

the Availability Digest™

Volume 13
Issue 7

--- achieving 100% uptime

July 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.

Follow us



[@availabilitydig](https://twitter.com/availabilitydig)

www.availabilitydigest.com

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](#)

[Amazon Downed By Overload](#)

[Sutter Health Loses Electronic Health Records](#)

[Best Practices](#)

[What is Failover?](#)

[Availability Topics](#)

[Sick But Not Dead](#)

[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](#).

Partial Failures

A partial failure in a system or application can go undetected, allowing erroneous processing to proceed without the user noticing. We call this 'sick but not dead.' In our article in this issue by the same name, we explore the consequences of partially failed components.

We give examples of actual instances of the impact of partial failures on companies, including Lloyd's Bank, Visa, and Los Angeles airport.

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

Never Again

Amazon Downed by Overload

One of the worst things that can happen to an application is to underestimate the load that it will have to handle. If the application becomes overloaded, it will slow down and may eventually fail. And that is exactly what happened to Amazon during its Prime Day sales event.

Amazon's Prime Day, its biggest non-holiday sales event, is a thirty-six-hour orgy of buying that spans over a dozen countries. Consumers bought over 100 million products during Prime Day 2018, which took place on July 16th. Amazon said that in the first hour of Prime Day 2018, customers had ordered more items than during the first Prime Day hour of the previous year.

However, there was a spike of issues when Prime Day kicked off. Prime Day suffered multiple outages as shoppers were left unable to access deals. The problem was with the millions of users trying to access the site at the same time. Shoppers reacted with anger as they struggled to find deals that Amazon typically offers only for a limited time.

Amazon said that it was trying to fix the problems, telling customers there were still 34 hours to shop Prime Day.

[--more--](#)

Sutter Health Loses Electronic Health Records

We have reported on several instances of data center outages caused by fire suppression systems even when there was no fire. In our November 2017 article "Fire Extinguishers Can Cause Data Center Outages," we described several such outages, including those at West Host (a web-hosting service), ING Bank, and Microsoft Azure. In all cases, disks in the servers were impacted.

Siemen's, a world leader in fire safety and fire suppression systems, performed a study to determine how a fire suppression system could damage disks in a data center. Their tests determined that it was not the sudden increase in pressure by the discharge of Inergen, a gas that is released when a fire suppression system is activated. Rather, the damage was caused by the load noise of the alarm sirens.

Now a fire-suppression system has shut down another major data center. In this case, it was that of Sacramento-based Sutter Health, one of the largest health-care providers in Northern California.

[--more--](#)

Best Practices

What Is Failover?

Failover is the process of switching to a backup or redundant system if something goes wrong with the primary system or if it is shut down for servicing. If a system has a failover mechanism, it can continue operating even if something goes wrong by switching over to its redundant system.

In order to provide high availability, businesses must have a failover mechanism in place.

Performed properly, the users don't notice that there has been a problem when a failover occurs. The switch to the backup system happens automatically and seamlessly. The backup system may be local, or it may be geographically separated from the primary system to prevent a common catastrophe (such as a flood or earthquake) from taking down both systems. Failover processes are programmed to operate automatically.

[--more--](#)

Availability Topics

Sick But Not Dead

Partial failures of a device are among the most difficult to detect and diagnose. A partially failed device is known as 'sick but not dead' – it is a 'soft failure.' It may generate invalid outputs that appear to be valid, causing havoc in the system. Or it may lead to unreasonably long latencies that can bring a system to its knees.

The fact that the device is operating improperly is often not clear to system operators or management systems. Therefore, failover to an operating backup may not be attempted; and the partially failed device remains in production.

Researchers are warning firms that they should make sure they are prepared for partial failures in their IT infrastructures. Systems cope with outages quite well. If one component stops working, the system will switch over to a backup system or to an alternate path. However, a marginal component showing only intermittent errors can have a severe impact on the overall infrastructure. Such errors can lead to performance degradation and loss of access to data or processes.

[--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>

To cancel your subscription, send an email to editor@availabilitydigest.com with the subject 'Cancel.'

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.
© 2018 Sombers Associates, Inc., and W. H. Highleyman