

the Availability Digest™

Volume 13
Issue 11

--- achieving 100% uptime

November 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:

[Best Practices](#)

[Swapping Replication Engines with ZDM – 2](#)

[Google Cloud Adds Dual Regional Support](#)

[Availability Topics](#)

[Object Storage](#)

[The Geek Corner](#)

[The Pancake Problem](#)

[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](#).

Object Storage – the New Kid on the Block

When I first started in Information Technology several decades ago, data was stored in files. A file had a file name by which it was accessed.

Then came relational databases. These SQL databases allowed data items that had a relationship of some sort with each other to be stored together in the database and to be accessed based on that relationship.

Now there is a new guy in town – objects. Object storage is a computer data storage architecture that manages data as objects. An object is a blob of data with an ID by which it can be retrieved quickly. Read about objects in our article “Object Storage.”

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

Swapping Replication Engines with Zero Downtime – Part 2

In Part 1 of this article, we described how an existing data replication engine could be replaced in a two-node scenario with a new version or with a different replication engine without requiring an outage. As pointed out in that article, it is important to avoid requiring different replication engines to communicate with each other. Interoperation between different versions is very dangerous, error-prone, and can produce faults.

In Part 1, we showed how a replication engine could be changed without interoperating between differing versions. Let us first review these procedures, and then in Part 2 we will look at eliminating the versioning problem by using additional nodes. We also will look at protecting the standby system while the replication engines are being switched.

[--more--](#)

Google Cloud Adds Dual Regional Support

Google wants to provide flexibility in its cloud services to differentiate it from AWS and Microsoft Azure.

One of the issues that companies face as they move data to the cloud is making sure that the data stays close to the application to reduce latency. Companies also require redundancy in the event of a catastrophic failure. Moreover, access to this data must be with low latency.

To meet these requirements, Google has introduced a new dual-regional storage option. Data is written to a single dual-regional bucket, and Google automatically replicates this data to a backup bucket. In the event of a region failure, Google transparently handles the failover to the data in the backup bucket.

Customers can pair specific Google dual regions, and data will be automatically replicated between these regions. The data is distributed redundantly across a multi-region that can span the U.S., the E.U., and/or Asia.

[--more--](#)

Availability Topics

Object Storage

When I first started in Information Technology several decades ago, data was stored in files. A file had a file name by which it was accessed. Each record in the file typically had a unique field that could be used as an index to access the record.

Then came relational databases. These SQL databases allowed data items that had a relationship of some sort with each other to be stored together in the database, and they could be accessed based on that relationship.

Now there is a new guy in town – objects. Object storage is a computer data storage architecture that manages data as objects, as opposed to other storage architectures like file systems that manage data as a file hierarchy and block storage that manages data as blocks within sectors and tracks.

An object is a blob of data with an ID. It can be stored in a flat address space and retrieved quickly by its ID. In this article, we will look at objects and how they are stored and retrieved.

[--more--](#)

The Geek Corner

The Pancake Problem

Jacob Goodman was a mathematician at the City College of New York. One day, at home folding towels for his wife, Goodman decided to restack the towels in the order of size, smallest folded towel on top. The problem was there was no room for a second stack. He was forced to flip over a few towels at a time until they were ordered.

A curious problem crossed his mind. How many flips would it take in the worst case to order the towels according to size? Thus, the so-called pancake sorting problem was born.

Two pancakes would take at most one flip. But if the pancakes were burnt on one side, and the resulting stack must have all pancakes with their burnt sides down, then it could take four flips.

The problem is simple to pose, yet tough to solve. No one is close to developing a general formula that predicts the number of flips required for any given number of pancakes.

[--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