ff²
the Availability Digest
--- achieving 100% uptime
Volume 2
Issue 12
December 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
A Journey from DR to Active/Active
Never Again
Another Hosting Service Bites the Dust
Availability Topics
Time Sync for Distributed Systems - Part 2
The Geek Corner
Failure State Diagrams - Recovery Following Repair

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

Can we ever achieve 100% uptime as our banner suggests? Never! We can come arbitrarily close, but perfection is elusive.

As we have seen in our *Never Again* stories, many companies believe that they have achieved such perfection only to find that they instead have "almost" achieved it. Getting rid of the "almost" is a very big step.

In our *Case Study* this month, we see an example of a company that seems to be doing everything right to achieve "almost" continuous availability. Our *Never Again* story is about a company that believed they had met this standard. They publicly proclaimed zero downtime, but they got caught by a failure that they never imagined.

Learn from these stories what you should and shouldn't do to achieve zero downtime - almost. And let us know about systems that you think are on either side of "almost."

Dr. Bill Highleyman - Managing Editor

Case Studies

Payment Authorization – A Journey from DR to Active/Active

A major provider of merchant services to over four million small to medium-sized merchants throughout the world provides, among other services, payment authorization for Visa, MasterCard, American Express, Diner's Club, Discover, and other credit and debit cards. Card transactions made at merchant point-of-sale (POS) devices and ATMs are verified to ensure that they are proper.

The company's payment authorization services are extremely critical. Should these services become unavailable, shoppers all over the world will not be able to use their credit or debit cards processed by the company. Therefore, the company has worked diligently to guarantee that its authorization services will always be available. It turned to data replication technology to satisfy this need.

Over a period of several years, the company has expanded its use of data replication technology from disaster recovery to active/active systems to application integration. This has been a slow and careful process as the company learned the benefits and pitfalls of data replication. Today, this effort has resulted in a system that indeed can be said to be continuously available.

.

--more--

Never Again

Rackspace – Another Hosting Service Bites the Dust

As one blogger said, "Our Internet infrastructure … is as fragile as a fine porcelain cup on the roof of a car zipping across a pot-holed goat track." This observation was made after yet another Internet infrastructure failure when Rackspace, a major hosting service for thousands of web sites, went down for reasons that would be hard to anticipate – a truck hitting a transformer that powered its data center.

In spite of triply-redundant power backup, this incident started a sequence of events that resulted in many of the web sites that Rackspace hosted going down for hours. The one faint glimmer of success in this disaster is that the company was completely open and honest with its customer base and worked hard to keep everyone informed.

--more--

Availability Topics

Time Synchronization for Distributed Systems – Part 2

Distributed systems often require that their nodes and the clients that access them all have the same view of time. In Part 1 of this three-part series, we showed how NTP calculates the time offset of a client relative to a time server. However, coordinating a client with a single time server leaves a lot of room for error. Here in Part 2, we describe the NTP facilities that allow us to reduce this error significantly.

--more--

The Geek Corner

Failure State Diagrams – Recovery Following Repair

In our December, 2006, article entitled Calculating Availability – The Three Rs, we derived the relationships for the availability of an n-node single-spared system that required recovery following a repair. These expressions were developed intuitively, but are they actually correct?

In our September, 2007, Geek Corner article, Calculating Availability – Failure State Diagrams, we introduced failure state diagrams as a method to formally derive availability relationships.

In this article, we use failure state diagrams to accurately derive the availability of a system that must be recovered after it has gone down and has subsequently been repaired. We show that our intuitive relationships are indeed approximations, but they are valid so long as the nodal availability is high and so long as there is only a modest number of nodes in the system. This is, in fact, the case for the redundant systems in which we are interested.

--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.