the Availability Digest

Parallel Sysplex – Fault Tolerance from IBM

April 2008

IBM's Parallel Sysplex, HP's NonStop server, and Stratus' ftServer are today the primary industry fault-tolerant offerings that can tolerate any single failure, thus leading to very high levels of availability. The Stratus line of fault-tolerant computers is aimed at seamlessly protecting industry-standard servers running operating systems such as Windows, Unix, and Linux. As a result, ftServer does not compete with the other two systems, Parallel Sysplex and NonStop servers, which do compete instead in the large enterprise marketplace. In this article, we will explore IBM's Parallel Sysplex and its features that address high availability.

IBM's Parallel Sysplex

IBM's Parallel Sysplex systems are multiprocessor clusters that can support from two to thirty-two mainframe nodes (typically S/390 or zSeries systems).[1] A Parallel Sysplex system is nearly linearly scalable up to its 32-processor limit.

A *node* may be a separate system or a logical partition (LPAR) within a system. The nodes do not have to be identical. They can be a mix of any servers that support the Parallel Sysplex environment.

The nodes in a Parallel Sysplex system interact as an active/active architecture. The system allows direct, concurrent read/write access to shared data from all processing nodes without sacrificing data integrity. Furthermore, work requests associated with a single transaction or database query can be dynamically distributed for parallel execution based on available processor capacity of the nodes in the Parallel Sysplex cluster.

Parallel Sysplex Architecture

All nodes in a Parallel Sysplex cluster connect to a shared disk subsystem. In this way, any node has access to any table or file in the cluster. For reliability, disks are organized either as mirrored pairs or as RAID arrays.

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, including locks, cache, and queues. It is the CF that allows shared access by all processing nodes to all tables and files in the database.

The CF also monitors the nodes and contains the status of the entire system. It is itself a mainframe system such as a zSeries system.

---

[1] "Parallel Sysplex Cluster Technology," IBM White Paper, www-03.ibm.com/servers/eserver/zseries/pso/sysover.html.
J. M. Nick, B. B. Moore, J. -Y. Chung, N. S. Bowen, S/390 Cluster Technology: Parallel Sysplex, *IBM Systems Journal*; November 2, 1997.

In addition, a Sysplex Timer connects to all nodes in the system. It provides a common clock for time synchronization between all nodes.

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

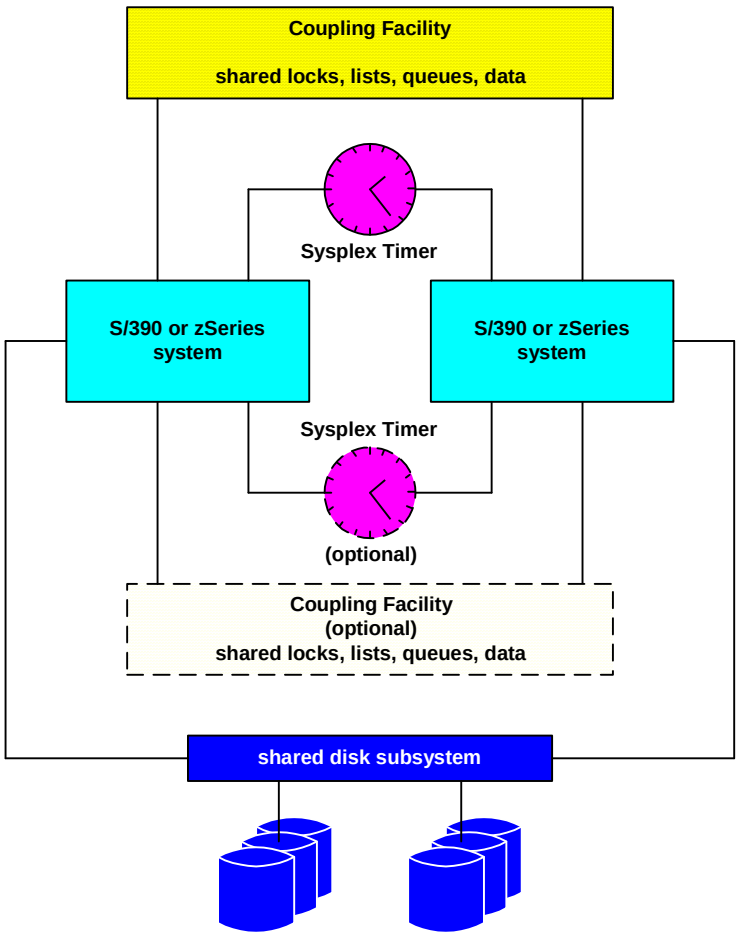

The Coupling Facility

The Parallel Sysplex architecture is a shared-data model. It enables multiple systems to cache the same data concurrently in local processor memory with full read/write access control and globally managed cache coherency and with high-performance and near-linear scalability.

The Coupling Facility (CF) is the key to this capability. The CF enables high-performance read/write sharing of data by applications running on each node of the cluster through global locking and cache coherency management mechanisms. It also provides cluster-wide queuing mechanisms for workload distribution and for message passing between nodes.

The CF is itself a processing system such as an IBM zSeries system. The CFs are attached to the other processors in the cluster via high-speed coupling links. The coupling links use specialized protocols for highly optimized transport of commands and responses to and from the CF. The coupling links are fiber-optic channels providing 100 megabyte per second data transfer rates.

If a Coupling Facility fails, critical data contents can be "rebuilt" into an optional alternate CF under cluster control.

The Coupling Facility architecture provides three mechanisms to enable efficient clustering protocols – globally-managed locks, globally-managed cache, and globally-managed queues.

Globally-Managed Locks

The CF supports high-performance, fine-grained global locking and contention protocols. Stored in the CF is a hashed lock table. Each entry in the lock table contains the system identifier of the first system to register an exclusive interest in any of the lock resource names that hash to that lock table entry and the identities of other systems that share an interest in that hash class.

Parallel Sysplex provides locking services to obtain, release, and modify lock ownership state information for program-specified lock requests. To request lock ownership, a program passes the software lock resource name, the hash class value (to use as the index to the coupling facility lock-table entry), and the shared or exclusive interest in the resource. If the system does not already have a registered compatible interest in the specified lock-table entry, Parallel Sysplex will issue a command to the CF to perform the registration.

Through use of efficient hashing algorithms, false lock resource contention within a hash class is kept to a minimum. This allows the majority of requests for locks to be granted

Globally Managed Cache

The global cache supported by the CF is, in essence, a second-level cache. The CF provides global coherency controls for distributed local processor caches. Any data that is not subject to shared access is maintained in a processor's local cache. However, if the data represents a shared resource, the processor will write that data to the CF shared cache.

Writes to CF cache can either be store-in-cache, in which the data is written to cache and only flushed to disk periodically, or store-through-cache, in which the data is also written to shared disk storage.

The CF cache structure contains a global buffer directory that tracks multisystem interest in shared-data items cached in the CF. A separate directory entry is maintained in the CF structure for each uniquely named data item. A directory entry is created the first time a command requests registration of interest in the data item or a write operation is performed to place the data item into the shared cache.

Whenever a data item in the CF global cache is updated, an invalid signal is sent to all other systems that have registered an interest in that data item. Each of these systems will mark its current copy in its local cache as being invalid.

When a program wants to read a shared data item, it attempts to read that data from its local cache. It the local cache copy is valid, it is returned to the program. If the local cache copy is invalid, the data item is fetched from the CF cache, marked as being valid in the local cache, and returned to the program.

Globally-Managed Queues

The CF provides a set of queuing constructs in support of workload distribution, message passing, and sharing of state information. The queues are event-driven for processing efficiency.

A process connects to a queue by registering an interest in that queue. Next, whenever a message is placed in the queue that causes it to go from empty to non-empty, a signal is sent to all registered processes and informs them of this change in state. The registered processes can then each fetch the message from the queue. When a message has been read by all registered processes, it is deleted from the queue.

Further messages are added to the tail of the queue. When the last message is deleted from the queue, a signal indicating that the queue is now empty is sent to all registered processes.

The Sysplex Timer

The sysplex timer serves as a common time-reference source for systems in the cluster. The timer distributes synchronizing clock signals to all nodes. This enables local processor time stamps to be used reliably on each node and to be synchronized with respect to all other cluster nodes without requiring any software serialization or message passing to maintain global time consistency.

Dynamic Load Balancing

The entire Parallel Sysplex cluster can be viewed as a single logical resource to end users and applications. Work can be directed to any cluster node having available capacity. This avoids the need to partition data or applications among individual nodes in the cluster or to replicate databases across multiple collocated servers.

During initial connection to the cluster, clients can be dynamically distributed and bound to server instances across the set of cluster nodes to effectively spread the workload. Subsequently, work requests submitted by a given client (such as transactions) can be executed on any system in the cluster based on available processing capacity.

The work requests do not have to be directed to a specific system node due to data-to-processor affinity. Rather, work will normally execute on the system on which the request is received; but in case of overutilization on a given node, work can be redirected for execution on other less-utilized system nodes in the cluster.

The *Workload Manager* (WLM) maintains the response levels for the various applications according to their individual SLAs (Service Level Agreements). The WLM automatically balances the workloads across all nodes in the cluster to meet these agreements.

Application Compatibility

A design goal of the Parallel Sysplex system is that no application changes are required to take advantage of the technology. For the most part, this is true, though some CICS attributes may need to be tuned to get the maximum advantage from the cluster.

Single System Image

A Parallel Sysplex cluster provides simplified system management by presenting a persistent single system image across failures to the operators, end users, database administrators, and others. The single system image is provided from several perspectives:

- Data access, allowing dynamic workload balancing and improved availability.
- Dynamic transaction routing, also for dynamic workload balancing and improved availability.

- End user interface, allowing logon to a logical network entity rather than to a physical console.

- Operational interfaces for easier systems management.

Fault Tolerance

Processor heartbeat monitoring is provided to monitor the health of the nodes in the cluster. In addition, functions are also provided to automatically terminate a failing node and to disconnect the node from its externally attached devices.

In the event of a hardware or software outage, either planned or unplanned, workloads can be dynamically redirected to available servers, thus providing near continuous application availability. In addition, servers can be dynamically removed from or added to a cluster. This allows installation and maintenance activities to be performed while the remaining systems continue to process work. Hardware and software upgrades can also be rolled through the system.

During the unavailability of an application subsystem, new work requests can be redirected to other data-sharing instances of the subsystem on other cluster nodes to provide continuous availability across the outage and subsequent recovery. This provides the ability to mask planned as well as unplanned outages from the end user.

Automatic Restart Manager

The *Automatic Restart Manager* (ARM) enables fast recovery of the application subsystems that might hold critical resources at the time of failure. If other instances of the failed subsystems in the cluster need any of these critical resources, the ARM will make these resources available quickly.

The ARM provides the following functions:

- It detects the failure of a critical task.
- It automatically restarts the failed task.
- It automatically redistributes work to appropriate surviving instances following a failure.

The ARM utilizes the shared-state support provided by the CF so that at any given point in time, the ARM is aware of the state of processes on all systems (even of processes that "exist" on failed nodes). The ARM monitors the processor heartbeats so that it is immediately made aware of node failures. Furthermore, the ARM is integrated with the Workload Manager (WLM) so that it can provide a target restart system based on the current resource utilization across the available nodes.

The ARM contains many other features to provide improved restarts, such as affinity of related processes, restart sequencing, and recovery when subsequent failures occur

Disaster Tolerance

IBM offers Parallel Sysplex as a geographically distributed system for disaster tolerance. This configuration is called Geographically Dispersed Parallel Sysplex (GDPS).[2] GDPS is a multisite application and data availability solution designed to provide the capability to manage the remote copy configuration and storage subsystem(s), to automate Parallel Sysplex operational tasks, and to perform failure recovery from a single point of control, thereby helping to improve application availability.

[2] "GDPS: The e-business Availability Solution," IBM White Paper; February, 2008.

The GDPS system allows remote Parallel Sysplex systems to back up each other. In addition to providing remote backup copies of databases, GDPS provides automated failover and system error recovery. All that is necessary is for an operator to authorize the failover.

Database synchronization between the nodal 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; but the impact on application responsiveness is minimal.

The main focus of GDPS data replication is to maintain the data consistency of the backup site. Its database must contain all updates made to the primary site up to a given point in time.

IBM benchmarks recorded a failover time of less than 20 seconds for a system with over 6,000 volumes and 20 terabytes of data using PPRC (synchronous) replication.

GDPS replication is unidirectional only, from the primary system to the backup system. Therefore, the backup system cannot participate in applications executing on the primary system. However, the backup system can be used for other processing activities.

Summary

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]

Furthermore, the complexity of the mainframe systems needed to implement and manage a Parallel Sysplex system is understandably expensive. 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]

However, Parallel Sysplex provides a valuable upgrade option with no reprogramming requirements for mainframe applications if fault tolerance is required.

[3] "Parallel Sysplex Availability Checklist," IBM Corporation; May, 2003.

[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.