

the *Availability Digest*[™]

Volume 5
Issue 10

- - - *achieving 100% uptime*

October 2010

Thanks to This Month's Availability Digest Sponsor



Multifaceted messaging middleware. Enables integration of diverse applications across the enterprise. Simple [API](#). Supports multiple messaging models. Fast, light, trouble-free, low-latency environment. [Solutions](#) include Secure RFT, Web Enablement, Secure Communications, Interactive & Queued Messaging. HP NonStop & OpenVMS, IBM z/OS CICS & iSeries, Stratus, Linux, Windows, HP-UX, IBM/AIX, Oracle/Solaris

In this issue:

[Never Again](#)
[The State of Virginia - Down for Days](#)
[Best Practices](#)
[2010 NonStop Availability Award](#)
[Recommended Reading](#)
[Megaplex Modeling: Future of NonStop](#)
[The Geek Corner](#)
[Simplifying Failover Analysis](#)

Browse through our [Useful Links](#).
Check our [article archive](#) for complete articles.
Sign up for your [free subscription](#).
Join us on our [Continuous Availability Forum](#).

Is Data Replication a Sufficient Backup?

One of our readers asked a very pertinent question. If you are creating a backup database via data replication, does that mean that you don't have to back up your database to tape? After all, you have two copies of the application database. Isn't that the same as backing up files on a single system?

I asked several others at the NonStop Symposium in September what they thought, and the general reaction that I received was the same as mine - "Arrgh!" But when asked why, there was often a pause.

Everyone agreed, though, that a major reason to back up is data corruption. If the source database becomes corrupted, that corruption is replicated to the target database; and you have lost both copies of your database. Recovering from a magnetic tape or a virtual tape backup is your only salvation.

In our Never Again story in this issue, the State of Virginia shut down twenty-six agencies when its fault-tolerant EMC database crashed due to a maintenance error and lost the agency's databases. Only the magnetic tape backup saved the state, and even then it took seven days to recover.

If you have views on this issue, please post them on the [LinkedIn Continuous Availability Forum](#). I plan to write about this topic in the next Digest.

Dr. Bill Highleyman, Managing Editor

Never Again

The State of Virginia – Down for Days

The Commonwealth of Virginia lost dozens of its computer systems for over a week, bringing the activities of over two dozen of its agencies to a halt. Tens of thousands of the state's citizens were affected, some seriously. The web site of the Division of Emergency Management went dark just as Hurricane Earl was approaching.

How could so many major systems go down simultaneously? Why did it take so long to restore services? As with any major outage, there was a chain of events. If any of these events had been avoided, the outage might not have happened. It started with a maintenance error, was made worse by poor backup procedures, and evidently was exacerbated by a lack of recovery-procedure testing.

[-- more --](#)

Best Practices

2010 NonStop Availability Award

Every year, the NonStop Availability Award is given to the NonStop user that has demonstrated superior high-availability practices. The award is determined based on four criteria:

- Outage Minutes
- Configuration and Operational Complexity
- Use of Best Practices for Outage Prevention
- Quality of Outage Data

Congratulations to this year's winner, Bank-Verlag of Germany. Congratulations also to the two runner-ups – Belgacom of Belgium and VocaLink of the U.K.

The NonStop Availability Award is a user-group sponsored award started years ago by ITUG, the International Tandem User Group. It is now sponsored by Connect, the HP Business Technology Community.

The NonStop community is focused on high availability – that is what HP NonStop systems are all about. The NonStop Availability Award is a recognition of those companies that have carried this technology to the extreme – often zero downtime with availability best practices.

[-- more --](#)

Recommended Reading

Megaplex Modeling: The Future of NonStop Demand

In its earlier papers, [Megaplex: An Odyssey of Innovation](#) and [Roadmap to the Megaplex](#), the Standish Group traced the history of the Tandem computer from its development to its current incarnation as an HP NonStop server. The papers described the innovations brought to computer technology by the Tandem systems with their fault-tolerant multiprocessor architecture, and Standish envisioned a new NonStop architecture that it coined the "Megaplex."

In its third paper, [Megaplex Modeling: The Future of NonStop Demand](#), Standish homes in on a specific architecture for the Megaplex and compares its cost with that of more traditional approaches. Standish envisions systems of HP blades that can run any operating system. For the Megaplex, Standish focuses on NonStop systems and Linux systems running on common blades. This means that any blade in the Megaplex can run highly critical applications in a NonStop environment and can run everyday applications in a Linux environment.

The Megaplex is sized to accommodate the median load. If the load increases, less critical services give up their blades to critical services to maintain critical service performance. Capacity is on-demand, and pricing is based on actual usage.

[-- more --](#)

The Geek Corner

Simplifying Failover Analysis

A system is down if it is not providing service to its users. To improve service availability, it is common to provide redundancy in the system. Typical system configurations that are used to provide high- or continuous availability through redundancy include active/passive systems and active/active systems.

A redundant system is certainly down if all redundant components required to provide service fail. However, it is also down for all users that are in the process of being failed over following a single-node failure. Until those users are once again connected to a properly operating system, they must sit idle. Furthermore, failover is not always successful. If failover fails, users are down until one of the systems is returned to service. This is called a *failover fault*.

In this article, we look at the impact of failover on the availability of services to the user. The analysis leads to a surprisingly simple technique for computing the net availability of a redundant system when failover is considered. It also shows by example how reasonably fast and reliable failovers can still have a dramatic impact on system availability.

[-- 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 is published monthly. It may be distributed freely. Please pass it on to an associate.
Managing Editor - Dr. Bill Highleyman editor@availabilitydigest.com.
© 2010 Sombers Associates, Inc., and W. H. Highleyman