



Volume 4
Issue 2

Availability Digest

--- achieving 100% uptime

February 2009

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

In this issue:

[Never Again](#)

[More Never Agains II](#)

[Availability Topics](#)

[Worsing on Worsening](#)

[Product Reviews](#)

[Stratus' Avance Brings Availability to the Edge](#)

[The Geek Corner](#)

[Meeting a Performance SLA – Part 2](#)

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>

The Quest for Higher Availability Is Over Four Decades Old

In 1967, Dr. Worsing, director of Boeing's computer center, scolded IBM's Field Service staff for the poor availability of IBM's System 360, though all he was looking for was 90% uptime. We have reprinted excerpts of his speech in this issue of the Availability Digest because much of what he had to say then still holds true today. As Dr. Worsing said:

“I'm still uneasily suspicious that, to the manufacturers, a better computer is a faster CPU.”

Over these four decades, server power has increased by factors of thousands. However, we have improved industry-standard server availability only by a factor of one hundred, from one 9 to three 9s.

Perhaps Dr. Worsing had great foresight. Availability is so low on the totem pole that today there are not even any accepted availability benchmarks for comparison purposes. In earlier “Availability Issues” articles, we have promoted the need for availability benchmarks and pointed out how they could be implemented. But to date, there has been no progress. If you have any thoughts on this issue, let us know at editor@availabilitydigest.com.

Dr. Bill Highleyman, Managing Editor

Never Again

More Never Agains II

Despite its title, this is the fourth in our semi-annual series of brief recaps of some of the many computing-system failures that have occurred over the last six months. Unlike prior recaps, power outages do not lead the list of failures this time. Network outages were the most predominant, accounting for over a third of all failures. Operator errors accounted for almost 20% of faults, ranging from one that destroyed a company to Google's disabling of its search engine.

[--more--](#)

Availability Topics

Worsing on Worsening

Ever since the first removal of a moth from the relay contacts of the Mark II computer in 1947, bugs have plagued computing. The search for high availability in computing systems is markedly historic. Over four decades ago, in a 1967 scolding given by Dr. R. A. Worsing, director of Boeing's computer center, to IBM Field Service management, he forcibly attacked what was then the current state of system availability.

Although at the time he would have been happy with a downtime of 2 hours per day (one 9 of availability), many of his observations hold today. We have improved availability in our industry-standard servers by a factor of one hundred (one 9 to three 9s). However, processor speeds have increased by a factor of thousands. One telling comment of his:

"I'm still uneasily suspicious that, to the manufacturers, a better computer is a faster CPU."

The history of availability improvement lends credence to that statement today.

Dr. Worsing's speech is fascinating and entertaining reading for anyone involved in availability and system support. In fact, legend has it that for many years, UNIVAC required all product development managers to read this speech yearly and to sign an annual declaration that they had read it.

[--more--](#)

Product Reviews

Stratus' Avance Brings Availability to the Edge

Business continuity has not yet been extended to the Edge. What is the Edge? It is everything outside of the corporate data center upon which the IT services of a company rely. These are the branch offices of an enterprise – the bank branches, the retail stores, the sales offices as well as small to medium businesses.

Avance™, from Stratus Technologies, brings high availability to the Edge. Avance also brings an added capability – virtualization. Not only can the servers sitting in the computer closet be highly available, but the various applications can also run on virtual machines hosted by a single, highly-available server, perhaps reducing branch IT costs significantly.

Avance provides an out-of-the-box, fault-tolerant virtualization solution supporting up to eight virtual machines with over four 9s of availability. Requiring no special hardware, it runs on a pair of standard x86 servers interconnected by an Ethernet link. One server acts as the primary node and the other as its backup node. From a deployment and management perspective, Avance creates a single-system image so that the operations staff sees only a single server.

If the cost of downtime in an Edge application is as little as \$1,000 per hour, Avance can pay for itself very quickly, perhaps in a year or so.

[--more--](#)

The Geek Corner

Configuring to Meet a Performance SLA – Part 3

Many applications carry with them a performance Service Level Agreement (SLA) that specifies the response times that they must achieve. The performance requirement is often expressed as a probability that the system's transaction-response time will be less than a given interval. For instance, "When handling 50 transactions per second, 98% of all transactions must complete within 500 milliseconds."

In Part 1 of this series, we derived the basic average response-time expression for a single-server system. In Part 2, we extended that result to a multiserver system in which multiple servers work off a common work queue.

We now show how to size a system to meet a performance SLA. If service time is exponentially distributed, the solution to this question is straightforward. If service time is not exponentially distributed, the solution is more complex. In this part, we explore exponentially-distributed service times. In Part 4, we will extend this to servers with general service-time distributions.

Though the derivation of SLA sizing is involved, and the calculations are complex, SLA sizing is reduced to simple graphics supported by an easy-to-use spreadsheet.

[--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.
Managing Editor - Dr. Bill Highleyman editor@availabilitydigest.com.
© 2009 Sombers Associates, Inc., and W. H. Highleyman