

IBM Parallel Sysplex clustering technology

The flexible solution to your enterprise computing requirements

Highlights

- **Designed for application availability of 99.99999%** — Removes single points of failure of server, LPAR, subsystems.
- **Resource Sharing** — Value today in a single server and multiple server environment. IBM Z brings even more dynamic real time efficient system resource utilization for "on demand" environment by Intelligent Resource Director. The combination of Parallel Sysplex® and Intelligent Resource Director enables your e-business application to take full utilization of all the servers' resources without affecting the datasharing environment.
- **Balance multiple workloads** — Let all your applications share system resources in order to meet business goals you define.
- **Dynamic workload balancing** — Maximize system performance and balance white space
- **Single system image (Increased productivity)** — Manage multiple systems as a single system from a single point of control and intelligently manage resources of multiple images within a single server.
- **Elimination of Batch window** — Parallel Sysplex enables technology to eliminate the batch window so your online transactions can continue 24x7.

An industry-leading solution

The IBM Z® platform is the defining standard in enterprise computing, particularly as enterprise computing is aligned with meeting strategic business objectives. The platform's scalability, ability to manage mixed workloads, availability and low total cost of ownership are the pillars on which enterprise-wide solutions can be built. [Parallel Sysplex](#) clustering technology builds on and extends the base strengths of the [IBM Z](#) platform.

Parallel Sysplex technology is a highly advanced commercial clustered system. With IBM's Parallel Sysplex technology, you can harness the power of up to 32 [z/OS](#)® systems, yet make these systems behave like a single, highly-scalable, highly-available logical computing platform. What's more, the underlying structure of the Parallel Sysplex remains virtually transparent to users, networks, applications, and even operations.

Each node can concurrently cache shared data in local processor memory through hardware-assisted cluster-wide serialization and coherency controls. As a result, work requests that are associated with a single workload, such as business transactions or data base queries, can be dynamically distributed for parallel execution on nodes in the sysplex cluster based on available processor capacity. This, in turn, facilitates dynamic workload balancing. In business terms, these capabilities can mean the difference between a system that supports high availability and one that doesn't.

What Parallel Sysplex clustering technology provides is flexibility — a means to match system performance to business requirements on an as-needed basis. It permits rapid response to unpredicted growth.

“At US Bank we look to expand through acquisitions. With our Parallel Sysplex environment, we know that we can quickly add processing power and grow databases while still maintaining consistent service levels even if we more than double the size of the business.”

Mark Sobotka, VP technical services US Bancorp (05/99)

Support for an “Always On” environment

Within a Parallel Sysplex cluster it is possible to construct an environment with practically no single points of failure. In a data sharing environment, all systems in the Parallel Sysplex can have concurrent access to all critical applications and data, the loss of a system due to either a planned or unplanned hardware or software outage does not necessitate loss of application availability. Peer instances of a failing subsystem executing on remaining healthy system nodes can take over recovery responsibility for resources held by the failing instance.

Alternatively, the failing subsystem can be automatically restarted on still-healthy systems using automatic restart capabilities to perform recovery for work in progress at the time of the failure. While the failing subsystem instance is unavailable, new work requests can be redirected to other data-sharing instances of the subsystem on other cluster nodes to provide continuous application availability across the failure and subsequent recovery. This provides the ability to mask planned as well as unplanned outages to the end user. Even though they work together and present a single image, the nodes in a Parallel Sysplex cluster remain individual systems, making installation, operation and maintenance nondisruptive.

Multitenancy

The Parallel Sysplex supports multitenancy between nodes in the cluster. This is true for both the IBM Z hardware, with support for a server to communicate to another server up to 2 levels back (n-2). z/OS images can be in the same Parallel Sysplex as another z/OS one level back. This is also true for middleware such as [CICS®](#), [Db2®](#), and [IMS™](#). This allows you to introduce changes such as hardware or software upgrades one system at a time. Remaining systems continue to process work. You can roll changes through your systems at a pace that makes sense for your business.

99.99999% availability

IBM has taken availability to a new level with the introduction of GDPS® Metro which provides end-end automated recovery even in the event that an entire data center becomes inoperative. A fully configured Parallel Sysplex data sharing environment together with GDPS Metro support for HyperSwap® for disk events is designed to provide 99.99999% availability. A stark contrast with the only 99.5% availability that many cloud providers guarantee.

More information on GDPS can be found at www.ibm.com/products/gdps or the Redbook®, “[IBM GDPS Family: An Introduction to Concepts and Capabilities](#)”.

Resource Sharing

As businesses grow and corporate data centers become consolidated under the same organization, the number of configured operating system images (or LPARs) increase as a natural result. Each image, regardless of whether it is on a single footprint or multiple footprints, requires basic hardware/software resources to function, such as files, tape drives, consoles, log data, system catalog, access to the JES2 checkpoint, etc. As the number of images increase, the requirement for these redundant hardware resources becomes very costly and difficult to manage. Many customers running with a single footprint with logical partitions are realizing that they too face the complexities presented by multiple images.

Resource Sharing uses Coupling technology to allow multiple z/OS images to manage and share a resource as if it were a single resource attached to a single image. Being able to efficiently share resources between multiple images enables IT organizations to utilize existing resources more efficiently and reduces or eliminates the need to purchase additional hardware and software to support multiple systems. In addition, z/OS Resource Sharing can reduce the overall cost and complexity in managing multiple system images.

This is NOT to be confused with Parallel Sysplex data sharing by the database subsystems. Resource Sharing delivers immediate value even for customers who are not leveraging data sharing, through native system exploitation delivered with the base z/OS software stack.

See more examples of Resource sharing - "[IBM Z - Parallel Sysplex: Coupling Facility Information](#)".

If you are already getting the benefits of resource sharing by managing multiple partitions on a single server, you are a perfect candidate for getting even greater value from Intelligent Resource Director.

Intelligent Resource Director

The IBM Z servers also automatically directs processor and I/O resources to priority work through Intelligent Resource Director (IRD). The IRD combines the strengths of three key technologies: z/OS Workload Manager, Logical Partitioning, and Parallel Sysplex clustering. It is the logical next step in resource sharing.

This powerful combination of IBM Z servers and z/OS provides the ability to intelligently manage numerous operating system images executing in the same Parallel Sysplex on a single server as a single compute resource, with dynamic workload management and physical resource balancing such as processors and I/O across logical partitions. In other words the system can dynamically allocate LPAR weights, logical CPs, channel paths, and channel controller work across multiple virtual servers in order to ensure that the unpredictable demands of the online business environment can be intelligently managed according to business goals.

Keeping work in balance

The entire Parallel Sysplex cluster can be viewed as a single logical resource to end users and business applications. That's critical if your business depends on sharing the most up-to-date information without delay.

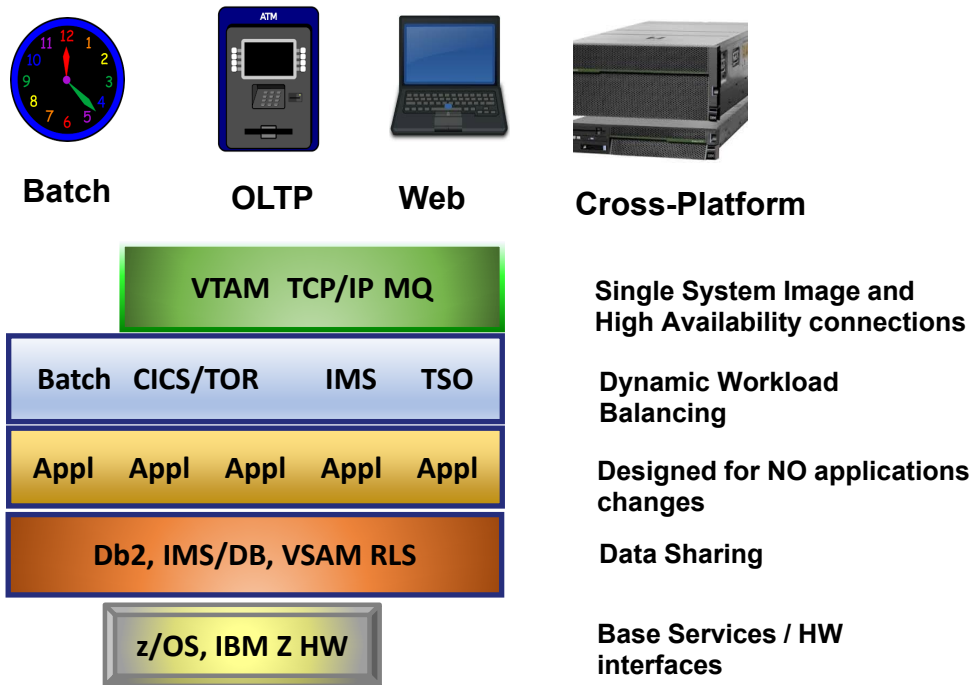
Just as work can be dynamically distributed across the individual processors within a single server, so too can work be directed to any node in a Parallel Sysplex cluster having available capacity.

"The Parallel Sysplex system gives us scalability and cost savings while adding to our ability to maintain service levels. We merged four banks and our workload was too big for any single processor. Capacity can be added to a Parallel Sysplex system when needed, as needed. Datasharing permits dynamic workload balancing—the ability to shift work as the situation demands. If a processor is off-line, for either an unscheduled outage or scheduled maintenance, dynamic workload balancing further enhances availability by transferring work to other processors according to user-defined priorities—an important point when you have a nationwide network."

Cassie Smith, Technical Manager Software Services, Absa Bank (04/00)

Workload balancing also permits you to run diverse applications across a Parallel Sysplex cluster while maintaining the response levels critical to your business. You select the service level agreements required for each workload, and the z/OS Workload Manager (WLM), along with the subsystems such as CICSplex®/System Manger (CP/SM), IMS shared queues automatically balances tasks across all the resources of the Parallel Sysplex cluster to meet your business goals.

Since work will naturally flow to the server with the most available resources to get the best response times, workload balancing also helps manage “white space,” or unused CPU resource, simplifying capacity planning management.



Session balancing

TCP/IP networks are the backbone of the communications infrastructure for the Internet. Support for TCP/IP session balancing and across the Parallel Sysplex cluster supports seamless fail over and recovery for planned or unplanned outages. Dynamic Virtual I/P Addressing (DVIPA) allows the Parallel Sysplex to be perceived as a single large server node, independently of which images in the Sysplex the server applications execute on. In addition, VIPAs can survive an outage by moving to another stack via “VIPA Takeover.” Another application instance can then pick up the workload. The work can then be moved back to the primary owner again when it is available. This helps hides planned and unplanned outages from the end users.

What further differentiates z/OS is the ability to combine real-time factors with information from the Workload Manager. The communication server has the unique ability to ensure that the best destination server instance is chosen for a particular client connection while helping to ensure that client/server specific Service Level Agreements are maintained. All this is done by Sysplex Distributor with existing IBM Z hardware and software.

Whether the work is coming from batch, SNA, TCP/IP, DRDA®, or IBM MQ® messages, dynamic workload balancing, getting the business requests into the system best able to process the transaction provides the performance and flexibility you need to give the responsiveness your customers demand, and it is invisible to the users. The result is a system that's well suited to new "front-office" applications, such as Business Intelligence or online business, application development or running traditional "back-office" online transaction processing and batch jobs.

To do real online business, you need flexible, scalable, continuously available parallel systems — that's what Parallel Sysplex technology provides. You can supply Web services on your existing Parallel Sysplex platform without any extra investments.

"Our Parallel Sysplex strategy is crucial for our major, IMS based, workload to allow continuous availability and workload growth. Performance is maintained at acceptable levels, while system and resource management are simplified. This positions us very well for accepting the new workload that ecommerce will bring us."

Hans Baken — IT Architect
ABN AMRO Bank,
The Netherlands (04/00)

[WebSphere® Application Server](#), based on the strengths of z/OS, provides scalability, availability, security and communication power. Workload Manager creates Web servers as required to meet user needs; individual Web work requests are dynamically balanced for optimum performance and have access to your existing business-critical applications and data that reside on IBM Z servers, including IMS, Db2 and VSAM.

Optimizing online transactions You won't find a better computing environment to handle your critical online transaction processing (OLTP) than a Parallel Sysplex cluster.

The benefit of processing multiple transactions across many processors instead of processing the same transactions on two large processors is that the more processors in your cluster, the smaller the transaction queue. Transactions don't have to wait — and neither do users. It's that simple.

Low cost of ownership with simplified system management

Numerous industry studies have shown that Parallel Sysplex clustering technology offers a low total cost of ownership. The cost advantage for Parallel Sysplex technology comes, in part, from the ability to manage the entire cluster as a single system from a single point of control and to build on existing assets.

Managing hundreds or even thousands of distributed servers is a time-consuming and expensive process with the amount of people needed to manage it grows with the number of distributed servers. By contrast, a Parallel Sysplex cluster can present itself as a single system image to Operations and to System Programmers. A single operator on a single console can view the entire complex and send commands to one or multiple images at the same time. Similarly, through the use of system symbolics, a System Programmer need only make a single definition change for it to have effect across the Parallel Sysplex.

By consolidating your standalone servers on a Parallel Sysplex cluster, you can drastically lower the number of footprints that you need to manage and maintain. With many databases and business applications from hundreds of independent software vendors enabled for Parallel Sysplex technology — as well as key applications and middleware from IBM — you can consolidate your critical business applications and gain the benefits of Parallel Sysplex technology as well.

More work in less time

Batch jobs remain an integral part of today's computing environment and that's putting many companies in an uncomfortable position. On the one hand, pressures to keep the OLTP environment up-and-running around the clock are cutting into the time once available for batch processing. On the other hand, increases in transaction volume are often accompanied by increased batch demands. The question is how do you fit more work into a shrinking window? The answer is Parallel Sysplex clustering technology with CICS Transactional VSAM.

Transactional VSAM allows VSAM to act as a database manager. Batch jobs can now perform updates with backout capabilities to the same recoverable files used by CICS online transactions, effectively removing the batch window. Together with WLM managed initiators, the high priority jobs can get an initiator to run on with the workload getting balanced automatically within the Sysplex.

Easy starting

How difficult is it to implement a Parallel Sysplex cluster? Toyota UK implemented a single-system cluster within a matter of weeks. Many other customers are finding the initial steps easy enough, and the benefits so clear, they've also made the move themselves. Yet we keep making it even easier.

If you prefer, IBM Implementation Services for Parallel Sysplex can help you make a transition to Parallel Sysplex clustering. IBM Lab Services have the experience to quickly implement Parallel Sysplex and set up your data sharing environment.

What about managing the Parallel Sysplex cluster and applications? z/OS Management Facility (z/OSMF) allows you to view sysplex resources. You can view sysplexes and systems in a sysplex. You also can view physical configurations, such as coupling facilities and LPARs, and logical resources, such as couple data sets and coupling facility structures. Many of the pages in the Sysplex Management task include a graphic view and a table view. You can drag the divider that separates the views to expand or reduce each section. In addition, IBM Z Automation can help automate the environment.

Compelling advantages

Does your business demand continuous computing you can't outgrow, at a low incremental cost? If it does, your IBM representative can design a comprehensive package consisting of hardware, software, services and maintenance support you'll need to get started with Parallel Sysplex clustering technology.

What you'll need:

To take advantage of the full power of Parallel Sysplex clustering technology, you need certain hardware and software elements

- z/OS
- IBM Z server. You need a minimum of two coupled LPARs for availability. Optimally, two IBM Z servers, although value can be obtained from a single server Parallel Sysplex.
- Two Coupling Facilities for availability
- Coupling links connecting the CFs to the servers hosting the z/OS systems. This can be internal (IC) links or external links.
- Server Time Protocol (STP) timing network. Multiple processors must be attached to a common time source, which maintains time synchronization in the Parallel Sysplex cluster.

To learn more

For additional information, see our home page at: www.ibm.com/products/zos/parallel-sysplex or call IBM DIRECT at 1 800 IBM-CALL in the U.S. and Canada