ff² the Availability Digest
--- achieving 100% uptime
Volume 2
Issue 3
March 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:

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

Important Notice for ITUG Members:

You have been receiving the Availability Digest via an ITUG mailing. This is the last issue that will be distributed in this manner. If you would like to continue receiving the free Digest, please sign up for it at www.availabilitydigest.com/signup.htm.

Many of the articles in the Digest are for subscribers only. If you would like to subscribe to the Availability Digest, you may do so at www.availabilitydigest.com/subscribe.htm.

We are always looking for new Case Studies and Never Again stories. We have only a limited number from our own customer base. If you have a story that you think would be suitable for publishing in the Digest, you can earn a free subscription. Just visit http://availabilitydigest.com/reporter.htm for details.
.
Dr. Bill Highleyman, Managing Editor

Case Studies

Telecom Italia's Active/Active Mobile Services

The Telecom Italia Group provides the bulk of mobile cell phone services in Italy and Brazil via its TIM (Telecom Italia Mobile) network. It cooperates with other mobile service providers to provide seamless cell phone services to 300 million subscribers in 28 countries.

The TIM network uses HP's Open Call Intelligent Network Servers (INS) running on HP NonStop servers to provide many of its services, such as SMS (Small Message Service), its text messaging service, and its Universal Messaging Service (UMS), which provides voice mail, email, and fax messaging services.

To provide disaster tolerance and capacity expansion, Telecom Italia has configured its INS system as an active/active system. TIM currently operates on two INS nodes, one in Rome and one in Milan. The two nodes are synchronized by using Shadowbase bidirectional replication. Data collisions are resolved with relative replication.

Telecom Italia plans to upgrade its INS operating system with no interruption to subscriber services. The ability to do zero downtime migrations such as this is a hallmark of active/active systems. TIM's active/active configuration also positions it to be able to add capacity easily by simply adding additional nodes and then redistributing its cell tower traffic.

--more for subscribers--

Never Again

The Great 2003 Northeast Blackout and the $6 Billion Software Bug

On August 14, Northeast North America went dark. Was this a continuation of the Blaster worm cyber attack that had occurred just three days earlier?

No. It turned out that the cause of the great 2003 Northeast Blackout was anything but sinister. The Blackout was, in fact, triggered on a hot day by a sagging transmission line contacting an untrimmed tree in Ohio and was aided by a hung alarm system. The failed transmission line imposed heavier loads on other transmission lines, which then began to fail. As each transmission line failed, it overloaded others, which then failed themselves. This cascade of failures led to the blackout. However, power controllers at Ohio's FirstEnergy were unaware of what was going on because there were no alarms being generated to alert them to the escalating problems.

Why did the GE Energy monitoring system fail after millions of hours of successful field experience? If alarms had been generated in the normal course of operations, there would have been ample opportunity to take corrective action and to prevent the blackout.

It took two months for a team of experts searching through four million lines of code to find the problem. It was a race condition with a very narrow window of opportunity. It was the six-billion dollar programming error.

--more for subscribers--

Best Practices

Microrebooting for Fast Recovery

If a system recovers in a time that is short enough so that users don't notice the failure, then no failure has occurred so far as the user is concerned. It is the purpose of the Recovery-Oriented Computing project being undertaken by a joint effort between Stanford University and UC Berkeley to study ways in which recovery from operator and software errors can be done this quickly.

A major contribution of the ROC project has been a technique known as microrebooting. With microrebooting, a first attempt is made to correct a failure by rebooting at the finest-grain level, typically an object suspected of causing the problem. Only if that is not successful are coarser levels of reboot attempted. This reboot escalation continues until the problem has been corrected or has been referred to an operator.

A prototype using an application running in a JBoss environment has shown a 98% reduction in user-perceived errors when microrebooting is used compared to full-system rebooting.

--more for subscribers--

Active/Active Topics

Migrating Your Application to Active/Active

In production today are many 24x7 mission-critical applications that are candidates for migrating to an active/active architecture. The cost of downtime for these systems is very expensive; and too often downtime can be excessively damaging to a company's business, to its reputation, and even to its market value.

Is this migration simply a matter of installing a data replication engine, bringing up a second node with the applications that are to run active/active, synchronizing the new database with the existing one, and then routing transactions to both? Probably not. There are many other factors to consider.

There are some applications whose nature does not allow migration to active/active, such as those which must process events in exact time sequence. There are others running on legacy systems that cannot be economically decomposed so that their databases can be replicated.

A suitable application running in an acceptable environment may contain functions that cannot be distributed and still work properly, such as unique number generators. Once these problems have been corrected and a suitable data replication engine installed, the application is ready to be moved into a multinode active/active network.

--more for subscribers--

Recommended Reading

Migrating Legacy Systems: Gateways, Interfaces, & the Incremental Approach

What does legacy migration have to do with continuous processing architectures? The answer is another question: "How do I get to there from here?" For instance, how do I migrate my current legacy system to an active/active system?

There are still in service many legacy applications that provide mission-critical services but are burdened with the inflexibility, high cost, and brittleness that is characteristic of such systems. If we want to move such a system to, say, an active/active architecture, is it as simple as replicating its database to a like system? Generally not. The legacy system must, in general, be migrated to an architecture in which its database is decomposable from its applications. This is not a simple process.

In their book, Michael Brodie and Michael Stonebraker detail an incremental migration approach that they dub Chicken Little. As opposed to the Cold Turkey approach, which attempts a massive cutover on one fateful day,

the Chicken Little approach decomposes the migration effort into small pieces that can be individually planned and executed over a period of time.

--more for subscribers--

Product Reviews

Shadowbase – The Active/Active Solution

Shadowbase from Gravic, Inc., (www.gravic.com) is a product set that maintains synchronism between geographically distributed, heterogeneous databases.

The proper implementation of an active/active system requires that multiple geographically distributed database copies be kept in synchronism so that any processing node in the application network has access to at least two database copies should one fail. Proper database synchronization requires the ability to

- replicate data changes from one database to the other database copies in the application network so that all database copies maintain the same application state,
- copy a database that is being actively updated in order to create or recover a remote database copy,
- compare two databases to verify that they are identical, and
- bring two databases into synchronism if necessary.

The Shadowbase suite of data replication tools performs all of the above functions. In addition to active/active systems, these products have many other uses, such as providing a hot standby; integrating disparate systems in heterogeneous applications; offloading query, backup, and extract activities; restoring corrupted databases online; and eliminating planned downtime.

--more for free--

The Geek Corner

Calculating Availability – Failover Faults

Redundant systems survive failures by transferring the functions of a failed component to another operating component. This transfer of functions is known as a failover.

Failover is a very difficult process to test. As a result, there is some probability that a failover will itself fail. Such a failure is known as a failover fault. Experience with some high-availability systems has indicated failover fault rates to be in the order of one in one hundred failover attempts.

Failover faults can have a serious impact on system availability. For systems with modest availability, failover faults are not terribly serious so far as overall system availability is concerned. However, as the inherent reliability of a system improves, that is, as the components become more reliable and as failover time decreases, the impact of failover faults can increase dramatically.

In the limit, once one component in a single-spared system has failed, the system availability is determined by the probability of a failover fault rather than by the probability that a second node will fail.

--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: _____

_____

_____

_____