new Vice President of Sales Joel Kocher was very happy. He called to find out why our relatively small company was buying so many PCs from them.

I told him we liked their reliable, low-cost product and their nationwide service. We had started reselling their computers to our customers. "Why don't they buy directly from us?" asked Joel. "Well, we sell them at our cost," said I. "How do you make any money doing that?" asked Joel. "We put $25,000 worth of software on them and make money on that," I replied.

Joel wanted to confer with his boss, the owner, and later got back to me with an offer to be their first reseller. Not only would they give our support calls special attention, but also they would give us a discount. Since there were no purchasing quotas, it sounded good to me.

To this day, we are still very satisfied with the company. PCs Limited became Dell Inc.

SCS welcomes working with suppliers and customers who choose to innovate with us.