ff² the Availability Digest
Volume 2
Issue 5
- - - achieving 100% uptime
May 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.

In the Availability Digest, we talk a great deal about active/active architectures that can achieve incredible availabilities - six 9s or better. However, there is another very important high-availability technology that is far more mature - clusters. Cluster technology has been with us for three decades, It is very mature with excellent product support. However, today's clusters can achieve availabilities of only four to five 9s.

In this issue, we focus on clusters. Though clusters are the established technology for high availability, the new guy on the block - active/active systems - can provide an improvement in reliability of over an order of magnitude at little additional cost.

If you are a cluster user but require higher availability than you have with cluster technology, let us come and talk to you about how a reconfiguration of your system to an active/active architecture can cut your downtime by a factor of ten or more.

Dr. Bill Highleyman, Managing Editor

Never Again

BlackBerry Gets Juiced

At about 5 pm Pacific Daylight Time on Tuesday, April 17, 2007, BlackBerry messages stopped flowing across North America and elsewhere. Millions of BlackBerry subscribers were without service until Wednesday morning when queued-up emails from the huge backlog of messages started to trickle through. It wasn't until Thursday that service was returned to normal.

How did this happen to a service that has come to be depended upon by such a large group of people? It was caused by a software upgrade that clearly was not tested as thoroughly as it should have been.

--more for subscribers--

Best Practices

Interview with Ron LaPedis on NonStop with XP Storage

Prior to his retirement, Ron LaPedis was product manager for business continuity in HP's NonStop Enterprise Division. He holds many certifications in business continuity.

During an interview with the Availability Digest, Ron offers some tips on improving the availability of NonStop servers. He discusses the use of HP's XP storage on a storage area network (SAN) to replace direct-connected disks on HP servers. Done properly, logical volumes on the SAN can be switched from a primary NonStop server to its backup in seconds.

However, certain rules must be followed with respect to node names and numbers, the communication network interconnecting the nodes, and TMF (NonStop's Transaction Monitoring Facility). In addition, certain application modifications can enhance failover times.

--more for free--

Active/Active Topics

Active/Active Versus Clusters

In our previous issues of the Availability Digest, we focused heavily on active/active architectures. But there is another, very important high-availability architecture, one, in fact, that is far more mature and predominant than active/active systems. That architecture is clusters. In this article, we describe the cluster architecture and compare it to active/active systems.

The bottom line is that cluster technology is very mature and is supported by a plethora of products. Active/active technology is the new guy on the block. Its product support is growing, and there are a number of very successful implementations. Most significantly, active/active technology can provide an order of magnitude or more improvement in availability.

--more for subscribers--

Recommended Reading

Blueprints for High Availability: Designing Resilient Distributed Systems

Evan Marcus and Hal Stern are well-known in the cluster community. Their book, *Blueprints for High Availability: Designing Resilient Distributed Systems*, is one of the most referenced books on this topic.

Easy to read and thorough in its content, it is an excellent reference for anyone wanting to learn about clustering technology. Also included are many insights into what is required outside of the cluster to achieve high availability.

High availability is insurance for your business. The amount that you are willing to spend on it depends on what downtime will cost you. The bottom line is that the requirements for high availability and its worth are different for every business. This book provides the background required to effectively evaluate the availability/cost tradeoff for your business.

--more for subscribers--

Product Reviews

HP's ServiceGuard Clustering Facility

HP's ServiceGuard is a cluster management facility. It allows a company to customize and control its high availability clusters. With ServiceGuard, the business can organize its applications into packages. In the event of a hardware or software fault, the company can designate that control of specific packages be transferred to another processing node in the cluster or that communications be transferred to a standby LAN.

ServiceGuard provides all of the services needed to efficiently manage an HP-UX or Linux cluster. With over 150,000 installations around the world, ServiceGuard is clearly a major contribution to the quest to achieve high availabilities. It is a very important element of HP's stated goal to achieve 5 9s:5 minutes of reliability.

--more for free--

The Geek Corner

Cluster Availability

In our companion article, *Active/Active versus Clusters*, we hinted at an availability difference between active/active systems and clusters. In this article, we look at that difference in more detail.

Though cluster and active/active architectures have a great deal in common, there is a tremendous difference in the availability that they offer. An active/active system can be an order of magnitude more reliable than a cluster. This is primarily due to the rapid failover times that can be achieved by active/active systems as compared to clusters.

A second contributor is that only the users at a failed node in an active/active system are affected by a failover, but all of the users in a cluster are affected by a failover.

--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 paid subscription only at
www.availabilitydigest.com/subscribe.

Name: _____

Email Address: _____

Company: _____

Title: _____

Telephone No. _____

Address: _____

_____

_____

_____