

The digest of current topics on Continuous Processing Architectures. More than Business Continuity Planning.

BCP tells you how to **recover** from the effects of downtime.
CPA tells you how to **avoid** the effects of downtime.

In this issue:

Never Again

So You Think Your System is Robust

Best Practices

HP Blows Up Data Center

Availability Topics

Availability versus Performance

Recommended Reading

Towards Zero Downtime

The Geek Corner

Estimating Data Collision Rates

Complete articles may be found at
<http://www.availabilitydigest.com/articles.htm>.

Welcome to our inaugural issue of the FREE Availability Digest. Access to all of our articles, past and present, no longer requires a subscription. Simply click on the “—more--” link at the end of the article synopsis in the Digest below to read the full article. Or go to www.availabilitydigest.com, and click on the Article Archive button to see all of our past articles.

In this issue, we look at a series of computer glitches that happened to major enterprises in the first six months of 2007 and learn some valuable lessons from these experiences. We explore ways to improve system availability by compromising performance to some extent. And be sure to check out HP’s explosive (literally) demonstration of fast failover.

We stand ready to advise and guide you in your search for higher availability. Please contact us about our availability services at editor@availabilitydigest.com.

Dr. Bill Highleyman, Managing Editor

Never Again

So You Think Your System is Robust?

A lot of people are convinced that their systems are robust. Sometimes, however, their best laid business continuity plans can go astray. Many sad tales tell of those who learned the hard way that business continuity planning is harder than they thought.

In this article, we relate the adventures of seven major enterprises that experienced debilitating computer glitches – Dow Jones, US Airways, the Canadian Revenue Agency, the FAA, M&T Bank, All Nippon Airways, and BlackBerry.

All the snafus occurred to these major enterprises in the first half of this year, 2007. It is interesting to note that four of the snafus had to do with failover faults to backup systems, and five occurred following a reconfiguration of one sort or another. It seems that testing was not at the top of the priority list for these enterprises.

[--more--](#)

Best Practices

HP Blows Up Data Center

As a demonstration of the rapid failover capabilities of HP systems, HP set up a data center with a mix of HP systems and backed up the data center with a remote data center. Populating the data center were NonStop, HP-UX, OpenVMS, Linux and Windows systems.

It then blew up the primary data center (yes, actually blew the data center up with explosives) to demonstrate fast recovery in a simulated natural gas line explosion.

Within two minutes, the backup data center had taken over the entire processing load.

A video describing this demonstration is available on the HP Web site.

[--more--](#)

Availability Topics

Availability versus Performance

Increased availability does not come for free. There are hardware approaches that increase cost, and there are software techniques that reduce performance.

Cost-based approaches for improving availability range from single-system architectures that employ redundant components to active/active systems and clusters.

Techniques for improving availability at the expense of performance are not typically found in the lower tiers of today's industry standard servers because these servers are highly optimized for performance. Availability is a secondary consideration.

However, because of the tremendous improvements in system performance over the years as compared to the modest improvements in system availability, it is often now desirable to trade off some of these performance gains for improved availability. This is especially true for applications which are involved in the 24x7 operations of today's enterprises.

The techniques for availability improvement at the expense of performance are substantially software-based and include database recovery, shared-nothing architectures, recursive rebooting, and software rejuvenation. These techniques are aimed at improving the availability of a single system. In addition, improved single-system availability translates to much higher availability for multinode architectures such as active/active systems and clusters.

[--more--](#)

Recommended Reading

Towards Zero Downtime: *High Availability Blueprints*

In his book entitled *Towards Zero Downtime: High Availability Blueprints*, Vishal Rupani focuses on the use of Windows clustering techniques and products. He introduces the topic by covering a broad range of availability issues such as storage and processing redundancy, and he highlights the need for an extensive discovery process to understand the client's current systems and future needs.

Rupani then details the installation, validation, and test procedures for Microsoft clustering and for several Microsoft cluster-aware applications such as SQL Server, Internet Information Server (IIS), Network Load Balancing, and clustered file servers. He follows this with a brief discussion of geographically-distributed fault-tolerant architectures and concludes with an actual in-depth case study that applies the concepts covered in his book.

[--more--](#)

The Geek Corner

Estimating Data Collision Rates

Data collisions are an unfortunate fact of life for active/active systems using asynchronous bidirectional data replication. There are several techniques that can be used to avoid or to automatically resolve data collisions. However, if these techniques cannot be used, collisions must be resolved manually. Manual collision resolution is a time consuming task and should be understood before initiating an active/active project.

This article provides a simple technique for estimating the rate of data collisions. It is shown that collisions are linearly proportional to replication latency and to the square of the system's update rate. They are inversely proportional to the size of the database and are a monotonically increasing function of the number of database copies in the network.

Special care must be taken in this estimation process if there are hot spots in the database or if some types of updates will not cause collisions.

[--more--](#)

Would you like to Sign Up for the free Digest by Fax?

Simply print out the following form, fill it in, and fax it to:

Availability Digest

+1 908 459 5543

Name: _____

Email Address: _____

Company: _____

Title: _____

Telephone No. _____

Address: _____

The Availability Digest may be distributed freely. Please pass it on to an associate.

To be a reporter, visit www.availabilitydigest.com/reporter.

Managing Editor - Dr. Bill Highleyman editor@availabilitydigest.com.

© 2007 Sombers Associates, Inc., and W. H. Highleyman