the Availability Digest™

- - - achieving 100% uptime

Volume 13
Issue 9

September 2018

The digest of topics on Continuous Availability. More than Business Continuity Planning.
BCP tells you how to *recover* from the effects of downtime.
CA tells you how to *avoid* the effects of downtime.

www.availabilitydigest.com

Follow us
@availabilitydig

In this issue:

Browse through our useful links.

See our article archive for complete articles.

Visit our Continuous Availability Forum.

Check out our seminars.

Check out our writing services.

Check out our consulting services.

Active/Active Takes a Leap Forward with AWS

Amazon Web Services (AWS) has now gone to an active/active architecture to ensure availability. Amazon has eighteen regions spread across the Americas, Asia Pacific, Europe, the Middle East, and Africa. It is providing a high-speed, low latency communication network connecting its regions.

Objects stored in Amazon S3 buckets are replicated in multiple regions to provide high availability and to allow local access of objects in the buckets. Asynchronous replication is used to keep the buckets synchronized.

An application can access data objects from the region to which it is closest. Furthermore, an application can update any bucket, and the update is replicated to the buckets in all other regions. Amazon supports a multi-master capability so that if a master fails, another instance will take over immediately.

Read about Amazon's architecture in our article "Active/Active AWS."

This article and our other stories in this issue are examples of what we write for the Digest and for others. If you have an article, a case study, or a white paper that you would like written, we encourage you to contact us. We also provide consulting services and seminars on high- and continuous availability.

- Dr. Bill Highleyman, Managing Editor

Best Practices

Reducing Data Center Failures

I have written extensively on data center failures. See my seven-part series “Help! My Data Center is Down!” published in the October 2011 through April 2012 issues of the Availability Digest. In Part 4 (Intranet Outages), I give over a dozen examples of data centers that were taken down by internal network outages that prevented servers from communicating with each other.

In data centers, several methods are used to communicate between servers:

- Message switching – Each switch in the network stores the entire message and forwards it to the next switch.
- Circuit switching – The switching equipment seeks a physical path through the network from the calling party to the called party. While it is being used to communicate, no other calls can use this path.
- Packet switching – A message is broken up into packets. Each packet is sent to its destination as soon as possible. The packet header provides the destination address and the information to assemble the packets in the proper order to reconstitute the message.

A data center can suffer an outage of some of its services if a server fails or if the network connecting the servers has a failure. A server failure is generally corrected by moving in a backup server. This article is concerned with network failures.

--more—

Leasing Dark Fiber

In our article in the previous issue of the Availability Digest, we discussed the use of dark fiber for communicating between end-points. Companies or municipalities will often run fiber cable to support high-speed, low-latency communications.

A fiber cable contains many strands of fiber. As fiber is laid in a community, many of these strands are initially unused. These ‘dark fibers’ provide a powerful means for additional communication capabilities in the community. The unused fibers are called ‘dark fibers’ because they are ‘unlit.’ They may be leased to other operators.

Once the fiber has been laid, all that is required to increase its bandwidth is to upgrade the equipment powering the fiber. There is no need for concern about sharing the bandwidth with other customers as in normal networks.

Optical fiber communications is the wave of the future for data communications. It is capable of incredibly high speeds and is absolutely secure. There is no way for a hacker to gain access to the data that is being transmitted over the fiber.

--more--

Availability Topics

Active/Active AWS

Amazon is investing heavily in multi-region active/active architectures. Why are multi-region architectures important?

- They reduce latency for end-users.
- They provide for disaster recovery.
- They satisfy certain business requirements.

Networks like those of Amazon have been used to speed up delivery of content, especially if the data is static and is cached in local servers. However, if the data is far away, latency can be significant:

- The latency between the U.S. and Europe is 140 msec.
- The latency between Europe and Australia is 300 msec.

Lower latency engages users to a greater extent:

- Amazon found that 100 msec. of extra latency caused a 1% drop in sales.
- Google found that 500 msec. of extra load time caused 20% fewer searches.
- Yahoo! found that 400 msec. of extra load time caused a 9% increase in the number of people who clicked 'back' before the page loaded.

Low latency is now becoming a requirement. Therefore, having locally available applications and content is becoming more important.

--more--

HP Launches Bug Bounty Program for Printers

HP is offering a bug bounty of up to $10,000 for flaws found in its printers. The vast majority of printers are attached to the same networks as local machines that could contain sensitive data. HP is concerned that a flaw in its printers could open the door to a company's entire network.

Historically, Chief Security Information Officers (CSIOs) have not been involved in the purchase of printers. Unpatched or vulnerable printers have not been considered a serious threat. However, CISOs are now expecting more secure printers.

HP has made cybersecurity a priority for its printers. In doing so, HP is helping to support the valuable role that CISOs play in their organizations.

However, as more networking capabilities and cloud functions are introduced, printers are presenting a larger attack surface for hackers. Attackers have started to focus more on targeting endpoint devices such as printers in the past year. HP doesn't want one of it printers serving as a conduit for a larger attack on a company.

The bug bounty currently covers the HP Laser Jet Enterprise printers and the HP Page-Wide Enterprise edition printers. HP quietly launched the program in May 2018.

--more--

Tweets

@availabilitydig – The Twitter Feed of Outages

A challenge every issue for the Availability Digest is to determine which of the many availability topics out there win coveted status as Digest articles. We always regret not focusing our attention on the topics we bypass.

Now with our Twitter presence, we don't have to feel guilty. This article highlights some of the @availabilitydig tweets that made headlines in recent days.

--more--

Sign up for your free subscription at http://www.availabilitydigest.com/signups.htm

The Availability Digest is published monthly. It may be distributed freely. Please pass it on to an associate.
Managing Editor - Dr. Bill Highleyman editor@availabilitydigest.com.