



Volume 2
Issue 6

the Availability Digest

--- achieving 100% uptime

June 2007

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:

Case Studies

HP's OpenCall INS Goes Active/Active

Never Again

VoIP PBX Succumbs to Overconfiguration

Best Practices

Katrina - The Harsh Teacher

Active/Active Topics

Let's Get an Availability Benchmark

Recommended Reading

Fire in the Computer Room, What Now?

Product Reviews

ETI-NET's Virtual Tape Server

The Geek Corner

Nodes, Subsystems, and Systems

Complete articles may be found at
www.availabilitydigest.com.

This is the month for the ITUG annual Summit meeting, held with the HP Technology Forum (HPTF) in Las Vegas from June 19 to June 21. As usual, this conference will be focused heavily on achieving very high system availability.

I will be very active at the conference and will be giving papers and hosting a panel. My paper, Active/Active in Action, presents several case histories of production active/active systems that are today achieving availabilities of six 9s and beyond. I am also presenting Active/Active versus Clusters, a comparison of the availability attributes of these two very important technologies.

In addition, I am hosting a panel on Active/Active Replication for SQL/MX tables and am co-moderating the Business Continuity SIG (Significant Interest Group).

I look forward to seeing those of you who will be at the conference with me. Please be sure to say "hello."

Dr. Bill Highleyman, Managing Editor

Case Studies

HP's OpenCall INS Goes Active/Active

A major telecommunications company is now providing an extended version of HP's OpenCall INS cell phone application to cell phone service providers. Running on NonStop servers in an active/active configuration, this

product is highly scalable. Furthermore, this configuration can survive multiple node failures and will virtually never be out of service.

The system configuration is a hierarchical active/active system, with one node acting as Master to the other Slave nodes.

System capacity can be easily modified by adding or removing Slave nodes. When a node is added or removed, the Master node is notified of the new configuration.

From an availability viewpoint, any number of Slave nodes may fail; and the system will continue to be operational within the capacity capabilities of the remaining nodes. Should a Master system fail, one of the Slave nodes is promoted to be the new Master; and the system continues in operation.

Operating the HP OpenCall INS applications in an active/active system configuration such as this provides unparalleled scalability and availability for cell phone service providers. Service availability is commensurate with what one expects from the telephone network. Downtime is virtually a thing of the past.

[--more for subscribers--](#)

Never Again

VoIP PBX Succumbs to Overconfiguration

A software product development firm specializing in high-availability solutions received a real-life lesson in the very principles which it teaches. A fundamental precept of high-availability systems is that the bigger they are, the easier they break. The company's new Voice over Internet (VoIP) PBX was overconfigured by the PBX vendor. The result – an early failure that didn't need to happen.

The system used an open source product to provide PBX functions for local and remote corporate telephones over its IP network. The PBX server was configured with four CPUs when only two were needed to provide the required capacity. Since the server was incapable of configuring around a failed CPU, any CPU failure would take down the server. Using twice as many CPUs as necessary increased the probability of server failure by a factor of two.

[--more--](#)

Best Practices

Katrina – The Harsh Teacher

The year 2005 brought with it the largest and most destructive hurricane season on record in U. S. history. The queen of the season was hurricane Katrina, which was the most devastating storm to hit the Gulf Coast Region in over 200 years. Its cost to the national economic infrastructure, estimated to exceed \$200 billion dollars, was greater than any other storm in the history of the United States.

Katrina taught us many painful lessons, ranging from evacuation to communications, emergency management, tracking missing people, family reunion, property security, protection of individuals, municipal infrastructure, and a whole lot more. Buried in these massive problems were the corporate problems of business continuity.

When we look at the problems that companies in the area faced as a result of Katrina, the IT infrastructure, though terribly important, was just one aspect of the multifaceted business continuity problem faced by these companies. Shortly after the business effects of the storm began to dissipate, IBM shared their learning experiences in an in-depth Web seminar. In this article, we review these lessons.

[--more for subscribers--](#)

Active/Active Topics

Let's Get an Availability Benchmark

Benchmarks have been around almost as long as computers, Benchmarks are immensely valuable to buyers of data processing systems for comparing cost and performance of the systems which they are considering. They are of equal value to system vendors to compare the competitiveness of their products with others.

Virtually all of today's accepted benchmarks focus on performance and cost. But what good is a super fast system if it is down and not available to its users? It seems that system availability is just as important today to system purchasers as is raw performance. Shouldn't availability be part of these benchmarks?

There are many problems to solve in order to achieve meaningful availability benchmarks, primarily due to the infrequency of failures in the field. However, a careful analysis of availability attributes can lead to a good start toward obtaining useful availability metrics.

[--more for subscribers--](#)

Recommended Reading

Fire in the Computer Room, What Now?

Heading to work, you hear on the radio that your office building has just burned down. Your cell phone rings. It's the CEO asking what your plans are to recover. How soon will it be before you have recovered the company records and are providing data processing services again? You *have* planned for this, haven't you?

The recovery from a disaster such as this requires extensive disaster recovery planning long before a disaster strikes. This book, *Fire in the Computer Room, What Now?*, walks us through the creation of a Disaster Recovery Plan to handle just this sort of situation. Following such a disaster, if you have to ask, "What Now?," then it is already too late to create this plan.

[--more for subscribers--](#)

Product Reviews

Virtual Tape for NonStop Servers with ETI-NET's EZX-BackBox

The EZX-BackBox virtual tape system from ETI-NET provides a powerful virtual tape robot solution for all HP NonStop servers and allows NonStop systems to use existing corporate storage environments.

EZX-BackBox emulates one or more native tape devices attached to HP NonStop servers. It is seen by the operating system as standard tape drives connected to SCSI or fiber channel ports. The standard NSK tape process is used. Virtual volumes are managed using unmodified Guardian media manager software such as DSM/TC or TMFCOM. Enterprise storage management products from IBM/Tivoli, Veritas, Legato, and others can easily be integrated for archive tape management and cross-platform backup consolidation.

Data deduplication is used to significantly reduce disk storage and bandwidth requirements. Data deduplication monitors the content of backup data streams being generated by the host tape process and writes to virtual storage only the actual data elements that have changed (for instance, only a record or a row) rather than an entire file or table.

[--more for free--](#)

The Geek Corner

Calculating Availability – Nodes, Subsystems, and Systems

In our previous articles, we have described multinode architectures that can provide very high availabilities. Predominant among these architectures are active/active and clustered systems. These *systems* are made up of *nodes* that are themselves computing systems.

In our analyses of system availability in our Geek Corner articles, we have talked throughout of *systems* and the *subsystems* that make up these systems. It is now time to bring the terms *system*, *subsystem*, and *node* together into a consistent whole.

The equivalence is that a *node* in a cluster or in an active/active system is a *subsystem* of that *system*. Sometimes, a node itself is a system and comprises multiple subsystems of its own. In these cases, the availability attributes of the nodal system are first calculated based on its subsystem attributes. These nodal availability attributes are then carried over as subsystem parameters for the calculation of the availability of the full active/active or clustered system.

[--more for subscribers--](#)

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

The free Digest, published monthly, provides abbreviated articles for your review.

Access to full article content is by paid subscription only at

www.availabilitydigest.com/subscribe.

Name: _____

Email Address: _____

Company: _____

Title: _____

Telephone No. _____

Address: _____

The Availability Digest may be distributed freely. Please pass it on to an associate.
Access to most detailed article content requires a paid subscription,
To sign up for the free Availability Digest or to subscribe, visit www.availabilitydigest.com/subscribe.
To be a reporter (free subscription), visit www.availabilitydigest.com/reporter.
Managing Editor - Dr. Bill Highleyman editor@availabilitydigest.com.
© 2006 Sombers Associates, Inc., and W. H. Highleyman