

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

[Major ISP Migrates from Sybase to NonStop](#)

#### Never Again

[Sydney's M5 Tunnel Closed Again](#)

#### Availability Topics

[Choosing a Database of Record](#)

#### Product Reviews

[Scaling MySQL with Continuent's uni/clusters](#)

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

### A Slow System is a Down System

The November 4<sup>th</sup> U.S. election may well set a participation record. That's good. The large turnouts resulted in long queues of voters waiting for over five hours. That's bad. For those voters who gave up and left the line, so far as they were concerned, the polling station was down.

This leads me to ponder the availability of overloaded servers in a computing infrastructure. If one second response times are expected, but a heavily loaded system takes twenty seconds to respond, so far as most users are concerned, the system is down.

System response time is often specified by a Service Level Agreement in the form of "95% of all transactions will complete within one second." If this performance is violated, the availability of the system suffers.

How do we size a system to meet a performance SLA? We will devote the next three issues of The Geek Corner to addressing this issue. Though the math is complex, we will provide easy-to-use spread sheets to answer this important question for you. Look for our introduction to this solution in next month's Digest.

Dr. Bill Highleyman, Managing Editor

---

## Case Studies

### **Major ISP Migrates from Sybase to NonStop with No Downtime**

The rapid growth of a major Internet Service Provider (ISP) led to capacity strains in its IT infrastructure. One such issue recently occurred in its login subsystem. Though comprising a large farm of redundant servers running many instances of Sybase on Linux, this login subsystem had reached the limits of its capacity.

The ISP therefore decided to architect and build an entirely new login subsystem using NonStop servers in an active/active configuration. Not only would the single-system image presented by the active/active system make the login complex significantly more manageable, but failover time could be reduced to seconds; and capacity could be added easily.

The problem became how to migrate from the old Sybase/Linux login complex to the new NonStop system without impacting the ISP's customers. The goal was to perform an online migration with minimal application downtime. By the judicious use of the Shadowbase data-replication engine from Gravic, Inc., the ISP was able to gracefully migrate all of its customers over a period of time to the new NonStop system with no interruption in service.

[--more--](#)

---

## Never Again

### **Sydney's M5 Tunnel Closed Again by Computer Glitch**

The M5 Motorway is a major link connecting Sydney, Australia, with points southwest. Completed in 2002, the 22-kilometer motorway is partly free and partly tolled. It currently carries over 100,000 motorists per day.

A 4.5-kilometer tunnel near Sydney airport has proven to be the Achilles heel of the motorway. Costing AU\$800 million, it has been closed six times in the six years since the tunnel opened, causing major backups and creating political outrage. Three of these outages were due to the failure of the safety-monitoring computer system.

The latest affront to Sydney motorists happened during the morning rush hour of September 22, 2008, when a circuit board failed and when the backup did not take over. Motorists were left stranded for three hours. This latest failure has spurred an effort to replace the system with technology that has been proven in other tunnel systems around the world.

[--more--](#)

---

# Availability Topics

## Choosing a Database of Record

In this age of corporate expansions, mergers, and acquisitions, companies find themselves with a hodgepodge of new corporate databases that overlap each other. What is needed is a master database that reflects the “single version of truth.” This database is referred to as the *database of record*. A database of record is not only necessary for the efficient operation of a company but is often required by government regulations.

To create a database of record, companies often build a data warehouse or an operational data store (ODS) to hold all of their data. Such a system contains the one “version of the truth,” and its database is the database of record.

The database of record must be continuously available. Therefore, redundant primary/backup systems are generally implemented. However, in the event of a primary system failure, there is typically a failover time interval during which data is not available; and some data may be lost.

An active/active system largely solves these twin problems. However, in an active/active system, there are multiple copies of the application database. Which one should be the database of record? We address that question in this article.

[--more--](#)

---

## Product Reviews

### Scaling MySQL with Continuent’s uni/cluster

A company’s data is its life blood in these days of 24x7 operations. Its data must always be available, and it must be protected against loss.

Continuent, Inc., of San Jose, California, meets these requirements with *Continuent uni/cluster for MySQL*. Continuent uni/cluster virtualizes two or more MySQL databases to make them appear to the application as a single, highly scalable MySQL database with extremely high reliability. Continuent uni/cluster protects against data loss by synchronously replicating updates across all copies of the database. Furthermore, it improves query performance by load-sharing read queries across the copies. The virtual database can be easily scaled by simply adding additional MySQL database servers.

There is no single point of failure in a uni/cluster. Recovery from a component failure anywhere in the cluster is transparent to the cluster’s clients.

*Continuent uni/cluster for PostgreSQL* provides the same features for PostgreSQL databases.

Continuent uni/cluster started out in the MySQL open-source community. It is available today as the Sequoia open-source implementation. Continuent uni/cluster is the commercially supported version of Sequoia.

[--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](http://www.availabilitydigest.com/reporter).

Managing Editor - Dr. Bill Highleyman [editor@availabilitydigest.com](mailto:editor@availabilitydigest.com).

© 2008 Sombers Associates, Inc., and W. H. Highleyman