



“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

Is IBM's Parallel Sysplex a NonStop Competitor?

June/July, 2006

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

HP's NonStop servers and IBM's Parallel Sysplex are the primary industry offerings that can tolerate any single failure, thus leading to very high levels of availability. But is Parallel Sysplex really a serious NonStop competitor? To explore this, we first describe the Parallel Sysplex architecture. We then look at a couple of important differences that answer this question.

IBM's Parallel Sysplex systems are multi-processor clusters that can support from two to thirty-two mainframe nodes (typically S/390 or zSeries systems as of this writing). A *node* may be a separate system or a logical partition (LPAR) within a system.

A Parallel Sysplex system¹ presents not only a single system image to its users, applications, and networks but also a single point of control for the operations staff. The following figure shows a high level view of a Parallel Sysplex system. It comprises two or more mainframe nodes (partitions or systems) and a shared disk subsystem.

The "brain" of the Parallel Sysplex system is the Coupling Facility, or CF. It maintains links to each node in the system and contains all shared resources such as locks, lists, queues, and shared data. It also monitors the nodes and contains the status of the entire system. It is itself a mainframe system such as a zSeries system.

In addition, a Sysplex Timer connects to all nodes in the system. It provides the time-of-day synchronization for all nodes.

Both the CF and the Sysplex Timer are single points of failure in the network. However, a second CF and/or Sysplex Timer can be provided for redundancy.

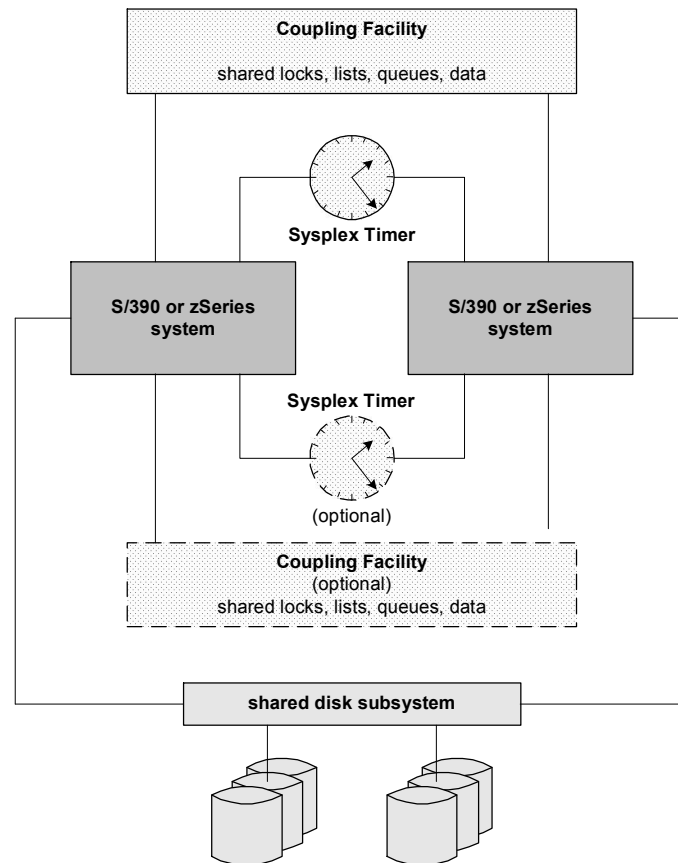
Since 1996, IBM has offered Parallel Sysplex as a geographically distributed system for disaster tolerance. This configuration is called Geographically Dispersed Parallel Sysplex (GDPS).² It is basically IBM's analog to RDF.

The GDPS system allows remote Parallel Sysplex systems to backup each other. In addition to providing remote backup copies of databases, GDPS provides automated failover and system error recovery. Database synchronization between the nodal

¹ "Parallel Sysplex Cluster Technology," IBM White Paper, www-03.ibm.com/servers/eserver/zseries/pso/sysover.html

² "GDPS: The e-business Availability Solution," IBM White Paper, March, 2005.

databases may be done either by synchronous replication over distances up to 100 km (Peer-to-Peer Remote Copy – PPRC – recently renamed Metro Mirror) or by asynchronous replication (eXtended Remote Copy – XRC – recently renamed Global Mirror). If PPRC is used, there is no data loss as a result of a node failure; but performance will generally be affected. If XRC is used, there may be seconds to minutes of data loss.



Now for the differences.

Unlike NonStop servers, Parallel Sysplex systems are not “out-of-the-box.” They cannot be ordered as a product from IBM. Rather, a Parallel Sysplex system comprises hardware products, software products, and extensive analysis services from IBM. IBM’s documentation states that “Continuous application availability for zSeries applications cannot be achieved without Parallel Sysplex. However, Parallel Sysplex on its own cannot provide a continuous application availability environment. Continuous or near-continuous application availability can only be achieved by properly designing, implementing, and managing the Parallel Sysplex systems environment.”³

³ “Parallel Sysplex Availability Checklist,” IBM Corporation; May, 2003.

Furthermore, the complex of mainframe systems needed to implement and manage a Parallel Sysplex system is understandably much more expensive than a NonStop system. TCO (total cost of ownership) studies by The Standish Group have indicated that an equivalent IBM Parallel Sysplex system has a five-year cost that can be twice that of an equivalent HP NonStop server.⁴

So unless you are committed to IBM for corporate policy reasons, NonStop remains the choice. It is out of the box so it is much easier to implement, and it is less expensive than an equivalent Parallel Sysplex system.

⁴ See “Digging the TCO Trenches,” 2004 Research Note from The Standish Group, at h20219.www2.hp.com/NonStopComputing/downloads/DiggingTCOTrenches.pdf. Also see “Dollars to Cents: TCO in the Trenches 2002,” 2002 Research Note from The Standish Group, at h20219.www2.hp.com/NonStopComputing/downloads/TCOTrenches02.pdf.