



IBM Parallel Sysplex clustering technology

Highlights

■ **Rapid response to unpredicted growth** — Quickly add computing power to match your requirements without disrupting your business.

■ **Resource Sharing** — Value today in a single server and multiple server environment.

The z/Architecture brings even more dynamic real time efficient system resource utilization for "on demand" e-business, by the introduction of 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.

■ **Advanced programming** — WebSphere™ Application Server permits the quick development of new applications using industry-standard object technology. The environment enables new applications that are accessing databases to receive the benefits of Parallel Sysplex scale and availability transparently.

■ **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.

■ **Continuous application availability** — Leading-edge application availability.

■ **Investment protection** — Build on your current investments in hardware, software, applications and skills, all at a reduced cost of computing.

Respond rapidly to changing conditions

The IBM® @server zSeries and S/390® server platforms are 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 zSeries and S/390 platform. Parallel Sysplex technology is a highly advanced commercial processing clustered system. It supports resource sharing and high-performance, multisystem, read/write data sharing, enabling the aggregate capacity of multiple z/OS and OS/390® systems to be applied against common workloads. This, in turn, facilitates dynamic workload balancing. In business terms, these capabilities can mean the difference between a system that supports rapid growth and one that doesn't.

Data sharing and workload balancing enable work to be directed to available processors "on the fly" and also permit servers to be dynamically added to the cluster without requiring costly downtime. This can be accomplished without splitting applications or databases across multiple servers,

a time-consuming — and expensive — process, often requiring significant new investments in staff or systems management.

Essentially, 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 — up to 32 zSeries or S/390 servers can be linked with near linear scalability.

“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)

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 datasets, 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 is the ability to share these common system resources and therefore manage them as a single entity. Redundant hardware (i.e., Tapes, DASD, consoles, catalogs, logs.) can be eliminated, providing reduced cost, simplified systems management along with increased performance, scalability and availability with GRS Star, channel constraint relief/reduced cost by elimination of CTC's with XCF Star, and increased performance with the Enhanced Catalog Sharing. With all that said, it is just a drop-in! In terms of implementation, over 1,400 customers are already using one or more resource sharing techniques in both multi- and single footprint environments to provide simplified systems management, reduced cost, increased performance and scalability. (See GF22-5115 “*Value of Resource Sharing*” at ibm.com/s390/pso)

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 z900 family of servers also automatically directs processor and I/O resources to priority work through Intelligent Resource Director (IRD). The z900 IRD combines the strengths of three key technologies: Workload Manager, Logical Partitioning and Parallel Sysplex clustering. It is the logical next step in resource sharing.

This powerful combination of z900 servers and z/OS provides the ability to intelligently manage numerous operating system images executing on a single server, as a single compute resource, with dynamic workload management and physical resource balancing (processors and I/O) across

logical partitions. In other words the system can dynamically allocate processors, channel paths and channel controller work across multiple virtual servers in order to ensure that the unpredictable demands of the e-business environment can be intelligently managed according to business goals.

Keeping work in balance

Equally important in today's economy, 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 SMP server, so too can work be directed to any node in a Parallel Sysplex cluster having available capacity. This avoids the need to partition data or applications among individual nodes in the cluster or to replicate databases across multiple servers.

“The Parallel Sysplex system gives us scalability and cost savings while adding to our ability to maintain service levels. We merged 4 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 or OS/390 Workload Manager (WLM), along with the subsystems, automatically balances tasks across all the resources of the Parallel Sysplex cluster to meet your business goals. Whether the work is coming from batch, SNA, TCP/IP, DRDA®, or MQSeries® (non-persistent) messages, dynamic session 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 e-business, application development or running traditional "back-office" online transaction processing and batch jobs.

"The implementation of VSAM Record Level Sharing and exploitation of OS/390's release migration/coexistence capabilities has enabled Bank of Montreal to reduce out-of-service time for one of its key applications by over 50%. At the same time, CPU capacity concerns have been eliminated. This application is involved in round-the-clock provision of direct customer access to chequing and savings accounts and bill payment services, including telephone banking, PC banking, point-of-sale debit cards, and ATMs. As such, its performance and availability contribute significantly to the Bank's image as a reliable service provider, a key factor in growing and retaining personal and commercial business."

Malcolm Sanderson,
Infrastructure Planning Department,
Bank of Montreal (04/00)

Programming advantages

Today's IT world is e-business oriented. The Web is an unavoidable issue, albeit a positive one, and the question is when — not if — businesses will join it. The Web changes many aspects, including the nature of applications and the role and scope of application development organizations. Most existing applications were conceived before the Web, as were the technologies and infrastructure they rely upon. Thus an up-to-date AD strategy should consider:

- *Entirely new e-business applications*
- *Evaluating how existing applications could become accessible through Web technologies*

This evaluation enables the grouping of existing applications as Web-accessible through front-end changes (for example, by translating between the 3270 and HTML protocols), reusable (applications that have presentation logic separate from business logic and, therefore, need only a new presentation or client layer), and not Web-accessible (needing to be rewritten). IBM has evolved all of its strategic technology offerings to enable customers to follow any of these options. CICS® and IMS™ support Web-enabling of existing applications. For new applications, Object Oriented technologies — in particular component - based development — are mandatory.

IBM's WebSphere Application Server (WAS) plays a fundamental role in the e-business environment. It is a major piece of the supporting infrastructure, a central hub that seamlessly blends with the network infrastructure and foundation services to enable the rapid development and deployment of applications on a distributed network environment. WebSphere's capabilities

span both the "classic" Web serving side (and associated application serving) and the new, component-based model. Thus, WebSphere is able to exploit the Enterprise JavaBean™ services available with CICS and Component Broker containers. Users can get the best of both worlds — dynamic e-business enabled applications with reliable, secure, transactional infrastructure provided by the CICS Transaction Server and Component Broker for OS/390.

Last, but certainly not least, WebSphere Application Server actively exploits Workload Manager and is fully Parallel Sysplex enabled. This allows for a highly scalable and reliable environment, unsurpassed in the industry.

Net benefits

Are you ready for e-business? If you're not, you could be at a serious competitive disadvantage.

Many customers have moved to deploy TCP/IP networks as the backbone of their communications infrastructure for the Internet, and so too has IBM moved rapidly to support robust native TCP/IP attachment. Ongoing support for TCP/IP includes such functions as a rewrite of the TCP/IP stack on OS/390 for significantly improved performance, and with enhancements in OS/390 Rel 10 to enable TCP/IP session balancing and workload balancing without additional hardware and increased availability across the Parallel Sysplex cluster for planned or unplanned outages, with seamless fail over and recovery support.

"Our 401K Web-based application was implemented prior to the supporting mainframe application being converted to PS. Implementation of PS allowed us to eliminate a daily outage, reducing it to weekly and eventually monthly. After conversion we measured a significant improvement in availability, which we expected. What we did not expect was the letters of thanks which came pouring in from our customers expressing their appreciation for the improved availability of Web application."

Large Financial Corporation (04/00)

What further differentiates z/OS and OS/390 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 zSeries and/or S/390 hardware and software.

To do real e-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 and OS/390, 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 zSeries and S/390 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.

Parallel Sysplex technology is the optimal environment for OLTP because of its high availability features and its ability to avoid processing bottlenecks by handling a large volume of transactions that can run anywhere across the cluster. The OLTP environment just keep getting better. With enhancements to support on-line reorgs and cache duplexing in DB2, elimination of CICS ENQ and data table affinities in CICS TS and Shared Message Queue support in IMS, data sharing and availability continue to be a priority across all of IBM. Additionally, new function continues to be made available between releases. These are highlighted on the DB2 and IMS Web sites (ibm.com/software/data/db2 and ibm.com/software/data/ims).

"Parallel Sysplex has enabled us to dramatically improve our cost/efficiency ratio. External audits have confirmed a 51 percent improvement, which helps us deliver better levels of customer service and still recognize significant cost savings of \$27.6 million over 5 years (from 1993 - 1997)."

Grant Kennedy,
Deputy General Manager
of Operations,

Standard Bank
of South Africa (05/99)

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.

Outsmarting the competition

According to conservative estimates, the typical business doubles its information volume every five years. That information — details on customers, on suppliers, on marketplace strategies that worked and those that didn't — represents an incredible resource that simply can't be duplicated by your competitors.

But with information, it's not only what you have, it's what you do with it. Business Intelligence systems can help identify new opportunities as well as create strategies and products that anticipate customer demands. IBM's Intelligent Miner™ product, an example of a Business Intelligence application, is the integrated solution for large-scale, sophisticated analysis of data. Turn IBM Intelligent Miner loose on your database, and its sophisticated algorithms will find patterns and relationships hidden in your data to uncover new revenue opportunities.

How powerful is the combination of DB2 and Parallel Sysplex clustering technology? In a demonstration on two S/390 G3 servers, six billion rows of data totalling 750 gigabytes (the largest data table at that time) was scanned in a mere 45 minutes using DB2 Version 5 data sharing on a Parallel Sysplex cluster (IBM Teraplex Benchmarks, 1997 ibm.com/solutions/

businessintelligence/teraplex/s390.htm). To put this in perspective, this query without any parallelism would have taken over 80 hours to complete. On today's technology, z900 servers and I/O improvements in the Enterprise Storage Server™, one can expect a 5 times improvement, finishing the query in about 9 minutes, or allow five times more data to be processed with the same response.

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 and the Workload Manager.

WLM managed initiators helps guarantee that the high priority jobs get an initiator to run on with the workload getting balanced automatically within the sysplex. (see HotTopics on ibm.com/s390/pso)

There when you need it

If these applications are critical to your business, Parallel Sysplex clustering provides another significant benefit — the closest thing to continuous computing possible today. Data sharing and workload balancing enable the best design point for application availability — in the event of a planned or unplanned outage, hardware or software — work can be redirected with little effect on end users. By comparison, a Gallup poll indicates that the

typical corporate PC server is down 1.6 hours per week.

“In the aftermath of a natural disaster, we've seen our claims workload increase by as much as 50 percent. Parallel Sysplex has enabled us to increase our claims application availability by over 32 percent and dynamically leverage the surplus capacity within the data center. This means we can plan for less total capacity to support a disaster and still consistently ensure that every client continues to receive the superior, personalized service they expect from Allstate.”

Allstate Insurance
Company (11/98)

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. You can introduce changes, such as software upgrades, one system at a time — remaining systems continue to process work. This allows you to roll changes through your systems at a pace that makes sense for your business.

Boeing's capability to maintain only one production line for the 777 is an excellent example of the business advantage of Parallel Sysplex. Boeing was able to reduce their planned outages from 104 hours per year to five hours per year when they implemented their Parallel Sysplex data sharing environment.

(see application brief GK20-6624)

IBM has taken availability to a new level with the introduction of the Geographically Dispersed Parallel Sysplex™ (GDPS™), which provides business continuity even in the event that an entire data center becomes inoperative. A GDPS consists of a Parallel Sysplex cluster spread across

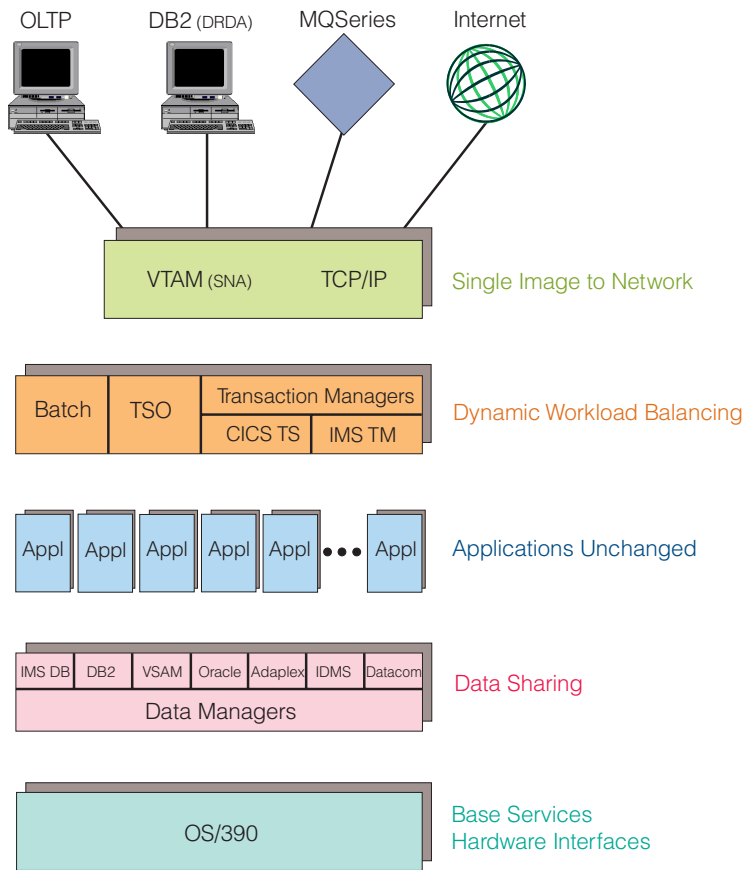
two sites separated by up to 40 kilometers based on Peer to Peer Remote Copy (PPRC) solution. For support of greater than 40 km distances between data centers GDPS also provides a service based on the open IBM Extended Remote Copy (XRC) architecture. All critical data is remote copied between the two sites. The GDPS provides the capability to manage the remote copy configuration, automate Parallel Sysplex operations, automate planned and unplanned reconfigurations and provides the fastest failure recovery available in the industry, using a single automated point of control. (See GDPS executive white paper GF22-5114 at ibm.com/s390/pso)

The price is right

Numerous industry studies have shown that Parallel Sysplex clustering technology offers a low total cost of ownership. For example, International Technology Group has demonstrated that the total five-year costs per user for Walker's e-2 (previously called Tamaris) financial suite are between 2.6 and 3.9 times lower with a Parallel Sysplex cluster than on comparable UNIX®-based systems. ("Strategies for Scalability", 12/97)

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. 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 en-



Parallel Sysplex
High level view of the software stack

abled 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.

For example, benchmarks have shown that Parallel Sysplex clusters provide the highest level of performance, availability and scalability yet achieved with a SAP R/3 application server using Parallel Sysplex DB2 data sharing as the database server.

Investment protection

Parallel Sysplex technology protects your investments. Your existing skills and training, as well as your backup and disaster recovery processes, can transfer to the Parallel Sysplex envi-

ronment. In addition, with WebSphere Application Server, you can leverage your existing data with new applications that more closely match your business model and processes.

The Parallel Sysplex computing system is helping Bell Atlantic to save money in another way; by protecting investments in their IMS applications:

“We’re getting 15 years of new life for TIRKs; it’s like getting a new application for free,” Higgins notes. “We’ll be able to move all our S/390 heritage applications into the Parallel Sysplex system, and we’ll spend zero on development to do this.”

Gerry Higgins,
 Assistant Vice President,
 Data Center, Network and
 Distribution Resources
 Bell Atlantic

Other savings come through the use of resource sharing, Intelligent Resource Director, the software pricing model, valuable professional services, and technology integration that already began with S/390 G3 servers.

Integrating key Parallel Sysplex technologies into z900 servers and S/390 enterprise servers provides much greater configuration flexibility and makes it easier to move into Parallel Sysplex clustering at reduced cost. Examples include the Internal Coupling Facility, introduced already with the S/390 G3 servers, and the Internal Coupling (IC) Channels and Integrated Cluster Bus (ICBs) introduced with the S/390 G5 servers. IC Channels are designed to provide very efficient, high-speed internal coupling links with lowest cost and highest performance. The ICB is designed to provide low-cost, high-bandwidth, and short-distance links.

To ensure a balanced system and an efficient Parallel Sysplex environment with the introduction of the z900 servers, new and faster Coupling links and ICBs are available.

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. With the Web-based Parallel Sysplex Configuration Assistant to set up your JCL, procedures and parameters, and the CF Structure Sizer to help with your tuning, setting up a Parallel Sysplex is a snap. Want to manage your environment with Resource Sharing? Just run z/OS Managed System Infrastructure for Setup (msys) to set up XCF, GRS Star, or Intelligent Resource Director, or run the SDSF Con-

figuration Assistant wizard for SDSF exploitation. Want to convert to Workload Manager Goal mode? Not a problem with the downloadable WLM Goal Mode Migration Aid to help.

If you prefer, IBM Global Services (IGS) can help you make a transition to Parallel Sysplex clustering. IGS has the experience to quickly implement Parallel Sysplex. What about managing the Parallel Sysplex cluster and applications? Not a sweat with System Automation for OS/390 V2.1 to automate the environment (ibm.com/s390/sa) and OS/390 R10 "AutoAlter" support to dynamically tune many of the CF structures size and object ratios.

Compelling advantages

How compelling are the advantages of Parallel Sysplex clustering technology? So compelling that competitors have adopted it as part of their strategic vision. Consultants are referring to it as the future of large-scale computing. But the clearest evidence of the benefits of Parallel Sysplex clustering is how rapidly customers have been adopting it. More than 1,400 of the world's leading companies have implemented a Parallel Sysplex Resource Sharing environment; more than 600 have operational production systems taking advantage of full data sharing.¹

"IBM Parallel Sysplex technology continues to be the best route for achieving Hewitt's three major goals: improved price/performance, horizontal/scalable growth and continuous availability."

Dan Kaberon, Manager,
Hewitt Associates (11/98)

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

Processors

You need a minimum installation of two coupled LPARs, which can run on any combination of the following:

- *IBM @server zSeries model 900 processor*
- *IBM S/390 Parallel Enterprise Server 9672 G3 or higher*

Coupling Facility

Each Parallel Sysplex configuration must have one or more of these Coupling Facilities and links:

- *IBM z900 - model 100 Coupling Facility*
- *IBM 9674 Coupling Facility Models C04, C05 or 9672 R06*
- *IBM z900 Server or S/390 G3, G4, G5 or G6 Server with Internal Coupling Facility (ICF)*
- *IBM z900 Server or S/390 Parallel Enterprise Server with a logical partition functioning as a Coupling Facility*
- *Coupling links or Integrated Coupling (IC) Channels and/or Integrated Cluster Bus (ICB)*

Sysplex Timer®

Multiple processors must be attached to a Sysplex Timer, which sets the time-of-day clocks and maintains time synchronization in a Parallel Sysplex cluster.

Software

Software needed for a Parallel Sysplex environment:

- *z/OS or OS/390*

Intelligent Resource Director

Requirements for Intelligent Resource Director are:

- *IBM z900 processor running z/OS*
- *z/OS Workload Manager in goal mode (for full benefits)*
- *Coupling Facility WLM structure*

To learn more

For additional information or a copy of this, see our home page at: ibm.com/eserver/zseries or call IBM DIRECT at 1 800 IBM-CALL in the U.S. and Canada.

¹ Progress of Parallel, ibm.com/s390/products/psa



© Copyright IBM Corporation 2000

IBM Corporation
Integrated Marketing Communications,
Server Group
Route 100
Somers, NY 10589

Produced in the United States of America
10-00
All Rights Reserved

References in this publication to IBM products or services do not imply that IBM intends to make them available in every country in which IBM operates. Consult your local IBM business contact for information on the products, features, and services available in your area.

IBM, the IBM logo, the e-business logo, CICS, DB2, DRDA, Enterprise Storage Server, GDPS, Geographically Dispersed Parallel Sysplex, IMS, Intelligent Miner, MQSeries, OS/390, Parallel Sysplex, S/390, S/390 Parallel Enterprise Server, ServerPac, Sysplex Timer, VTAM, WebSphere, z/Architecture, z/OS and zSeries are trademarks or registered trademarks of the International Business Machines Corporation.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States or other countries.

Tivoli is a registered trademark of Tivoli Systems Inc. in the United States, other countries or both.

UNIX is a registered trademark of The Open Group.

Windows NT is a registered trademarks of Microsoft Corporation.

LINUX is a registered trademark of Linus Torvalds.

Other trademarks and registered trademarks are the properties of their respective companies.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

Photographs shown are engineering prototypes. Changes may be incorporated in production models.

This equipment is subject to all applicable FCC rules and will comply with them upon delivery.

Information concerning non-IBM products was obtained from the suppliers of those products. Questions concerning those products should be directed to those suppliers.

All customer examples described are presented as illustrations of how these customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.