

the Availability Digest™

Volume 13
Issue 1

--- achieving 100% uptime

January 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

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:

[Case Studies](#)

[The Bogus Missile Launch](#)

[Never Again](#)

[Meltdown and Spectre Security Threats](#)

[Best Practices](#)

[Serverless Computing](#)

[Availability Topics](#)

[Availability in the Cloud](#)

[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](#).

Turn to the Cloud for High Availability

Increasingly, businesses are turning to the cloud to provide high availability for their critical applications. The cloud provides anywhere and anytime access to services, tools and data.

In our article exploring the high availability that clouds can provide, we point out that an appropriate cloud provider guarantees a level of availability. The guarantee compensates the provider's customer if the provider misses the documented availability metrics. The better the guarantee and availability, the more reliable and expensive is the service.

This article and our other stories in the January 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

Case Studies

The Bogus Missile Launch

On Saturday morning, January 13, 2018, Hawaiian residents were greeted by a terrifying Emergency Alert System message that turned out to be entirely false:

“The U.S. Pacific Command has detected a missile threat to Hawaii. A missile may impact on land or sea within minutes. This is not a drill. If you are indoors, stay indoors. If you are outdoors, seek immediate shelter in a building. Remain indoors well away from windows. If you are driving, pull safely to the side of the road and seek shelter in a building or lay on the floor. We will announce when the threat has ended. Take immediate action measures. This is not a drill.”

Emergency alert messages were posted on message signs along the highways. Many residents and visitors in Hawaii thought they were going to die.

As it turned out,, there was no missile launch. This was a false alarm. The Hawaiian governor knew it was false after two minutes but didn't know his Twitter password in order to announce this.

[--more--](#)

Never Again

Meltdown and Spectre Security Threats

A new class of security threats has just reared its ugly head. These threats, known as Meltdown and Spectre, exploit a common feature of modern microprocessors in order to steal sensitive data from memory. This feature is known as ‘speculative execution.’

In order to speed up application processing, today's CPUs employ speculative execution. When a program reaches a conditional branch (such as ‘If FLAG is true, do Process A; else do Process B’), the CPU decides in advance which branch to take. A ‘branch predictor’ determines the likely result of the condition and executes the branch that is most likely to run before the test is actually completed.

If the guess is correct, the chip appears to be running faster than had it awaited the results of the test. If the guess is wrong, the chip has to throw away any speculative results. These results are stored in cache memory, from which Meltdown and Spectre can access sensitive data.

[--more--](#)

Best Practices

Serverless Computing

A new technology is about to become one of IT's most valuable tools. It is serverless computing, which allows one to build and run applications and services without thinking about servers. With serverless computing, you don't need to provision, scale, and manage any servers. Everything required to run and scale a high-availability application is handled by a cloud service provider.

Serverless computing is a new way of hosting applications on infrastructure that end users do not manage. Serverless architectures let developers and users execute tasks without worrying about servers, virtual machines, or the underlying compute resources.

With serverless computing, applications can be executed without any effort being put forth to provision or manage servers. The serverless computing environment handles all of that for you. This gives developers the ability to focus on what they are supposed to be doing – managing applications or creating new applications.

[--more--](#)

Availability Topics

Availability in the Cloud

Businesses and their employees demand 100% availability of their applications. Such high availability is an enabler of worker productivity and provides differentiation from the business' competitors.

One route businesses are employing to achieve high availability is to take advantage of the cloud. High-availability is the holy grail of the cloud. The cloud embodies the idea of anywhere and anytime access to services, tools and data.

A typical cloud SLA (Service Level Agreement) will guarantee an availability of 99.99% or higher. This is less than one hour of downtime per year and is definitely considered high availability.

Cloud services are complex. Troubles can arise from the interactions of multiple components and automated systems over distributed networks and data centers. This results in issues that can take a long time to resolve.

The availability of applications can be improved by using multiple cloud providers. Services can be hosted and replicated to various providers at several locations. Even downtime related to maintenance can be reduced by spreading a service over multiple providers.

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