



the Availability Digest

Volume 2
Issue 11

--- achieving 100% uptime

November 2007

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

Asymmetric Active/Active at Banco de Credito

Never Again

Triple Redundancy Failure on Space Station

Availability Topics

Distributed Time Synchronization - Part 1

Product Reviews

Time Synchronization for NonStop Servers

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

An Availability Digest subscriber from Singapore writes:

"Just cannot fathom and accept why in today's advanced IT world, system/database upgrades need to bring down critical national services (which provide citizenship and PR applications, birth registrations, collections for passports, entry visas, and various foreign visitor passes such as student passes, long-term social visit passes and professional visit passes, etc.) for days! Surely, there has to be ways to reduce the system downtime so as to minimize the disruption to the public."

There is no reason at all with today's technology that such a situation should occur. This sounds like an old mainframe system. I would guess that one problem is the legacy problem. There is so much investment in old legacy systems, and their software is most likely not portable, that it is just too much of an investment to come into the modern world - especially for a government agency that is under tight budgets. It would take a public outcry to improve such a situation.

The technology is there to correct problems such as this. The world will move in this direction – eventually.

Dr. Bill Highleyman, Managing Editor

Case Studies

Asymmetric Active/Active at Banco de Credito

Banco de Credito e Inversiones (Bci) is the third largest bank in Chile and serves 10% of Chile's population of 16,000,000 people. It had been running a pair of NonStop servers in an active/standby configuration. However, to improve availability and to gain additional capacity, the bank decided to reconfigure its systems as an active/active pair.

In order to minimize application modifications, it has implemented an asymmetric active/active network in which one of the NonStop servers is the master node and the other is the slave node. All write transactions are routed to the master node, which performs the database updates. Modifications to the master's database are replicated to the slave node's database so that either node may process read-only transactions.

By using a master/slave asymmetric configuration, data collisions are avoided. This has resulted in minimizing the need for application code changes.

[--more--](#)

Never Again

Triple Redundancy Failure on the Space Station

In June, 2007, a triply-redundant attitude and environmental control computer provided by Russia failed on the International Space Station (ISS). Had this been a mission to Mars, it would have been fatal. Only the space station's proximity to Earth, which put it in range of support and resupply missions, prevented a tragedy.

Though the problem was circumvented in a few days by the space station crew, it took weeks for the station crew and ground engineers to determine the source of the problem. It turned out to be caused by a single point of failure in the otherwise triply-redundant system – a failure that was highly unlikely but one that occurred anyway. The culprit was a power surge monitor that could command all three computers to shut down if a power surge were detected.

During this experience, many technological and diplomatic lessons were learned. Understanding the interaction between the crew members was imperative since the U.S. and Russia are bound to be partners in a Mars mission.

[--more--](#)

Availability Topics

Time Synchronization for Distributed Systems – Part 1

In many distributed systems, it is imperative that each node march to the same time. In fact, this requirement is often extended to include client systems that access nodes in the distributed system.

Time synchronization is a complex problem across distributed systems. Each node in a distributed system necessarily has its own clock. Communication between nodes can take tens of milliseconds and, especially over the Internet, can be quite variable and unpredictable. Consequently, tight (e.g., submillisecond) time synchronization cannot be achieved simply by having one master timekeeping server send timing messages to the other nodes in the network.

In this series of articles, we discuss current technologies for maintaining time synchronization across the network. The most common approach today is NTP (Network Time Protocol), used extensively over the Internet. However, some synchronization problems go beyond the scope of NTP and require additional technologies.

In Part 1, we discuss the basics of how NTP adjusts time between systems. Part 2 discusses additional NTP features that make these adjustments more precise. Part 3 discusses a totally different approach, called Lamport logical clocks, which has the potential of virtually eliminating data collisions in active/active systems without reverting to synchronous replication.

[--more--](#)

Product Reviews

Time Synchronization for NonStop Servers

In our companion article, [Time Synchronization for Distributed Systems – Part 1](#) (November, 2007, *Availability Digest*), we described the Network Time Protocol (NTP) so ubiquitous on the Internet for maintaining time synchronization between nodes in a distributed system and a civil time reference source. NTP is an open-source facility that is available on a wide variety of platforms, including Windows, Unix, and Linux.

Unfortunately, there is not an NTP open-source port available for HP NonStop servers. Rather, these services are provided by proprietary products that are compatible with NTP running on other systems. In this article, we review two of these products. One is from Bowden Systems, and the other is from HP.

[--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.

© 2007 Sombers Associates, Inc., and W. H. Highleyman