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

Sneak Peek - Active/Active Systems: Theory and Practice

If you're planning to bake in the Las Vegas heat by attending the HP Technology Forum (HPTF) this June, I hope to see you at my preconference seminar on "Active/Active Systems: Theory and Practice." I'll be teaching on the first day of the conference, June 15th.

The all-day presentation will include the defining of active/active systems, why they work, and how to implement them. Included will be several active/active examples that are in production today and that have yet to fail.

To get a taste of what my Las Vegas seminar will include, I encourage you to participate in my Thursday, May 21st, "sneak peek" webinar. Sponsored by HP, it's one-hour in length and will begin at 11:00 am EDT. Best of all, the webinar is free. If you're interested, you can register at
https://www1.gotomeeting.com/register/117299321.

Dr. Bill Highleyman, Managing Editor

Never Again

Why Back Up?

The blogging platform JournalSpace suddenly went out of business when it lost its database.

Apparently, the database's demise was the malicious act of a disgruntled employee – even worse, the IT manager. JournalSpace claims that it had caught the IT manager stealing from the company. They summarily fired him; but he did a slash-and-burn on his way out, overwriting the entire database with garbage.

It was then that JournalSpace management realized that there was no backup. Their valiant attempts to recover the database failed, and they closed their doors days later. Thousands of bloggers lost years of their work.

--more--

Availability Topics

Active/Active Full-Day Seminar at HPTF

On June 15th, Dr. Bill Highleyman, Managing Editor of the Availability Digest, will teach a full-day preconference seminar at HPTF&E, the HP Technology Forum and Expo. Sponsored by HP and Connect, the HP Business Technology Community, the conference will be held from June 15th through June 18th in Las Vegas.

Entitled "Active/Active Systems: Theory and Practice," the seminar will describe active/active systems and why they provide such high availability. It will conclude with several examples of NonStop, OpenVMS, and Linux active/active systems that have been in production for years without a failure.

Dr. Bill will present an overview of his seminar during a one-hour "sneak peek" webinar on Thursday, May 21st, at 11 am Eastern Daylight Time. You can register at https://www1.gotomeeting.com/register/117299321.

--more--

Recommended Reading

High Availability Network Fundamentals

High Availability Network Fundamentals is a practical guide to predicting network availability, complete with many examples using Cisco products in complex networks.

Though calculating availability of complex systems is a highly-mathematical exercise, the author takes great pains to provide a detailed and accurate approach to availability analysis without requiring intensive mathematics – just a little simple arithmetic. There are only four equations in the entire book, yet he shows the reader how to do extensive availability analysis through a process that he dubs "divide and conquer." He shows the reader how to account for hardware and software faults, failover time, power outages, human error, and network device redundancy.

The book includes a CD containing the SHARC System Hardware Availability and Reliability Calculator spreadsheet. This spreadsheet does all the arithmetic work for the analyst as well as performs calculations deemed too complex for the reader.

Though the book focuses on network availability, its techniques are directly applicable to redundant systems such as clusters, primary/standby configurations, and active/active systems.

--more--

The Geek Corner

Configuring to Meet a Performance SLA – Part 5:
Multiple Servers with General Service Times

An SLA often contains a performance specification in the form of "99% of all transactions will complete in less than 50 msec." Whether or not we can meet this specification is the SLA question.

In Part 4 of this series on meeting a performance SLA specification, we used the Gamma distribution to provide a powerful method for determining the probability that the response time of a server with general service times will be less than a specified amount. Alternatively, we showed how to calculate the response time that could be achieved with a certain probability. We reduced these rather complex calculations to a set of charts and spreadsheets.

However, most systems comprise a set of servers acting in tandem through which an event must pass before a response to that event is generated. Each of these servers could have a different distribution of service times and could be carrying different loads. How do we calculate the performance SLA parameters for such a complex system?

Here in Part 5, we show how to use the results of the previous articles to solve this problem.

--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: _____

_____

_____

_____