



**“Achieving Century Uptimes”
An Informational Series on Enterprise
Computing**

**As Seen in *The Connection*, An ITUG Publication
December 2006 – Present**

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.

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

Achieving Century Uptimes
Part 25: Is the Active/Active Topic Getting Stale?
November/December 2010

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

Active/active systems¹ don't just provide *high availability* – they provide *continuous availability*. Active/active systems don't just provide *disaster recovery* – they provide **disaster tolerance**. Faults ranging from server failures to entire data-center outages can be transparent to the users of the applications. Active/active systems can eliminate all forms of application downtime and can lessen or reduce to zero data loss when catastrophic system or site failures occur.

But has the active/active topic become stale? To hear comments from some quarters, one might think so. Comments have included objections that we are hearing too much about active/active, that the costs of active/active systems limit them to just a handful of highly critical applications with million-dollar-per-hour downtime costs, that active/active technology is pie-in-the-sky,² and that systems today are so reliable that this seemingly overly complicated technology is not needed anyway.

But we think that the active/active message is just getting out. Let's look at these and other objections, at where active/active is today, and where active/active will be going in the future.

Progression from Pie-in-the-Sky

Is active/active a pie-in-the-sky technology? Let's look at its history.

In the late 1980's the first pioneers dipped their toes into the active/active waters. Wolfgang Breidbach guided Bank-Verlag into what may be the first NonStop active/active system ever built³ as the German banks introduced debit cards. Wells Fargo implemented an active/active ATM network.⁴ Both of these systems used home-grown data-replication engines to synchronize the nodal databases because active/active replication engines had not yet appeared in the marketplace. A decade would lapse before this happened.

Around the same time, in 1984, Digital Equipment Corp. introduced VAX clusters. These systems were intended to provide continuous availability, but technical problems tended to limit their introduction; and they did not compete well against Tandem in the early days. However, now known as HP's OpenVMS Split-Site Clusters,⁵ they offer support for up to 96 nodes in an

¹ What is Active/Active, *Availability Digest*; October 2006.

² An unattainable promise of a great return.

³ Bank-Verlag – The Active/Active Pioneer, *Availability Digest*; December 2006.

⁴ Wells Fargo's Pioneering Active/Active ATM Network, *Availability Digest*; September 2010.

⁵ OpenVMS Active/Active Split-Site Clusters, *Availability Digest*; June 2008.

active/active environment. These systems are mature and are installed in hundreds of applications providing the benefits of active/active.⁶ OpenVMS Clusters and their predecessors were probably the first commercial offering of active/active technology.

IBM's Parallel Sysplex⁷ systems were introduced in 1994. Though these systems require considerable consulting support from IBM, they, too, provide active/active capabilities across up to 32 IBM mainframe nodes. They are found primarily in classic IBM shops that had to squeeze more availability out of their IBM mainframes.⁸

Also in the mid-90s, commercial off-the-shelf, bidirectional database-replication engines necessary for active/active systems began to appear, primarily in the NonStop world. It is the NonStop user community that is the most demanding of application availability, and so it is natural that such products first appeared in this arena. Today, some of these products include DRNet from Network Technologies, Inc. (NTI), GoldenGate from Oracle,⁹ and Shadowbase from Gravic.¹⁰ An important point to make, which we shall address later, is that some of these products are heterogeneous and extend far beyond the NonStop world to Windows, Linux, UNIX, IBM systems, and others. Thus, the benefits of active/active now extend past the worlds of NonStop, OpenVMS, and IBM mainframes to the world of industry-standard servers.

No longer is active/active a pie-in-the-sky architecture. It is a mature technology in use worldwide in many applications, including finance, securities trading, ATM/POS, telecom, Internet, healthcare, transportation, process control, and even gaming.

Is Active/Active Overly Complicated?

There is no representation that active/active is a simple technology. Some key issues are resolving data collisions and potential data loss if asynchronous replication is used or avoiding them by using performance-impacting synchronous replication. However, there are commercial solutions to these and other active/active issues available today, primarily in the many off-the-shelf data-replication engines.

Another critical issue is modifying existing applications to run in a distributed environment. If this is a critical impediment, an application can be run in a "sizzling-hot" environment in which it runs on only one node in an active/active system, thus still achieving most of the availability benefits of active/active. As new applications designed for distributed processing come online, this problem will disappear.

⁶ UK National Health Service – Blood and Transport, *Availability Digest*; October, 2008.

⁷ Parallel Sysplex – Fault Tolerance from IBM, *Availability Digest*; April 2008.

⁸ Handelsbanken Turns to Parallel Sysplex, *Availability Digest*; October 2009.

⁹ Flexible Availability Options with GoldenGate's Transactional Data Management Platform (TDM), *Availability Digest*; February 2007.

¹⁰ Shadowbase, the Active/Active Solution, *Availability Digest*; March 2007.

Is Active/Active *Too Expensive* for Ordinary Applications?

Perhaps. At least with today's capabilities. Going active/active is not inexpensive.¹¹ However, if you are already running a backup site (which you probably are if you are considering active/active), then you are already covering a major portion of the cost.

Additional costs include possibly needing to modify your application to run in a distributed environment, redundant independent networks, and additional license fees. Your low-cost backup system and application licenses may have to be upgraded to full operational licenses, and you will need to license a bidirectional data-replication product and perhaps distributed application management tools.

However, if your downtime costs more than cover these additional expenses, if you are focused on customer satisfaction, if you are in an intensely competitive environment where application availability is a valuable differentiating factor, or if you are under regulatory restrictions, active/active systems could be your best answer to achieve continuous system availability.

Active/active costs are going to become less of an impediment as time goes on. We have come a long way in commoditizing active/active technology beyond the NonStop, OpenVMS, and IBM mainframe worlds. It is now time to move into or integrate commodity servers. Today's bidirectional replication products support Windows, Linux, UNIX, and even the smaller IBM iSeries systems. They can work heterogeneously with systems from different vendors in a Megaplex, a fabric of resources for efficiently and economically implementing new applications.¹²

Multiple expensive data centers will not be needed for commodity systems. Though one node should be actively managed, there is no reason why redundant nodes cannot be in a closet somewhere and managed remotely from the primary node's site. Should a node fail and the closet node needs to take over the failed node's load, this can be an unattended function. Certainly, the license costs for additional nodes will be a fraction of those for mainframe nodes.

A major impediment today is the problem of modifying applications to run in a distributed active/active environment.¹³ However, companies are now designing new applications to run in these environments. VocaLink, the company that has just introduced real-time payment services for the U.K. banking system,¹⁴ designed its new Faster Payments Service system specifically to run active/active. Introduced in 2008, this NonStop system has run for more than two years without an outage due either to unplanned downtime or planned downtime.

¹¹ Achieving Century Uptimes – Part 24: Is It Worth the Effort to Move to Active/Active?, *The Connection*; September/October, 2010.

¹² Megaplex – An Odyssey of Innovation, *The Connection*; May/June 2010.

¹³ Achieving Century Uptimes – Part 25: Is Your Application Active/Active Ready?, *The Connection*; January/February, 2010.

¹⁴ Faster Payments – Bringing Payment Processing Into the 21st Century, *Availability Digest*; June 2010.

As time goes on, there is no reason not to design new critical applications running on commodity servers that will operate in a distributed environment. This is not a particularly difficult discipline – it is only sometimes difficult after the fact. As this trend takes hold, active/active technology will become much less expensive and will begin to penetrate those markets in which downtime costs are thousands or tens of thousands of dollars an hour, not hundreds of thousands of dollars per hour, multiplying the number of applications that are candidates for active/active technology manyfold.

Is Active/Active Overly Promoted?

To our knowledge, there is only one publication devoted to continuous availability, and that is the *Availability Digest*.¹⁵ Even the respected *Disaster Recovery Journal*, which focuses on the next best topic, Business Continuity Planning (BCP), just published its first article on continuous availability using active/active architectures.¹⁶

There are no conferences promoting active/active best practices as there are for BCP. There are no accreditations offered for active/active expertise. The best you can find today are the few active/active sessions given at HP NonStop meetings and conferences and the occasional webinar out of the NonStop world. But these sessions show that active/active is still of great interest.

No, we don't think that active/active technology suffers from over-promotion. In fact, the problem is to spread understanding of this technology not only within the NonStop world but also beyond into the world of commodity servers.

Is Active/Active Needed?

Most companies today run their critical applications with a backup that is usually remotely located to support disaster recovery. Even though these systems typically have recovery times measured in several hours, perhaps down to several minutes, they are considered “good enough.” But “good enough” is the enemy of “great” preventing operations from acquiring the resources from management to make the application great. Increasing pressures for round-the-clock continuous operation may soon render these systems “not so good.” This is perhaps the time for many companies to start reevaluating their availability needs, especially with regard to the costs and consequences of any downtime that they either now have or may experience in the future given their current approach.

One argument often put forth is that today's hardware and operating-system software are so reliable that failures are rare. NonStop Integrity DMR (dual modular redundancy) systems exhibit perhaps six 9s of availability, and the availability of TMR (triple modular redundancy) systems probably cannot even be measured. The NonStop operating system has been hardened over several decades and is rock-solid.

¹⁵ www.availabilitydigest.com.

¹⁶ W. H. Highleyman, *Achieving Continuous Availability with Active/Active Networks*, *Disaster Recovery Journal*; Summer 2010.

However, these arguments ignore the primary causes of faults. The vast majority of system faults today are not caused by hardware failures or by operating system bugs. They are caused by a host of other problems. System operator errors and maintenance activities can take a system down. Network faults can take a system down. Power outages and cooling failures can take a system down. Natural and man-made disasters such as hurricanes, fires, floods, earthquakes, disgruntled operator malfeasance, and terrorist attacks can take a system down – the infamous 9/11 attack destroyed dozens of data centers. Systems have been taken down by battery-room explosions, lightning strikes, and even confiscation of critical servers by law enforcement.

All told, these causes tend to reduce the availability of commodity servers to about three 9s and NonStop servers to about four 9s. Will this be *good enough* as competitive pressures push for continuous availability? We think not. Active/active technology is the next obvious step to convert “good enough” to “great.”

Where is Active/Active Going?

One inhibitor of achieving continuous availability today is the current economic meltdown. There is not much that we can do about that. Companies are understandably cautious in their investments that do not result in an immediate benefit to their bottom line. A system that is “good enough” may be acceptable. But as the recession comes to an end, the need for continuous availability will have progressed by several years. We think that we will see a resurgence in interest in continuous availability as this occurs in order for companies to remain competitive. This is especially true as companies implement the latest trends for matching their availability needs to their business needs and deploy applications across a range of platforms and environments in the Megaplex.

Another inhibitor is the problem of modifying an application to make it active/active ready. This may be quite difficult for many applications, especially older ones for which the source code is lost and is no longer available. It may be impossible to modify third-party products whose source code is proprietary. Going forward, active/active technology may be more likely applied to new applications that can be designed from the start to work properly in a distributed environment. This trend has already begun, as we have previously noted with respect to VocaLink.

Is active/active technology limited to large mainframe systems because of its cost? We know this is not the case. Today’s products support active/active configurations for commodity servers. We expect that as the active/active word gets spread, interest in it will begin to build in the commodity server community. As this happens, active/active technology will be positioned to explode.

What’s Next

All this having been said, the NonStop community has perhaps been drenched with the active/active story. We are therefore going to broaden the scope of our monthly column. Look for articles in the coming issues of *The Connection* that deal with more general topics in the

realm of availability, integrity, scalability, reliability, and security. If you have a topic that you would like us to explore, please contact us.