

the Availability Digest™

Volume 13
Issue 10

--- achieving 100% uptime

October 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](https://twitter.com/availabilitydig)

Technical
Writing

The articles you read in the Availability Digest result from years of experience in researching and writing a variety of technical documents and marketing content. It's what we do best, and we provide our services to others who value high-quality content created by IT specialists. [Ask us](#) about

- articles
- white papers
- case studies
- web content
- manuals
- specifications
- patent disclosures

In this issue:

[Never Again](#)

[Facebook Breach of 30 Million Users](#)

[Best Practices](#)

[California Kills Power to Prevent Wildfires](#)

[Hyper-Availability](#)

[Swapping Replication Engines with ZDM – Part 1](#)

[Tweets](#)

[The Twitter Feed of Outages](#)

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](#).

There Is Now Hyper-Availability for Critical Apps

Decades ago, applications were configured for high availability by deploying them with backup applications to which the production applications could failover should they become compromised. Then continuous availability became the mantra with active/active systems, in which two or more applications processed transactions simultaneously, keeping their databases synchronized via bidirectional data replication.

Now there is hyper-availability. Hyper-availability is meant to be a step forward from continuous availability. It depends upon fast failover, testing without disruptions, visibility and monitoring, and orchestration and automation.

Read about the architecture of hyper-availability in our article "Hyper-Availability."

This article and our other stories in the October 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

Never Again

Facebook Breach of 30 Million Users

Facebook disclosed that it has detected the biggest data breach in its history, one which affected 30 million users.

Still reeling from the Cambridge Analytics scandal, in which a UK-based digital consultancy harvested the personal data of 87 million Facebook users, Facebook is under intense scrutiny for its ability to keep the data of its more than two billion users safe. CEO Mark Zuckerberg testified in front of the U.S. Senate to explain the Cambridge Analytics scandal.

Recently, threat actors were able to steal access tokens belonging to 30 million user accounts. An access token is a set of code granted to a user after logging in for the first time so that the user does not need to log back in every time he goes to a new page. Facebook uses access tokens for logins. They allow for secure access without needing a password.

Facebook says the breach started when hackers with some access tokens exploited a combination of three bugs related to its 'View As' privacy feature. 'View As' allows a user to view his profile from the perspective of someone else.

[--more--](#)

Best Practices

California Kills Power to Prevent Wildfires

California has a history of power lines sometimes sparking wildfires during high wind conditions. In October, 2018, strong winds swept California, knocking down trees. One woman was killed when a Eucalyptus tree fell on her car as she was coming out of the car port at her apartment building. Wind gusts up to 55 miles per hour were reported in Northern California.

Utilities are responding in elevated fire risk situations by turning off the power in those areas. Two major California utilities cut power to customers following concerns about downed power lines sparking wildfires. In Northern California, Pacific Gas & Electric (PG&E) began cutting power to tens of thousands of customers after the National Weather Service warned of extreme fire danger due to high winds, low humidity, and dry vegetation.

San Diego Gas & Electric did the same. It was the second time in a few months that it had killed power. October saw about 84,000 customers impacted. Schools in the affected areas cancelled classes.

[--more--](#)

Hyper-Availability

When I first started writing mission-critical applications decades ago, normal applications had an availability of about three 9s (about eight hours of downtime per year). High-availability was achieved by using redundant systems with fast failover from a failed production system to its backup. This level of high availability achieved about four 9s of availability (less than one hour of downtime per year).

Then came continuous availability with active/active systems. These systems comprised two or more systems actively processing transactions. Their databases were kept synchronized via bidirectional replication. Whenever a change was made to one database, that change was immediately replicated to the other systems in the application network. Should a system fail, all that was required to recover was to route all transactions to surviving systems. Active/active systems provided what we called 'continuous availability.' This was equivalent to five 9s of availability (about five minutes of downtime per year).

Now we have hyper-availability. What is hyper-availability? We explore that concept in this article.

[--more--](#)

Swapping Replication Engines with Zero Downtime – Part 1

A mission-critical application often runs in redundant systems to ensure that it is always available to its users. Such a system may be configured as an active/passive pair, in which one system runs the production workload while the other system is standing by, ready to take over application processing should the production system fail.

Alternatively, the system may be configured as an active/active architecture, in which both systems are processing transactions. The databases are kept synchronized via bidirectional replication.

Sometimes, companies may decide to change data replication engines or to upgrade to a new version of the existing data replication engine. With mission-critical applications, it is necessary to do so without taking down the applications – a zero downtime migration (ZDM). Furthermore, it is imperative that a backup copy of the database is always available, ready to take over if the production database fails. The backup database must be kept synchronized with the production database while the data replication is being migrated.

In this three-part series of articles, we describe how a data replication engine can be changed without taking down either applications or the backup database.

[--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.

© 2017 Sombers Associates, Inc., and W. H. Highleyman