

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

Cellular Provider Goes Active/Active

Never Again

The Planet Blows Up

Availability Topics

Active/Active Systems - A Taxonomy

The Geek Corner

Why Are Active/Active Systems So Reliable?

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

When Will We Ever Learn?

Systems Fail. Data centers fail. The Internet fails. Yet businesses foolishly hang on to a blind faith that such failures will not impact the Internet services upon which they so desperately depend.

In our Never Again articles, we have described over and over again the many failures of systems that provide services via the Internet for thousands of customers, systems run by such major service providers as Rackspace, Amazon, Hostway, Salesforce, PayPal, and, in this issue, The Planet. When these systems fail, they are often down for days.

Yet there is an awesome absence of backup plans throughout the entire service chain, from the dedicated server providers to the hosting services that use these servers to the end customers trying to run their online stores. The net result – millions of dollars in lost revenues for online retailers whenever such an outage occurs.

Backup plans must be in place for a failure anywhere along the service chain. It's appropriate if high availability is sold as an add-on service, such as Amazon's newly announced Availability Zones. What is inappropriate is for the online store to have no fallback when its web site fails.

Here in the Availability Digest, we cover topics that can lead to a solution to this growing problem. We encourage you to talk to us if the loss of your Internet services can throw you into a financial crisis.

Dr. Bill Highleyman, Managing Editor

Case Studies

Cellular Provider Goes Active/Active for Prepaid Calls

The premier cellular service provider in greater Africa serves over 20 million subscribers in South Africa, Mozambique, Tanzania, Lesotho, and the Democratic Republic of the Congo.

Prepaid calling cards represent a major portion of its subscriber activity. Should the system that authorizes calling card calls go down, a major part of Africa's cellular service is lost.

The system that provides the company's online authorization and tracking of calling card calls is its Prepaid Front End (PPFE). To ensure the continued availability of this system, the cellular provider has implemented the PPFE as a two-node active/active system using NonStop nodes. The nodal databases are kept in synchronism via asynchronous data replication.

The PPFE provides information via data replication to several other ancillary systems that provide billing, web access to card information, card recharge and merge functions, and a data warehouse of subscriber activity.

[-- more --](#)

Never Again

The Planet Blows Up

On Saturday, May 31, 2008, an explosion blew out three walls in the Houston data center of The Planet, one of the world's largest providers of dedicated servers for more than 22,000 businesses. The resulting damage disabled 9,000 servers used by 7,500 web-hosting companies, taking down web sites serving millions of customers. It was days before service was completely restored.

An important lesson to be learned from this incident is that it is impossible to protect a data center against all disasters. You can't anticipate what you don't know. The only option is to ensure that there is a backup plan in place that will allow the company's services to continue at an alternate site should the data center be taken out of service.

More importantly, it is imperative for every link in the computing chain to have a backup plan, from the primary service provider to the wholesaler who is reselling capacity and to the retailer who is the end user.

[--more--](#)

Availability Topics

Active/Active Systems – A Taxonomy

Active/active systems are characterized by the attributes of availability, survivability, scalability, and consistency. They eliminate planned downtime and have expected mean time between failures that can be measured in centuries. They can survive the failure of an entire data center. They are scalable by simply adding computing resources with no impact to the users. They will consistently execute operations in a predictable manner across the application network.

A common way to implement active/active systems is via context-free virtualized pools of compute resources. Each pool is virtualized to appear as a single resource to external users. If a member of a pool fails, work that it had in progress is simply resubmitted to another member. Thus, “failover” becomes “resubmission.”

At least one of the resource pools must maintain application state, usually by managing the application database. Since the database is a virtualized pool, there will be two or more copies of the application database in the database pool. These database copies must be kept in synchronism, which is typically done via asynchronous or synchronous replication.

In this article, we develop a taxonomy that can be used to classify various active/active architectures. The taxonomy is applied to several examples of active/active systems in production today.

[--more--](#)

The Geek Corner

Why Are Active/Active Systems So Reliable?

Active/active systems can achieve availabilities of six 9s and beyond. Six 9s is an average of just 30 seconds of downtime per year. These systems achieve such high availabilities by providing very rapid recovery from faults – recovery times measured in seconds or subseconds. In fact, if the recovery time is fast enough, users will not realize that there has been a fault. In effect, no fault has occurred.

In addition, failover faults, which can plague active/backup systems, are almost nonexistent in active/active systems because the other nodes in the system are known to be operational. After all, they are actively processing transactions.

In this article, we analyze the downtime probability of a system that is subject to failover times and failover faults and use this analysis to compare active/active systems to other active/backup redundant configurations. It becomes clear that active/active systems achieve their high availabilities via the philosophy of “*Let it fail, but fix it fast.*” “Fix it fast” is achieved via the technique of “*Resubmit rather than fail over.*”

[--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.

© 2008 Sombers Associates, Inc., and W. H. Highleyman