"The Availability Corner"
Advice and Solutions for Enterprise Computing

As Seen in *The Connection*, An ITUG Publication
November 2004 – September 2006

About the Authors:

Dr. Bill Highleyman, Paul J. Holenstein, and Dr. Bruce Holenstein, have a combined experience of over 90 years in the implementation of fault-tolerant, highly available computing systems. This experience ranges from the early days of custom redundant systems to today's fault-tolerant offerings from HP (NonStop) and Stratus.

Series Topics:

Testing Your System Recovery Plan (09/06)
Is IBM's Parallel Sysplex a NonStop Competitor? (06/06)
Grid Computing (03/06)
The Net Present Value of Active/Active Systems (01/06)
TCO for Active/Active Systems (11/05)
Fault Tolerance vs. High Availability (09/05)
The Great Tape Backup Paradigm Shift (07/05)
The Language of Availability (05/05)
What Reliability Do We Really Need? (01/05)
Let's Measure System Reliability in Centuries (11/04)

Gravic, Inc.
Shadowbase Products Group
17 General Warren Blvd.
Malvern, PA 19355
610-647-6250
www.ShadowbaseSoftware.com

The Availability Corner

TCO for Active/Active Systems

November/December 2005

Dr. Bill Highleyman
Dr. Bruce Holenstein
Paul J. Holenstein

As with any system, we can optimize availability, cost, and performance. Pick any two. When it comes to cost, understanding and optimizing the total cost of ownership (TCO) of an active/active system is a lot more involved than it is for standard monolithic systems.

Of course, active/active systems have all of the cost components of monolithic systems – initial hardware and software costs, licenses, maintenance contracts, networks, personnel, facilities – and let's not forget the cost of downtime. But with active/active systems, many of these areas involve many more decisions than we are used to making.

Active/active systems are networks of peer nodes, each engaging in a common application, and each applying transactions simultaneously to a common redundant application database. The first decision to be made is what system to use for the nodes. We talked in our last column about using highly available systems (UNIX or PC systems) versus fault-tolerant systems such as NonStop servers. In order to achieve high availabilities of six nines or greater, it is important to have additional nodes in the system over and above those required for the application load so that capacity is maintained in the event of a node failure. It can be shown[1] that to achieve a given level of availability (say six 9s or better), fewer fault-tolerant nodes may be required than highly available nodes. Thus, hardware costs for the system could in fact be less if fault-tolerant nodes are used.

Active/active systems may incur additional software licensing costs if licenses are provided on a node basis in which each node must incur its full share of software licenses. If licenses are offered based on node size or application load, then this difference becomes less significant.

Active/active systems certainly have more extensive networking requirements in order to provide communication between nodes, and these networks must be very reliable. This generally means that the communication links between nodes should be redundant. These links are needed not only for database synchronization but also for heartbeats to monitor the health of nodes and for switching user traffic in the event of a node failure. An important area for compromise between cost and availability is the degree of redundancy provided for each type of communication link.

---

[1] W. H. Highleyman, Paul J. Holenstein, Bruce D. Holenstein, "*Breaking the Availability Barrier; Survivable Systems for Enterprise Computing*," Author House; 2004.

Because the nodes in an active/active network are likely to be geographically distributed for disaster tolerance or for data locality reasons, there will be duplication of facilities and to some extent personnel depending upon the degree to which remote sites can be configured to be "lights-out" operations.

To manage this complex network, a very good network management facility such as HP's OpenView is required. The cost of this tool could well be offset by personnel savings if it enables "lights-out" operations at remote nodes.

Another cost savings to be chalked up to fault-tolerant systems is again "lights-out" operations. Since a fault does not immediately take a node down as it may for a UNIX or PC system, there is less of a need for personnel resident at each node in the network to quickly repair the node and return it to service.

And finally there is the cost of downtime that must be considered when comparing an active/active architecture to a monolithic three or four 9s architecture. A recent advocacy survey[2] showed that the cost of downtime for half of the respondents was $10,000 per hour and over $100,000 per hour for 20% of the respondents. Things could be a lot worse. A 22-hour outage in 1999 at eBay not only cost them an estimated $3 million to $5 million in lost business, but it also sent their stock plummeting by 25%, representing a market capitalization hit measured in the billions of dollars. AOL suffered a similar fate of a 15% stock loss in 1996 due to an outage. Even worse for the corporate image, these outages became "CNN Moments."

Want to know more about computing a meaningful TCO for active/active systems? That will be the topic of our next Availability Corner.

[2] See www.hpuseradvocacy.org, Instapoll Archive dated December 18, 2003.