

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](#)

[BlackBerry Takes Another Dive](#)

[Best Practices](#)

[Rules of Availability – Part 1](#)

[Availability Topics](#)

[Fault Tolerance for Virtual Environments](#)

[The Geek Corner](#)

[Heterogeneous Systems - Part 1](#)

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

Schedule an Active/Active Seminar

I just returned from South Africa, where I presented our seminar, "Achieving Century Uptimes: The Theory and Practice of Active/Active Systems," to a large and enthusiastic group of IT professionals. Africa, you say? Are they into active/active systems? You bet. Some of the largest companies in South Africa are running active/active configurations to achieve the extreme reliabilities that they need. Companies such as Vodacom, the largest cellular telephone service provider in Africa, and BANKSERV, which provides interbank electronic transaction switching services to the South African banking sector.

Are your costs of downtime concerning you? Why not consider our one-day, three-day, or five-day seminars on this exciting and useful topic? Contact us at editor@availabilitydigest.com for our syllabuses and pricing information.. We'd love to educate your IT people.

Dr. Bill Highleyman, Managing Editor

Never Again

BlackBerry Takes Another Dive

Déjà vu. It happened again. It was less than a year ago that the BlackBerry email network of RIM (Research in Motion) went down, and it took days to work through the backlog of email that had built up during the half-day outage.

Just last month, email and Internet service again disappeared from the screens of millions of BlackBerrys in North America and Canada. Again, it was over a day before RIM's email service worked off the resulting backlog and returned to normal.

The BlackBerry outage in April of 2007 was caused by a software upgrade that had not been properly tested. RIM management said that this would never happen again. Guess what? The outage in February, 2008, was caused by a software upgrade that had not been properly tested. Will they ever learn?

[--more--](#)

Best Practices

Rules of Availability – Part 1

There are many ways in use today to achieve high availabilities. Predominant among these techniques are lockstepped processors, checkpointed or persistent processes, clusters, and active/active systems. All use some form of redundancy to recover quickly from faults, and all are subject to a common set of principles.

Many of these principles are set forth in the book series entitled *Breaking the Availability Barrier*, which I coauthored. The principles are presented as sixty-four "Rules of Availability." The rules focus on active/active architectures, but many of the rules are applicable in a broader sense. In this article, which is the first part in a series on availability principles, we review some of these rules.

[--more--](#)

Availability Topics

Fault Tolerance for Virtual Environments – Part 1

Not only do businesses today depend upon information technology (IT) for their very existences, but IT costs have become a major part of an enterprise's budget. As corporate data centers become bigger and bigger, often supporting thousands of servers, their costs for hardware, space, administration, and energy are rapidly increasing.

Virtualization lets one physical server do the work of many. It does so by creating *virtual machines (VMs)*. A single physical server can host several virtual machines. Through virtualization, a data center can significantly reduce the number of physical machines that it requires and can enjoy all of the savings that go along with that reduction.

However, since a physical server is now hosting multiple guest operating systems and their applications, redundancy of physical servers in a virtual environment is necessary. Should a server fail, there must be a failover mechanism in place to rapidly move the failed virtual machines to functioning servers.

In this multipart article, we describe today's virtualization techniques. We then look at the redundancy mechanisms that are available today to provide fault tolerance. We finally review some products that offer fault-tolerant virtualization.

[--more--](#)

The Geek Corner

Calculating Availability – Heterogeneous Systems - Part 1

In all of our availability analyses to date, we have assumed that the nodes in a system are identical. But what if the nodal availabilities are not the same? What if one node is in a safe area, and the other is in Hurricane Alley in Florida? The Florida node will have an availability less than the other node because it stands to be destroyed by a hurricane at some time. What, then, is the availability of the redundant system?

In this article, we review some simple probability relationships necessary for analyzing this situation and others like it. In our next article, we apply these relationships to heterogeneous systems.

[--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