



Volume 4
Issue 8

Availability Digest

--- achieving 100% uptime

August 2009

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.
www.availabilitydigest.com

In this issue:

[Never Again](#)

[Twitter Taken Down by DDoS Attack](#)

[Best Practices](#)

[Achieving Fast Failover - Part 1](#)

[Availability Topics](#)

[The IPv4 Doomsday](#)

[Product Reviews](#)

[Replicating Windows, Linux with Double-Take](#)

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

Sign up for your free subscription at
<http://www.availabilitydigest.com/signups.htm>

Be an Active/Active Emissary

I recently gave a pair of three-day courses, "Active/Active: Theory and Practice," to a major bank. It was well-received as there is a lot of excitement in the industry about active/active technology. The course covered the theory behind the continuous availability provided by these systems, how to build them, and several real-life case studies.

The primary criticism from the attendees was that it was too NonStop-focused, with not enough detail on Linux and Windows systems, where clustering is the current technology to achieve high (though not continuous) availability. Unfortunately, the criticism is well-founded. Active/active has come out of the NonStop (formerly Tandem) world because that is where the highest needs for availability are. Though the technology has not yet made it into the mainstream of other systems, it is equally applicable to these systems – an observation that was a major focus of the course.

The good news is that over half of the attendees were from other disciplines and got the message. Hopefully, they will dip their toes in the water to explore this exciting technology. In the meantime, I feel like a missionary trying to spread the word. You can help. Please share with me your experience and insights into active/active and other continuous availability topics so that I can share them with others. Thanks.

Dr. Bill Highleyman, Managing Editor

Never Again

Twitter Taken Down by DDoS Attack

On Thursday, August 6, 2009, the Twitter social networking site went down. It suffered repeated outages, timeouts, and serious slow-downs for at least three days. What caused this failure?

To add to the mystery, Facebook and Live Journal simultaneously had similar problems. Were these outages somehow related? They occurred at about the same time as the 2009 Defcon 17 hackers conference held from July 30th to August 2nd. Could this have been some misguided mischief?

Working in concert, it took but a few hours for these sites to determine the reason for their problems. They were under a distributed denial of service attack (DDoS). Spam messages were flooding their sites, all being queries against the blog of a single user who went by the user name Cyxymu. Clearly, someone was out to silence Cyxymu. But why?

[--more--](#)

Best Practices

Achieving Fast Failover in Active/Active Systems – Part 1

Active/active systems provide continuous availability not because they avoid faults but because they can recover from faults so quickly that users don't notice that there has been an outage. This capability requires not only that failover to a backup component be rapid but that it be reliable.

An active/active system comprises two or more geographically separated processing nodes that cooperate in a common application. This requires that each node has access to a local copy of the application database. The databases are kept synchronized via data replication.

Should a node fail, all that is required to recover from the fault is to move the users or to reroute transactions to one or more surviving nodes. Failover can be accomplished in seconds or even in subseconds. Furthermore, failover is risk-free and reliable because it is known that the surviving nodes are operational. After all, they are actively processing transactions.

But how can users or transactions be moved so quickly between processing nodes? That is the subject of this two-part article.

[--more--](#)

Availability Topics

The IPv4 Doomsday

Remember the Year 2000 Doomsday? Though potentially a worldwide disaster, the Y2K problem only affected individual systems; and massive efforts by system developers avoided most problems.

But what if instead the problem had affected the entire Internet? Not only might these individual systems have been taken down, but *all* systems except those few that lived in isolation might have been lost. Without the Internet, much of today's global commerce would come to a halt.

Well, we are facing just that problem. The issue is the Internet Protocol (IP) that interconnects systems and users all over the world via a global packet-switching network. The current version of IP, IPv4, has an address space that seemed inexhaustible when the Internet was first being specified. After all, IPv4 used a 32-bit address field that provided over four billion unique addresses.

However, various estimates now indicate that the IPv4 address space will be exhausted in the next two years or so. What now? The planned answer is a massively extended address space that is made available in the next version of IP, IPv6. But is IPv6 ready for prime time? We look at the issue in this article.

[--more--](#)

Product Reviews

Replicating Windows and Linux Environments with Double-Take

Double-Take is a unidirectional asynchronous data-replication engine that replicates file changes in Windows and Linux environments over unlimited distances. It intercepts file and directory changes and replicates them in real time to one or more target servers, where the changes are applied to the target file systems. Likewise, multiple source servers can replicate their files to a single target server, where an image of each source file system is maintained.

For data writes, Double-Take replicates only byte changes. Entire files need not be replicated, thus minimizing network load and replication latency. Either upon command or automatically, Double-Take will compare a target file system to its source and correct any errors that have occurred.

With Double-Take's mirroring capability, files can be migrated to other servers by copying the files first and then keeping them updated via replication. Users are unaffected during the migration process.

Should a source server failure occur, Double-Take can automatically fail over to the target server and fall back if desired when the source server is returned to service.

Any file can be included in the list of files to be replicated. Thus, an entire system can be replicated to the target server. Should the source server fail, the entire operational server can be quickly restored onto the target server.

[--more--](#)

Sign up for your free subscription at <http://www.availabilitydigest.com/signups.htm>

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.

© 2009 Sombers Associates, Inc., and W. H. Highleyman