



the  
Volume 5  
Issue 1

# the Availability Digest

--- achieving 100% uptime

January 2010

---

Current topics on Continuous Processing Architectures. More than Business Continuity Planning.  
BCP tells you how to **recover** from the effects of downtime.  
CPA tells you how to **avoid** the effects of downtime.  
[www.availabilitydigest.com](http://www.availabilitydigest.com)

## In this issue:

### [Never Again](#)

[BlackBerry – OMG, It's Déjà Vu!](#)

### [Best Practices](#)

[Roll-Your-Own Replication Engine - 1](#)

### [Availability Topics](#)

[Defining Active/Active - Revision 1](#)

### [Product Reviews](#)

[Stratus Bets \\$50K You Won't Be Down](#)

---

Complete articles may be found at  
<http://www.availabilitydigest.com/articles.htm>

Sign up for your free subscription at  
<http://www.availabilitydigest.com/signups.htm>

## What's Active/Active, and What's Not? Join us on LinkedIn for a discussion.

Does this sound like a repeat from our last issue? It is. In that issue, we made a first pass at defining active/active. We set up the LinkedIn *Continuous Availability Forum* so that you could critique our definition. Almost one hundred of your peers have joined us on this forum. Based on the many constructive suggestions we received, we have revised our definition; and we post it in this issue as "Defining Active/Active – Revision 1."

We intend for this to be a living document and to keep it updated based on your feedback. Please review it and join us on our *Continuous Availability Forum* at <http://www.linkedin.com/groupsDirectory?results=&size=1260921605283&pplSearchOrigin=GLHD&keywords=continuous+availability+forum>.

We want your comments so that we can develop a general agreement on what this thing called active/active really is.

The *Continuous Availability Forum* has been so active since it started that we are anxious to ensure that it continues to provide a platform for the exchange of ideas, experiences, and help for those in the continuous and high-availability communities. So please feel free to start any new discussion thread that you think would be helpful to you. I am doing so with this issue, asking for your experiences in achieving continuous or high availability, starting with one of my own.

Dr. Bill Highleyman, Managing Editor

---

## Never Again

### BlackBerry – OMG, It's Déjà Vu!

BlackBerry's email service went down twice in December, 2009, with each outage lasting for hours.

The RIM BlackBerry smartphone has become a perceived necessity for anyone on the move. Consequently, one would think that RIM (Research in Motion) would consider its services mission-critical. But in the last five years, it has achieved only about three 9s of availability – down for hours per year.

Interestingly, these problems have had a common thread – upgrades. Over the past three years, RIM has suffered five major outages, each lasting for hours, after RIM tried to upgrade its Internet browsing and email services.

In the U.S., BlackBerry is the number one smartphone, with 20% of the market. However, BlackBerry faces stiff competition from Apple's iPhone (the number two smartphone that is rapidly gaining on BlackBerry), Palm's new Pre, Motorola's Droid, Google's forthcoming Nexus One, and Verizon/Microsoft's Sidekick. As users rely more and more on smartphones as an all-in-one communication device for voice, email, texting, and Internet access, outages won't win fans.

We talk often about the cost of downtime in terms of dollars, safety, stock value, and publicity. In RIM's case, its outage history could relegate it to an "also ran."

[--more --](#)

---

## Best Practices

### Roll-Your-Own Replication Engine – Part 1

Active/active systems depend upon synchronized distributed copies of the application database. The predominant technology for keeping database copies in synchronism is data replication. With data replication, changes made to any one database copy are immediately replicated to the other database copies in the application network. Though it is certainly feasible to incorporate data replication within the application, it is more common to utilize a replication engine that can serve the needs of multiple applications.

There are many excellent, commercially-available data-replication engines available today serving a wide variety of server systems and databases. However, it is always tempting to consider building your own replication engine so that you can save all of those license fees while at the same time ensuring that it will meet your requirements.

Organizations have built their own data-replication engines. However, it is quite a complex task. The purpose of this two-part series of articles is to ensure that you have thought out all of your active/active data-replication requirements so that you don't get caught in an embarrassing or untenable situation during production.

[--more --](#)

---

## Availability Topics

### Defining Active/Active – Revision 1

If we ask a dozen people what the term “active/active” means, the general consensus will probably be that it is a technique for building extremely reliable computing systems. In fact, this is true. There are many examples of such systems that have been in production for a decade or more with never a failure.

But when we probe deeper, we start to find the caveats – the limitations imposed by one technology or another to achieve such high reliability. How far can these caveats reach before we must conclude that a particular approach is really not suitable for our applications because of the approach’s limitations?

In our previous article, we made a first pass at defining the ideal active/active system. We categorized active/active architectures and listed their advantages and disadvantages relative to our ideal system. We started a LinkedIn thread to discuss just this issue. Several excellent and constructive suggestions were made, and we incorporate them into this Revision 1 of the document. As more suggestions come in, we will continue to keep the document updated.

To join us, search on Continuous Availability Forum under “Groups” on LinkedIn ([www.linkedin.com](http://www.linkedin.com)). We want your comments as well.

[--more --](#)

---

## Product Reviews

### Stratus Bets \$50,000 That You Won’t Be Down

Stratus Technologies has been providing fault-tolerant servers for the last three decades. It claims that its current fault-tolerant servers provide in excess of six 9s availability (up 99.9999% of the time, corresponding to 32 seconds of downtime per year).

Stratus is now putting its money where its mouth is with its Zero Downtime \$50K Guarantee. If you buy a Stratus ftServer® 6300 running Microsoft Windows Server 2008 (downgradable to Windows Server 2003 if required for application compatibility) before February 26, 2010, and if the system fails in production during the first six months of deployment, Stratus will pay you \$50,000. In many cases, this is more than the cost of the system.

Continuous availability is no longer a technological problem. It is an exercise in balancing system cost with downtime cost. Stratus’ ftServer is an affordable starting point to achieve extreme availabilities. Stratus says so – with its wallet.

[--more --](#)

Sign up for your free subscription at <http://www.availabilitydigest.com/signups.htm>

**Would You Like to Sign Up for the Free Digest by Fax?**

Simply print out the following form, fill it in, and fax it to:

Availability Digest  
+1 908 459 5543

**Name:** \_\_\_\_\_

**Email Address:** \_\_\_\_\_

**Company:** \_\_\_\_\_

**Title:** \_\_\_\_\_

**Telephone No.:** \_\_\_\_\_

**Address:** \_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

The Availability Digest may be distributed freely. Please pass it on to an associate.  
To be a reporter, visit [www.availabilitydigest.com/reporter](http://www.availabilitydigest.com/reporter).  
Managing Editor - Dr. Bill Highleyman [editor@availabilitydigest.com](mailto:editor@availabilitydigest.com).  
© 2010 Sombers Associates, Inc., and W. H. Highleyman