



the

# Availability Digest

Volume 1  
Issue 1

--- achieving 100% uptime

October 2006

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

Tackling Switchover Times

Do You Know Where Your Train Is?

### Never Again

The Case of the Flying Cable

Software Bug Causes Train Wreck

Bear Terrorizes Squirrels

### Best Practices

With 100% Uptime, Do I Need a

Business Continuity Plan?

### Active/Active Topics

All About Continuous Processing Architectures

What is Active/Active?

### Recommended Reading

Business Continuity Planning: IT

Examination Handbook

### Product Reviews

How Much Will Active/Active Cost Me?

### The Geek Corner

Calculating Availability

What Is That Nerd Logo?

Complete articles may be found at  
[www.availabilitydigest.com](http://www.availabilitydigest.com).

Dear Reader:

Welcome to the inaugural issue of the Availability Digest. It is devoted to the search for 100% uptime in our data processing systems, with particular focus on active/active technology.

The Digest is freely distributed monthly and contains a synopsis of articles of interest relative to 100% uptime. The full articles may be viewed by clicking on the -- more -- button at the end of each synopsis, or they may be found on our web site at [www.availabilitydigest.com](http://www.availabilitydigest.com). Access to the full article content requires a subscription to the Availability Digest except for this inaugural issue, in which we are happy to provide you with the full content so that you can see the quality of our articles.

The Digest will be distributed on an unsolicited basis again in November. However, to avoid spamming, we will stop unsolicited distribution shortly. Therefore, if you would like to continue to receive your free Digest, please sign up now at our web site. There you may also subscribe to the full article content.

We are always looking for good content and welcome you to become an AD Reporter in return for a free subscription. See our web site for details.

Dr. Bill Highleyman, Managing Editor

---

## Case Studies

### ***Tackling Switchover Times***

For many companies, going to active/active is a major effort. In order to improve the availability of their data processing services in the short term, many are deciding to focus instead on reducing the switchover time for planned outages. Current switchover processes make the systems unavailable, often for hours.

This is what Sabre Holdings and Mayo Clinic are doing with their NonStop systems. Both have reduced their switchover times from several hours to a few minutes. Sabre got rid of long-running transactions and prevented batch jobs from being scheduled during the switchover process, among other initiatives.

Mayo upgraded its backup system to be a full production system so that once switched over, there is no need to have to switch back to the old active system. Among Mayo's other improvements was the modification of its file copy procedures to significantly reduce file copying activity during switchover.

Come hear the presentations by Sabre and Mayo at the ITUG 2006 Summit in October. Their presentations describe these efforts and successes. Sabre's paper is entitled "Continuous Availability – Reducing Switchover from Hours to Minutes." Mayo's paper is "Use of NonStop Technology for an Electronic Medical Record: A Journey." [-- more --](#)

### ***Do You Know Where Your Train Is?***

A transit authority in a major city has gone to an active/active system to monitor its train traffic. It is using a pair of Sun servers to receive track occupancy, switch, and signal indications from the field; to display train locations on track maps on the controllers' consoles and a large overhead display; and to send controller-generated commands to the field.

Schedules and event logs are recorded on disk and are replicated to the other system via Oracle's data replicator. However, system state information such as train positions and switch and signal states have only short-term but immediate value and are therefore replicated over a redundant LAN connecting the two systems.

Since all field indications must be processed in exact time sequence, the system cannot be run in true active/active mode with both systems independently processing indications. Therefore, the system is run in "sizzling hot standby" mode, with the standby system being ready to take over the control functions in less than a second. The switchover of users is done via a unique technique which uses the gratuitous ARP capability of IP. [-- more --](#)

---

## Never Again

### ***The Case of the Flying Cable***

A fault-tolerant system depended upon communications with the outside world. Therefore, its communication interface was totally redundant. It had dual communication processors connected to dual LANs driving redundant communication lines. Each set of equipment was powered by an independent external power source.

The communication subsystem hadn't failed for over a decade and seemed solid. Not quite! One day, a technician was pulling cables through the false flooring when a cable end came loose and went flying. It collided with the power strip supplying one side of the redundant network equipment. The on/off switches on the power strips had been disabled so that they would not be accidentally turned off. However, the recessed circuit breakers were active for safety reasons.

As luck would have it, the corner of the cable connector went into the recess of one of the circuit breakers and tripped it. No problem, right? After all, the other side of the communications system would carry on. Wrong! Long ago, with no one noticing, all of the communication equipment had been plugged into only one power strip; and that is the one that got hit. That ended a decade of perfect availability. [-- more --](#)

### ***Software Bug Causes Train Wreck***

An international long-haul railroad uses a fault-tolerant system to track its trains. Controllers monitor train movements and can control the trains by positioning switches and changing signal states.

The system had been in operation for over a decade with its normal set of problems, all of which had been corrected. However, unbeknownst to the operations staff, an obscure bug lurked inside. There was one particular set of parallel tracks, one northbound and one southbound, that would lose the southbound train display if trains were on each set of tracks. The southbound train would become a ghost train. Since traffic was light, this hardly ever happened.

But one day it did while the controller had diverted his attention. When he returned to his console, he saw that the track was unoccupied and cleared the next train onto that track. Unfortunately, the track was not unoccupied – the ghost train was still there.

Also unfortunately, at this time the engineer of the trailing train was “otherwise occupied” and did not see the train in front of him. The resulting crash caused significant damage to the trains and the track. Luckily, no one was seriously injured because of the low speed of the collision. [-- more --](#)

### ***Bear Terrorizes Squirrels***

For fifteen years in our home in northern New Hampshire, we've been putting out peanuts to feed the squirrels, chipmunks, blue jays, crows, and sometimes even - yes - fox. We do this to entertain our dogs.

This summer, a giant intruder decided that he liked peanuts, too. It didn't take long for him to decimate the entire stash. Sorry, little guys, but I guess we are going to have to wait for hibernation time before we can treat you again.

What does this have to do with availability? Well, what occurred was an unanticipated event that denied the little critters their peanut service for several months. Don't let this happen to you. Subscribe to the Availability Digest. *-- a paid advertisement --*



# Best Practices

## ***With 100% Uptime, Do I Need a Business Continuity Plan?***

As much as we would like to achieve 100% uptime, we never can. There will always be the possibility of some unforeseen event that could take down our active/active application network, like the 2004 Northeast blackout in North America or the worldwide SARS epidemic (which fortunately, so far, has fizzled).

Even the theory behind the extreme availability of active/active systems predicts that they can fail (see our article entitled "[Calculating Availability](#)"). Though the probability of the simultaneous failure of enough nodes to take down the system may equate to centuries of uptime, that does mean that such a failure could occur in the next few centuries. And the day that unlikely event occurs could be tomorrow.

If you are not prepared to cope with that "impossible" event, you are simply playing Russian Roulette with your business. Every business needs an up-to-date and tested Business Continuity Plan. [-- more --](#)

---

## Active/Active Topics

### ***All About Continuous Processing Architectures***

We can never achieve 100% uptime, but we can come arbitrarily close. There are several techniques in common use today to achieve extreme availabilities, including active/active systems, lock-stepped processors, and synchronized process pairs.

An active/active system is one in which a set of independent nodes in an application network cooperate in a common application. Should a node fail, its users can be immediately switched to a surviving node. Active/active systems are in use by many financial institutions, among others, especially in the NonStop community.

Lock-stepped processor pairs are an old technology still in use today but primarily as a building block in fault-tolerant systems. Both NonStop servers and Stratus ftServers use lock-stepped processors that provide fast-fail, with other processors available to pick up the load of a failed processor pair.

Synchronized process pairs are used in NonStop systems to provide services to processes in the event of the failure of their assigned processors.

Sometimes, all three of these technologies are in use in one system, such as in NonStop active/active systems. [-- more --](#)

### ***What Is Active/Active?***

An active/active system is a network of independent processing nodes, each having access to a common replicated database such that all nodes can participate in a common application. Should a node fail, its users can be immediately switched to a surviving node. Therefore, an active/active system achieves its extreme availability through the philosophy of *let it fail, but fix it fast*.

The database copies may either be network-attached or attached directly to two or more nodes in the network. The distributed database copies in the application network are kept in synchronism by replicating changes made in any one database to all of the other database copies in real time. There are at least three ways in which data can be replicated – by network transactions, by asynchronous replication, or by synchronous replication. Each has its advantages and disadvantages.

In addition to extreme availabilities, active/active networks can also provide disaster tolerance, data locality, elimination of planned downtime, capacity expansion, load balancing, risk-free failover testing, and lights-out operation. With an active/active system, all purchased capacity is used. There are no standby systems waiting idly by to pick up the load of a failed processor. [-- more --](#)

---

## Recommended Reading

### ***Business Continuity Planning: IT Examination Handbook***

Business continuity planning is not really a focus of the Availability Digest. It is a massive subject in itself for which there is a wealth of literature.

However, if you are interested in approaching 100% uptime, you had better be very interested in business continuity planning. No matter how reliable you have built your system, there is still a possibility, no matter how small, that your system will fail someday. At that time, you had better be prepared with an up-to-date and well-rehearsed plan to cope with that outage.

What better way to learn about what should go into a good Business Continuity Plan than to peek at the handbook that the auditors use. That handbook is the Business Continuity Planning: IT Examination Handbook.

The Handbook clearly discusses everything that should go into a good Business Continuity Plan. It presents the four main stages of business continuity planning. The Business Impact Analysis identifies the potential impact of unanticipated disruptions on business processes. The Risk Assessment phase prioritizes these threats. The Risk Management phase is the actual preparation of the plan. The Risk Monitoring phase is the practice and testing of the plan. [-- more --](#)

---

## Product Reviews

### ***How Much Will Active/Active Cost Me?***

Moving to an active/active architecture is a big step. One of the biggest questions is how much will it cost?

On the one hand, there is additional cost for multiple systems, for multiple software licenses, for expanded network capabilities, for additional facilities and staff if disaster tolerance is to be achieved, and so on.

On the other hand, there is the savings of downtime costs, a savings which in some cases can be significant (some firms report costs in excess of six figures per hour). In addition, no unused spare processing capacity need be purchased – all purchased capacity is used.

The total cost of ownership (TCO) and the return on investment (ROI) are complex tasks. Fortunately, there is a service available from The Standish Group that will help you do this. It is called VirtualADVISOR, and there is no software required to be purchased or installed. You work with a Standish advisor to review your various parameters. Using large databases of over 50,000 projects, VirtualADVISOR estimates what your TCO and ROI will be based on different approaches and will provide you with a detailed report that you can share with your management.

[-- more --](#)

---

# The Geek Corner

## Calculating Availability

The availability of an active/active application network is a function of the availability of the individual nodes, the number of nodes in the system, and the number of those nodes that are spares - that is, the number of nodes that can fail and leave the system still operational. Though there are several complicating factors, the basic availability equation is

$$A = 1 - f(1 - a)^{s+1}$$

where  $A$  is the system availability;  $a$  is the availability of a node;  $s$  is the number of spare nodes; and  $f$  is the number of ways that  $s+1$  nodes in a network of  $n$  nodes may fail, thus taking down the system..

The derivation of this relationship yields some interesting and perhaps nonintuitive rules:

- ♦ Adding more nodes to an active/active network makes it less reliable if the spare count remains unchanged.
- ♦ Adding an additional spare node to the system increases the availability of the system by the number of 9s which represents the node's availability – almost.

[-- more --](#)

## What's That Nerd Logo?

It's the availability equation, of course, though admittedly somewhat stylized. It gives the probability of failure of a redundant system with one spare. The system will survive the failure of any one component; but if two components fail, the system is down.

$ff^2$

The first "f" represents the number of ways that two components can fail. The second "f" represents the probability that any one of the components will fail. Therefore, the probability that two specific components will fail is  $f^2$ . The probability that the system will fail is the probability that two specific components will fail multiplied by the number of ways that two components can fail, or  $ff^2$ , if you will excuse our stylization.

[-- 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**

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

*After this issue, access to full article content will be by subscription only at*

[www.availabilitydigest.com/subscribe](http://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 subscription,  
To sign up for the free Availability Digest or to subscribe, visit [www.availabilitydigest.com/subscribe](http://www.availabilitydigest.com/subscribe).  
To be a reporter (free subscription, visit [www.availabilitydigest.com/reporter](http://www.availabilitydigest.com/reporter).  
Managing Editor - Dr. Bill Highleyman [editor@availabilitydigest.com](mailto:editor@availabilitydigest.com).  
© 2006 Sombers Associates, Inc., and W. H. Highleyman