GRAVIC®
Shadowbase

“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

Testing Your System Recovery Plan

September/October, 2006

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

A major railroad turned to its diesel generators to continue power for its train control systems. Guess what? The UPS would not come up, and train control came to a virtual stop. Why wouldn't it start? We don't know, but even more to the point is that the railroad didn't know. In another case, an electric utility tried to start their diesel generators during an outage, and these generators also would not start – in this case because the fuel was old and had turned into jelly. These are clear cases of recovery plans that weren't periodically tested.

From our perspective as IT professionals, ensuring that there are satisfactory backup computing facilities is high on the list of business continuity planning.

Of course, like the diesel generators, a backup site is of little use if it is not operational, if it has not been configured to support the critical processing activities required during the crisis, and if the operations people are not thoroughly versed in the minutiae of the recovery procedures. People training is especially important as this will be a high-stress time. Things will go wrong, and management will be demanding answers. Remember, to paraphrase Wendy Bartlett of HP, when things go wrong, people get stupider. This is not the time for discovering how something should be done. This is the time for rote reacting according to established and well-rehearsed procedures.

However, here we come to a big roadblock. Today's backup sites are typically comprised of systems doing other (or no) work. They can take hours to bring online as the database and the applications are brought up. If your services are 24x7 mission-critical, how do you take down your perfectly running system for a few hours to switch over to the backup system as a test and then take another few hours to switch back to the primary system? The answer often is that you don't, and the recovery plan collects dust. It gets out of date, and operations personnel forget all the little details of a successful recovery. They are left to figure out a lot of the details in the stress of the crisis – a certain recipe for further disaster.

How can we solve this problem? One way that is being proven today is to move to an active/active system. Like an active/backup system, an active/active system also has multiple processing nodes that can be geographically separated for disaster tolerance. However, in an active/active system, all nodes are actively processing transactions against a distributed common copy of the application database. These database copies are kept in synchronism via some technology such as data replication.

Thus, should one node fail, whether by hardware or software fault or by natural or man-made disaster, all that must be done is to switch the users from the failed node to a surviving node or nodes. This can be done in seconds.

It must be ensured that the surviving node or nodes have sufficient capacity to handle the load during this time of failure. Perhaps some noncritical load can be shed. Otherwise, some additional capacity can be provided in the application network. Since all purchased capacity is available for the application (as opposed to only half of it in an active/backup configuration), it is often possible to achieve the required excess capacity with less purchased capacity than that required for an active/backup configuration.[1]

If running active/active should be seen not to be practical for a certain application, the current active/backup configuration can be upgraded to an active/active configuration in which the "backup" system is ready to process transactions; but no transactions are, in fact, routed to it during normal operation. Should the primary system fail, users can be switched to the backup system, again within seconds. This is called "sizzling hot takeover."

The use of active/active technology solves the recovery plan testing syndrome. If recovery from a node failure can be accomplished almost transparently to the users, node failures can be simulated at will to test the recovery plan. This technique is, in fact, being used today. A major financial firm in the mid-western United States keeps their people on their toes with unannounced simulated node failures in a sizzling hot takeover configuration. A bank service provider with processing centers in Florida and the Midwest will close down its Florida center with approaching hurricanes. A major healthcare facility is taking steps to go to active/active to facilitate recovery testing and system upgrades.

Remember that an untested recovery plan is no better than the dust it collects. Active/active technology provides the methods to painlessly test your recovery procedures.

[1] Dr. Bill Highleyman, P. J. Holenstein, Dr. Bruce Holenstein, *Breaking the Availability Barrier: Survivable systems for Enterprise Computing*, AuthorHouse; 2004.