

# the Availability Digest™

Volume 12  
Issue 8

--- achieving 100% uptime

August 2017

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](http://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](#)

[Marketo Goes Down, Revived By Customer](#)

### [Best Practices](#)

[Dark DR – Avoid Its Costs With Active/Active](#)

### [Availability Topics](#)

[Should Cyber Victims Be Fined?](#)

### [The Geek Corner](#)

[Dr. Bill's Doctoral Thesis](#)

### [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](#).

## The Inefficiency of Backup Systems

A common strategy to protect a production system is to back it up with a similar system that is just sitting there idly, ready to take over. The backup system is kept synchronized with the production system via data replication so that transactions can be rerouted to it should the production system fail. Switching to the backup system is called 'failover.'

Unfortunately, a failover can take minutes to hours depending upon how much work has to be done on the backup system to bring it up-to-date. Even worse, the failover may fail. This is often caused by 'configuration drift' in which updates made to the production system are not made to the backup system.

Running an active/active architecture solves these problems, as discussed in our article "Dark DR – Avoid Its Costs With Active/Active."

These articles are an example of the stories 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, come talk to us. We also provide consulting services and seminars on high- and continuous availability. We will be glad to help you.

Dr. Bill Highleyman, Managing Editor

---

## Never Again

### Marketo Goes Down, Revived By Customer

Marketo Inc. is a marketing tech giant that makes automation software, sold to companies that plan to promote themselves via multiple digital channels (email, social media, websites, etc.).

On July 25, 2017, Marketo's website suddenly went down. Its applications crashed. Marketo's customers lost access to their content on the Marketo platform, scheduled promotional emails were sent with broken links, and customers had no way to interrupt their compromised online campaigns.

What happened? It turned out that Marketo had forgotten to renew its domain name, an oversight that seems to be incredibly ludicrous but happens to more companies than you might think. In Marketo's case, it was a customer who figured it out and re-registered the domain name on behalf of Marketo.

[--more--](#)

---

## Best Practices

### Dark DR – Avoid Its Costs With Active/Active

Traditional disaster recovery (DR) methods utilize an active/passive architecture, in which there are two redundant servers. One server is actively processing the application, and the other acts as a backup system ready to take over should the production system fail. Therefore, this architecture requires the cost of two systems, even though only one is providing processing capacity.

The two systems typically are located remotely from each other to avoid a dual failure due to some local disaster. From an operations viewpoint, the backup server remains 'dark' until it is needed. Thus, this architecture is commonly called "dark DR."

In contrast, the servers in an active/active system are both processing transactions. Should one server fail, all transactions are simply routed to the surviving server. Of course, both servers must have sufficient capacity to accommodate the full operational load; but in normal operation, the reduced workload on each server results in performance improvement.

[--more--](#)

---

---

## Availability Topics

### Should Cyber Victims Be Fined?

The title of this paper seems absolutely absurd. Should cyber victims be fined? They already have suffered from a cyberattack. Why should they be made to suffer further? But this is exactly what the United Kingdom is now doing by threatening to fine essential infrastructure operators who experience a cyberattack. Essential providers include those in the utilities, health, transport, and communications sectors.

The U.K.'s seemingly extreme response is in part due to the £60 million cyberattack on TalkTalk in 2015. TalkTalk is a major U.K. communications company. In October 2015, TalkTalk suffered a devastating cyberattack. A consequence was its loss of 95,000 subscribers. As a result, profits slumped from £72 million in 2015 to just £2 million in 2016.

It will be interesting to follow this measure to see if it encourages companies to make further investments in combating cybercrime.

[--more--](#)

---

## The Geek Corner

### Dr. Bill's Doctoral Thesis

During my tenure at Bell Telephone Laboratories in the early 1960s, my work at Bell Labs involved an effort to build a system that would read hand-written numbers. At the time, long distance calls were recorded by an operator on a hand-written ticket and then transcribed manually onto punch cards for processing by a billing program. This required a great deal of labor, and Bell Labs sought a solution that would automate the process. My assignment was to build a system that could automatically read the operators' tickets and create the punch cards.

This was an ideal task for a doctoral thesis and led to my 211-page dissertation "Linear Decision Functions, With Application to Pattern Recognition." Recognize that back then, there were no PCs. The only computer the Labs had was a monstrous IBM 7090. I had to build a system that used no computer.

My thesis was one of the early attempts at recognizing hand-printed text. It failed because of a lack of computer resources and reliable categorization algorithms. As time has marched on, and as a great deal of research has been accomplished, these limitations have been overcome. Now, even machine reading of alphanumeric hand-written script is common.

[--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--](#)

---

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](mailto:editor@availabilitydigest.com).

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