



Volume 2
Issue 2

the Availability Digest

--- achieving 100% uptime

February 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

[College Learns from SAN disaster](#)

Never Again

[Don't Wait for the Other Shoe to Drop](#)

Best Practices

[Recovery-Oriented Computing](#)

Active/Active Topics

[Transaction Replication](#)

Recommended Reading

[Unix Backup and Recovery](#)

Product Reviews

[GoldenGate's TDM](#)

The Geek Corner

[Calculating Availability - Failover](#)

Complete articles may be found at
www.availabilitydigest.com.

It's not the backup. It's the recovery that counts.

We have been through the painful act of recovery at many of our client sites. In this issue, this is the topic on which we focus. Our recommended reading is a leading book on backup and recovery of Unix systems. Both our Case Study and our Never Again articles describe the agonies of a failed recovery. The Geek Corner looks at the impact of failover time on system availability.

As you know, access to some of the Digest articles requires a paid subscription. If you would like to be able to access these articles, please subscribe at <http://availabilitydigest.com/subscribe.htm>.

And don't forget our AD Reporter program. We can only draw so many stories from our own customer base. If you can suggest articles which we use, you are entitled to a free subscription. Read the details and report your story at <http://availabilitydigest.com/reporter.htm>.

Thank you.
Dr. Bill Highleyman, Managing Editor

Case Studies

Community College Learns From SAN Disaster

Cuesta College, a large community college in California, consolidated all of its heterogeneous data processing systems by linking them together with a fiber-channel SAN. All of the mission-critical data that had been stored on the individual servers was now stored on redundant SAN storage, where it could be secured and made available across the enterprise.

The College was satisfied that the SAN's redundancy would protect the data in the event of a component failure. What it did not anticipate was a failover fault. The SAN controller failed, and the failover was faulty. The result – two days of hard downtime and several weeks cleaning up stray problems.

Now a dually redundant SAN provides the data protection the College originally sought. Two subsequent catastrophic failures of SAN arrays have proved the resiliency of the new configuration.

[--more for subscribers--](#)

Never Again

Don't Wait for the Other Shoe to Drop

Redundant systems are great for protecting against a failure. But once a failure occurs, fix it fast before a second failure turns a problem into a disaster.

Such a delay cost a small company dearly. What should have been a weekend ordeal turned into a two-month disaster. The company provides online services via a SQL Server system running on Windows 2000 Server. To ensure data integrity, the system uses RAID storage, which is backed up nightly to a remote service.

However, due to management inaction, a failed drive in a RAID array was left unattended. Then the other shoe dropped. A second drive failed, and the system upon which the company's services depended went down.

The lack of a tested and documented recovery plan exacerbated the situation and led to a recovery time measured in weeks rather than in days from a failure that shouldn't have happened in the first place.

[--more for subscribers--](#)

Best Practices

Recovery-Oriented Computing

Users tend to perceive system availability more in terms of recovery time than in terms of failure rate. Much effort has been put into the improvement of the performance of computing systems over the last several decades as well as into improving their availability. However, little effort has been made to improve their recovery time.

The need to achieve rapid recovery times has been hastened by the emergence of large, heterogeneous systems, especially with regard to Internet services. The complexity of these systems has led to high failure rates, difficulty in locating and correcting faults, and long recovery times.

If recovery time can be made small enough, users will perceive a faultless system. This is the goal of the Recovery-Oriented Computing project, staffed by researchers from UC Berkeley and Stanford University. The ROC project is focused on reducing and containing faults, automatically locating faults, and recovering rapidly from faults.

A key component of their research is microbooting for fast recovery. This technique is described in a companion article to be published next month. Microbooting prototypes have demonstrated a 50:1 reduction in user-perceived faults.

[--more for subscribers--](#)

Active/Active Topics

Transaction Replication

Active/active systems depend upon the applications at different nodes having access to distributed copies of the application database. The database copies must all be synchronized so that they present the same application state to their local application instances.

In our previous articles, we have discussed software and hardware replication of changes to achieve this goal. In this article, we discuss database synchronization by transaction replication – the application of each transaction independently to all database copies.

The primary issue with transaction replication is database corruption should a transaction be processed differently in two different nodes. There are certain types of applications that are corruption-safe, such as insert-only applications.

Scalability is another concern since each node must process all transactions. However, in two-node active/active applications, this may not be a problem if the nodes have to be sized anyway to carry the entire transaction load so that one node can provide the required capacity in the event of the failure of the other node.

Replicated transactions may provide an easier implementation path than data replication. It is an ideal technique to support a sizzling-hot standby system for immediate and assured failover.

[--more for subscribers--](#)

Recommended Reading

Unix Backup and Recovery

Backing up is a pain. But it is the *restore* that counts.

This is the message that Curtis Preston delivers in his book, *Unix Backup and Recovery*. Preston has been involved in backup and recovery for much of his professional career. In his book, aimed at heavy-duty business Unix systems and the databases they run, he passes on all of the knowledge that he wishes he had when he first started out as a System Administrator.

This book reviews in detail both commercial and freely-available file system and database backup and recovery utilities. It is applicable to the small shop with no money to spend and to large shops with hundreds of servers. It provides full examples of the use of each utility, with significant effort spent on the nuances of the syntax of each.

Preston sprinkles his book with vignettes of actual recovery horror stories experienced by him and his cohorts. These stories are as entertaining as they are educational.

[--more for subscribers--](#)

Product Reviews

Flexible Availability Options with GoldenGate's TDM

The Transaction Data Manager from GoldenGate is a data replication engine that is used to synchronize two databases in near real time. It finds application in active/active systems which achieve extreme availabilities by using multiple nodes to run a common application against a distributed database comprising synchronized database copies. TDM can furthermore be used to maintain a hot standby system ready to take over in the event of a primary system failure for disaster recovery, to distribute the database to other systems for browsing, querying, or reporting, or to consolidate data from multiple databases into a common master database.

TDM can provide heterogeneous database synchronization across disparate databases, operating systems, and computing platforms. It maintains the referential integrity of the target databases so that they can be used even while they are being updated with the most recent changes.

The GoldenGate data replication suite also includes GoldenGate Director for managing TDM in an application network and GoldenGate Veridata for verifying that two databases are equivalent.

[--more for free--](#)

The Geek Corner

Calculating Availability – Failover

Failover time plays a very important and sometimes dominant role in system availability. This article explores the effect of failover time on system availability.

For some system configurations such as an active/standby system, failover times in the order of hours completely mask the system down time due to dual system failures. In these cases, the resulting system cannot really be considered a high-availability system. It is a disaster-tolerant system in that it can recover from failures of the active system but only at the cost of seriously reduced availability.

Clusters fare significantly better. Failover times contribute a significant but not an overwhelming contribution to the failure probability of a cluster.

Active/active systems still retain their attribute of extreme availability in the presence of failover times so long as these times can be kept short, measured in seconds rather than in minutes.

[--more for subscribers--](#)

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

The free Digest, published monthly, provides abbreviated articles for your review.

Access to full article content is by subscription only at

www.availabilitydigest.com/subscribe.

Name: _____

Email Address: _____

Company: _____

Title: _____

Telephone No. _____

Address: _____

The Availability Digest may be distributed freely. Please pass it on to an associate.
Access to most detailed article content requires a subscription,
To sign up for the free Availability Digest or to subscribe, visit www.availabilitydigest.com/subscribe.
To be a reporter (free subscription, visit www.availabilitydigest.com/reporter.
Managing Editor - Dr. Bill Highleyman editor@availabilitydigest.com.
© 2007 Sombers Associates, Inc., and W. H. Highleyman