

IBM @server zSeries 990

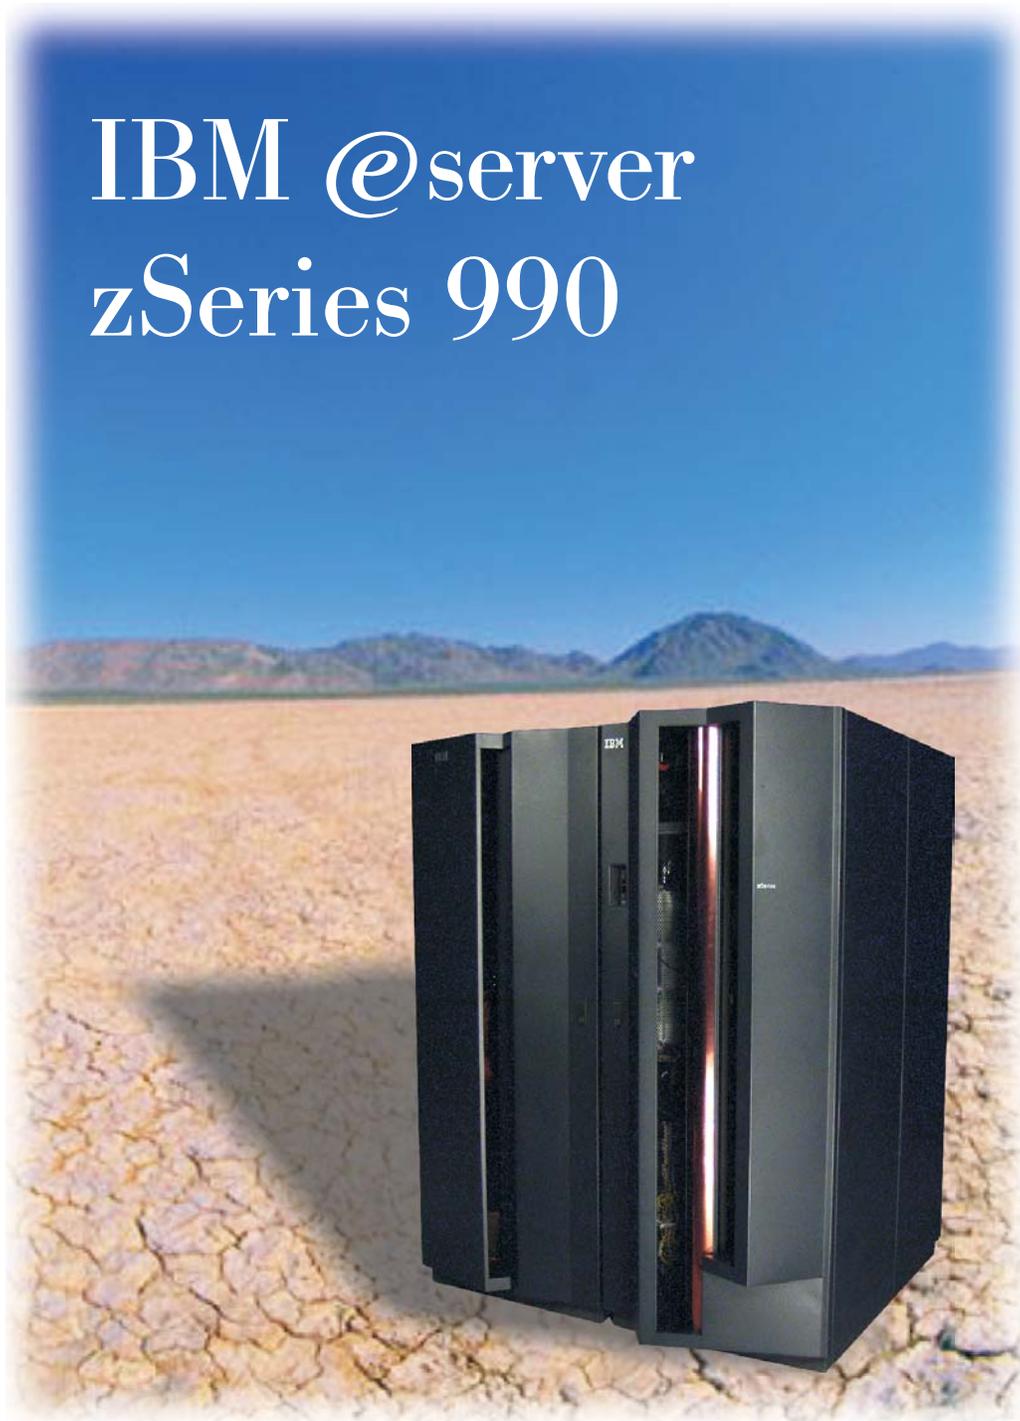


Table of Contents

<i>z/Architecture</i>	<i>page 5</i>
<i>IBM @server zSeries 990</i>	<i>page 6</i>
<i>z990 Family Models</i>	<i>page 9</i>
<i>z990 and z900 Performance Comparison</i>	<i>page 11</i>
<i>z990 I/O SubSystem</i>	<i>page 12</i>
<i>z990 Channels and I/O Connectivity</i>	<i>page 14</i>
<i>Fibre Channel Connectivity</i>	<i>page 16</i>
<i>Open Systems Adapter-Express Features (OSA-Express)</i>	<i>page 19</i>
<i>HiperSockets</i>	<i>page 22</i>
<i>Cryptography</i>	<i>page 23</i>
<i>Availability</i>	<i>page 24</i>
<i>Advanced Availability Functions</i>	<i>page 26</i>
<i>Parallel Sysplex Cluster Technology</i>	<i>page 27</i>
<i>z990 Support for Linux</i>	<i>page 36</i>
<i>zSeries 990 Family Configuration Detail</i>	<i>page 40</i>
<i>Physical Characteristics</i>	<i>page 43</i>
<i>Coupling Facility - CF Level of Support</i>	<i>page 44</i>
<i>z/OS</i>	<i>page 46</i>
<i>z/VM</i>	<i>page 69</i>
<i>To Learn More</i>	<i>page 74</i>

zSeries Overview

Technology has always accelerated the pace of change. New technologies enable new ways of doing business, shifting markets, changing customer expectations, and redefining business models. Each major enhancement to technology presents opportunities. Companies that understand and prepare for changes can gain advantage over competitors and lead their industries.

Customers of every size, and in every industry are looking for ways to make their businesses more resilient in the face of change and uncertainty. They want the ability to react to rapidly changing market conditions, manage risk, outpace their competitors with new capabilities and deliver clear returns on investments.

Welcome to the on demand era, the next phase of e-business, in which companies move beyond simply integrating their processes to actually being able to sense and respond to fluctuating market conditions and provide products and services to customers on demand. While the former notion of on demand as a utility capability is a key component, on demand companies have much broader capabilities.

What does an on demand company look like?

- *Responsive: It can sense and respond in real time to the changing needs of customers, employees, suppliers and partners*
- *Variable: It employs variable cost structures to do business at high levels of productivity, cost control, capital efficiency and financial predictability.*
- *Focused: It concentrates on its core competencies – areas where it has a differentiating advantage – and draws on the skills of strategic partners to manage needs outside of these competencies.*

- *Resilient: It can handle the ups and downs of the global market, and manage changes and threats with consistent availability, security and privacy – around the world, around the clock.*

To support an on demand business, the IT infrastructure must evolve to support it. At its heart the data center must transition to reflect these needs; the data center must be responsive to changing demands, it must be variable to support the diverse environment, it must be flexible so that applications can run on the optimal resources at any point in time, and it must be resilient to support an always open for business environment.

The on demand era plays to the strengths of the IBM **@server™** zSeries™. The IBM **@server** zSeries 900 (z900) was launched in 2000 and was the first IBM server "designed from the ground up for e-business." The latest member of the family, the IBM **@server** zSeries 990 (z990), brings enriched functions that are required for the on demand data center.

The "responsive" data center needs to have systems that are managed to the quality of service goals of the business; they need systems that can be upgraded transparently to the user and they must be adaptable to the changing requirements of the business. With the zSeries you have a server with high levels of reliability and a balanced design to ensure high levels of utilization and consistently high service to the user. The capacity on demand features continue to evolve ensuring that upgrading the servers is timely and meets the needs of your business, and it's not just the capacity of the servers that can be changed on demand, but also the mix of workload and the allocation of resources to reflect the evolving needs and priorities of the business.

The variable data center needs to be able to respond to the ever changing demands that occur when you support multiple diverse workloads as a single entity. It must respond to maintain the quality of service required and the cost of utilizing the resources must reflect the changing environment. The zSeries Intelligent Resource Director (IRD), which combines three key zSeries technologies, z/OS™ Workload Manager (WLM), Logical Partitioning and Parallel Sysplex® helps ensure that your most important workloads get the resources they need and constantly manages the resources according to the changing priorities of the business. With Workload License Charging (WLC), as the resources required by different applications, middle ware and operating systems change over time, the software costs change to reflect this. In addition, new virtual Linux® servers can be added in just minutes with zSeries virtualization technology to respond rapidly to huge increases in user activity.

The flexible data center must be adaptable to support change and ease integration. This is achieved through a combination of Open and industry standards along with the adaptability to direct resources where they are required. The zSeries along with other IBM servers have been investing in standards for years. Key is the support for Linux, but let's not forget Java™ and XML and industry standard technologies, such as FCP, Ethernet and SCSI.

Finally the on demand data center must be resilient. The zSeries has always been renowned for reliability and availability. The zSeries platform will protect from both scheduled and unscheduled outages, and GDPS™ enables protection from loss of complete sites.

Tools for Managing e-business

The IBM @server product line is backed by a comprehensive suite of offerings and resources that provide value at every stage of IT implementation. These tools can help customers test possible solutions, obtain financing, plan and implement applications and middleware, manage capacity and availability, improve performance and obtain technical support across the entire infrastructure. The result is an easier way to handle the complexities and rapid growth of e-business. In addition, IBM Global Services experts can help with business and IT consulting, business transformation and total systems management services, as well as customized e-business solutions.

z/Architecture

The zSeries is based on the z/Architecture™, which is designed to reduce bottlenecks associated with the lack of addressable memory and automatically directs resources to priority work through Intelligent Resource Director. The z/Architecture is a 64-bit superset of ESA/390.

z/Architecture is implemented on the z990 to allow full 64-bit real and virtual storage support. A maximum 256 GB of real storage is available on z990 servers. z990 can define any LPAR as having 31-bit or 64-bit addressability.

z/Architecture has:

- *64-bit general registers.*
- *New 64-bit integer instructions. Most ESA/390 architecture instructions with 32-bit operands have new 64-bit and 32- to 64-bit analogs.*
- *64-bit addressing is supported for both operands and instructions for both real addressing and virtual addressing.*
- *64-bit address generation. z/Architecture provides 64-bit virtual addressing in an address space, and 64-bit real addressing.*
- *64-bit control registers. z/Architecture control registers can specify regions, segments, or can force virtual addresses to be treated as real addresses.*
- *The prefix area is expanded from 4K to 8K bytes.*
- *New instructions provide quad-word storage consistency.*
- *The 64-bit I/O architecture allows CCW indirect data addressing to designate data addresses above 2 GB for both format-0 and format-1 CCWs.*
- *IEEE Floating Point architecture adds twelve new instructions for 64-bit integer conversion.*
- *The 64-bit SIE architecture allows a z/Architecture server to support both ESA/390 (31-bit) and z/Architecture (64-bit) guests. Zone Relocation is expanded to 64-bit for LPAR and z/VM™.*

- *64-bit operands and general registers are used for all Cryptographic instructions*
- *The implementation of 64-bit z/Architecture can reduce problems associated with lack of addressable memory by making the addressing capability virtually unlimited (16 Exabytes).*

z/Architecture Operating System Support

The z/Architecture is a tri-modal architecture capable of executing in 24-bit, 31-bit, or 64-bit addressing modes. Operating systems and middleware products have been modified to exploit the new capabilities of the z/Architecture. Immediate benefit can be realized by the elimination of the overhead of Central Storage to Expanded Storage page movement and the relief provided for those constrained by the 2 GB real storage limit of ESA/390. Application programs can run unmodified on the zSeries family of servers.

Expanded Storage (ES) is still supported for operating systems running in ESA/390 mode (31-bit). For z/Architecture mode (64-bit), ES is supported by z/VM. ES is not supported by z/OS in z/Architecture mode.

Although z/OS does not support Expanded Storage when running under the new architecture, all of the Hiperspace™ and VIO APIs, as well as the Move Page (MVPG) instruction, continue to operate in a compatible manner. There is no need to change products that use Hiperspaces.

Some of the exploiters of z/Architecture for z/OS include:

- *DB2 Universal Database™ Server for z/OS*
- *IMS™*
- *Virtual Storage Access Method (VSAM)*
- *Remote Dual Copy (XRC)*
- *Tape and DASD access method*

IBM @server zSeries 990

Operating System	ESA/390 (31-bit)	z/Arch (64-bit)	Compatibility	Exploitation
OS/390® Version 2 Release 10	Yes	Yes	Yes	No
z/OS Version 1 Release 1	No	No	No	No
z/OS Version 1 Release 2	No*	Yes	Yes	No
z/OS Version 1 Release 3	No*	Yes	Yes	No
z/OS Version 1 Release 4	No*	Yes	Yes	Yes
z/OS Version 1 Release 5	No	Yes	Included	Included
Linux for S/390®	Yes	No	Yes	Yes
Linux for zSeries	No	Yes	Yes	Yes
z/VM Version 3 Release 1	Yes	Yes	Yes	No
z/VM Version 4 Release 2, 3	Yes	Yes	Yes	No
z/VM Version 4 Release 4	Yes	Yes	Included	Included
VSE/ESA™ Ver. 2 Release 5, 6, 7	Yes	No	Yes	No
TPF Version 4 Release 1 (ESA mode only)	Yes	No	Yes	No

* ESA/390 mode permitted for migration purposes only

Note: All operating systems must run in LPAR mode. Basic mode is not supported on the z990.

IBM @server zSeries is the enterprise class e-business server optimized for integration, transactions and data of the next generation e-business world. In implementing the z/Architecture with new technology solutions, the zSeries models are designed to facilitate the IT business transformation and reduce the stress of business-to-business and business-to-customer growth pressure. The zSeries represents an advanced generation of servers that feature enhanced performance, support for zSeries Parallel Sysplex clustering, improved hardware management controls and innovative functions to address e-business processing.

The z990 processor, the latest addition to the zSeries family, enhances performance by exploiting new technology through many design enhancements. With a new super-scalar microprocessor and the CMOS9S-SOI technology the z990 is designed to further extend and integrate key platform characteristics such as dynamic flexible partitioning and resource management in mixed and unpredictable workload environments, providing scalability, high availability and Quality of Service to emerging e-business applications such as WebSphere®, Java and Linux.

The z990 has 4 models available as new build systems and as upgrades from the z900.

The four new z990 models are designed with a multi-book system structure which provides up to 32 Processor Units (PUs) that can be characterized prior to the shipment of the machine as either Central Processors (CPs), Integrated Facility for Linux (IFLs), or Internal Coupling Facilities (ICFs).

With the introduction of the z990 customers can expect to see the following performance improvements:

Number of CPs	Base	Ratio
1	z900 2C1	1.52 - 1.58
8	z900 2C8	1.48 - 1.55
16	z900 2C16	1.45 - 1.53
32	z900 2C16	2.4 - 2.9

Note: Greater than 16 CPs requires a minimum of two operating system images

The Large System Performance Reference (LSPR) should be referenced when considering performance on the z990. Visit: ibm.com/servers/eserver/zseries/lSpr/ for more information on LSPR.

To support the new scalability of the z990 a new improvement to the I/O Subsystem has been introduced to “break the barrier” of 256 channels per Central Electronic Complex (CEC). This provides “horizontal” growth by allowing the definition of two Logical Channel SubSystems each capable of supporting up to 256 channels giving a total of up to 512 channels per CEC with a Statement of Direction (SOD) to support up to four Logical Channel SubSystem and 1024 channels. The increased scalability is further supported by the increase in the number of Logical Partitions available from the current 15 LPARs to a new 30 LPARs (SOD up to 60 LPARs). There is still a 256 channel limit per operating system image.

These are some of the significant enhancements in the zSeries 990 processor that bring improved performance, availability and function to the platform. The following sections highlight the functions and features of the server.

z990 Design and Technology

The z990 is designed to provide balanced system performance. From processor storage to the system’s I/O and network channels, end-to-end bandwidth is provided to deliver data where and when it is needed.

The z990 provides a significant increase in system scalability and opportunity for server consolidation by providing four models, from one to four MultiChip Modules (MCMs), delivering up to a maximum 32-way configuration. The MCMs are configured in a book package, with each book comprised of a MultiChip Module (MCM), memory cards and Self-Timed Interconnects. The MCM, which measures approximately 93 x 93 millimeters (43% smaller than the z900), contains the processor unit (PU) chips, the cache structure chips and the processor storage controller chips. The MCM contains 101 glass ceramic layers to provide interconnection between the chips and the off-module environment. In total, there is approximately 0.4 kilometer of internal copper wiring on this module. This new MCM packaging delivers an MCM 42% smaller than the z900, with 23% more I/O connections and 133% I/O density improvement. Each MCM provides support for 12 PUs and 32 MB level 2 cache. Each PU contains 122 million transistors and measures 14.1 mm x 18.9 mm. The design of the MCM technology on the z990 provides the flexibility to configure the PUs for different uses, two of the PUs are reserved for use as System Assist Processors (SAPs), two are reserved as spares. The remaining inactive 8 PUs on the MCM are available to be characterized as either CPs, ICF processors for Coupling Facility applications, IFLs for Linux applications or as optional SAPs, providing the customer with tremendous flexibility in establishing the best system for running applications. Each model of the z990 must always be ordered with at least one CP, IFL or ICF. The z990 MCM is the world’s densest logic package with over three billion transistors.

The PU, which uses the latest chip technology from IBM semiconductor laboratories, is built on CMOS 9S-SOI with copper interconnections. The 14.1 mm x 18.9 mm chip has a cycle time of 0.83 nanoseconds. Implemented on this chip is the z/Architecture with its 64-bit capabilities including instructions, 64-bit General Purpose Registers and translation facilities.

Each book can support up to 64 GB of Memory, delivered on two memory cards, and 12 (STIs) giving a total of 256 GB of memory and 48 STIs on the D32 model. The memory is delivered on 8 GB, 16 GB or 32 GB memory cards which can be purchased in 8GB increments. The two memory cards associated with each book must be the same size. Each book has 3 MBAs and each MBA supports 4 STIs.

All books are interconnected with a super fast bi-directional redundant ring structure which allows the system to be operated and controlled by PR/SM™ operating in LPAR mode as a symmetrical, memory coherent, multiprocessor. PR/SM provides the ability to configure and operate as many as 30 logical partitions which may be assigned processors, memory and I/O resources from any of the available books. The design of the ring structure also provides the ability to upgrade from model to model concurrently (stage delivery). The z990 supports LPAR mode only (basic mode is no longer supported).

The MultiChip Module is the technology cornerstone for flexible PU deployment in the z990 models. For most models, the ability of the MCM to have inactive PUs allows such features as Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrades (CIU), and the ability to add CPs, ICFs and IFLs dynamically providing nondisruptive

upgrade of processing capability. Also, the ability to add CPs lets a z990 with spare PU capacity become a backup for other systems in the enterprise; expanding the z990 system to meet an emergency outage situation. This is called Capacity BackUp (CBU). The greater capacity of the z990 offers customers even more flexibility for using this feature to backup critical systems in their enterprise.

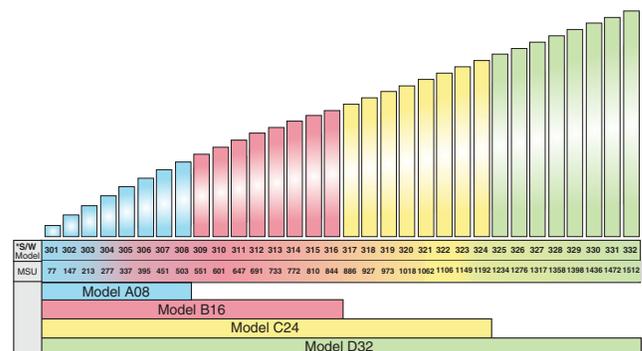
In order to support the highly scalable multi-book system design the I/O SubSystem has been enhanced by introducing a new multiple Logical Channel SubSystem (LCSSs) which provides the capability to install up to 512 (SOD 1024) channels across three I/O cages (256 channels per operating system image). I/O improvements in the Parallel Sysplex Coupling Link architecture and technology support faster and more efficient transmission between the Coupling Facility and production systems. HiperSockets™ provide high-speed capability to communicate between Logical Partitions; this is based on high-speed TCP/IP memory speed transfers and provides value in allowing applications running in one partition to communicate with applications running in another. The z990 Gigabit Ethernet feature is one of the first adapters in the industry capable of achieving line speed: one Gigabit per second. Industry standard and openness are design objectives for I/O in z900. The improved I/O subsystem is delivering new horizons in I/O capability and has eliminated the 256 limit to I/O attachments for a mainframe.

z990 Family Models

The z990 offers 4 new models, the A08, B16, C24 and D32, which can be configured to give customers a highly scalable solution to meet the needs of both high transaction processing applications and the demands of e-business. The new model structure provides between 1-32 configurable Processor Units (PUs) which can be characterized as either CPs, IFLs, or ICFs. A new easy-to-enable ability to “turn off” CPs is available on z990 (a similar offering was available via RPQ on z900). The objective is to allow customers to purchase capacity for future use with no impact on software billing. An MES feature will enable the CPs for use where the customer requires the increased capacity. There are a wide range of upgrade options available which are indicated in the z990 Models chart.

Unlike other zSeries processor offerings, it is no longer possible to tell by the hardware model (A08, B16, C24, D32) the number of PUs that are being used as CPs. For software billing purposes only, there will be a “software” model associated with the number of PUs that are characterized as CPs. This number will be reported by the Store System Information (STSI) instruction for software billing purposes only. There is no affinity between the hardware model and the number of CPs. For example, it is possible to have a model B16 which has 5 PUs characterized as CPs, so for software billing purposes, the STSI instruction would report 305. The more normal configuration for a 5-way would be an A08 with 5 PUs characterized as CPs. The STSI instruction would also report 305 for that configuration.

z990 Models



* S/W Model refers to number of installed CPs. Reported by STSI instruction.
Model 300 does not have any CPs.

z990 and IBM @server On/Off Capacity on Demand

IBM @server On/Off Capacity on Demand is offered with z990 processors to provide a temporary increase in capacity to meet customers peak workload requirements. The scope of On/Off Capacity on Demand is to allow customers to temporarily turn on unassigned/unowned PUs available within the current model for use as CPs. Temporary use of IFLs, CFs, memory and channels is not supported.

Before customers can order temporary capacity, they must have a signed agreement for Customer Initiated Upgrade (CIU) facility. In addition to that agreement, they must agree to specific terms and conditions which govern the use of temporary capacity.

Typically, On/Off Capacity on Demand will be ordered through CIU, however there will be an RPQ available if no RSF connection is present.

The use of On/Off Capacity on Demand is mutually exclusive with Capacity BackUp (CBU) and no physical hardware upgrade will be supported while On/Off Capacity on Demand is active.

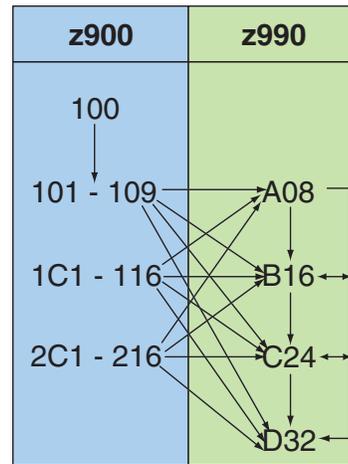
This important new function for zSeries gives customers greater control and ability to add capacity to meet the requirements of an unpredictable on demand application environment.

Customers can also take advantage of the Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU), and Capacity BackUp (CBU) which are described later in the document (page 24).

The z990 has also been designed to offer a high performance and efficient I/O structure. All z990 models ship with two frames the A-Frame and the Z-Frame, which supports the installation of up to three I/O cages. Each I/O cage has the capability of plugging up to 28 I/O cards. When used in conjunction with the software that supports multiple Logical Channel SubSystems, it is possible to have up to 420 ESCON® channels in a single I/O cage and a maximum of 512 ESCON (SOD 1024) channels across 3 I/O cages. Alternatively, three I/O cages will support up to 120 FICON™ channels and up to 360 ESCON channels. Each book will support up to 12 STIs for I/O connectivity. Seven STIs are required to support the 28 channel slots in each I/O cage so in order to support a fully configured three I/O cage system 21 STIs are required. To achieve this maximum I/O connectivity requires at least a B16 model which provides 24 STIs.

The following chart shows the upgrade from z900 to z990. There are any to any upgrades from any of the z900 general purpose models. A z900 Coupling Facility model 100 must first be upgraded to a z900 general purpose model before upgrading to a z990. There are no direct upgrades from 9672 G5/G6 or IBM @server zSeries 800 (z800).

Model Upgrades



z990 and z900 Performance Comparison

The performance design of the z/Architecture enables the entire server to support a new standard of performance for all applications through expanding upon a balanced system approach. As CMOS technology has been enhanced to support not only additional processing power, but also more engines, the entire server is modified to support the increase in processing power. The I/O subsystem supports a greater amount of bandwidth through internal changes, thus providing for larger and quicker data movement into and out of the server. Support of larger amounts of data within the server required improved management of storage configurations made available through integration of the software operating system and hardware support of 64-bit addressing. The combined balanced system effect allows for increases in performance across a broad spectrum of work. However, due to the increased flexibility in the z990 model structure and resource management in the system, it is expected that there will be larger performance variability than has been previously seen by our traditional customer set. This variability may be observed in several ways. The range of performance ratings across the individual LSPR workloads is likely to have a larger spread than past processors. There will also be more performance variation of individual LPAR partitions as the impact of fluctuating resource requirements of other partitions can be more pronounced with the increased number of partitions and additional CPs available on the z990. The customer impact of this increased variability will be seen as increased deviations of workloads from single-number-metric based factors such as MIPS, MSUs and CPU time chargeback algorithms. It is important to realize the z990 has been optimized to run many workloads at high utilization rates.

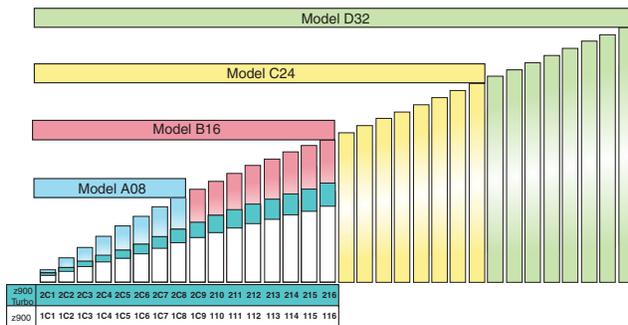
It is also important to notice that the LSPR workloads for z990 have been updated to reflect more closely our customers' current and growth workloads. The traditional TSO LSPR workload is replaced by a new, heavy Java technology-based online workload referred to as Trade2-EJB (a stock trading application). The traditional CICS®/DB2® LSPR online workload has been updated to have a Web-frontend which then connects to CICS. This updated workload is referred to as WEB/CICS/DB2 and is representative of customers who Web-enable access to their legacy applications. Continuing in the LSPR for z990 will be the legacy online workload, IMS, and two legacy batch workloads CB84 and CBW2. The z990 LSPR will provide performance ratios for individual workloads as well as a "default mixed workload" which is used to establish single-number-metrics such as MIPS, MSUs and SRM constants. The z990 default mixed workload will be composed of equal amounts of five workloads, Trade2-EJB, WEB/CICS/DB2, IMS, CB84 and CBW2. Additionally, the z990 LSPR will rate all z/Architecture processors running in LPAR mode and 64-bit mode. The existing z900 processors have all been re-measured using the new workloads — all running in LPAR mode and 64-bit mode.

Using the new LSPR "default mixed workload," and with all processors executing in 64-bit and LPAR mode, the following results have been estimated:

- *Comparing a one-way z900 Model 2C1 to a z990 model with one CP enabled, it is estimated that the z990 model has 1.52 to 1.58 times the capacity of the 2C1.*
- *Comparing an 8-way z900 Model 2C8 to a z990 model with eight CPs enabled, it is estimated that the z990 model has 1.48 to 1.55 times the capacity of the 2C8.*

z990 I/O SubSystem

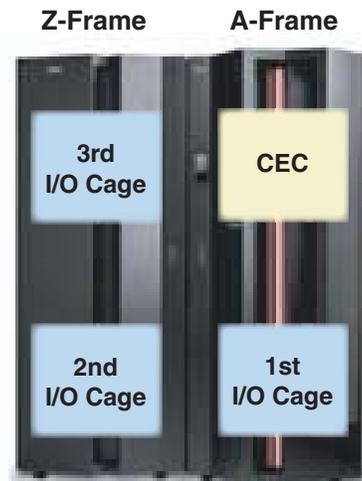
- Comparing a 16-way z900 Model 216 to a z990 model with sixteen CPs enabled, it is estimated that the z990 model has 1.45 to 1.53 times the capacity of the 216.
- Comparing a 16-way z900 Model 216 to a z990 model with thirty-two CPs enabled, and the workload executing on the z990 executes in two 16-way LPARs, it is estimated that the z990 model has 2.4 to 2.9 times the capacity of the 216.



Note: Expected performance improvements are based on hardware changes. Additional performance benefits may be obtained as the z/Architecture is fully exploited.

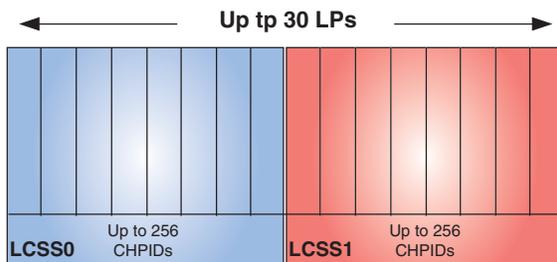
The z990 contains an I/O subsystem infrastructure which uses an I/O cage that provides 28 I/O slots and the ability to have one to three I/O cages giving a total of 84 I/O slots. Each I/O cage has seven x 4 I/O slot domains each of which has one STI for connectivity giving the requirement of seven STIs to support each I/O cage. ESCON, FICON Express™ and OSA-Express cards plug into the z990 I/O cage. All I/O cards and their support cards can be hot-plugged in the I/O cage. Installation of an I/O cage remains a disruptive MES, so the Plan Ahead feature remains an important consideration when ordering a z990 system. The A08 model has 12 available STIs with connectivity to a maximum of 12 I/O domains, i.e. 48 I/O slots, so if more than 48 I/O slots are required a Model B16 is required. Each model ships with one I/O cage as standard in the A-Frame (the A-Frame also contains the processor CEC cage) and any additional I/O cages are installed in the Z-Frame. The z990 provides a 400 percent increase in I/O bandwidth provided by the STIs.

z990 Cage Layout



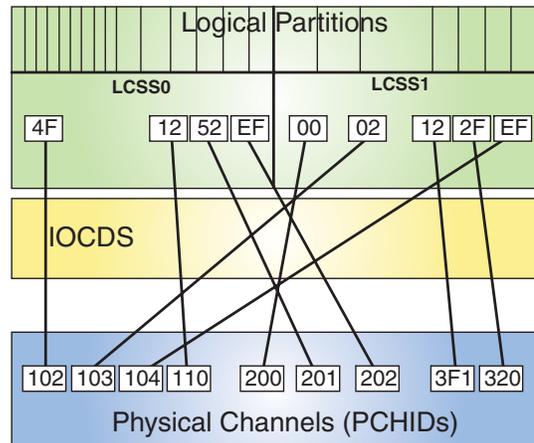
z990 multiple Logical Channel SubSystems (LCSSs) and support greater than 15 Logical Partitions (LPARs)

In order to provide the increased channel connectivity required to support the scalability of the z990, the z990 channel I/O SubSystem gives a breakthrough in connectivity, by providing up to 2 LCSS per CEC (SOD for up to 4 LCSS), each of which can support up to 256 channels with exploitation software installed. This support is provided in such a way that is transparent to the programs operating in the Logical Partition (LPAR). Each Logical Channel SubSystem may have from 1 to 256 channels and may in turn be configured with 1 to 15 logical partitions. Each Logical Partition runs under a single LCSS. As with previous zSeries servers, Multiple Image Facility (MIF) channel sharing as well as all other channel subsystem features are available to each Logical Partition configured to each Logical Channel SubSystem.



Physical Channel IDs (PCHIDs) SubSystem

In order to accommodate the new support for 512 channels (SOD 1024) introduced with the multiple Logical Channel SubSystem (LCSSs), a new Physical Channel ID (PCHID) is being introduced. The PCHID represents the physical location of an I/O feature in the I/O cage. CHPID numbers are no longer pre-assigned and it is now a customer responsibility to do this assignment via IOCP/HCD. CHPID assignment is done by associating a CHPID number with a physical location, the PCHID. It is important to note that although it is possible to have multiple LCSSs, there is still a single IOCDS to define the I/O subsystem. There is a new CHPID mapping tool available to aid in the mapping of CHPIDs to PCHIDs. The CHPID Mapping tool is available from Resource Link™, at ibm.com/servers/resourcelink.



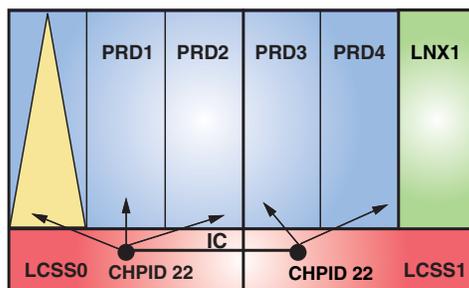
Note: Crypto no longer requires a CHPID

z990 Channels and I/O Connectivity

Logical Channel SubSystem (LCSS) Spanning

The concept of spanning channels provides the ability for a channel to be configured to multiple logical channel subsystems and therefore they may be transparently shared by any/all of the logical partitions in those LCSSs. Normal Multiple Image Facility (MIF) sharing of a channel is confined to a single LCSS. Initially, only IC and HiperSockets are supported for spanning. There is an SOD which states IBM intends to support channel spanning between appropriate channel types (e.g. ESCON architecture prevents the spanning of ESCON channels) on the z990.

IC Channel Spanning



A z990 with all three I/O cages installed has a total of 84 I/O slots. These slots can be plugged with a mixture of cards providing the I/O connectivity, networking connectivity, coupling and cryptographic capability of the machine.

Up to 512 ESCON Channels

The high density ESCON feature has 16 ports, 15 of which can be activated for customer use. One port is always reserved as a spare which is activated in the event of a failure of one of the other ports. For high availability the initial order of ESCON features will deliver two cards and the active ports will be distributed across those cards. After the initial install the ESCON features are installed in increments of one. ESCON channels are available in four-port increments and are activated using IBM Licensed Internal Code, Control Code (LIC CC).

Up to 120 FICON Express Channels

The z990 supports up to 120 FICON Express channels. FICON Express is available in long wavelength (LX) and short wavelength (SX) features. Each FICON Express feature has two independent ports which support two channels per card. The LX and SX cannot be intermixed on a single feature. The maximum number of FICON Express features is 60 which can be installed across three I/O cages with a maximum of 20 features per I/O cage.

The z990 supports up to 120 FCP channels for Linux. The same two-port FICON Express feature card used for FICON Express channels is used for FCP channels. FCP channels are enabled on these existing features via a microcode load with a new mode of operation and new CHPID definition. FCP is available in long wavelength (LX) and short wavelength (SX) features, though the LX and SX cannot be intermixed on a single feature. Note, the maximum combined native FICON (FC), FICON Bridge (FCV), FCP channels (FCP), OSA-E, PCICA, and PCIXCC is 120.

InterSystem Channel-3 (ISC-3)

A four port ISC-3 card is provided on the z990 family of servers. It consists of a mother card with two daughter cards which have two ports each. Each daughter card port is capable of operation at 1 gigabits per second (Gbps) in Compatibility Mode or 2 Gbps in peer mode up to a distance of 10 km. The mode is selected for each port via the CHPID type in the IOCDs. The ports are orderable in one-port increments.

An RPQ card (8P2197) is available to allow ISC-3 distances up to 20 km. This card runs in Peer Mode at 1 Gbps and/or Compatibility Mode at 1 Gbps. The ports are orderable in two-port increments.

Integrated Cluster Bus-2 (ICB-2)

The ICB-2 feature is a coupling link used to provide high-speed communication between a 9672 G5/G6 server and a z990 server over a short distance (less than 7 meters). The ICB-2 is supported via an STI-2 card which resides in the I/O cage and converts the 2 GigaBytes per second (GBps) input into two 333 MegaBytes per second (MBps) ICB-2s. ICB-2 is not supported between z990 and other zSeries servers.

Integrated ClusterBus-3 (ICB-3)

The ICB-3 feature is a coupling link used to provide high-speed communication between a z990 server and a z900 General Purpose Server or Model 100 Coupling Facility over a short distance (less than 7 meters). The ICB-3 is supported via an STI-3 card which resides in the I/O cage and converts the 2 GBps input into two 1 GBps ICB-3s.

Integrated ClusterBus-4 (ICB-4)

The ICB-4 feature is a coupling link used to provide high-speed communication between a z990 servers over a short distance (less than 7 meters). The ICB-4 consists of one link that attaches directly to a 2 GBps STI port on the server and does not require connectivity to an I/O cage.

Internal Coupling (IC)

IC links emulate the coupling links between images within a single server. IC links are defined in the IOCP. There is no physical channel involved. A z/OS image can connect to a coupling facility on the same server using IC capabilities.

Fibre Channel Connectivity

The on demand e-business world requires fast data access, continuous data availability, and improved flexibility all with lower cost of ownership. The new increased number of FICON Express cards available on the z990 helps distinguish this new server family, further setting it apart as enterprise class in terms of the number of simultaneous I/O connections available for these FICON Express features.

FICON Express Channel Card Features

Performance

With its 2 Gigabit per second link speed capability, the FICON Express channel card feature (feature codes 2319, 2320) is the latest zSeries implementation for the Fibre Channel Architecture. The FICON Express card has two ports and can achieve improved performance over the previous generation FICON channel card. For example, attached to a 100 MBps link (1 Gbps), a single FICON Express feature configured as a native FICON channel is capable of supporting up to 7,200 I/O operations/sec (channel is 100% utilized) and an aggregate total throughput of 120 MBps on z990.

With 2 Gbps links, customers may expect up to 170 MBps of total throughput. The 2 Gbps link speeds are applicable to native FICON and FCP channels on zSeries only and for full benefit, require 2 Gbps capable devices as well. Customers can leverage this additional bandwidth capacity to consolidate channels and reduce configuration complexity, infrastructure costs, and the number of channels that must be managed. Please note, no additional hardware or code is needed in order to obtain 2 Gbps links. The functionality was incorporated in all zSeries with March 2002 LIC. The link speed is auto-negotiated between server and devices.

Flexibility - Three channel types supported

The FICON Express features support three different channel types: 1) FCV Mode for FICON Bridge Channels, 2) FC mode for Native FICON channels (including the FICON CTC function), and 3) FCP mode for Fibre Channels (FCP

channels). Support for FCP devices means that zSeries servers will be capable of attaching to select fibre channel switches/directors and FCP/SCSI devices and may access these devices from Linux for zSeries.

Distance

All channels defined on FICON Express LX channel card features at 1 Gbps link speeds support a maximum unrepeated distance of up to 10 km (6.2 miles, or up to 20 km via RPQ, or up to 100 km with repeaters) over nine micron single mode fiber and up to 550 meters (1,804 feet) over 50 or 62.5 micron multimode fiber through Mode Conditioning Patch (MCP) cables. At 2 Gbps link speeds FICON Express LX channel card features support up to 10 km (6.2 miles, or up to 12 km via RPQ, or up to 100 km with repeaters) over nine micron single mode fiber. At 2 Gbps link speeds Mode Conditioning Patch (MCP) cables on 50 or 62.5 micron multimode fiber are not supported. The maximum unrepeated distances for 1 Gbps links defined on the SX FICON Express channel cards are up to 500 meters and 300 meters (984 feet) for 50 and 62.5 micron multimode fiber, respectively. The maximum unrepeated distances for 2 Gbps links defined on the SX FICON Express channel cards are up to 300 meters and 120 meters for 50 and 62.5 micron multimode fiber, respectively. The FICON Express channel cards are designed to reduce the data droop effect that made long distances not viable for ESCON. This distance capability is becoming increasingly important as customers are moving toward remote I/O, vaulting for disaster recovery and Geographically Dispersed Parallel Sysplex™ for availability.

Shared infrastructure

Both FCP and FICON (FC-SB-2) are adopted by INCITS (International) as standards to the Fibre Channel Architecture. Using open connectivity standards leads to shared I/O fiber cabling and switch infrastructures, facilitated data sharing, storage management and SAN implementation, and integration between the mainframe and UNIX®/Intel® technologies.

Native FICON Channels

Native FICON channels and devices can help to reduce bandwidth constraints and channel contention to enable easier server consolidation, new application growth, large business intelligence queries and exploitation of e-business.

Currently, the IBM Enterprise Storage Server™ (ESS) models F10, F20 and 800 have two host adapters to support native FICON. These host adapters each have one port per card and can either be FC 3021 for long wavelength or FC 3032 for short wavelength on the F10/F20 or FC 3024 for long wavelength and 3025 for short wavelength on the 800. All three models can support up to 16 FICON ports per ESS. The Model 800 is 2 Gb link capable. The IBM TotalStorage™ Enterprise Tape Controller® 3590 Model A60 provides up to two FICON interfaces which can coexist with ESCON on the same box. Customers can utilize IBM's highest capacity, highest performance tape drive to support their new business models.

Many Fibre Channel directors provide dynamic connectivity to native FICON control units. The IBM 2032 models 001, 064 and 140 (resell of the McDATA ED-5000, ED-6064 and ED-6140) are 32-, 64- and 140-port high availability directors. The IBM 2042 models 001, 128 and 256 (resell of the INRANGE Intrepid FC/9000) are 64-, 128- and 256-port high availability directors. All have features that provide interface support to allow the unit to be managed by System Automation for OS/390. The McDATA ED-6064, ED-6140 and INRANGE Intrepid FC/9000-001/-128/-256 support 2 Gbps links as well.

Wave Division Multiplexor and Optical Amplifiers that support 2 Gbps FICON Express links are: Cisco ONS 15540 ESP (LX, SX) and optical amplifier (LX, SX), Nortel Optera Metro 5200 and 5300E and optical amplifier, and the IBM 2029 Fiber Saver.

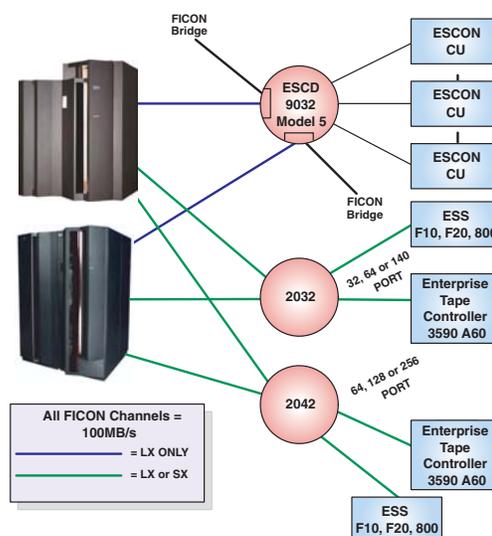
Although native FICON channels don't exploit zSeries IRD Dynamic Channel Path Management, the raw bandwidth and distance capabilities that native FICON end-to-end connectivity has to offer makes them of interest for anyone with a need for high performance, large data transfers or enhanced multi-site solutions.

FICON CTC function

Native FICON channels support CTC on the z990, z900 and z800. G5 and G6 servers can connect to a zSeries FICON CTC as well. This FICON CTC connectivity will increase bandwidth between G5, G6, z990, z900, and z800 systems.

Because the FICON CTC function is included as part of the native FICON (FC) mode of operation on zSeries, FICON CTC is not limited to intersystem connectivity (as is the case with ESCON), but will also support multiple device definitions. For example, ESCON channels that are dedicated as CTC cannot communicate with any other device, whereas native FICON (FC) channels are not dedicated to CTC only. Native can support both device and CTC mode definition concurrently, allowing for greater connectivity flexibility.

FICON Connectivity

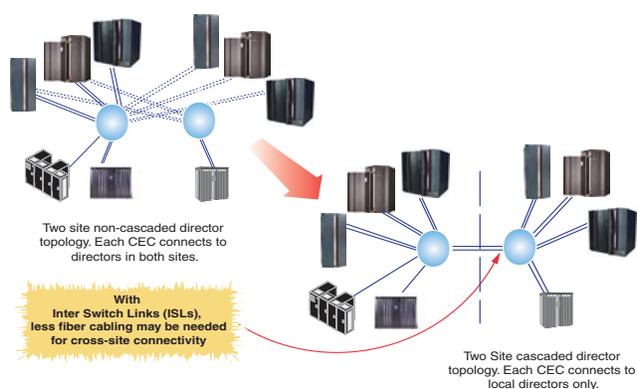


FICON Support for Cascaded Directors

Native FICON (FC) channels now support cascaded directors. This support is for a two director, single hop configuration only. This means that a Native FICON (FC) channel or a FICON CTC can connect a server to a device or other server via two (same vendor) FICON Directors in between. This type of cascaded support is important for disaster recovery and business continuity solutions because it can provide high availability, extended distance connectivity, and (particularly with the implementation of 2 Gbps Inter Switch Links), has the potential for fiber infrastructure cost savings by reducing the number of channels for interconnecting the 2 sites.

FICON cascaded directors have the added value of ensuring high integrity connectivity. New integrity features introduced within the FICON Express channel and the FICON cascaded switch fabric to ensure the detection and reporting of any miscabling actions occurring within the fabric and prevents data from being delivered to the wrong end point.

FICON cascaded directors are offered in conjunction with INRANGE and McDATA switch vendors.



FICON Bridge Channel

Introduced first on the 9672 G5 processors, the FICON Bridge (FCV) channel is still an effective way to use FICON bandwidth with existing ESCON control units. FICON Express LX channel cards in FCV (FICON Converted) Mode of operation can attach to the 9032 Model 005 ESCON Director through the use of a director bridge card. Up to 16 bridge cards are supportable on a single 9032 Model 005 with each card capable of sustaining up to eight concurrent ESCON data transfers. 9032 Model 005 ESCON Directors can be field upgradeable at no charge to support the bridge cards, and bridge cards and ESCON cards can coexist in the same director.

FCP Channels

zSeries supports FCP channels, switches and FCP/SCSI devices with full fabric connectivity under Linux for zSeries. Support for FCP devices means that zSeries servers will be capable of attaching to select FCP/SCSI devices and may access these devices from Linux for zSeries. This expanded attachability means that customers have more choices for new storage solutions, or may have the ability to use existing storage devices, thus leveraging existing investments and lowering total cost of ownership for their Linux implementation.

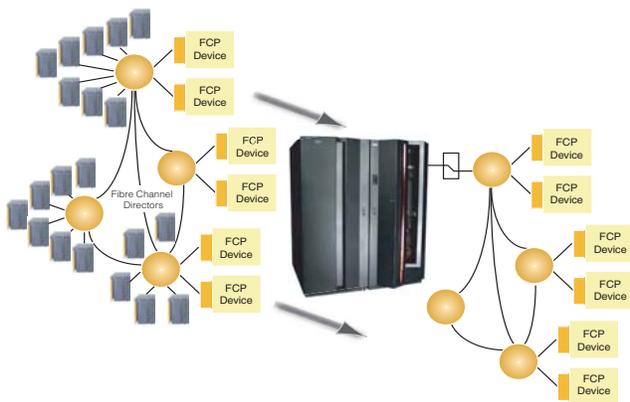
The support for FCP channels is for Linux only. Linux may be the native operating system on the zSeries server (note z990 runs LPAR mode only), or it can be in LPAR mode or, operating as a guest under z/VM 4.3 or later. Please note, FCP device support is not available for native z/VM; rather, z/VM makes the FCP function available to the Linux guest.

The 2 Gbps capability on the FICON Express channel cards means that 2 Gbps link speeds are available for FCP channels as well.

Open Systems Adapter-Express Features (OSA-Express)

FCP Full fabric connectivity

FCP full fabric support means that any number of (single vendor) FCP directors/ switches can be placed between the server and FCP/ SCSI device thereby allowing many “hops” through a storage network for I/O connectivity. This support along with 2 Gbps link capability is being delivered together with IBM switch vendors Brocade, INRANGE, and McDATA. FCP full fabric connectivity enables multiple FCP switches/ directors on a fabric to share links and therefore provide improved utilization of inter-site connected resources and infrastructure. Further savings may be realized in the reduction of the number of fiber optic cabling and director ports.



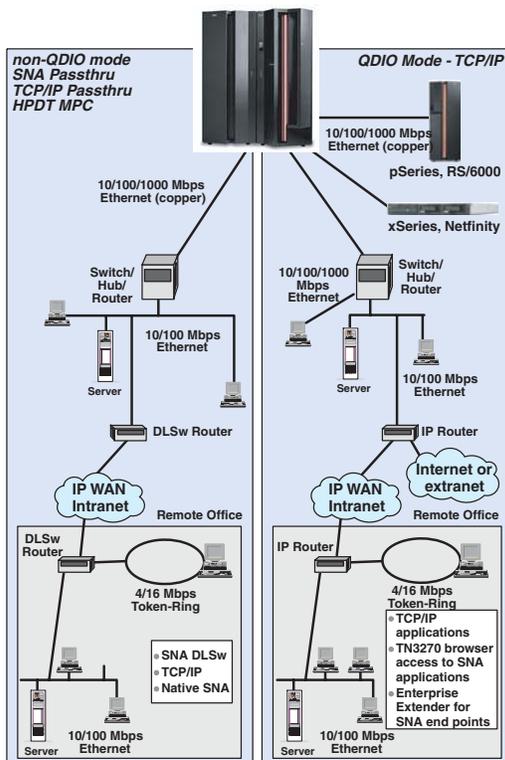
With the introduction of the z990 its increased processing capacity, and the availability of multiple Logical Channel Subsystems, the OSA-Express Adapter family of Local Area Network (LAN) adapters is also expanding by offering a maximum of 24 features per system, versus the maximum of 12 features per system on prior generations. This ensures a balanced solution to maximize throughput and responsiveness in an on demand application environment. These features combined with z/OS, z/OS.e, or OS/390, z/VM or VM/ESA®, Linux and VSE/ESA deliver a balanced system solution to maximize throughput and minimize host interrupts to continue to satisfy your business goals.

Each of the OSA-Express features offers two ports for connectivity delivered in a single I/O slot, with up to a maximum of 48 ports, each port uses a single CHPID. You can choose any combination of OSA-Express features: the new OSA-Express Gigabit Ethernet LX or SX, the new OSA-Express 1000BASE-T Ethernet or OSA-Express Token-Ring. The new OSA-Express GbE features have a new connector type, LC Duplex, replacing the current SC Duplex connectors. The prior OSA-Express Gigabit LX and SX, the OSA-Express Fast Ethernet, and the OSA-Express Token-Ring can be carried forward on an upgrade from z900.

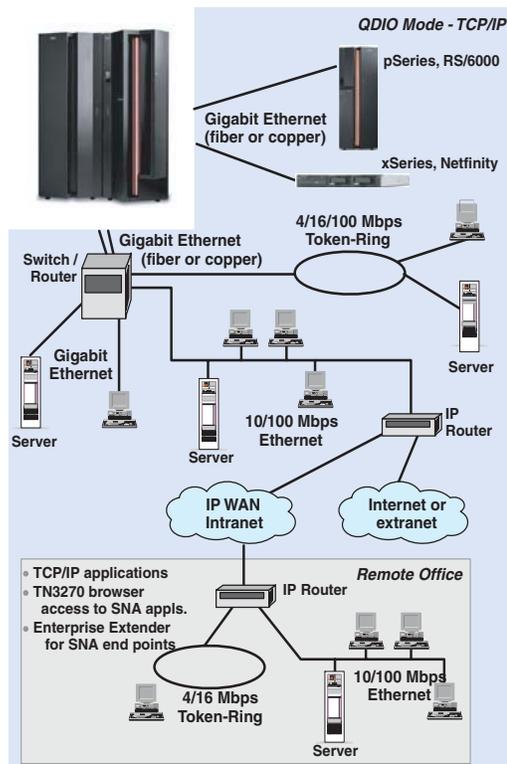
The new OSA-Express 1000BASE-T Ethernet feature replaces the current Fast Ethernet (10/100 Mbps) feature. This new feature is capable of operating at 10,100 or 1000 Mbps using the same copper cabling infrastructure as Fast Ethernet and it supports Auto-negotiation, QDIO and non-QDIO environments allowing you to make the most of your TCP/IP and SNA/APPN®/HPR environments at up to gigabit speeds.

OSA-Express ATM and OSA-2 FDDI are no longer supported. If ATM or FDDI support are still required, a multiprotocol switch or router with the appropriate network interface for example, 1000BASE-T Ethernet, GbE LX or GbE SX can be used to provide connectivity between the LAN and the ATM network or FDDI LAN.

z990 OSA-Express 1000BASE-T Ethernet



z990 OSA-Express Gigabit Ethernet



Queued Direct Input/Output (QDIO)

The OSA-Express Gigabit Ethernet, 1000BASE-T Ethernet and Token-Ring features support QDIO, which was first introduced in Communication Server for OS/390 2.7.

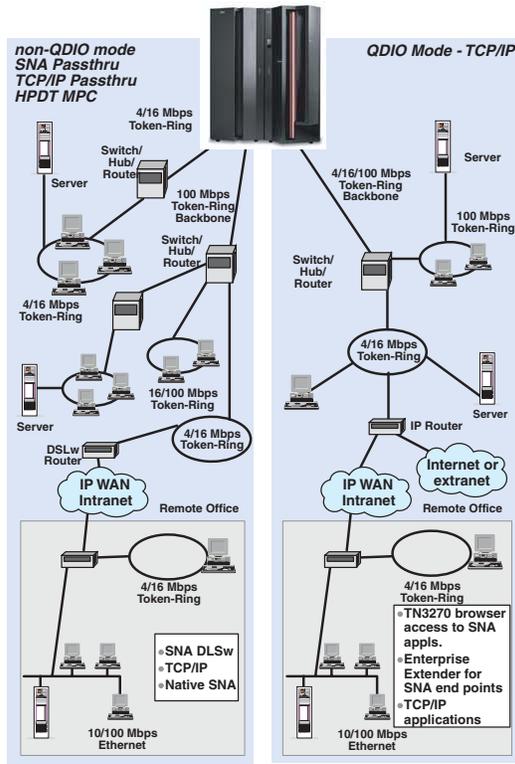
Queued Direct Input/Output (QDIO), a highly efficient data transfer architecture, breaks the barriers associated with the Channel Control Word (CCW/ESCON) architecture increasing data rates and reducing CPU cycle consumption. QDIO is designed to dramatically reduce system overhead and improves throughput using S/390 memory queues and a signaling protocol to directly exchange data between the OSA-Express microprocessor and CS for OS/390 2.10 and z/OS.

Full Virtual Local Area Network (VLAN) support is planned to be available in z/OS 1.5 Communications Server (CS) for the OSA-Express 1000BASE-T Ethernet, Fast Ethernet and Gigabit Ethernet feature when configured in QDIO mode. Full VLAN support in a Linux on zSeries environment was delivered for QDIO mode in April 2002 for z800 and z900.

NON-QDIO operational mode

The OSA-Express 1000BASE-T Ethernet, Fast Ethernet and Token-Ring also support the non-QDIO mode of operation. The adapter can only be set (via the CHPID type parameter) to one mode at a time. The non-QDIO mode does not provide the benefits of QDIO. However, it does provide for protocol support similar to the OSA-2 adapter, but at higher levels of performance. This support includes native SNA, APPN, High Performance Routing, TCP/IP Pass-through, and HPDT MPC. The new OSA-Express 1000BASE-T Ethernet provides support for TCP/IP and SNA/APPN/HPR up to 1 gigabit per second over the copper wiring infrastructure.

z990 OSA-Express Token-Ring



LPAR Support of OSA-Express

For z990 customers or customers who use the Processor Resource/Systems Manager™ (PR/SM) capabilities of the z900 and the S/390 servers, IBM offers the Multiple Image Facility (MIF), allowing the sharing of physical channels by any number of LPARs. Since a port on an OSA-Express feature is like a channel, sharing of an OSA-Express port is done using MIF. The LPARs are defined in the Hardware Configuration Definition (HCD). Depending upon the feature, and how it is defined, SNA/APPN/HPR and TCP/IP traffic can flow simultaneously through any given port.

Server to User connections

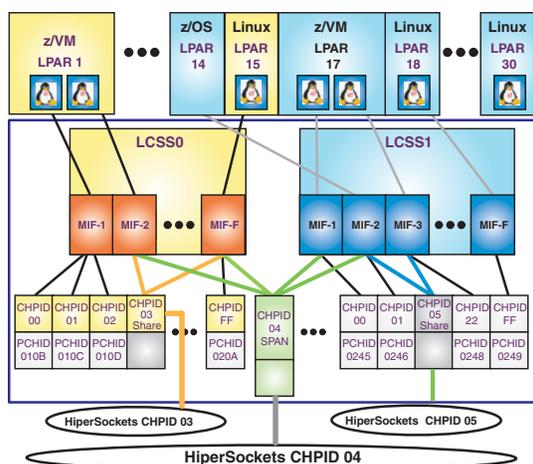
A key strength of OSA-Express and associated Communications Server protocol support is the ability to accommodate the customer's attachment requirements, spanning combinations of TCP/IP and SNA applications and devices. Customers requiring TCP/IP connections from the remote site to the TCP/IP or SNA applications on zSeries and S/390 can use OSA-Express with QDIO and either direct TCP/IP access or appropriate SNA to IP integration technologies, such as TN3270 Server and Enterprise Extender. Customers who require the use of SNA-based connections from the remote site can use a TCP/IP or SNA transport to the data center and connect into zSeries and S/390 using appropriate OSA-Express features in non-QDIO mode.

HiperSockets

HiperSockets, a feature unique to the zSeries, provides a “TCP/IP network in the server” that allows high-speed any-to-any connectivity among virtual servers (TCP/IP images) within a zSeries system without any physical cabling.

HiperSockets minimizes network latency and maximizes bandwidth between combinations of Linux, z/OS and z/VM virtual servers. These OS images can be first level (directly under an LPAR), or second level images (under z/VM).

With new support for up to sixteen HiperSockets per LPAR on the z990 (four on the z800 and z900), one could separate traffic to different HiperSockets for security (separation of LAN traffic, no external wire-tapping, monitoring) and performance and management reasons (separate sysplex traffic Linux or non-sysplex LPAR traffic).



HiperSockets does not use an external network, therefore, it can free up system and network resources, eliminating attachment cost while improving availability and performance. HiperSockets can have significant value in server consolidation, for example, by connecting multiple Linux virtual servers under z/VM to z/OS machines. Furthermore, HiperSockets will be utilized by TCP/IP in place of XCF for sysplex connectivity between images which exist in

the same server, thus z/OS TCP/IP uses HiperSockets for connectivity between sysplex images in the same server and uses XCF for connectivity between images in different servers. Management and administration cost reductions over existing configurations are possible.

HiperSockets acts like any other TCP/IP network interface, so TCP/IP features like IP Security (IPSec) in Virtual Private Networks (VPN) and SSL can be used to provide heightened security for flows within the same CHPID. HiperSockets supports multiple frame sizes, which is configured on a per HiperSockets CHPID basis. This support gives the user the flexibility to optimize and tune each HiperSockets to the predominant traffic profile, for example to distinguish between “high bandwidth” workloads such as FTP versus lower bandwidth interactive workloads.

The HiperSockets function provides many possibilities for improved integration between workloads in different LPARs, bound only by the combinations of operating systems and their respective applications. HiperSockets is intended to provide the fastest zSeries connection between e-business and Enterprise Resource Planning (ERP) solutions sharing information while running on the same server. WebSphere http and Web Application Servers or Apache http servers can be running in a Linux image (LPAR or z/VM guest) and will be able to use HiperSockets for very fast TCP/IP traffic transfer to a DB2 database server running in a z/OS LPAR. System performance is optimized because this allows you to keep your Web and transaction application environments in close proximity to your data and eliminates any exposure to network related outages, thus improving availability.

The z/OS HiperSockets Accelerator function can improve performance and cost efficiencies when attaching a high number of TCP/IP images via HiperSockets to a “front end” z/OS system for shared access to a set of OSA-Express adapters.

Cryptography

In the on demand era security will be a strong requirement. The zSeries products will continue to address security with announcements and deliveries of products and features.

The main focus in cryptography will continue to be very high and scalable performance for SSL algorithms, secondly, to provide competitive secure, symmetric performance for financial and banking applications using PIN/POS type encryption. As in the past zSeries will deliver seamless integration of the cryptography facilities through use of ICSF. Use of ICSF will assure that applications will work without change regardless of how and where the cryptographic functions are implemented, and also assure that the cryptography work is load balanced across the hardware resources. Finally we will be focused on required certifications and open standards.

The existing PCICA card feature will continue to be available on the z990 – for SSL acceleration / clear key operations. To support the increased number LPARs available on z990 the configuration options for the crypto PCICA adapter — introduced with the z900 — will be extended to allow sharing of a PCICA over the whole range of LPARs with a max of 16 LPARs sharing one PCICA adapter.

In addition to the PCICA, a new crypto adapter (PCIXCC) will be introduced as a functional replacement for the CMOS Cryptographic Coprocessor and the PCI Cryptographic Coprocessor. The PCIXCC adapter design introduces a breakthrough concept which supports high security demanding applications requiring a FIPS 140-2 Level 4 certified crypto module, also as an execution environment for customer written programs and a high performance path for Public Key / SSL operations. The PCIXCC HW and

microcode design will support almost all of the past cryptographic functions which were provided on the zSeries 900 via the CMOS Cryptographic Coprocessor (CCF), the PCI Cryptographic Coprocessor (PCICC). At the system software level the SSL related operations will be directed to the PCICA adapter and the Secure Crypto operations to the PCIXCC adapter.

The zSeries cryptography is further advanced with the introduction of the Cryptographic Assist Architecture implemented on every z990 processor (CPU). With unprecedented scalability and data rates the z990 processor provides a set of symmetric cryptographic functions, synchronously executed, which enormously enhance the performance of the en/decrypt function of SSL, VPN (Virtual-Private-Network) and data storing applications which do not require FIPS 140-2 Level 4 security. The on-processor crypto functions run at z990 processor speed, an order of magnitude faster than the CMOS Crypto Coprocessor in the zSeries 900. As these crypto functions are implemented in each and every CPU the affinity problem of pre-z990 systems (which had only two CMOS Crypto Coprocessors) is eliminated. The Crypto Assist Architecture includes DES and T-DES data en/decryption, MAC message authentication and SHA-1 secure hashing; all of these functions are directly available to application programs (zSeries Architecture instructions) and so will reduce programming overhead. To conform with US Export and Import Regulations of other countries a SE panel is provided for proper enable/disable of "strong" cryptographic functions.

Availability

z990 Capacity Upgrade on Demand (CUoD)

Capacity Upgrade on Demand allows for the nondisruptive addition of one or more Central Processors (CPs), Internal Coupling Facilities (ICFs) and/or Integrated Facility for Linux (IFLs). Capacity Upgrade on Demand can quickly add processors up to the maximum number of available inactive engines. This provides customers with the capacity for much needed dynamic growth in an unpredictable e-business world. The Capacity Upgrade on Demand function, combined with Parallel Sysplex technology, enables virtually unlimited capacity upgrade capability.

The CUoD functions are:

- *Nondisruptive CP, ICF, and IFL upgrades within minutes*
- *Dynamic upgrade of all I/O cards in the I/O Cage*
- *Dynamic upgrade of spare installed memory*
- *Plan Ahead and Concurrent Conditioning*

Concurrent Conditioning configures a system for hot plugging of I/O based on a future specified target configuration. Concurrent Conditioning of the zSeries I/O is minimized by the fact that all I/O cards plugging into the zSeries I/O cage are hot pluggable. This means that the only I/O to be conditioned is the I/O cage itself. The question of whether or not to concurrently condition a cage is a very important consideration, especially with the rapid change in the IT environment (e-business) as well as the technology. Migration to FICON Express or additional OSA-Express networking is exceptionally easy and nondisruptive with the appropriate microcode load and if the cage space is available.

The z990 supports concurrent memory upgrade. This capability will allow a processor's memory to be increased without disrupting the processor operation. To take advantage of this capability, a customer should not plan processor storage on the 16 or 32 GB increments. A customer with a Model A08, for example, with 8 GB of storage will be able to concurrently upgrade to 16 GB but will not be able to get to the next increment of 24 GB without a disruption.

Plan Ahead for PUs is done by ordering a "more book" model. For example, if a customer needs 5 PUs initially, but plans to grow to need 10 PUs, he should not order an A08, but a model B16 with only 5 PUs initially active.

The Plan Ahead process can easily identify the customer configuration that is required to meet future needs. The result of concurrent conditioning is a flexible IT infrastructure that can accommodate unpredictable growth in a low risk, nondisruptive way. Depending on the required Concurrent Conditioning, there should be minimal cost associated with dormant z990 capacity. This creates an attractive option for businesses to quickly respond to changing environments, bringing new applications online or growing existing applications without disrupting users.

z990 Server Capacity BackUp (CBU)

Capacity BackUp (CBU) is offered with the z990 processors to provide reserved emergency backup CPU capacity for situations where customers have lost capacity in another part of their establishment and want to recover by adding reserved capacity on a designated z990 system. A CBU system normally operates with a "base" CP configuration and with a preconfigured number of additional Processor Units (PUs) reserved for activation in case of an emergency.

The z990 technology is suited for providing capacity backup. The reserved CBU processing units are on the same technology building block, the MCM, as the regular CPs. Therefore, a single processor can support two diverse configurations with the same MCM. For CBU purposes, the Models A08, B16, C24 and D32 can scale from a 1-way to a 32-way; with the purpose of providing capacity backup.

The “base” CBU configuration must have sufficient memory and channels to accommodate the potential needs of the larger CBU target machine. When capacity is needed in an emergency, the primary operation performed is activating the emergency CBU configuration with the reserved PUs added into the configuration as CPs.

Upon request from the customer, IBM can remotely activate the emergency configuration. This is a fast electronic activation that eliminates time associated with waiting for an IBM CE to arrive onsite to perform the activation. A customer request through the Hardware Master Console and Remote Support Facility could drive activation time down to minutes; a request by telephone (for customers without RSF) could drive activation to less than an hour.

The z990 supports concurrent CBU downgrade. This function enables a Capacity BackUp Server to be returned to its normal configuration without an outage (i.e. PowerOn Reset (POR)).

Automatic Enablement of CBU for Geographically Dispersed Parallel Sysplex (GDPS)

The intent of GDPS support for CBU is to enable automatic management of the reserved PUs provided by the CBU feature in the event of a processor failure and/or a site failure. Upon detection of a processor failure, site failure or planned disaster test GDPS will activate CBU to dynamically add PUs to the processors in the takeover site to acquire processing power required to restart mission-critical production workloads.

z990 Server Customer Initiated Upgrade (CIU)

Customer Initiated Upgrade (CIU) is the capability to initiate a processor and/or memory upgrade when spare PUs/installed unused memory, are available via the Web using IBM Resource Link. Customers can download and apply the upgrade using functions on the Hardware Management Console via the Remote Support Facility.

This unique and important function for zSeries gives the customer greater control and ability in adding capacity to the System to meet resource requirements for unpredictable e-business workloads and for applications which are difficult to size.

Advanced Availability Functions

Transparent Sparing

z990 offers 12 PU MCMs with 2 PUs reserved as Spares. In the case of processor failure, these Spares are used for transparent sparing.

Enhanced Dynamic Memory Sparing

The z990 has enhanced this robust recovery design with 16 times more chips available for sparing. This will virtually eliminate the need to replace a memory card due to DRAM failure.

Enhanced Storage Protect Keys: z990 has enhanced the memory storage protect key design by adding a third key array to each memory card. The arrays are parity checked and employ a Triple Voting strategy to assure accuracy. This will reduce the need for memory card replacement due to key array failure.

ESCON Port Sparing: The ESCON 16-port I/O card is delivered with one unused port dedicated for sparing in the event of a port failure on that card. Other unused ports are available for nondisruptive growth of ESCON channels.

Concurrent Maintenance

- *Concurrent Service for I/O Cards:* All the cards which plug into the new I/O Cage are able to be added and replaced concurrent with system operation. This eliminates any need to schedule outage to service or upgrade the I/O subsystem on this cage.
- *Upgrade for Coupling Links:* z990 has concurrent maintenance for the ISC-3 adapter card. Also, Coupling links can be added concurrently. This eliminates a need for scheduled downtime in the demanding sysplex environment.
- *Cryptographic Cards:* The PCIXCC and PCICA cards plug in the I/O cage and can be added or replaced concurrently with system operation.

- *Redundant Cage Controllers:* The Power and Service Control Network features redundant Cage Controllers for Logic and Power control. This design enables non-disruptive service to the controllers and eliminates customer scheduled outage.
- *Auto-Switchover for Service Element:* The z990 has two Service Elements. In the event of failure on the Primary SE, the switchover to the backup is handled automatically. There is no need for any intervention by the customer or Service Representative.

Concurrent Capacity Backup Downgrade (CBU Undo)

This function allows the customer to downgrade the disaster backup machine to its normal configuration without requiring the PowerOn Reset (POR).

Concurrent Memory Upgrade

This function allows adding memory concurrently, up to the maximum amount physically installed.

Customer Initiated Upgrade (CIU)

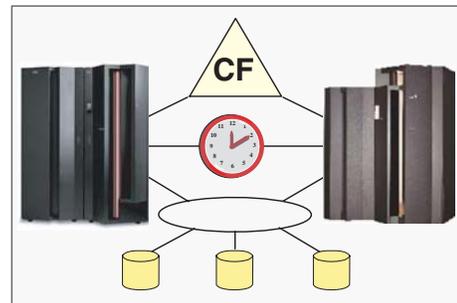
Customer Initiated Upgrade (CIU) is the capability to initiate a processor and/or memory upgrade when spare PUs/installed unused memory are available via the Web using IBM Resource Link.

- *z/Architecture*
 - 64-bit Architecture
- *Intelligent Resource Director*
 - LPAR CPU Management
 - Dynamic Channel Path Management
 - Channel Subsystem Priority Queuing
- *HiperSockets*
- *Internal System Channel-3*
- *Integrated Cluster Bus-4*
- *Internal Coupling Channel*

Parallel Sysplex Cluster Technology

- *Nondisruptive replacement of I/O*
- *Two port OSA-Express adapters at line speed*
- *ESCON Port Sparing*
- *1 Gbps/ 2 Gbps Auto-negotiated FICON Express*
- *FICON CTC*
- *FCP channels for Linux*
- *Concurrent Maintenance for ISC-3 adapter card*
- *Concurrent upgrade for Coupling Links*
- *Concurrent Service for all I/O cards*
- *63k devices per LCSS may be defined*
- *Redundant Cage Controllers*
- *Auto-Switchover for Service Element*
- *More Dense Logic Modules (MCMs)*
- *Enhanced Storage Protect Keys*
- *Enhanced Dynamic Memory Sparing*
- *Enhanced Hardware Compression*
- *Integrated Facilities for Linux (IFLs)*
- *Nondisruptive CBU downgrade*
- *Concurrent Memory Upgrade*
- *Customer Initiated Upgrade*
- *Hybrid Cooling*

Parallel Sysplex clustering was designed to bring the power of parallel processing to business-critical zSeries and S/390 applications. A Parallel Sysplex cluster consists of up to 32 z/OS and/or OS/390 images coupled to one or more Coupling Facilities (CFs or ICFs) using high-speed specialized links for communication. The Coupling Facilities, at the heart of the Parallel Sysplex cluster, enable high speed, read/write data sharing and resource sharing among all the z/OS and OS/390 images in a cluster. All images are also connected to a Sysplex Timer[®] to ensure time synchronization.



Parallel Sysplex Resource Sharing enables multiple system resources to be managed as a single logical resource shared among all of the images. Some examples of resource sharing include Automatic Tape Switching (ATS star), GRS “star,” and Enhanced Catalog Sharing; all of which provide simplified systems management, increase performance and/or scalability. For more detail, please see *S/390 Value of Resource Sharing White Paper – GF22-5115-00* on the Parallel Sysplex home page at ibm.com/servers/eserver/zseries/pso.

Although there is a significant value in a single footprint and multi-footprint environment with resource sharing, those customers looking for high availability must move on to a database data sharing configuration. With the Parallel Sysplex environment, combined with the Workload Manager and CICS TS or IMS TM, incoming work can be dynamically routed to the z/OS or the OS/390 image most capable of handling the work. This dynamic workload balancing, along with the capability to have read/write access data from anywhere in the Parallel Sysplex cluster, provides the scalability and availability that businesses demand today. When configured properly, a Parallel Sysplex cluster has no single point of failure and can provide customers with near continuous application availability over planned and unplanned outages. For detailed information on IBM's Parallel Sysplex technology, visit our Parallel Sysplex home page at ibm.com/servers/eserver/zseries/psa.

Coupling Facility Configuration Alternatives

IBM offers different options for configuring a functioning Coupling Facility:

- *Standalone Coupling Facility: z900 Model 100, z800 Model 0CF and 9672-R06 models provide a physically isolated, totally independent CF environment. There is no unique standalone coupling facility model offered with the z990. Customers can achieve the same physically isolated environment as on prior mainframe families by ordering a z990 with PUs characterized as ICFs. There are no software charges associated with such configuration. An ICF or CF partition sharing a server with any operating system images not in the sysplex acts like a logical standalone CF.*
- *Internal Coupling Facility (ICF): Customers considering clustering technology can get started with Parallel Sysplex technology at a lower cost by using an ICF instead of purchasing a standalone Coupling Facility. An ICF feature is a processor that can only run Coupling Facility Control Code (CFCC) in a partition. Since CF LPARs on ICFs are restricted to running only CFCC, there are no IBM software charges associated with ICFs. ICFs are ideal for Intelligent Resource Director and resource sharing environments as well as for data sharing environments where System Managed CF Structure Duplexing is exploited.*
- *Coupling Facility partition on a z990, z900, z800 or 9672 server using standard LPAR: A CF can be configured to run in either a dedicated or shared CP partition. IBM software charges apply. This may be a good alternative for test configurations that require very little CF processing resource or for providing hot-standby CF backup using the Dynamic Coupling Facility Dispatching function.*

A Coupling Facility can be configured to take advantage of a combination of different Parallel Sysplex capabilities:

- *Dynamic CF Dispatch: Prior to the availability of the Dynamic CF Dispatch algorithm, shared CF partitions could only use the "active wait" algorithm. With active wait, a CF partition uses all of its allotted time-slice, whether it has any requests to service or not. The optional Dynamic CF Dispatch algorithm puts a CF partition to "sleep" when there are no requests to service and the longer there are no requests, the longer the partition sleeps. Although less responsive than the active wait algorithm, Dynamic CF Dispatch will conserve CP or ICF resources when a CF partition has no work to process and will make the resources available to other partitions sharing the resource. Dynamic CF Dispatch can be used for test CFs and also for creating a hot-standby partition to back up an active CF.*

- *Dynamic ICF Expansion: Dynamic ICF expansion provides value by providing extra CF capacity when there are unexpected peaks in the workload or in case of loss of CF capacity in the cluster.*
 - *ICF Expansion into shared CPs. A CF partition running with dedicated ICFs needing processing capacity beyond what is available with the dedicated CP ICFs, can “grow” into the shared pool of application CPs being used to execute S/390 applications on the same server.*
 - *ICF Expansion into shared ICFs. A CF partition running with dedicated ICFs can “grow” into the shared pool of ICFs in case the dedicated ICF capacity is not sufficient. The resulting partition, an “L-shaped” LPAR, will be composed of both shared ICF and dedicated ICF processors, enabling more efficient utilization of ICF resources across the various CF LPARs.*

System-Managed CF Structure Duplexing

System-Managed Coupling Facility (CF) Structure Duplexing provides a general purpose, hardware-assisted, easy-to-exploit mechanism for duplexing CF structure data. This provides a robust recovery mechanism for failures such as loss of a single structure or CF or loss of connectivity to a single CF, through rapid failover to the backup instance of the duplexed structure pair.

Benefits of System-Managed CF Structure Duplexing include:

- *Availability*
Faster recovery of structures by having the data already there in the second CF eliminating the time and processing required for structure rebuilds. System-Managed Duplexing also provides basic recovery for many structures that have no simple means to recover data for failed structures, failed CFs, and losses of CF connectivity.

- **Manageability and Usability**

A consistent procedure to duplexing does not “set up” structures and manage structures across multiple exploiters.

- **Reliability**

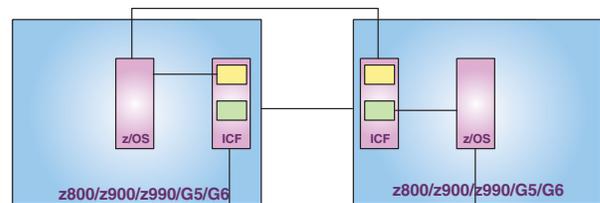
A common framework provides less effort on behalf of the exploiters, resulting in more reliable subsystem code.

- **Cost Benefits**

Facilitates the use of non-standalone CFs (e.g. ICFs) for data sharing environments in addition to resource sharing environments.

- **Flexibility**

The diagram below represents creation of a duplexed copy of the structure within a System-Managed CF Duplexing Configuration.



A robust failure recovery capability

Note: An example of two systems in a Parallel Sysplex with CF Duplexing

Customers who are interested in testing and/or deploying System-Managed CF Structure Duplexing in their sysplex should review GM13-0103, *System-Managed CF Structure Duplexing* at ibm.com/server/eserver/zseries/pso or ibm.com/server/eserver/zseries/library/techpapers/gm130103.html to understand the performance and other considerations of using this feature.

Parallel Sysplex Coupling Connectivity

The Coupling Facilities communicate with z/OS and OS/390 images in the Parallel Sysplex environment over specialized high-speed links. For availability purposes, it is recommended that there be at least two links connecting each z/OS or OS/390 image to each CF in a Parallel Sysplex cluster. As processor performance increases, it is important to also use faster links so that link performance does not become constrained. The performance, availability and distance requirements of a Parallel Sysplex environment are the key factors that will identify the appropriate connectivity option for a given configuration.

Parallel Sysplex coupling links on the zSeries have been enhanced with the introduction of Peer Mode. When connecting a zSeries server (z990/z900/z800) to a z800 Model OCF, a z900 Model 100 or a zSeries ICF, the links can be configured to operate in Peer Mode. This allows for higher data transfer rates to and from the Coupling Facilities. In Peer Mode, the fiber-optic single mode coupling link (ISC-3) provides 200 Gbps capacity, up to 10 km, 100 Gbps up to 20 km, the ICB-3 link with 1 GBps peak capacity, the ICB-4 for z990 to z990 connection at 2 GBps, and the IC-3 link with 1.25 GBps capacity. Additional Peer Mode benefits are obtained by enabling the link to be MIFed between z/OS (or OS/390) and CF LPARs. The peer link acts simultaneously as both a CF Sender and CF Receiver link, reducing the number of links required. Larger and more data buffers and improved protocols also improve long distance performance. For connectivity to 9672s, zSeries ISC-3 CF links can be configured to run in Compatibility Mode with the same characteristics as links on the 9672 of 100 Gbps. The z900 and z990 also support ICB-2 links for connectivity to 9672s. The ICB coupling link speeds described above are theoretical maximums.

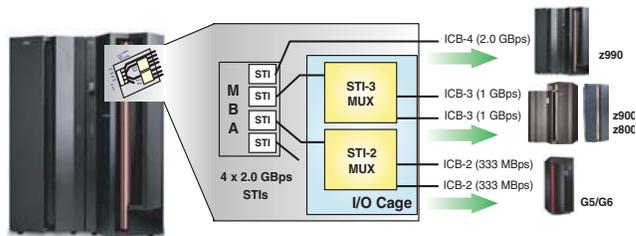
z990 Theoretical Maximum Coupling Link Speed

Connectivity Options	z990 ISC-3	z990 ICB-2	z990 ICB-3	z990 ICB-4
G5/G6 ISC	1 Gbps Compatibility	n/a	n/a	n/a
z800/z900 ISC-3	2 Gbps Peer Mode*	n/a	n/a	n/a
z990 ISC-3	2 Gbps Peer Mode	n/a	n/a	n/a
G5/G6 ICB	n/a	333 MBps Compatibility	n/a	n/a
z900 ICB-2	n/a	Not Supported	n/a	n/a
z990 ICB-2	n/a	Not Supported	n/a	n/a
z800/z900 ICB-3	n/a	n/a	1 GBps Peer Mode	n/a
z990 ICB-3	n/a	n/a	Not Recommended	n/a
z990 ICB-4	n/a	n/a	n/a	2 GBps Peer Mode

- *ISC-3. InterSystem Coupling Facility-3 channels provide the connectivity required for resource or data sharing between the Coupling Facility and the systems directly attached to it.*

ISC-3 channels are point-to-point connections that require a unique channel definition at each end of the channel. ISC-3 channels operating in Peer Mode provide connection between zSeries (z800/z900/z990) general purpose models and zSeries Coupling Facilities. ISC-3 channels operating in Compatibility Mode provide connection between z990 and HiPerLink (ISC-2) channels on 9672 G5 and G6 and the 9674 R06 Models. A four port ISC-3 card structure is provided on the z990 processors. It consists of a mother card with two daughter cards which have 2 ports each. Each daughter card is capable of operation at 1 Gbps in Compatibility Mode or 2 Gbps in peer mode up to a distance of 10 km. From 10 to 20 km, an RPQ card which comes in two port increments is available which runs at 1 Gbps in both Peer and Compatibility Modes. The mode is selected for each port via CHPID type in the IOCDs. The ports are activated in one port increments.

- ISC-2. HiPerLinks, based on single-mode CF links, are available on 9672s (G3-G6) and 9674s only. ISC-3s replace HiPerLinks on zSeries 900 and z990 models.
- ICB-2. The Integrated Cluster Bus-2 is used to provide high-speed coupling communication between a zSeries server or CF and a 9672 G5/G6 server or CF over short distances (~7 meters). For longer distances, ISC links must be used. The z990 features the STI-2 card which resides in an I/O cage and provides 2 ICB-2 ports each capable of up to 333 MegaBytes per second. The ports are activated in one port increments. Up to 4 STI-2 cards, 8 ICB-2 links are available on the z990.
- ICB-3. The Integrated Cluster Bus-3 is used to provide high-speed coupling communication between a z990 server or CF and a z800/z900 server or CF or between two z800/z900s over short distances (~7 meters). For longer distances, ISC-3 links must be used. The z990 features the STI-3 card which resides in an I/O cage and provides 2 ICB-3 ports each capable of up to 1 GB per second. The ports are activated in one port increments. Up to 8 STI-3 cards, 16 ICB-3 links are available on the z990. ICB-3 links operate in “peer mode.”
- ICB-4. The Integrated Coupling Bus-4 is a “native” coupling connection available for connecting a z990 server or CF to another z990 server or CF over short distances. Capable of up to 2 GB per second, the ICB-4 is the fast external coupling connection available for the z990. The ICB-4 connection consists of one link that directly attaches to an STI port on the system and does not require connectivity to a card in the I/O cage. One feature is required for each end of the link. Up to 16 ICB-4 features can be configured on a z990 depending on model selected.
- IC. The Internal Coupling channel emulates the Coupling links providing connectivity between images within a single server. No hardware is required, however a minimum of 2 CHPID numbers must be defined in the IOCDs. IC links provide the fastest Parallel Sysplex connectivity.



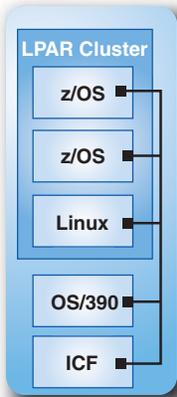
Intelligent Resource Director

Exclusive to IBM z/Architecture is Intelligent Resource Director (IRD), a function that optimizes processor and channel resource utilization across Logical Partitions (LPARs) based on workload priorities. IRD combines the strengths of the zSeries LPARs, Parallel Sysplex clustering, and z/OS Workload Manager.

Intelligent Resource Director uses the concept of an LPAR cluster, the subset of z/OS systems in a Parallel Sysplex cluster that are running as LPARs on the same zSeries server. On a z990, systems that are part of the same LPAR cluster may be in different LCSSs. In a Parallel Sysplex environment, Workload Manager directs work to the appropriate resources based on business policy. With IRD, resources are directed to the priority work. Together, Parallel Sysplex technology and IRD provide the flexibility and responsiveness to on demand e-business workloads unrivaled in the industry.

IRD has three major functions: LPAR CPU Management, Dynamic Channel Path Management, and Channel Subsystem Priority Queuing.

IRD Scope



LPAR CPU Management

LPAR CPU Management allows WLM working in goal mode to manage the processor weighting and logical processors across an LPAR cluster. CPU resources are automatically moved toward LPARs with the greatest need by adjusting the partition's weight. WLM also manages the available processors by adjusting the number of logical CPs in each LPAR. This optimizes the processor speed and multiprogramming level for each workload, reduces MP overhead, and gives z/OS more control over how CP resources are distributed to meet your business goals.

z/OS 1.2 enhances the LPAR CPU management capabilities and will allow the dynamic assignment of CPU resources to non-z/OS partitions outside the z/OS LPAR cluster such as Linux or z/VM.

Dynamic Channel Path Management

In the past, and on other architectures, I/O paths are defined with a fixed relationship between processors and devices. With z/OS and the zSeries, paths may be dynamically assigned to control units to reflect the I/O load. For example, in an environment where an installation normally requires four channels to several control units, but occasionally needs as many as six, system programmers must currently define all six channels to each control unit that may require them. With Dynamic Channel Path Management, the system programmer need only define the four channels to the control units, and indicate that DCM may add an additional two. As the control unit becomes more heavily used, DCM may assign channels from a pool of managed channels, identified by the system programmer, to the control unit. If the work shifts to other control units, DCM will unassign them from lesser utilized control units and assign them to what are now the more heavily used ones. DCM is for ESCON and FICON Bridge channels can help reduce the number of channels required to effectively run a workload. DCM can also reduce the cost of the fiber infrastructure required for connectivity between multiple data centers. On a z990 with multiple Logical Channel SubSystems (LCSSs), the scope of DCM management is within the Logical Channel SubSystems. Although an LPAR cluster can span multiple LCSS, when DCM is used it will only consider systems in the same LPAR cluster and the same LCSS.

Channel Subsystem Priority Queuing

The notion of I/O Priority Queuing is not new; it has been in place in OS/390 for many years. With IRD, this capability is extended into the I/O channel subsystem. Now, when higher priority workloads are running in an LPAR cluster, their I/Os will be given higher priority and will be sent to the attached I/O devices (normally disk but also tape and network devices) ahead of I/O for lower priority workloads. LPAR priorities are managed by WLM in goal mode.

Channel Subsystem Priority Queuing provides two advantages. First, customers who did not share I/O connectivity via MIF (Multiple Image Facility) out of concern that a lower priority I/O intensive workload might preempt the I/O of higher priority workloads, can now share the channels and reduce costs. Second, high priority workloads may even benefit with improved performance if there were I/O contention with lower priority workloads. Initially, Channel Subsystem Priority Queuing is implemented for Parallel OEMI and ESCON, FICON Bridge and native FICON channels.

On a z990, the scope of Channel Subsystem I/O Priority Queuing is a Logical Channel SubSystems.

Channel Subsystem Priority Queuing complements the IBM Enterprise Storage Server capability to manage I/O priority across CECs.

With IRD, the combination of z/OS and zSeries working in synergy extends the industry leading workload management tradition of S/390 and OS/390 to ensure that the most important work on a server meets its goals, to increase the efficiency of existing hardware, and to reduce the amount of intervention in a constantly changing environment.

Parallel Sysplex Professional Services

IBM provides extensive services to assist customers with migrating their environments and applications to benefit from Parallel Sysplex clustering. A basic set of IBM services is designed to help address planning and early implementation requirements. These services can help you reduce the time and costs of planning a Parallel Sysplex environment and moving it into production. An advanced optional package of services is also available and includes data sharing application enablement, project management and business consultation through advanced capacity planning and application stress testing. For more information on Parallel Sysplex Professional Services, visit IBM's Web site at ibm.com/servers/eserver/zseries/psso/services.html

Geographically Dispersed Parallel Sysplex

The Geographically Dispersed Parallel Sysplex (GDPS) complements a multisite Parallel Sysplex environment by providing a single, automated solution to dynamically manage disk and tape storage subsystem mirroring, processors, and network resources to allow a business to attain "continuous availability" and near transparent business continuity/disaster recovery without data loss. GDPS provides the ability to perform a controlled site switch for both planned and unplanned site outages, while maintaining full data integrity across multiple storage subsystems.

GDPS requires Tivoli® NetView® for z/OS or Tivoli NetView for OS/390, Tivoli System Automation for OS/390, and remote copy technologies. GDPS supports both the synchronous Peer-to-Peer Remote Copy (PPRC) as well as the asynchronous Extended Remote Copy (XRC) forms of remote copy. GDPS/PPRC is a continuous availability solution and near transparent business continuity/disaster recovery solution that allows a customer to meet a Recovery Time Objective (RTO) of less than an hour, a Recovery Point Objective (RPO) of no data loss, and protects against metropolitan area disasters (up to 40 km between sites).

GDPS/XRC is a business continuity/disaster recovery solution that allows a customer to meet an RTO of one to two hours, an RPO of less than a minute, and protects against metropolitan as well as regional disasters, since the distance between sites is unlimited. XRC can use either common communication links and channel extender technology or dark fiber as the connectivity between sites.

Note: Dark fiber refers to dedicated strands of fiber optic cable with no electronics between the ends (source and destination).

Geographically Dispersed Parallel Sysplex HyperSwap™

The GDPS/PPRC HyperSwap function is designed to broaden the continuous availability attributes of GDPS/PPRC by extending the Parallel Sysplex redundancy to disk subsystems. The HyperSwap function can help significantly increase the speed of recovering from planned and unplanned disk and site reconfigurations. The HyperSwap function is designed to be controlled by complete automation, allowing all aspects of the site switch to be controlled via GDPS.

The HyperSwap function delivered in GDPS 2.7 provides the ability to transparently switch all primary PPRC disk subsystems with the secondary PPRC disk subsystems for a planned switch configuration. The GDPS 2.7 planned HyperSwap capability enables disk configuration maintenance and planned site maintenance without requiring any applications to be quiesced. Large configurations can be supported, as HyperSwap has been designed to provide capacity and capability to swap large numbers of disk devices very quickly. The important ability to re-synchronize incremental disk data changes, in both directions, between primary/secondary PPRC disks is provided as part of this function.

GDPS 2.8 adds the capability to transparently switch all processing to the PPRC secondary disks in the event of a primary disk subsystem failure. With the GDPS 2.8 unplanned HyperSwap function, disk subsystem failures no longer constitute a single point of failure for an entire sysplex. A data sharing, application enables Parallel Sysplex, combined with the GDPS/PPRC 2.8 unplanned HyperSwap capability, to lay the foundation for continuous availability, even in the event of a complete site failure. In the event of a complete failure of the site where the primary disk resides, the systems in the site with the secondary disks can continue to remain active even though workload running on these systems needs to be restarted.

Geographically Dispersed Parallel Sysplex support for Peer-to-Peer Virtual Tape Server (PtP VTS)

The GDPS solution has been extended to include tape data in its management of data consistency and integrity across sites with the announced support of the Peer-to-Peer VTS configuration (IBM United States Hardware Announcement 101-215). The PtP VTS provides a hardware-based duplex tape solution and GDPS automatically manages the duplexed tapes in the event of a planned site switch or a site failure. At the present time, the GDPS PtP support is only available for a GDPS/PPRC (Peer to Peer Remote Copy) configuration.

A new I/O VTS selection option is provided especially for use with GDPS, so that all virtual volumes are processed from a primary VTS, and a copy is stored on the secondary VTS. Control capability has been added to allow GDPS to “freeze” copy operations, so that tape data consistency can be maintained across GDPS managed sites during a switch between the primary and secondary VTSs. Synchronization of system data sets such as catalogs, the tape control database, and tape management databases is also provided with the PtP VTS after an emergency switchover.

Operational data, data that is used directly by applications supporting end users, is normally found on disk. For the past several years, GDPS has provided continuous availability and near transparent business continuity for disk resident data. However, there is another category of data that “supports” the operational data, which is typically found on tape subsystems. Support data typically covers migrated data, point-in-time backups, archive data, etc. For sustained operation in the failover site, the support data is indispensable. Furthermore, several enterprises have mission-critical data that only resides on tape. By extending GDPS support to data resident on tape, the GDPS solution provides continuous availability and near transparent business continuity benefit for both disk and tape resident data. Enterprises will no longer be forced to develop and utilize processes that create duplex tapes and maintain the tape copies in alternate sites. For example, previous techniques created two copies of each DBMS image copy and archived log as part of the batch process and manual transportation of each set of tapes to different locations.

Automatic Enablement of CBU for Geographically Dispersed Parallel Sysplex

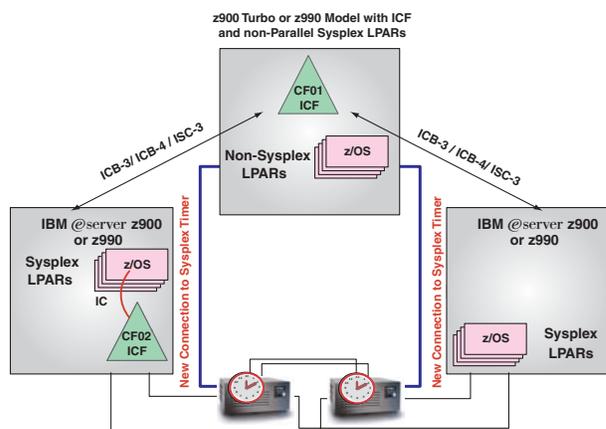
The intent of the GDPS (CBU) is to enable automatic management of the reserved PUs provided by the CBU feature in the event of a processor failure and/or a site failure. Upon detection of a site failure or planned disaster test, GDPS will dynamically add PUs to the configuration in the takeover site to restore processing power for mission-critical production workloads.

GDPS is discussed in a white paper available at ibm.com/server/eserver/zseries/psa/library.html. GDPS is a service offering of IBM Global Services. For IBM Installation Services for GDPS, refer to the IBM Web site.

Message Time Ordering

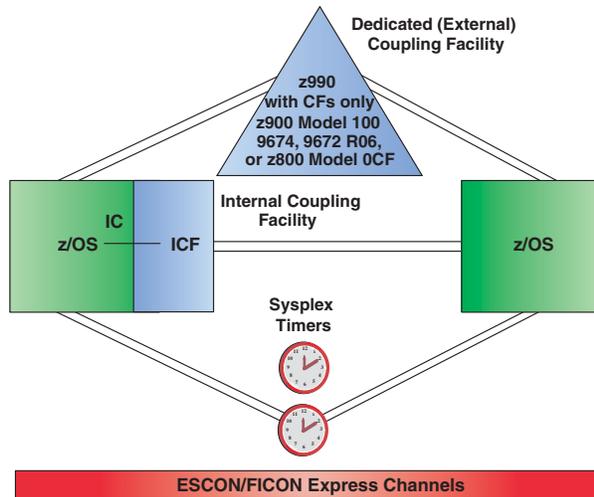
(Sysplex Timer Connectivity to Coupling Facilities)

As processor and Coupling Facility link technologies have improved over the years, the requirement for time synchronization tolerance between systems in a Parallel Sysplex environment has become ever more rigorous. In order to ensure that any exchanges of timestamped information between systems in a sysplex involving the Coupling Facility observe the correct time ordering, time stamps are now included in the message-transfer protocol between the systems and the Coupling Facility. Therefore, when a Coupling Facility is configured as an ICF on any z990 or z900 Models 2C1 through 216, the Coupling Facility will require connectivity to the same 9037 Sysplex Timer that the systems in its Parallel Sysplex cluster are using for time synchronization. If the ICF is on the same server as a member of its Parallel Sysplex environment, no additional connectivity is required, since the server already has connectivity to the Sysplex Timer. However, when an ICF is configured on any z990 or z900 Models 2C1 through 216 which do not host any systems in the same Parallel Sysplex cluster, it is necessary to attach the server to the 9037 Sysplex Timer.



z990 Support for Linux

Continuous Availability Recommended Configuration for Parallel Sysplex



Note: z990 will attach to 9037-001 or 9037-002. Service for 9037-001 will be discontinued at the end of 2003.

Components and assumptions

- *Two Coupling Facilities; at least one external*
- *Two Sysplex Timers*
- *Two z/OS or OS/390 servers with redundant backup capacity*
- *Two links from each CF to each image*
- *Two hardware management consoles*
- *Two ESCON or FICON Directors with cross-connected disks*
- *Dual electrical power grids*
- *Cloned OS/390 images, latest available software levels*
- *Automation capabilities for recovery/restart*
- *Critical data on RAID and/or mirrored disks*

Key attributes can include

- *No single point of failure*
- *Fast, automatic recovery*
 - *CF: rebuild in surviving CF*
 - *CEC, z/OS, OS/390: restart subsystems on surviving image*
 - *TM/DBMS: restart in place*
- *Surviving components absorb new work*
- *No service loss for planned or unplanned outages*
- *Near unlimited, plug-and-play growth capacity*

Linux for zSeries

Linux and zSeries are a great team. Linux has the same appearance and application programming interfaces no matter what platform it is running on, since it is designed to be platform-independent. When Linux is run on a zSeries server it inherits the legendary qualities of service that businesses world-wide rely on for hosting their most important workloads. Linux is open standards-based, supporting rapid application portability and can be adapted to suit changing business needs. The flexibility and openness of Linux make it very popular with developers, whose contributions endow Linux with an extensive and diverse application portfolio. zSeries servers enable massive scalability within a single server, either horizontally or vertically. Hundreds of Linux images can run simultaneously, providing unique server consolidation capabilities while reducing both cost and complexity.

Of course, no matter which Linux applications are brought to the zSeries platform, they all benefit from high-speed access to the applications and corporate data that resides on zSeries.

IBM developed the code that enables Linux to run on zSeries servers, and made it available to the Open Source community. The term used to describe this enabling code is "patches."

To eliminate the need for an external 2074 Console controller and associated consoles, an administrator may utilize the Hardware Master Console (HMC) functions "Integrated 3270 Console Support" for operating z/VM images, and "Integrated ASCII Console Support" to operate Linux images.

The support is implemented using an internal communications method — SCLP — which enables the operating system to communicate with the HMC. The software support will be made available for z/VM in Version 4 Release 4. An update for Linux will be made available to IBM Linux Distribution Partners.

Linux for zSeries supports the 64-bit architecture available on zSeries processors. This architecture eliminates the existing main storage limitation of 2 GB. Linux for zSeries provides full exploitation of the architecture in both real and virtual modes. Linux for zSeries is based on the Linux 2.4 kernel. Linux for S/390 is also able to execute on zSeries and S/390 in 32-bit mode:

IBM Middleware

- *Connectors*
 - *DB2 Connect™ Version 8.1*
 - *DB2 Connect Enterprise Edition Version 7.2*
 - *DB2 Connect Unlimited Edition Version 7.2*
 - *CICS Transaction Gateway Version 5.0*
 - *IMS Connect Version 1.1 and 1.2*
- *WebSphere Family*
 - *WebSphere Application Server Version 5.0*
 - *WebSphere Application Server for Developers Version 5.0*
 - *WebSphere Application Server Network Deployment Version 5.0*
 - *WebSphere Application Server Advanced Edition 4.0*
 - *WebSphere Application Server Advanced Single Server Edition Version 4.0*
 - *WebSphere Application Server Advanced Developer Edition Version 4.0*
 - *WebSphere Application Server Advanced Edition Version 3.5*
 - *WebSphere Commerce Business Edition Version 5.4*
 - *WebSphere Host On-Demand Version 7.0 and 6.0*
 - *WebSphere MQ Everyplace Version 2.0 and 1.2*
 - *WebSphere MQ Version 5.3*
 - *WebSphere Personalization Server for Multiplatforms Version 4.0*
 - *WebSphere Personalization Server Version 3.5*
 - *WebSphere Portal Server for Multiplatforms Version 4.1 and 4.2*
- *Data Management*
 - *DB2 Universal Database Enterprise Server Edition Version 8.1*
 - *DB2 Universal Developers Edition Version 8.1*
 - *DB2 Personal Developers Edition Version 8.1*
 - *DB2 Net.Data Version 8.1*
 - *DB2 Runtime Client Version 8.1*
 - *DB2 Spatial Extender Version 8.1*
 - *DB2 Intelligent Miner™ Modeling Version 8.1*
 - *DB2 Intelligent Miner Scoring Version 8.1*
 - *DB2 Intelligent Miner Visualization Version 8.1*
 - *DB2 Net Search Extender Version 8.1*
 - *DB2 Universal Database Enterprise Edition Version 7.2*
 - *DB2 Universal Database Developers Edition Version 7.2*
 - *DB2 Intelligent Miner Scoring Version 7.1*
 - *DB2 Net Search Extender Version 7.2*
- *Tivoli*
 - *Tivoli Access Manager for e-business Version 3.9 and 4.1*
 - *Tivoli Access Manager for Operating Systems Version 4.1*
 - *Tivoli Configuration Manager Version 4.2*
 - *Tivoli Decision Support for OS/390 Version 1.5.1*
 - *Tivoli Distributed Monitoring Version 4.1*
 - *Tivoli Enterprise Console Version 3.8 and 3.7.1*
 - *Tivoli Identity Manager Version 1.1*
 - *Tivoli Monitoring for Transaction Performance Version 5.1*
 - *Tivoli Monitoring Version 5.1.1 and 5.1*
 - *Tivoli NetView for z/OS Version 5.1*
 - *Tivoli Remote Control Version 3.8*
 - *Tivoli Risk Manager Version 4.1 and 3.8*
 - *Tivoli Software Distribution Version 4.0*
 - *Tivoli Storage Manager™ Version 5.1.5 and 5.1*
 - *Tivoli Storage Manager Client Version 4.2*
 - *Tivoli Switch Analyzer Version 1.2*
 - *Tivoli User Admin Version 3.8*
 - *Tivoli Workload Scheduler Version 8.1*

- *Informix*
 - *Informix C-ISAM*
- *U2*
 - *IBM UniData Version 5.2x*
- *Other IBM Software Products*
 - *IBM Application Workload Modeler Version 1.1 and Release 1*
 - *IBM Developer Kit Version 1.4 and 1.3.1*
 - *IBM Directory Integrator Version 5.1*
 - *IBM Directory Server Version 5.1 and 4.1*
 - *IBM HTTP Server Version 1.3.19.3*
 - *IBM Object REXX Version 2.2*
 - *IBM Screen Customizer Version 2.0.7 and 2.0.6*

Linux Distribution Partners

- *SuSE Linux Enterprise Server 7 for S/390 and zSeries*
Product Information at suse.de/en/produkte/susesoft/S390/
- *Turbolinux Server 6 for zSeries and S/390*
Product Information at turbolinux.com/products/s390
- *Red Hat Linux 7.2 for S/390*
Product Information at redhat.com/software/S390

z/VM Version 4

z/VM supports Linux on the mainframe. Within the VM environment, Linux images benefit from the ability to share hardware and software resources and use internal high-speed communications. While benefiting from the reliability, availability and serviceability of IBM **@server** zSeries, z/VM V4 offers an ideal platform for consolidating Linux workloads on a single physical server which allows you to run tens to hundreds of Linux images. z/VM V4 is priced on a per-engine basis (one-time charge) and supports IBM Integrated Facility for Linux (IFL) processor features for Linux based workloads, as well as standard engines for all other zSeries and S/390 workloads.

Integrated Facility for Linux (IFL)

The Integrated Facility for Linux feature of the zSeries servers provides a way to add processing capacity, exclusively for Linux workload, with no effect on the model designation. Operating systems like z/OS, TPF, and VSE/ESA will not execute on Integrated Facility for Linux engines. Consequently, these engines will not affect the software charges for workload running on the other engines in the system.

OSA-Express Ethernet for Linux

Driver support is provided for the functions of the new OSA-Express Gigabit Ethernet and 1000BASE-T Ethernet features.

OSA-Express Enhancements for Linux

A new function, Checksum Offload, offered for the new OSA-Express GbE and 1000BASE-T Ethernet features, is being introduced for the Linux for zSeries and z/OS environments. Checksum Offload provides the capability of calculating the Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Protocol (IP) header checksums. Checksum verifies the correctness of files. By moving the checksum calculations to a Gigabit or 1000BASE-T Ethernet feature, host CPU cycles are reduced and performance is improved. When checksum is offloaded, the OSA-Express feature performs the checksum calculations for Internet Protocol Version 4 (IPv4) packets.

HiperSockets

HiperSockets can be used for communication between Linux images and Linux or z/OS images, whether Linux is running in an IFL LPAR, natively or under z/VM.

Virtual Local Area Networks (VLANs), IEEE standard 802.1q, is now being offered for HiperSockets in a Linux for zSeries environment. VLANs can reduce overhead by allowing networks to be organized for optimum traffic flow; the network is organized by traffic patterns rather than physical location. This enhancement permits traffic to flow on a VLAN connection both over HiperSockets and between HiperSockets and an OSA-Express GbE, 1000BASE-T Ethernet, or Fast Ethernet feature.

Internet Protocol Version 4 (IPv4) broadcast packets are now supported over HiperSockets. TCP/IP applications that support IPv4 broadcast, such as OMPROUTE when running Routing Information Protocol Version 1 (RIPv1), can send and receive broadcast packets over HiperSockets interfaces. This support is exclusive to z990.

You can now transparently bridge traffic between a HiperSockets and an external OSA-Express network attachment. New Linux Layer 2 Switch (Linux L2S) support can simplify network addressing between HiperSockets and OSA-Express. You can now seamlessly integrate HiperSockets connected operating systems into external networks, without requiring intervening network routing overhead, thus increasing performance and simplifying configuration.

The currently available distributions; SuSE SLES 7, SuSE SLES 8, Red Hat 7.1 and Red Hat 7.2 support z990 compatibility and exploitation of 30 LPARs and 2 Logical Channel SubSystems. Support to further exploit z990 functions will be delivered as an Open Source Contribution via www.software.ibm.com/developerworks/opensource/linux390/index.shtml. A beta version is planned to be delivered on June 16, 2003 and a formal release will follow. IBM is working with its distribution partners to provide these functions in future distribution releases.

Fibre Channel Protocol (FCP channel) Support for Linux

Support for FCP channels enables zSeries servers to connect to select Fibre Channel Switches and FCP/SCSI devices under Linux for zSeries. This expanded attachability provides a larger selection of storage solutions for Linux implementations.

Cryptographic Support for Linux

Linux for zSeries running on standard z990, z900, and z800 engines is capable of exploiting the hardware cryptographic feature provided by the PCICA feature (PCI Cryptographic Accelerator). This enables customers implementing e-business applications on Linux for zSeries to utilize this enhanced hardware security.

zSeries 990 Family Configuration Detail

Linux Support

Environment

- z990, z900, z800 or S/390 single image
- zSeries or S/390 LPAR
- VM/ESA or z/VM guest

Block devices

- VM minidisks
- ECKD 3380 or 3390 DASDs
- VM virtual disk in storage

Network devices

- Virtual CTC
- ESCON CTC
- OSA-Express (Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet, Token-Ring) up to 24 features/48 ports on z990
- HiperSockets (up to 4096 TCP/IP stacks on up to HiperSockets on z990)
- 3172
- IUCV
- Character devices
- 3215 console
- Integrated console

Additional information is available at ibm.com/linux/ and at ibm.com/zseries/linux/.

Maximum of 512 CHPIDs (SOD 1024); 3 I/O cages (28 slots each) = 84 I/O slots

All features that require I/O slots are included in this matrix.

Feature	Per System			Increments, Ports/feature	Comments
	Minimum Features	Maximum Features	Maximum Ports		
ESCON, 16 port	0 ¹	35	512	4/16 LICCC controlled	1 spare port per feature
FICON Express ^{5,8}	0 ¹	60	120	2 ports	Configured as FC, FCV, max 20 features per I/O cage
STI-2 ICB-2 link ⁶	0	4	N/A	1/2	
	0 ¹	N/A	8	N/A	
STI-3 ICB-3 ⁶	0	8	N/A	1/2	
	0 ¹	N/A	16	N/A	
ISC-3 M ISC link ^{3,6}	0	8	N/A	4 ³	
	0 ¹	N/A	32	1	
OSA-Express ^{5,2}	0	24	48	2 ports	
PCICA (6) ^{4,5,7}	0	6 (5)	12	2 accelerator cards	total # can not exceed 2 features/ I/O cage
PCIXCC ^{4,5,7}	0	4	4	1 coprocessor	

- 1) At least one I/O or Coupling Link feature must be present; ESCON/ FICON or ICB/ISC.
- 2) OSA-Express any combination of GbE LX, GbE SX, 1000Base-T Ethernet, TR
- 3) ISC-3 M is the "mother" card which supports two "daughter" cards (ISC-D). Each daughter card has two ports. ISC-3 is offered in 1-port increments; ports are activated across all installed ISC-3 M features.
- 4) The total number of PCICAs and PCIXCCs cannot exceed eight features per system.
- 5) The total number of FICON Express, OSA-Express, PCICA and PCIXCC cannot exceed 20 features per I/O cage and 60 features per system.
- 6) The maximum number of Coupling Links combined (ICB-2, ICB-3, ICB-4, and active ISC-3 links) cannot exceed 32 per system.
- 7) PCICA and PCIXCC do not require a CHPID
- 8) Maximum of 48 FICON Express featured / 96 channels on Model A08 (needs 3 I/O cages)

Processor Unit Assignments

Model	Min. PU*	Max. PU	SAP Standard	Spares Standard
A08**	1	8	2	2
B16**	1	16	4	4
C24**	1	24	6	6
D32**	1	32	8	8

*PU can be characterized as CP, IFL, ICF, Optional SAPs, unassigned CPs, and/or unassigned IFLs up to a max number of PUs for the model
 **Customer will be required to purchase at least one CP, IFL or ICF feature for any model.

Processor Memory

z990 Model	Minimum	Maximum
A08	8 GB	64 GB
B16	8 GB	128 GB
C24	8 GB	192 GB
D32	8 GB	256 GB

Max two memory cards per z990 book. Memory cards 8 GB, 16 GB or 32 GB.

Channels

Model	A08	B16	C24	D32
ESCON Min	0	0	0	0
**ESCON Max	512	512	512	512
FICON *Min	0	0	0	0
FICON *Max	96	120	120	120

*FICON Express and FCP configured on the same FICON Express features. Max channels total 120.

**ESCON increments of 4 channels.

Coupling Links

Links	IC	ICB-2*	ICB-3**	ICB-4	ISC-3	Max External Links
	0-32	0-8	0-16	0-16	0-32	32

*requires STI-2 card

**requires STI-3 card

Note: At least one I/O channel (FICON, ESCON) or one coupling link (ISC, ICB) must be present.

Cryptographic Features

	PCICA ^{1,2}	PCIXCC ²
Minimum	0	0
Maximum	6 ³	4 ⁴

1. Max two PCICA features per I/O cage
2. Max eight PCICA and PCIXCC features per system
3. Two accelerator cards per PCICA feature
4. One coprocessor per PCIXCC feature

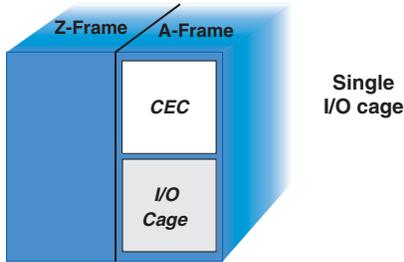
OSA-Express Features

	8
Minimum	0
Maximum	24

*Any combination of GbE LX, GbE SX, 1000BASE-T Ethernet, TR

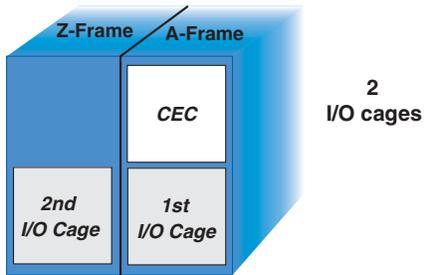
z990 Frame and I/O Configuration Content: Planning for I/O

The following diagrams show the capability and flexibility built into the I/O subsystem. All machines are shipped with two frames, the A-Frame and the Z-Frame and can have between one and three I/O cages. Each I/O cage has 28 I/O slots.



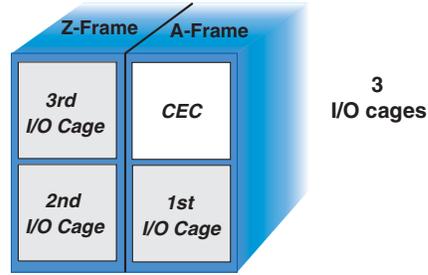
I/O Feature Type	Features	Maximum
ESCON	28 cards	420 channels
FICON Express	20	40 channels
OSA-Express	20	40 ports
PCIXCC	4	4
PCICA	2	4 cards

Maximum combined FICON Express, OSA-Express, PCICA/PCIXCC features is 20.



I/O Feature Type	Features	Maximum
ESCON	35 cards	512 channels
FICON Express	40	80 channels
O	S A	-
PCIXCC	4	4
PCICA	4	8 cards

Maximum combined FICON Express, OSA-Express, PCICA/PCIXCC features is 40.



I/O Feature Type	Features	Maximum
ESCON	35 cards	512 channels
FICON Express	60	120 channels
OSA-Express	24	48 ports
PCIXCC	4	4
PCICA	6	12 cards

Maximum combined FICON Express, OSA-Express, PCICA/PCIXCC features is 60.

General Information:

- ESCON configured in 4-port increments. Up to 28 channels in 2 cards, 60 channels in 4 cards, 88 channels in 6 cards, 120 in 8 cards, etc. up to a maximum 512 channels on 35 cards.
- OSA-Express can be Gigabit Ethernet, 1000BASE-T Ethernet or Token-Ring
- Total number of PCIXCC / PCICA is 8 per system
- If ICB-2 or ICB-3 are required on the system, these will use up a single I/O slot for each ICB-2 or ICB-3 to accommodate the STI-2 and STI-3 cards.

Physical Characteristics

Channels and channel adapters no longer supported on z990

The following channels and/or channel adapters are no longer supported:

- *Parallel channels - an ESCON converter is required; IBM 9034 or Optica 34600 FXBT*
- *OSA-2 adapters - use equivalent OSA-Express adapters, for FDDI use 1000BASE-T or Gigabit Ethernet with appropriate multi-protocol switch or router*
- *OSA-Express ATM - use 1000BASE-T or Gigabit Ethernet with appropriate multi-protocol switch or router*
- *4 Port ESCON cards - will be replaced with 16 port ESCON cards during upgrade*
- *FICON (pre FICON Express) - will be replaced with FICON Express during upgrade*
- *PCICC - replaced with PCIXCC for most functions*

The first ICB-2 or 3 required a slot. The second to the fourth required another slot. The fifth to the sixth required another slot. (STI - 2/3 cards each provide two ICBs)

z990 Power/Heating/Cooling

System Power Consumption (kW)

Model / Config	1 I/O Cage	2 I/O Cage	3 I/O Cage
A08	6.74	10.46	13.81
B16	9.57	13.27	16.98
C24	11.82	15.53	19.23
D32	13.98	17.68	21.39

Note: Assumes maximum configuration of I/O Cages 60 amp cords

System Cooling (Air Flow Rate - CFM)

Model / Config	1 I/O Cage	2 I/O Cage	3 I/O Cage
A08	1400	1800	2200
B16	1800	2200	3000
C24	2200	2600	3250
D32	2200	3000	3250

Note: Assumes chilled underfloor temperature of 24°C and maximum configuration of I/O cages

Heat Output (kBtu/hr)

Model / Config	1 I/O Cage	2 I/O Cage	3 I/O Cage
A08	22.92	35.56	46.95
B16	32.54	45.12	57.73
C24	40.19	52.80	65.38
D32	47.53	60.11	72.73

z990 Dimensions

	z990
# of Frames	2 Frames IBF contained within 2 frames
Height (w/ covers) Width (w/ covers) Depth (w/ covers)	194.1 cm / 76.4 in (40 EIA) 157.7 cm / 62.1 in (each frame 30.2 in) 157.7 cm / 62.1 in
Height Reduction Width Reduction	178.5 cm / 70.3 in (38 EIA) None
Machine Area Service Clearance	2.49 sq. meters / 26.78 sq. feet 5.45 sq. meters / 58.69 sq. feet (IBF contained within the frame)

Coupling Facility — CF Level of Support

CF Level	Function
12	64-bit support for Coupling Facility, CF Duplexing z990 Compatibility Support Toleration for > 15 LPAR ID on z990 Enhanced Storage Protect DB2 Performance Message Time Ordering
11	9672 G5/G6 CF Duplexing z990 Compatibility Support Toleration for > 15 LPAR ID on z990
10	z900 GA2 support
9	MQSeries® shared message queues WLM Multi-system Enclave Support Intelligent Resource Director** IC3/ISC3/ICB3 Peer Mode**
8	Dynamic ICF expansion into shared ICF pool Systems-Managed Rebuild
7	Shared ICF partitions on server models DB2 Delete Name optimization
6	ICB & IC TPF support
5	DB2 cache structure duplexing DB2 castout performance improvement Dynamic ICF expansion into shared CP pool*
4	Performance optimization for IMS & VSAM RLS Dynamic CF Dispatching Internal Coupling Facility* IMS shared message queue extensions
3	IMS shared message queue base
2	DB2 performance VSAM RLS 255 Connectors/1023 structures for IMS Batch DL1 (non-BMP)
1	Dynamic Alter support CICS temporary storage queues System Logger

Notes:

- G5 base CF level code is CF Level 6 and can be upgraded to CF Level 11
- G6 base CF level code is CF Level 8 and can be upgraded to CF Level 11
- z900 base CF level code is CF Level 9
- z800 and z990 base CF level code is CF Level 12
- Detailed information regarding CF Levels can be found in Coupling Facility Level (CF LEVEL) Considerations at ibm.com/s390/psocftable.html

*G3, G4, G5 and G6 only

**z900 required

Please note that although a particular back level machine may be updated to a more current CFCC level, NOT all the functions of that CFCC level may be able to run on that hardware platform, i.e., G3/G4 can be upgraded to CF Level 8 but it cannot use dynamic ICF expansion into shared ICF pool.

Fiber-Optic Cabling and System Connectivity

In the world of open systems and Storage Area Networks (SANs), the changing requirements for fiber optic cabling are directly related to the system hardware configuration. As industry-standard protocols and higher data rates continue to be embraced in these environments, the fiber-optic cabling options can become numerous and complex.

Today's marketplace is evolving towards new Small Form Factor (SFF) fiber optic connectors, short wavelength (SX) and long wavelength (LX) laser transceivers, and increasing link speeds from one Gigabit per second (Gbps) to 10 Gbps. New industry-standard SFF fiber optic connectors and transceivers are utilized on the zSeries ESCON and FICON Express features, on the ISC-3 feature, and on the zSeries ETR feature. These new features must coexist with the current infrastructure that utilizes a different “family” of fiber optic connectors and transceivers.

As a result of this complex and continually changing landscape, IBM is providing you with multiple fiber cabling services options to provide flexibility in meeting your fiber cabling needs.

IBM Network Integration and Deployment Services for zSeries fiber cabling (zSeries fiber cabling services) enables businesses to choose the zSeries configuration that best matches their computing environment without having to worry about planning and implementing the fiber-optic cabling. By teaming with IBM, businesses can receive a the right solution for their zSeries fiber connectivity requirements, including consulting and project management, as well as the fiber-optic jumper cables and installation to complete the zSeries integration.

zSeries fiber cabling now offers three options to ensure the right solution for your fiber cable jumper cable installation. Enterprise fiber cabling offers two additional options to meet your structured (trunking) environments requirements.

- *Fiber-optic jumper cabling package*
will analyze your zSeries channel configuration and your existing fiber-optic cabling to determine the appropriate fiber-optic jumper cables required, then supply, label and install the fiber-optic jumper cables and complete the installation with a detailed connection report
- *Fiber-optic jumper migration and reuse for a zSeries upgrade*
will plan, organize, re-label, re-rout and re-plug your existing fiber-optic jumper cables for reuse with the upgraded zSeries server
- *Fiber-optic jumper cables and installation*
will supply the fiber-optic jumper cables you specify, then label and install the fiber-optic jumper cables

Enterprise fiber cabling options

- *zSeries fiber optic trunk cabling package*
will analyze your zSeries channel configuration and your existing fiber-optic infrastructure to determine the appropriate fiber-optic harnesses, fiber-optic trunk cables and the fiber-optic patch panel boxes required, then supply, label and install the fiber-optic components to connect your new zSeries server to your existing structured fiber cabling infrastructure
- *Enterprise fiber cabling package*
will analyze your entire data center configuration and existing fiber-optic infrastructure to determine the appropriate end-to-end enterprise solution for connectivity. This is a customized offering that includes trunk cables, zone cabinets, patch panels and direct attach harnesses for servers, directors and storage devices.

These tailored zSeries fiber cabling options use the same proven planning and implementation methodologies as IBM's customized enterprise fiber cabling services, only focused on your zSeries fiber cabling needs.

Fiber Quick Connect (FQC): FQC, a zSeries configuration option, reduces the cable bulk associated with the installation of potentially 240 (z800) to 256 (z900) to 420 (z990) ESCON channels in one I/O cage. Fiber harnesses, which are factory-installed, enable connection to IBM's Fiber Transport System (FTS) direct-attach fiber trunk cables. Each trunk can have up to 72 fiber pairs. Four trunks can displace the 240 to 256 fiber optic cables on the z800 or z900.

In planning for zSeries systems, refer to *Planning for: S/390 Fiber Optic Links (ESCON, FICON, Coupling Links, and Open System Adapters)*, GA23-0367, and the *Installation Manual Physical Planning (IMPP)* manual. Refer to the services section of Resource Link for further details on the zSeries Fiber Cabling Service options and the Fiber Quick Connect configuration option.

Access Resource Link at ibm.com/servers/resourcelink

z/OS

While zSeries servers are supported by a number of different operating systems, their most advanced features are powered by z/OS. z/OS is the foundation for the future of zSeries, an integral part of the z/Architecture designed and developed to quickly respond to the demanding quality of service requirements for enterprise e-business.

z/OS is the robust operating system that is based on the 64-bit z/Architecture. It delivers the highest qualities of service for enterprise transactions and data, and extends these qualities to new applications using the latest software technologies. It provides a highly secure, scalable, high-performance base on which to deploy Internet and Java-enabled applications, providing a comprehensive and diverse application execution environment. z/OS takes advantage of the latest software technologies: new object-oriented programming models that permit the rapid design, development and deployment of applications essential to enterprise e-business. It protects your investment in your present mainframe applications by providing options for integrating existing applications within your e-business infrastructure. It provides a solid base for delivering on the benefits of industry-specific UNIX and e-business applications, supporting new technologies such as Enterprise JavaBeans™, XML, HTML, LDAP, Digital Certificates and Unicode. It supports such technological advances as Parallel Sysplex processing, Intelligent Resource Director and TCP/IP networking capabilities.

z/OS helps make critical data and processing functions accessible to end users regardless of their location in the heterogeneous e-business world. The z/OS base includes z/OS Communications Server, which enables: world-class TCP/IP and SNA networking support, including enterprise class dependability, performance, and scalability; highly secure connectivity; support for multiple protocols, and efficient use of networking assets.

The z/OS operating system combines many features that change the playing field of I/T infrastructure design:

- *Intelligent Resource Director expands the capabilities of z/OS Workload Manager to react to changing conditions and prioritize critical business workloads*
- *Support for 64-bit real memory and initial support of 64-bit virtual storage*
- *A new installation and configuration infrastructure that simplifies the installation and configuration of z/OS and related products*
- *Software pricing models designed to support e-business reality*

z/OS.e

Unique for the z800 is z/OS.e, a specially priced offering for z/OS, providing select function at an exceptional price. z/OS.e is intended to help customers by making the deployment of new applications on the z800 very attractively priced.

For more information about z/OS.e, see the *IBM @server zSeries 800 and z/OS, z/OS.e, z/VM and VSE/ESA Reference Guide*, GM13-0117.

Intelligent Resource Director

Intelligent Resource Director (IRD) is a key feature of the z/Architecture which extends the Workload Manager to work with PR/SM on zSeries servers to dynamically manage resources across an LPAR cluster. An LPAR cluster is the subset of the systems that are running as LPARs on the same zSeries server. Based on business goals, WLM can adjust processor capacity, channel paths, and I/O requests across LPARs without human intervention.

IRD assigns resources to the application; the application is not assigned to the resource. This capability of a system to dynamically direct resources to respond to the needs of individual components within the system is zSeries unique. It enables the system to continuously allocate resources for different applications. IRD is made up of three parts that work together to respond to the demands of e-business:

- *LPAR CPU Management*
- *Dynamic Channel Path Management*
- *Channel Subsystem Priority Queuing*
(see page 31)

The z/OS Intelligent Resource Director (IRD) Planning Wizard helps to plan your IRD implementation by asking questions about your enterprise setup, and produces a worksheet that describes the issues on each of your systems that you must consider before you can implement IRD. The z/OS IRD Planning Wizard supports z/OS 1.2 and higher, and is available at: ibm.com/servers/eserver/zseries/zos/wizards/

64-bit Real Storage Support

z/OS provides 64-bit real storage support for the zSeries servers. z/OS will continue to provide 31-bit real and expanded storage support for G5/G6 and Multiprise 3000 servers. The 64-bit real support eliminates expanded storage and is designed to help eliminate paging. The 64-bit real support may allow you to consolidate your current systems into fewer LPARs.

These z/OS functions are enhanced to exploit 64-bit real storage above 2 GB:

- *Traditional Access Methods (BSAM, QSAM, and others)*
- *VSAM for extended format data sets*
- *Hierarchical File System (HFS)*
- *Extended Remote Copy (XRC)*

These IBM products are enhanced to exploit 64-bit real storage above 2 GB:

- *DB2 Version 6 (with PTF), DB2 V7, DB2 V8*
- *IMS Version 7, IMS Version 8*

64-bit Virtual Storage Support

In addition, z/OS 1.2, 1.3 and 1.4 deliver basic 64-bit virtual storage management support. Assembler programs can obtain virtual storage above 2 GB for storing and manipulating data.

z/OS 1.5 delivers a 64-bit Application toolkit that provides assembler system services that enable middleware to share 64-bit virtual storage among multiple address spaces and also gives middleware and application developers the ability to compile programs with the new 64-bit C/C++ compiler.

For more information see IBM's 64-bit virtual storage roadmap at: ibm.com/servers/eserver/zseries/library/whitepapers/gm130076.html

Note: DB2 V8 takes advantage of the 64-bit virtual capability of z/Architecture and requires z/OS 1.3 or later as a prerequisite.

zSeries Bimodal Support for z/OS

The z/OS Migration Accommodation Offering is available to assist customers in migrating from OS/390 to z/OS. This addresses customer requests to have a fallback option to 31-bit mode when first migrating to z/OS in 64-bit mode on a zSeries server. It is very important that our OS/390 customers make the move to z/OS. This offering is available for 6 months starting from the registration of a z/OS license to a zSeries server. This offering can also assist customers running z/OS on a 31-bit machine who may in the event of a disaster have the requirement to recover to a 64-bit system.

msys for Operations Automation in z/OS Support

z/OS Managed System Infrastructure for Operations (msys for Operations) plays an important role in the area of outage avoidance. msys for Operations addresses certain self-healing, self-managing qualities of the IBM autonomic computing initiative on zSeries today. msys for Operations is an automation infrastructure within z/OS delivering functionality that targets system resources. It is based upon IBM's well-proven automation technology.

msys for Operations functions control and manage both hardware and software resources. The focus is on simplifying complicated operator interaction, detecting failure situations and reacting to them quickly and precisely. This is achieved through panel driven operator dialogs and automated recovery routines that run in the background.

IBM Health Checker for z/OS & Sysplex

The IBM Health Checker for z/OS and Sysplex is a tool that checks the current, active z/OS and Parallel Sysplex settings and compares their values to those best practices either suggested by IBM or defined by you as your criteria. The objective of the IBM Health Checker for z/OS and Sysplex is to identify potential problems before they impact your availability. This tool will be updated periodically with additional function and runs on all z/OS releases and on OS/390 2.10.

Download from: ibm.com/servers/eserver/zseries/zos/downloads/

z/OS Managed System Infrastructure for Setup — Self-Configuring Support

z/OS Managed System Infrastructure for Setup (msys for setup) is the strategic solution for product installation, configuration and function enablement. msys for Setup allows usage of consistent interfaces with wizard-like configuration dialogs. In z/OS 1.4, the msys for Setup Framework has been enhanced to provide multi-user capability. The workplace has been improved with the provision of the 'Control Panel,' which clarifies which parts of the sysplex or system are being updated, and provides clearer dialogs and a job progress indicator.

The dialogs use defaults and best practices values whenever possible and derive low-level values from answers to high-level questions. After the configuration parameters have been specified, msys for Setup can automatically update the system configuration directly. The user can see in detail what the changes will be before they are made.

System-Managed CF Structure Duplexing

System-Managed Coupling Facility (CF) Structure Duplexing is designed to significantly enhance Parallel Sysplex availability. It provides a robust failure recovery capability via CF structure redundancy, and it enhances Parallel Sysplex ease of use by reducing the complexity of CF structure recovery. These benefits are achieved by creating a duplicate (or duplexed) copy of a CF structure and then maintaining the two structure instances in a synchronized state during normal CF operation. In the event of a CF-related failure (or even a planned outage of a CF), failover to the remaining copy of the duplexed structures is initiated and quickly completed transparent to the CF structure user and without manual intervention.

TCP/IP Networking Enhancements

z/OS can provide near continuous availability for TCP/IP applications and their users with two key features in z/OS: Sysplex Distributor and VIPA nondisruptive Takeover and Takeback. This is a prime example of IBM innovation and integration in software and hardware to bring value added qualities — namely very high resiliency and availability — to the zSeries networking environment.

Virtual IP Address Nondisruptive Takeover

VIPA represents an IP address that is not tied to a specific hardware adapter address. The benefit is that if an adapter fails, the IP protocol can find an alternate path to the same software, be it the TCP/IP services on zSeries or an application.

VIPA Takeover supports movement to a backup IP stack on a different server in a Parallel Sysplex cluster in case of a failure of the primary IP stack. VIPA Nondisruptive Takeover provides VIPA takeback support. This allows the movement of workload back from the alternate to the primary IP stack.

Sysplex Distributor

Introduced in OS/390 2.10, Sysplex Distributor is a software-only means of distributing IP workload across a Parallel Sysplex cluster. Client connections appear to be connected to a single IP address, yet the connections are routed to servers on different zSeries or S/390 servers. In addition to load balancing, Sysplex Distributor simplifies the task of moving applications within a Parallel Sysplex environment.

e-business Security

z/OS Security Server is designed to address the demands of e-business enterprises. In z/OS 1.3, full function PKI (Public Key Infrastructure) support enables customers to create and manage large numbers of Digital Certificates. Many large e-business enterprises are in need of this large scale, full life-cycle capability for Digital Certificates. UNIX System Services provides greater security granularity for HFS and zFS file systems with support for Access Control Lists (ACLs).

z/OS provides industry-standard Lightweight Directory Protocol (LDAP) Version 3 services supporting thousands of concurrent clients.

Security Server (RACF®) in z/OS 1.4 supports Enterprise Identity Mapping (EIM) a cross platform IBM @server infrastructure that helps administrators and application developers to solve the problem of managing multiple user registries across their enterprise.

z990 Support Features

z/OS provides support for the z990 is provided in 2 stages: Compatibility and Exploitation.

Compatibility support provides OS/390 2.10, z/OS 1.2, z/OS 1.3 and z/OS 1.4 the ability to run on a z990 processor. Also, Compatibility Support is required on all members of a sysplex if any z/OS image of that sysplex is running on z990 in an LPAR ID>15, or if any Coupling Facility Image is running on a z990 in an LPAR ID >15. The z/OS 1.4 Compatibility Support is available via an optional, unpriced, orderable feature — z/OS 1.4 z990 Compatibility Support. Compatibility Support for z/OS 1.2, 1.3 and OS/390 2.10 is available via a Web download, called z990 Compatibility Support for Selected Releases from the z/OS home page at: ibm.com/eserver/zseries/zos/downloads.

z/OS.e support is available for z800 customers who intend to run z/OS.e in a sysplex which includes a z990 processor. The z/OS.e 1.4 support is available via an optional, unpriced, orderable feature — z/OS.e 1.4 z990 Coexistence is available via the same Web download as the z/OS support — z990 Compatibility for Selected Releases.

Exploitation support provides z/OS 1.4 the ability to run in a logical channel subsystem that has an id that is greater than zero and provides the ability to use an LPAR id greater than 15. The z/OS 1.4 support is planned to be available in 4Q2003 via an optional, unpriced, orderable feature — z/OS 1.4 z990 Exploitation Support. This new feature will include all the code provided by the z/OS 1.4 z990 Compatibility Support feature.

Note: There will be no exploitation support for any release prior to z/OS 1.4 Exploitation support will be integrated into z/OS 1.5 and later. An update to the z/OS.e 1.4 feature will also be provided to allow customers to maintain the same code base with z/OS. The z/OS.e 1.4 support is planned to be available in 4Q2003 via an optional, unpriced, orderable feature — z/OS.e 1.4 Coexistence Update.

ICSF support for z990 is described in the Security Services section of this document.

System Services

z/OS Version 1 Release 4 base elements

Base Control Program (BCP)
JES2
ESCON Director Support
MICR/OCR Support
Bulk Data Transfer base
DFSMSdftp™
EREP/MVS™ V3R5
High Level Assembler V1R4
ICKDSF R17
ISPF
TSO/E
3270 PC File Transfer Program V1R11
FFST™/ESA
TIOC

z/OS Version 1 Release 4 optional prices features

DFSMSdss™
DFSMSHsm™
DFSMSrmm™
DFSMSStvs
JES3
Bulk Data Transfer (BDT) File to File
BDT, SNA, NJE

z/OS Version 1 Release 4 optional non-priced features

z990 Compatibility Support
z990 Exploitation Support (available 4Q2003)
Console Enhancement (available 1Q2004)

The backbone of the z/OS system is the Base Control Program (BCP) with JES2 or JES3. These provide the essential integrity and the uninterrupted services that make z/OS the system of choice when workloads must be processed reliably, securely, and with complete data.

The existence of a completely integrated set of system services assures that a customer can focus on extracting the maximum business value from the z/OS installation. The system manages the workload, program libraries and I/O devices. Complexities are minimized and problem determination is facilitated with recovery and reporting facilities.

Highlights

Workload Manager (WLM) addresses the need for managing mixed workload distribution, load balancing and the distribution of computing resources to competing workloads. It does this while providing fewer, simpler system externals and performance management goals expressed in Service Level Agreement terms, automatic work and resource management. All this is done with a single policy that can be used across the sysplex to provide a single control point, eliminating the need to manage each individual image.

Dynamic balancing of JES2 batch initiators across sysplex has been enhanced in z/OS 1.4 to provide better utilization of processor resources. WLM checks every 10 seconds to see if re-balancing is required. WLM is more aggressive in reducing initiators on constrained systems and starting new ones on less utilized systems ensuring that processors are more evenly utilized.

WLM Improvements for WebSphere

Performance block reporting for enclaves and multi-period classes provides improved workload balancing for middleware applications such as WebSphere. WLM Enqueue Management establishes a new interface to allow reporting of resource contention. The priority of the task holding the enqueue can be increased to allow the resource to be released more quickly.

Data Management

DFSMSdfp provides a comprehensive set of functions to manage storage resources on the system. Data management functions support storage and retrieval of data on disk, optical and tape devices. Program management functions allow creation and retrieval of executable program libraries. Device management functions provide the means to define and control the operation of input and output storage devices. Distributed File Manager (DFM) supports access to remote data and storage resources.

In z/OS 1.3, DFSMS focuses on improving a customer's business continuance by adding features that minimize the occurrence of system outages and enhance a customer's disaster recovery capabilities. The new functions also improve a customer's business efficiency by providing better system performance and throughput and by providing usability enhancements that increase Storage Administrator productivity.

Benefits

- *Improved business continuity*
 - *Coupled Extended Remote Copy to enhance XRC disaster recovery for very large customers*
 - *SMS data set separation on separate storage controllers*
 - *SMS overflow storage group for peak workload handling*
 - *VSAM RLS Supports System Managed Coupling Facility structure*
- *Improved Performance*
 - *Large storage (64-bit) support for all VSAM record organizations (extended and non-extended format VSAM)*
 - *DFSMSHsm recall workload can be balanced across multiple hosts by placing recall requests in a common queue in the Coupling Facility*
 - *DFSMSDss one-step HFS data set copy*

DFSMSStvs

DFSMS Transactional VSAM Services (DFSMSStvs) is an optional, priced feature for z/OS 1.4 that enables running batch jobs concurrently with CICS online transactions to allow updates to the shared VSAM data sets. Multiple batch jobs and online transactions can be run against the same VSAM data sets and DFSMSStvs provide for data integrity for concurrent batch updates while CICS provides it for online updates.

DFSMSStvs offers a number of benefits, such as:

- *Contributes to the reduction or elimination of your batch window for CICS applications and other VSAM applications by allowing concurrent batch and online updating of VSAM recoverable data sets*
- *Increased system availability with simpler, faster, and more reliable recovery operations for VSAM storage structures*
- *Simplifies scheduling batch jobs because multiple batch jobs that access the same files can be run concurrently on one or more z/OS images in a Parallel Sysplex instead of serially on one image*
- *Provides the ability to share VSAM data sets at record level with integrity and commit and rollback functions for non-CICS applications*
- *Offers Backup-while-open to be taken using DFSMSdss and DFSMSHsm*
- *Enables batch applications to use the same forward recovery logs*
- *Helps enable 24x7 CICS TS applications*

For more detailed information on DFSMS, visit: storage.ibm.com/software/sms/index.html

Parallel Sysplex

zSeries Parallel Sysplex technology provides industry-leading availability by allowing z/OS workloads to be balanced across multiple servers to provide near continuous availability. This form of clustering, known as single system image, also provides tremendous scale. Up to 32 zSeries servers with z/OS can appear as a single image (see page 25, Parallel Sysplex Clustering Technology).

Base Control Program

- *The BCP has support for the IEEE floating point hardware, standard on all servers supported by z/OS*
- *DB2 can support customers with a requirement for more than 10,000 DB2 dynamically-allocated table spaces*
- *Increased number of system symbols*
- *Batched group buffer pool writes in single CF operation reduces overhead and improves performance in update-intensive DB2 data sharing*
- *CFRM performance has been enhancements to eliminate I/O bottlenecks that impact CFRM CDS*
- *Multi-system cascaded transactions subordinate failure notification allows IMS to exploit RRS multi-system cascaded transactions feature*

JES2 and JES3

In z/OS 1.2, JES2 and JES3 allow an installation to have up to 999,999 jobs. In addition, both provide the installation the ability to obtain (spinoff) their JESlog data sets prior to job completion.

The JES2 Health monitor, in z/OS 1.4, provides improved diagnostics. Even when JES2 is not responding to commands, the JES2 monitor, running in a separate address space, will be able to provide information about JES2's status. JES2 also provides enhanced recovery from corrupted checkpoint data to prevent multisystem outages

In z/OS 1.4, JES3 provides additional compatibility for initialization errors and the MAINPROC refresh function which provides dynamic addition of systems to the sysplex.

Console Availability Enhancement

The z/OS 1.4 Console Enhancements Feature extends the availability capabilities of z/OS 1.4, and represents the first phase of IBM's Consoles strategy to enhance the operator messaging architecture of z/OS. The focus for this first phase will be message production and consumption in a Parallel Sysplex environment and its restructuring to ensure high availability.

IBM plans to provide the z/OS 1.4 Console Enhancements Feature in 1Q2004.

Console Enhancement - Benefits in Detail: The z/OS 1.4 Console Enhancement Feature focuses on preventing the possibility of system outages due to exhaustion of system resources during message production and consumption between systems. To accomplish this, data space capabilities are employed to manage message buffering, and the possibility for bottlenecks is removed by instituting queuing independence. IBM also plans to enhance consoles so that there is always a console available to receive critical messages during system initialization. The hardware system console use would be disabled only when MCS, SMCS, or EMCS consoles become available.

System Management Services

z/OS Version 1 Release 4 base elements

HCD
SMP/E V3
Managed System Infrastructure for Setup
Managed System Infrastructure for Operations

z/OS Version 1 Release 4 optional prices features

RMF™
SDSF
HCM

z/OS provides systems management functions and features to manage not only host resources, but also distributed systems resources. These capabilities have a long, successful history of usage by S/390 customers. z/OS has enhanced many of these systems management functions and features to provide more robust control and automation of the basic processes of z/OS.

Highlights

- *msys for Setup can be used to define the following technologies*
 - *Parallel Sysplex (z/OS R1, R2, R3, R4)*
 - *Base Sysplex (R2)*
 - *TCP/IP (R2)*
 - *ISPF (R2)*
 - *UNIX System Services (R3)*
 - *Language Environment® (R3)*
- *msys for Operations enhancements in z/OS 1.3 include automation to handle ENQ resources contention and auxiliary storage shortages. msys for Operations can now interface with the Hardware Management Console (HMC) to provide Hardware functions such as de-activating LPARs.*
- *RMF can show the contention for Cryptographic Coprocessors, including a description of which workloads are using or are delayed in access to the cryptographic coprocessors*
- *WLM can show the contention for PCICC Cryptographic Coprocessors, including a description of which workloads are using or are delayed in access to the cryptographic coprocessors*
- *Application State Recording, a new feature of z/OS 1.4 provides more granular performance reporting for middleware such as WebSphere*

Benefits can include:

- *Increased system availability*
- *Improved productivity of system programmers*
- *A more consistent approach for configuring z/OS components or products*
- *System set-up and automation using best practices which can greatly improve availability*

msys for Setup

msys for Setup is part of IBM's advanced self-management technology and helps customers to enable functions as quickly and easily as possible. The value of msys for Setup is extended in z/OS 1.4 with the delivery of further msys for Setup exploitation and major enhancements to the framework. The msys for Setup framework improvements include:

- *Multiple User Support: With z/OS 1.4, multiple users can be logged on and work concurrently from different workstations. Furthermore, as part of the user enrollment process, the msys for Setup user administrator can control which msys for Setup workplace functions a user can access.*
- *Improved User Interface: The graphical user interface (msys for Setup workplace) has been redesigned and is now easier to learn and use. These valuable ease of use enhancements make working with msys for Setup more intuitive.*

In addition to these framework highlights, TCP/IP provides additional configuration capability through msys for Setup:

- *The TCP/IP plug-in provides customization of the TN3270 Server*
- *End-user definition of port reservation and port sharing is now supported*

msys for Operations

The objective of msys for Operations is to simplify day-to-day operational tasks associated with Parallel Sysplex and z/OS. This is achieved by reducing operator complexity, creating greater operational awareness of important indicators and improving system recoverability. All of these factors are essential to z/OS availability and in turn directly affect the performance and availability of every business application.

It is important to note that msys for Operations is not a product nor will it be displacing any product. It will however, remove the necessity for many customers to write their own automation exploiting code to achieve control over specific system events. Today many installations write and maintain their own versions of this type of code. For the less sophisticated customer this is virtually impossible. By supplying msys for Operations, each installation can decide to enable and make use of the new capabilities or leave them disabled, continuing to use existing automation routines.

Whether their interest is in the panel driven operator dialogs that assist in the control of couple data sets, coupling facilities and managing coupling facility structures OR in the background recovery routines that guard against console buffer shortages, long running enqueues or auxiliary storage shortages, z/OS customers will find something in msys for Operations that will benefit operations in their organizations and lead to improved availability.

SMP/E

SMP/E provides the ability to install software products and service either from DASD or tape, or directly from a network source, such as the Internet. By installing directly from a network source, SMP/E is enabling a more seamless integration of electronic software delivery and installation.

Security Services

z/OS Version 1 Release 4 base elements

Cryptographic Services (Integrated Cryptographic Service Facility, Open Cryptographic Services Facility, System SSL, and Data Confidentiality Services)

z/OS Version 1 Release 4 optional priced features

Security Server includes:

- RACF
 - Network Authentication Service (Kerberos services)¹
 - LDAP Server¹
 - Firewall Technologies²
 - DCE Security Server
 - Open Cryptographic Enhanced Plug-ins¹
 - Public Key Infrastructure Services¹
-

z/OS Version 1 Release 4 optional no charge features

Network Authentication Service Level 3

System SSL Security Level 3

Open Cryptographic Services Facility Security Level 3

1. These components are now a licensed part of the z/OS base and can be used without purchasing/enabling Security Server

2. Firewall technologies (IKE and GUI only) are now part of the z/OS base

The optional Security Server for z/OS combines the traditional benefits of RACF with an integrated set of security functions essential to e-business applications. It forms the basis for all security services from traditional applications, UNIX applications, and distributed systems. Access to resources can be selectively controlled, audited, and managed with appropriate centralized or decentralized control as required by each installation.

Highlights

- Full function PKI support
- New cryptographic algorithms and standards
- More granular access control for UNIX file systems (zFS and HFS)
- Integrated package of security and directory functions, with advanced security capabilities
- Tivoli administration interface

Benefits include:

- Extensive security controls over emerging e-business opportunities
- Flexible control of access to protected resources, including installation-defined items
- Choice of centralized or decentralized control of security profiles
- Choice of platform for DCE security server, with RACF for access control, and remote administration
- Support for the Kerberos V5 network authentication protocol
- Support for z800/z900/z990 cryptographic features

RACF

- Digital Certificates can be automatically authenticated without administrator action
- Administrative enhancements enable definition of profiles granting partial authority. Handling of new passwords and removal of class authority are simplified.
- Customers designing e-business applications need a way to associate more users under a RACF Group definition, so RACF allows you to create a new kind of group that can contain an unlimited number of users

SSL

- Ability for applications to create multiple SSL environments within a single process. An application can now modify environment attributes without terminating any SSL sessions already underway.
- IPv6 Support: This support allows System SSL to be used in an IPv6 network configuration. It also enables System SSL to support both IPv4 and IPv6 Internet protocol addresses.

- *Performance is improved with CRL Caching: Today, SSL supports certificate revocation lists (CRLs) stored in an LDAP server. Each time a certificate needs to be validated, a request is made to the LDAP server to get the list of CRLs. CRL Caching enables applications to request that the retrieved list of CRLs be cached for a defined length of time.*
- *Support for industry standards such as Advanced Encryption Standard (AES), Transport Layer Services (TLS) and Digital Certificate Standard (DCS) Certificates, including PKCS#12.*
- *Enhanced Environment Close: System SSL permits existing connections to remain active and run to completion after their SSL environment has been closed. This removes the requirement that the customer application must manage the SSL environment until all SSL connections have been closed.*
- *Sysplex Session ID Caching: The sysplex session cache support makes server session information available across the sysplex. An SSL session established with a server on one system in the sysplex can be resumed using a server on another system in the sysplex.*
- *System SSL of RSA Private Keys Stored in ICSF: With z/OS 1.4, support is introduced that will allow a certificate's private key to reside in ICSF thus lifting a restriction where the private key had to reside in the RACF database.*
- *LDAP Server now at V3 level. Bulk load utility for populating the LDAP Directory.*
- *The LDAP server supports thousands of concurrent clients, increasing the maximum number of concurrently connected clients by an order of magnitude.*
- *Mandatory Authentication Methods (required by IETF RFC 2829) supported in z/OS 1.4: The CRAM-MD5 and DIGEST-MD5 authentication methods have been added. The methods avoid flowing the user's password over the connection to the server.*
- *TLS: z/OS LDAP now provides support for TLS (Transport Layer Security) as defined in IETF RFC 2830 as an alternative to SSL support.*
- *Modify DN: Today the LDAP Server only supports modification of the most significant portion of a DN (Distinguished Name). z/OS 1.4 has expanded this support to allow any portion of the DN to be changed. This essentially means that relocation of an entry, and possibly a subtree, to a different location in the DIT (directory information tree) is possible.*
- *ACL Updates: Access control has been improved allowing finer-grained access control definitions to be established, providing greater interoperability with other IBM-provided LDAP directory products.*

LDAP

- *Fail-over LDAP provides greater availability: You can now specify a list of Security Server-LDAP servers to be used for storing certificate revocation lists (CRLs). When certificate validation is being performed, this list will be used to determine which LDAP server to connect to for the CRL information.*
- *Client access to information in multiple directories is supported with the LDAP protocol*

ICSF

Also featured as part of Security Management is the Integrated Cryptographic Service Facility, which provides cryptographic functions for data security, data integrity, personal identification, digital signatures, and the management of cryptographic keys. These functions are provided through the combination of secure cryptographic hardware, the ICSF cryptographic API, and the ICSF administration interface.

z/OS supports the PCI Cryptographic Coprocessor feature for the S/390 Parallel Enterprise Server™ G5/G6 and z800/z900 servers. These pluggable cryptographic coprocessor cards provide additional cryptographic capabilities, and more flexible and scalable cryptographic function. ICSF is able to route cryptographic requests to either the CMOS Cryptographic Coprocessor or to a PCI Cryptographic Coprocessor based on workload and capability of the coprocessors. These coprocessors therefore increase e-business capacity, as well as provide an opportunity for more rapid implementation of new cryptographic algorithms.

Initially, z/OS supports the unique cryptography capability of the z990 by providing Clear Key and SSL acceleration for z/OS 1.4 z990 Compatibility Support feature and for z/OS 1.3 via the z990 Cryptographic CP Assist Support for z/OS 1.3 Web download. The Web download is available at ibm.com/servers/eserver/zseries/zos/downloads/. IBM plans to provide Clear Key, SSL acceleration, and Secure Key cryptographic support in 4Q2003 for z/OS 1.4, z/OS.e 1.4, z/OS 1.3, z/OS.e 1.3, z/OS 1.2, and OS/390 2.10. This support will be made available via the z990 Cryptographic Support Web download via the Web at: ibm.com/servers/eserver/zseries/zos/downloads/.

z/OS enhances SSL performance through support of the z800/z900/z990 IBM PCI Cryptographic Accelerator (PCICA), a fast cryptographic processor designed to provide leading edge performance of the complex RSA cryptographic operations used in the Secure Sockets Layer (SSL) protocol. This will benefit functions that use System SSL, such as the z/OS HTTP Server (and WebSphere), TN3270 server, LDAP server and CICS Transaction Gateway server.

PKI Services

PKI Services is a z/OS component that provides a complete Certificate Authority (CA) package for full certificate life cycle management.

- *Component of the z/OS Security Server (RACF)*
- *User request driven via customizable Web pages for browser or server certificates*
- *Automatic or administrator approval process administered via same Web interface*
- *End user / administrator revocation process*
- *Certificate validation service for z/OS applications*
- *With z/OS 1.4 multiple instances can share database via VSAM record-level sharing within a Parallel Sysplex cluster*

Firewall

- *Enhanced Security Server-Firewall Technologies Provide Sysplex-wide Security Association Support: This function enables VPN (virtual private network) security associations to be dynamically re-established on a backup processor in a sysplex when a Dynamic Virtual IP Address (DVIPA) takeover occurs. When the Dynamic Virtual IP Address give-back occurs, the security association will be reestablished on the original processor in the sysplex. When used in conjunction with z/OS Communications Server's TCP/IP DVIPA takeover/give-back capability, this function provides customers with improved availability of IPsec security associations.*

Network Authentication Service

- A component of Security Server, Network Authentication Service, provides authentication, delegation and data confidentiality services which are interoperable with other industry implementations based on the MIT Kerberos V5 reference implementation. Network Authentication Service, administered with RACF commands, supports both the native Kerberos API functions as well as the GSS-API Kerberos security mechanism and does not require DCE.
- IPv6 supported by Kerberos with z/OS 1.4 for improved network security scalability
- Kerberos in z/OS 1.4 provides an alternative database to RACF by offering support for its own registry database using the UNIX System Services NDBM (New Database Manager) support. NDBM provides full Kerberos administration support.

Application Enablement Services

z/OS Version 1 Release 4 base elements

Language Environment C/C++ IBM Open Class® Library
DCE Application Support
Encina® Toolkit Executive

z/OS Version 1 Release 4 optional priced features

C/C++ with Debug Tool
C/C++ without Debug Tool
DFSORT™
GDDM®-PGF V2R1.31
GDDM-REXX V3R21
HLASM Toolkit V1R4

z/OS provides a solid infrastructure in which you can build new applications, extend existing applications, and run existing OLTP and batch processes.

Highlights

The IBM @server brand is about uncompromising flexibility in selecting, building and deploying the applications your business needs. Toward that end, IBM offers the industry's broadest range of platforms and operating systems. IBM is committed to industry-standard, cross-platform technologies — such as Java, XML, HTML, SOAP and UDDI — that are at the heart of a flexible e-business infrastructure. Support for these standards in our key middleware — including DB2 Universal Database, WebSphere Application Server and MQSeries — means you won't be locked in to a single platform as your business grows. As a result, you always have the flexibility to deploy applications in a cost-effective way.

C/C++

- z/OS 1.2 and 1.3 provide two levels of the C/C++ compiler as a migration aid. One is identical to the compiler provided with OS/390 2.10. The other is fully compliant with the ISO C++ standard, also known as the ANSI C++ standard. This includes support for ISO Standard C++ Libraries, and other language features of ISO C++ 1998. New application development involving C++ classes should make use of the C++ Standard Class Library, instead of the C/C++ IBM Open Class Library.
- Enhanced ASCII support provides the ability to produce code that contains ASCII string literals and character constants. This allows ASCII dependent logic to continue working on ASCII platforms, thus eliminating the need to find all such places in the code and converting them to EBCDIC when porting UNIX applications to z/OS.
- The C/C++ Compiler includes support for IEEE Floating Point, and 64-bit long integer format. Support is also added to the C/C++ Runtime Library. Extra Performance Linkage (XPLINK).

Extra Performance Linkage (XPLINK)

When you run a C or C++ application, there is overhead associated with each function call. The more highly functionalized a program, the more overhead. XPLINK cuts down on the overhead associated with these function calls and improves the performance of these applications. In order to exploit the bulk of "high-performance linkage" customers must recompile their C and C++ programs under the new XPLINK environment. The new IBM SDK for z/OS Java 2 Technology Edition V1.4 has been re-written to take advantage of this unique z/OS function.

Java

SDK for z/OS, Java 2 Technology Edition V1.4 provides a full-function Software Development Kit (SDK) at the Java 2 technology level, compliant with the Sun SDK 1.4 APIs.

With SDK for z/OS, Java 2 Technology Edition V1.4, you can:

- *Test and deploy Java applications at the Java 2 SDK 1.4 API level*
- *Continue the "write once, run anywhere" Java paradigm at the Java 2 API level*
- *Take advantage of the new Java 2 function, including XML*
- *Continue to take advantage of:*
 - *Security APIs introduced in the SDK1.3 product*
 - *Java Record Input/Output (JRIO) APIs*
 - *Persistent reusable technology*
- *Leverage traditional zSeries benefits: scalability, reliability, availability, and serviceability*

The program will be available via download from the IBM **@server** zSeries Java Web site and by tape from IBM Software Delivery and Fulfillment (SDF).

Note: SDK for z/OS Java 2 Technology Edition V1.4 requires z/OS 1.4 or later, however look out for release 1.4.1 in June 2003 which will run on z/OS 1.2 or above.

For additional information about zSeries Java products, go to: **ibm.com/servers/eserver/zseries/software/java/**

Unicode

z/OS provide Unicode Callable System Services - code page and case conversions from EBCDIC to Unicode:

- *DB2 V7 is the first exploiter*
- *New hardware instruction on zSeries processor has been implemented to provide superior performance*
- *Unicode Normalization Services allows programmers to decompose or compose characters from another code page and can apply normalization forms to have the same meaning*

REXX Functions

z/OS 1.4 extends the REXX language on z/OS when used in a UNIX System Services MVS REXX environment. It includes functions for standard REXX I/O. A number of I/O stream commands are provided to help control stream processing. Additional functions are also included to easily access some common file services and environments variables. Also included with this package is the ability to interrupt a running REXX program and enter an immediate REXX command such as TS to start interactive tracing.

Provides:

- *Greater interoperability of z/OS as compared to other platforms*
- *Enhancements for porting REXX CGIs and other programs*

Network Communication Services

z/OS Version 1 Release 4 base elements

z/OS Communications Server (Multiprotocol/HPR Services, TCP/IP Services, SNA/APPN Services)

OSA Support Facility

z/OS Version 1 Release 4 optional no charge features

z/OS Communications Server Security Level 3

The z/OS base includes z/OS Communications Server, which enables world class TCP/IP and SNA networking support, including enterprise class dependability; performance and scalability; highly secure connectivity; support for multiple protocols; and efficient use of networking assets.

Highlights

- IPv6 support provides increased scalability and simplified configuration
- HiperSockets, introduced in z/OS 1.2, provides very high-speed, low latency TCP/IP data communications across LPARs resident within the same zSeries server. HiperSockets acts like a TCP/IP network within the server.
- HiperSockets Accelerator provides an “accelerated routing path” which concentrates traffic between OSA-Express external network connections and HiperSockets connected LPARs. This function can improve performance, simplify configuration, and increase scalability while lowering cost by reducing the number of networking adapters and associated I/O cage slots required for large numbers of virtual servers.
- The Sysplex Distributor has been extended to control workload balancing, working in conjunction with Cisco MultiNode Load Balancing (MNLB).
- Host-based Intrusion Detection Services (IDS) complements network-based IDS sensors and scanners by providing defense mechanisms that discard attacking packets before they cause damage, discard packets exceeding established thresholds, and limit the number of connections from “greedy” users.

- OSA-Express Queued Direct Input/Output support (now includes IP Broadcast support, used by DHCP, for IPv4)
- LPAR-to-LPAR communication sharing a single OSA-Express ethernet or Token-Ring port
- Integration of SNA/APPN, TCP/IP and AnyNet®
- A single high-performance TCP/IP stack providing support for both IPv4 and IPv6 applications
- High Performance Native Sockets (HPNS) for TCP/IP applications
- Support for the latest security protocols - SSL & TLS
- Simple Network Time Protocol Support (SNTP) for client/server synchronization
- Multinode Persistent Sessions for SNA applications running in a Parallel Sysplex environment
- SNA 3270, Sockets, and APPC application support
- Native ATM support for high-speed networking
- Easy access to host applications from Java technology-enabled Web browser
- Greater performance and mobility for CICS Sockets

Benefits can include:

- Simplified deployment of client/server applications
- Function for new e-business Internet and intranet applications
- Multivendor, multiplatform connectivity
- High performance, high availability, network choice
- SNA class of service over IP networks
- World class TCP/IP services
- Dramatic improvements in TCP/IP performance include optimization of the TCP/IP stack, and inclusion of a number of performance related capabilities such as:
 - UNIX Sendmail is supported including the POP3 server
 - TN3270 sessions use Secure Sockets Layer, with increased number of supported ports (255)

- *Multi-Node Persistent Sessions (MNPS) includes recovery for Dependent LU Requester (DLUR)-owned LUs. This leads to full Parallel Sysplex exploitation with native IP networks, including nondisruptive session-switching*
- *Client/Server Affinity improvements to allow recovery over sync points in multiphase commit giving higher availability for these applications*
- *Simple Network Management Protocol (SNMP) Distributed Protocol Interface (DPI) instance level registration*
- *X-Motif uses Data Link Libraries (DLL) to reduce disk space*
- *Virtual Private Network (VPN) easier to establish*
- *Addition of Secure Sockets Layer (SSL) and Transport Layer Services (TLS) authentication to the TN3270e server to protect against unauthorized access to SNA applications from TCP/IP clients*
- *Support of Triple DES session encryption for SNA applications*
- *Service Policy Agent for IBM Communications Server Enhancements*
- *Simplified configuration and diagnostics for Enterprise Extender (EE)*

UNIX System Services

z/OS Version 1 Release 4 base elements

z/OS UNIX System Services

Further improvement to TCP/IP sysplex support to increase usability, availability and performance. Sysplex Sockets, XCF Dynamics, and System Symbolics continue to improve performance, provide nondisruptive growth and reduce definition effort. Dynamic Fast Response Web caching within the TCP/IP service stack will provide up to two times improvement in performance. Service Policy Agent offers Priority Networking for TCP/IP, Internet Security with enhanced, further improved Firewall technologies and inclusion of the latest security standards such as IPSec and SNMPv3. Enhanced addressing for TN3270 Server reduces consumption of resources and increase capacity.

Benefits can include:

- *Enhanced Multipath Load Balancing*
- *Support for dynamic updates to Service Policy*
- *Higher availability in Parallel Sysplex environments with Virtual Addressing Takeover (VIPA)*
- *Support of Internet Key Exchange (IKE) protocol to automatically create and distribute encryption keys for dynamic IP clients*

z/OS contains the UNIX application services (shell, utilities and debugger) and the UNIX system services (kernel and runtime environment). The shell and utilities provide the standard command interface familiar to interactive UNIX users. z/OS includes all of the commands and utilities specified in the X/Open Company's Single UNIX Specification, also known as UNIX 95 or XPG4.2.

The z/OS UNIX Services Debugger provides a set of commands that allow a C language program to be debugged interactively. The command set is familiar to many UNIX users. With Language Environment, z/OS supports industry standards for C programming, shell and utilities, client/server applications, and the majority of the standards for thread management and the X/Open Single UNIX Specification. The combination of open computing and z/OS allows the transparent exchange of data, easy portability of applications, cross-network management of data and applications, and the exploitation of traditional MVS system strengths in an open environment.

Highlights

- X/Open UNIX 95 Branded
- Permanent z/OS UNIX Kernel (restart capability provided in z/OS 1.3)
- Integrated/converged z/OS UNIX Sockets
- Web application and UNIX C program performance improvements
- Improved z/OS UNIX setup
- Web threading improvements
- Flexibility improvements — HFS files can be program controlled or APF authorized
- Addition of new UNIX commands
- Multiprocess/MultiUser Kernel Support

Benefits can include:

- Development and execution of UNIX applications — z/OS is a UNIX platform
- Portability of applications to and from other platforms
- Use of UNIX development skills in a z/OS environment
- Consolidate multiple UNIX systems
- Scalability for high growth UNIX applications

UNIX

- Performance enhancements include recompiled and optimized functions within the kernel, and shell and utilities; addition of Socket Functions; use of Communication Storage Management buffer transfer instead of data movement; and optimized NFS Logical File System
- Multiprocess/MultiUser allows faster process creation for customers and reduced storage usage for servers
- Semaphores without contention using the hardware Perform Locked Operation (PLO) instruction
- Shared memory (captured storage) reduces real storage when sharing large amounts of virtual storage
- UNIX System Services and UNIX debugger add support for IEEE-floating point

UNIX Parallel Environment

- Enhanced to support MPI 1.2 specification for C/C++ applications
- Support for Intelligent Miner data with the use of WLM facilities

USS

- Make porting from other platforms easier with support of the Magic Number (#!), double-square-bracket conditional testing, and other functions
- Parallel Sysplex support to share UNIX file systems benefits with Web server applications and others who access the hierarchical file system. This support can make your data and information that reside in the HFS available to your customers at any time, no matter where the applications are running in the Parallel Sysplex environment.
- Shared HFS availability has been improved in z/OS 1.4. If a system goes down, dead system recovery moves file systems that have been defined as AUTOMOVE=YES to another system in the sysplex. z/OS provides a method for customers to indicate where file systems should be moved to. The availability of this function will improve performance and workload balancing.

Distributed Computing Services

z/OS Version 1 Release 4 base elements

Network File System (NFS)

DCE Base Services

Distributed File Service (including zFS and SMB)

- Network File System (NFS) acts as a file server to workstations, personal computers, or other authorized systems in a TCP/IP network. It also provides a z/OS client. The remote data sets or files are mounted from the mainframe (z/OS) to appear as local directories and files on the client system. NFS also provides access to the Hierarchical File System (HFS).
- DCE enables data encryption using the data encryption standard (DES) algorithm and the commercial data masking facility (CDMF)

Highlights

- Remote Procedure Call (RPC) lets calls between programs running on different platforms appear as local procedure calls
- Directory Services allows resources to be found anywhere in an enterprise without the need to know local names
- Security Services helps solve security problems common in a distributed environment by handling identification and certification of users, clients, servers and systems. Distributed Time Services synchronizes clocks running on different nodes.
- Security based on Kerberos authentication with access to RACF data as well

Benefits can include

- Transparency of data and logic
- Distributed, consistent directory service
- Security for both clients and servers, including encryption if required
- Scalability of distributed applications
- Interoperability and portability

zSeries File System (zFS)

zFS is a UNIX file system that can be accessed with the z/OS hierarchical file system file Application Programming Interfaces (APIs). It provides significant performance gains in most customer environments requiring files 8K in size or greater. zFS introduces a new construct called an aggregate which allows space sharing between file systems. HFS does not support space sharing between file systems. zFS also provides a facility to take an almost instantaneous read-only copy of a file system. zFS has two modes of operation:

- *Compatibility mode aggregate: An aggregate defined to only contain a single zFS file system. The VSAM Linear Data Set is used for the name of the zFS file system. The result is an HFS compatibility mode aggregate or file system where the VSAM Linear Data Set cluster name = the zFS Aggregate name = the zFS File System name.*

- *Multi-File system aggregate: An aggregate defined to contain one or more zFS file systems. The extended function in zFS allows multiple zFS file systems to be defined in the same zFS aggregate for space sharing purposes. The zFS aggregate is a VSAM linear data set and the zFS aggregate name is the same as the VSAM linear data set cluster name. Note: It is recommended that compatibility mode aggregates be used for the initial deployment of zFS.*

zFS or zFS-related administration, system management, performance, configuration support, and scalability have been further enhanced in z/OS 1.4:

- *Dynamic reconfiguration for file system configuration options*
- *Dynamic use of secondary allocation for a zFS aggregate (data set) or file system*
- *Use of file systems in different zFS aggregates with the same name. This allows administrators to restore aggregate data and simultaneously mount a file system from the new and original aggregate to recover file system data.*
- *Use of Static system symbol substitutions providing system-specific information to be specified and enables the configuration to be shared between systems in a sysplex*
- *Improvements in the UNIX System Services automount support for zFS*
- *Addition of ISHELL support for zFS*
- *zFS aggregate awareness in UNIX System Services. This allows the automove of zFS file systems in the same zFS aggregate (data set) to the same system in a sysplex when the owning system goes down.*
- *Ability to perform I/O operations in parallel for a zFS aggregate that spans multiple DASD volumes. This provides improved performance when using multi-DASD volume aggregates.*
- *Support for 64-bit user virtual buffer addresses*

Distributed File Services (DFS) SMB

Microsoft® Windows® networking compatible file and print serving is available with Native SMB File and Print Serving for Windows Clients. SMB file serving enables z/OS to share HFS, zFS, Sequential files and Record Files Systems (RFS) such as PDS, PDS/E or VSAM data sets with Windows workstations. SMB can automatically handle the conversion between ASCII and EBCDIC making full use of USS file tagging Access Control Lists (ACLs) support. This allows development of applications on Windows for deployment on z/OS. z/OS also supports printing of SMB files without requiring that code be installed on the clients and without requiring unique printer setup on the workstations. OS/2® client file serving support is also available.

SMB Support has been further enhanced in z/OS 1.4 by:

- *Simplified User administration with Windows Domain ID mapping*
- *Performance improvements with RFS and large PDS or PDS/E file systems*
- *Exploitation of zFS performance*

e-business Services

z/OS Version 1 Release 4 base elements

IBM HTTP Server
Text Search

z/OS Version 1 Release 4 optional no charge features

IBM HTTP Server North America Secure

Businesses increasingly use the Internet to market products and conduct business with suppliers and customers. WebSphere Application Server for OS/390 (separately orderable) enables the use of z/OS as a Web server with the benefits of security, the utilization of large storage capacity, centralized skills, a single point of entry and control, consolidation of multiple Web sites, and secure Internet transactions.

Text Search is a powerful, full-text indexing and search server that supports high-speed searching of Web sites and other documents stored on the z/OS server. Free text searching is supported, as well as Boolean logic. Text Search returns a ranked list of hits. A full API set is provided so that Text Search functions can be used in other programs and products. It can also be used for Web-usage mining.

The IBM HTTP Server offers HTTP 1.1 compliance, support for Java technology, and the ability to manage Internet processing through the Workload Manager (WLM). In addition, it provides:

- *Automatic browser detection to respond to requests with the version of a Web page or document appropriate for that browser*
- *Dramatic improvements in single and multiple processor environments, coupled with significant CPU requirement reductions, provide higher throughput and shorter response time*
- *Platform for Internet Content Selection (PICS) provides a way for users to filter material they encounter on the Internet and accept or reject the material based on its rating*
- *Web usage statistics*

Benefits can include

- *e-business with security*
- *Utilization of large storage capacity*
- *Single point of entry and control*
- *Consolidation of multiple Web sites*
- *Secure Internet transactions*
- *Exploitation of z/OS WLM*

Print Services

z/OS Version 1 Release 4 optional priced features

Infoprint® Server

– IP PrintWay™

– NetSpool™

– z/OS Print Interface

Today, businesses are looking for better ways to get information to workers throughout the enterprise, and to suppliers, partners and customers around the globe. Application output that was once printed on centralized data center printers and mailed must now be electronically distributed and printed, or presented over the Web. Infoprint Server provides a reliable, highly available, secure and scalable foundation for your e-business output infrastructure. Infoprint Server and its companion product, Infoprint Server Transforms, include a print interface, printer inventory, VTAM® application output capture program, and print drivers and management tools that let you manage any print job to any printer defined to Infoprint Server.

These capabilities give you the flexibility to deliver the output of e-business anywhere you need it:

- *Legacy CICS and IMS applications that generate SNA Character String (SCS) or 3270 output formats can print to LAN-attached PCL printers, without changes to the application program*
- *Support for the output of e-business includes the ability to send output as e-mail instead of, or in addition to print*
- *A consolidated printer inventory lets you define all printers used with Infoprint Server, and printers driven by Print Services Facility (PSF) in one place. Printers can be defined and modified from a single easy-to-use interface.*

- *IP PrintWay provides support for printers attached to the network using TCP/IP, VTAM-controlled coax printers, and for printers and servers over the Internet using the industry standard Internet Printing Protocol (IPP). Easy to use ISPF menus also enable management of distributed printers.*
- *The Print Interface supports print submission from applications running in UNIX System Services (USS), from Windows users via native Windows SMB, from applications on other servers, and over the Internet using IPP.*
- *Data stream transforms let you print AFP™ applications on printers using PCL, PostScript or PDF. You can also print PCL, PostScript and PDF output on AFP printers.*
- *A transform from SAP to AFP and a certified SAP Output Management System lets you print SAP application output on your fast, reliable AFP printers, and receive print completion notification back at the SAP Application Server.*

Benefits of consolidating your enterprise printing onto z/OS using Infoprint Server can include:

- *Reduced total cost of ownership for distributed print operations*
- *Improved productivity with simplified print operations and management*
- *Investment protection and leverage for your AFP applications and printers*
- *Faster deployment of e-business initiatives with flexible output delivery options*

Softcopy Publications Support

z/OS Version 1 Release 4 base elements

BookManager® READ V3 BookManager Bookserver
GDDM (including PLCK and OS/2 Link and REXX code) V3R2

z/OS Version 1 Release 4 optional priced features

BookManager Build

Test

To further reduce migration time, z/OS is system integration-tested using a production-like environment. The z/OS environment includes subsystems such as CICS, IMS, DB2 and WebSphere. This additional testing supplements existing functional tests, with a focus on tasks performed by customers in the production environment, thus helping establishments move more quickly to new functions.

LookAt

A convenient way to look up the meaning for over 600,000 different messages. LookAt provides information for OS/390, z/OS, z/VM, VSE and is available via the Internet at ibm.com/servers/s390/os390/bkserv/lookat/lookat.html or on your z/OS or OS/390, VM, AIX or Windows systems. In addition there is support for Palm VIIx (applet) or PocketPC (uses LookAt site).

Publications

For a list of the publications available for z/OS, visit the z/OS library Web site at: ibm.com/servers/eserver/zseries/zos/bkserv.

Installation Considerations

z/OS 1.4 provides a strong foundation for exploitation of further zSeries architecture enhancements. Customers are encouraged to migrate to this release so that they are well positioned to take full advantage of what zSeries has to offer. Customer are able to upgrade from OS/390 2.10

or any previous release of z/OS and are able to order z/OS via Customized Offerings (ServerPac, SystemPac® and CBPDO) for a complete system. For more information about SystemPac, visit ibm.com/ca/custompac.

Customers are strongly encouraged to implement WLM goal mode before installing z/OS 1.4 rather than concurrent with their 1.4 installation. Goal mode is required for z/OS 1.3 and later.

z/OS 1.4 and 1.5 will be supported on the following IBM servers:

- *S/390 Parallel Enterprise Server G5/G6 (or compatible server) (31-bit mode)*
- *Multiprise 3000 Enterprise Server (or compatible server) (31-bit mode)*
- *IBM @server zSeries 800, 900 and 990 servers (or compatible servers) (64-bit mode*)*

(* 31-bit mode for z/OS 1.2, z/OS 1.3 and z/OS 1.4 is supported on zSeries servers for up to 6 months via the Bimodal accommodation offering. For more information see ibm.com/servers/eserver/zseries/zos/installation/bimodal.html.)

Statement of Direction: z/OS 1.6, planned for September 2004 establishes an architectural level set, requiring z/Architecture hardware.

Migration/Coexistence

IBM has converged on a consistent migration and coexistence policy. This consistent migration and coexistence policy is based on the current coexistence policy. Migration forward as well as backward should be made within the same releases supported by the coexistence policy. Four releases is the general migration and coexistence policy that should be assumed, except where special provisions have been provided.

This consistent coexistence, migration and fallback policy applies to release migrations for:

- *Single system configurations*
- *Individual systems within a multisystem configuration*
- *Cases where a simultaneous IPL is used to migrate all systems in a multisystem configuration at the same time.*

See section Planned Changes to the Migration/Coexistence Policy for more information on changes to the Coexistence-Migration-Fallback policy that IBM intends to make with z/OS 1.6 (release planned to be available in September 2004).

For additional information on the current coexistence and release migration information, refer to *z/OS Planning for Installation* (GA22-7504) at: ibm.com/servers/eserver/zseries/zos/bkserv/find_books.html

The following table describes the migration/coexistence supported releases. The release shown in column 1 is the highest release running in a multisystem configuration.

Migration/Coexistence Supported Releases (OS/390 & z/OS)

Release	Migration/Coexistence Release
R10	R10
z/OS R1 ²	{z/OS R1 OS/390 R10} ¹
z/OS R2	z/OS R2, {z/OS R1 ² , OS/390 R10} ¹
z/OS R3	z/OS R3, z/OS R2, {z/OS R1 ² , OS/390 R10} ¹
z/OS R4	z/OS R4, z/OS R3, z/OS R2, {z/OS R1 ² , OS/390 R10} ¹

1. OS/390 R10 and z/OS R1 treated as one coexistence level
2. z/OS R1 is not supported on z990 and cannot coexist in a sysplex which includes a z990

This information is being provided to you to facilitate release planning and to ensure appropriate positioning for future software migrations. z/OS 1.4 is expected to be orderable until 1Q2004.

Since z/OS allows the JES element to be separately staged, z/OS also supports the coexistence of certain lower-level JES releases with the JES release provided with z/OS.

Note: Specific functions might only be available on the up-level systems, or it might be necessary to up-level all systems to enable some functions.

For additional information on z/OS planning, refer to *z/OS Planning for Installation*, available on the z/OS Web site in the z/OS Library at: ibm.com/servers/eserver/zseries/zos

For the latest information regarding z/OS and OS/390 marketing and service withdrawal dates visit: ibm.com/servers/eserver/zseries/zos/support/zos_eos_dates.html

Planned Changes to Migration/Coexistence policy:

IBM intends to align the Coexistence-Migration-Fallback policy with the Service policy. The intention with this policy change is to simplify and provide greater predictability to aid in release migrations. IBM intends to continue with the practice of providing service support for each release of z/OS or z/OS.e for three years following its general availability (GA) date. IBM, at its sole discretion, may choose to leave a release supported for more than three years.

This change to the Coexistence-Migration-Fallback policy is planned to be effective starting with z/OS 1.6. In some cases, more than three releases may be coexistence, migration, and fallback supported, if IBM at its sole discretion chooses to provide service support for greater than three years for a release. However, as an exception, any z/OS or z/OS.e release having three or fewer months of service remaining at the time of GA of a new release will not be coexistence, migration, and fallback supported.

For z/OS 1.5 and 1.4, and z/OS.e 1.5 and 1.4, and lower releases, the existing Coexistence-Migration-Fallback policy (and any accommodations previously provided) remains in effect and is unchanged as documented above.

Wizards

By using your browser on the Internet you can get help to set up some of the more complex functions that you do infrequently.

Our wizards are interactive assistants that ask you a series of questions about the task you want to perform (for example, setting up a Parallel Sysplex). The wizards simplify your planning and configuration needs by exploiting recommended values and by building customized checklists for you to use. For configuration tasks, our wizards also generate outputs like jobs, policies, or parmlib members that you can upload to z/OS and use.

For more information about wizards, visit the Web site at: ibm.com/servers/eserver/zseries/zos/wizards/

Order z/OS through the Internet

ShopzSeries (formerly SHOPS390) provides an easy way to plan and order your z/OS ServerPac or CBPDO. It will analyze your current installation, determine the correct product migration, and present your new configuration based on z/OS. Additional products can also be added to your order (including determination of whether all product requisites are satisfied).

IBM Service is available through ShopzSeries as an Internet-based, strategic software tool available to fulfill orders 24x7 (365 days a year). Through Service in ShopzSeries, customers can quickly and easily order and receive corrective or preventive service electronically over the Internet or by standard physical media. For preventive service, customers can order just critical service (HIPERs and PTFs that resolve PE PTFs), the latest recommended service (which includes all critical service), or all available service. Service in ShopzSeries reduces customers' research time and effort by using their uploaded SMP/E Consolidated Software Inventory (CSI) so that all applicable service, including reach ahead service, for the installed FMIDs in the target zones is selected. ShopzSeries also uses the CSI information to limit the size of the service order, only sending applicable service that has not already been processed by SMP/E. IBM Technical Support is available through RETAIN for problems and Q&A.

ShopzSeries is available in the U.S. and several countries in Europe. In countries where ShopzSeries is not available yet, please contact your IBM representative (or Business Partner) to handle your order via the traditional IBM ordering process.

For more details and availability, visit the ShopzSeries Web site at: ibm.com/software/shopzSeries

z/VM

The infrastructure for e-business on demand™ environments can be deployed with visionary technology on z/VM, including advanced virtualization, support for open-source software, and autonomic computing enhancements. With virtualization technology as its foundation, z/VM continues to provide new function and technology exploitation on the mainframe that enables you to virtualize processor, communication, memory, storage, I/O, and networking resources, with the potential to help reduce the need to plan for, purchase, and install hardware to support new workloads.

z/VM supports Linux, one of the world's leading Open Source operating system, on the mainframe. Within the VM environment, Linux images can share hardware resources and use internal high-speed communications. While benefiting from the reliability, availability and serviceability of zSeries servers, z/VM V4 offers an ideal platform for consolidating select UNIX, Windows, and Linux workloads on a single physical zSeries server, which allows you to run tens to hundreds of Linux images. z/VM V4 is priced on a per-engine basis and supports IBM Integrated Facility for Linux (IFL) processor features for Linux-based workloads, as well as standard engines for all other zSeries and S/390 workloads.

z/VM V4 is the follow-on product for VM/ESA and z/VM V3. It provides additional support and exploitation opportunities for the thousands of users who have built enterprise-wide automation and infrastructure enhancements on the VM platform in support of their applications, database systems, and e-business solutions.

z/VM Version 3 (V3)

z/VM offers a new technology base for customers looking to use IBM Virtual Machine technology on one of the industry's best-of-breed server platform, zSeries. When z/VM is running on a zSeries server, it's possible to run 64-bit capable OS/390 2.10, z/OS, and Linux for zSeries as guest

systems of z/VM, in addition to ESA/390 guest operating systems such as OS/390, VSE/ESA, TPF, and Linux for S/390. To operate z/OS as a guest of z/VM on a zSeries server, z/VM must be operating in 64-bit mode. z/VM will allow customers to develop and test their 64-bit Parallel Sysplex applications in a guest environment before putting them into production. This may help reduce the need to invest in separate standalone configurations.

z/VM also reduces storage constraints by eliminating the 2 GB central storage limitation, providing plenty of headroom for increasing e-business demands and growing back-office applications within a single machine image. Customers experiencing real memory constraints can experience relief by running z/VM in 64-bit mode on a zSeries server. This constraint relief is provided for ESA/390 guest operating systems. Support for large real memory with z/VM may benefit customers running a large number of Linux for zSeries and S/390 guest systems.

Additional enhancements in z/VM V3 include:

- *Native FlashCopy for Enterprise Storage Server (ESS) for high-speed data copy*
- *Guest support enhancements for 3494 VTS and FICON attached 3590 A60 Tape Controller*

Connectivity enhancements for TCP/IP Feature for z/VM:

- *Improved security with the inclusion of a Secure Socket Layer (SSL) server*
- *Transparent data access to remote systems data with an NFS Client*
- *Capability and usability improvements to FTP server for Web browsers*
- *Reduced load on hosts with support for IP Multicasting*
- *Improved data transfer performance with QDIO supporting Gigabit Ethernet, Fast Ethernet, and 155 ATM (Ethernet LAN Emulation)*
- *Support for the DFSMS/MVS® Program Management binder and loader functionality to enhance application affinity between CMS and OS/390 or z/OS*

z/VM Version 4 (V4)

With z/VM and the IBM Integrated Facility for Linux (IFL), a low-cost, flexible environment is created to test and develop on Linux while running Linux production applications on IBM z990, z900, z800, S/390 Parallel Enterprise Server Generation 5 and 6, S/390 Multiprise 3000, or equivalent servers. z/VM V4 support for IFL processor features is designed to run Linux workloads without increasing the IBM software charges for z/OS, z/OS.e, OS/390, VM, VSE/ESA, or TPF operating systems and applications running on standard engines of the z990, z900, z800, and S/390 servers in other logical partitions.

Engine-based pricing for z/VM V4 and its optional features allows customers to exploit the zSeries and S/390 servers more cost effectively than a discrete server implementation. With engine-based pricing, customers pay a one-time software license charge (OTC) for each processor engine. This can be for standard processor engines or IFL processor features. Traditional operating systems such as z/OS, z/OS.e, OS/390, TPF, VSE/ESA, z/VM 3.1, or VM/ESA are not supported nor can they operate on IFL processor features. Only Linux workloads in an LPAR or Linux guests of z/VM V4 can operate on the IFL processor feature.

z/VM provides the capability to account for the use of system resources by virtual machines, including those running Linux. Accounting records are produced that track a virtual machine's use of processor, paging, I/O, and virtual network resources, including virtual channel-to-channel adapters (CTCAs), inter-user communication vehicle (IUCV) or advanced program-to-program (APPC) connections, and virtual (Guest LAN) network interface cards (NICs).

z/VM 4.4 extends its virtualization technology in support of Linux and other guests while providing some enhancements that enable z/VM to be self-optimized and self-managed:

- *Reducing contention for the z/VM Control Program (CP) scheduler lock may help increase the number of Linux and other guest virtual machines that can be managed concurrently.*
- *Enhancing the Virtual Machine Resource Manager (VMRM) to provide the infrastructure necessary to support more extensive workload and systems resource management features by providing:*
 - *Monitor data showing actual workload achievement*
 - *An interface to dynamically change users in workloads, workload characteristics, and goals*
 - *More flexibility using the VMRM configuration file when managing multiple users*
 - *Improvements in the reliability and performance of the VMRM service virtual machine's monitor data handling*
 - *Serviceability enhancements including improved messages, logfile entries, and new server options*
- *Simulating virtual FICON CTCA devices for guest operating systems enhances previous virtual-CTCA support by adding the FICON protocol as an option for guest operating systems. Guests use virtual CTCAs to communicate among themselves within a single z/VM system image, without the need for real FICON CTCAs.*
- *Supporting real and virtual integrated 3270 console devices. Real support enables this device, provided by the Hardware management Console, to be used as the system operator console. Virtual support enables testing of guest operating systems and utilities such as Stand-Alone Program Loader (SAPL) and stand-alone DASD Dump-Restore (DDR), that support the integrated 3270 console device.*
- *Delivering the Performance Toolkit for VM™ to process Linux performance data obtained from the Resource Management Facility (RMF) Performance Monitoring (PM) client application, rmfpm. Linux performance data obtained from RMF is presented on display screens and in printed reports similarly to the way VM data is presented.*

With corresponding function available in Linux on zSeries and S/390, z/VM 4.4 provides:

- *The attachment of Small Computer System Interface (SCSI) devices to guest Linux images using Fibre Channel Protocol (FCP) channels on zSeries processors*
- *IPL from FCP-attached disks for Linux and other guest operating systems with necessary SCSI support, when z/VM is running on a z990, z900, or z800 server equipped with this function*
- *Enhanced page-fault handling*
- *Clear-key RSA functions of the IBM PCI Cryptographic Coprocessor (PCICC) or the IBM PCI Cryptographic Accelerator (PCICA)*

Exploiting New Technology

z/VM provides a highly-flexible test and production environment for enterprises deploying the latest e-business solutions. Enterprises that require multi-system server solutions will find that z/VM helps them meet the demands of their businesses and IT infrastructures with a broad range of support for such operating system environments as z/OS, z/OS.e, OS/390, TPF, VSE/ESA, CMS, and Linux for zSeries and S/390. The ability to support multiple machine images and architectures enables z/VM to run multiple production and test versions of zSeries and S/390 operating systems, all on the same system. z/VM can help simplify migration from one release to another, facilitate the transition to newer applications, provide a test system whenever one is needed, and consolidate several systems onto one physical machine. z/VM can also be used to provide virtual access to the latest DASD and processor architecture for systems that lack such support. New technological enhancements in z/VM 4.4 provide:

- *Exploitation of the IBM @server zSeries 990 server*
 - *Extending Dynamic-I/O configuration support allows channel paths, control units, and devices to be dynamically added, changed, and deleted in a multiple Logical Channel Subsystems (LCSS) environment.*

- *Support for extended I/O-measurement facilities provides improved capacity planning and I/O performance measurement*
- *Handling I/O-configuration definition and dynamic-I/O configuration in an environment of up to 30 LPARs, in October 2003, an increase from the previous limit of 15*
- *Support for the zSeries capability to cascade two FICON directors within a Fibre-Channel fabric. z/VM and its guests can take advantage of this enhanced and simplified connectivity, which is particularly useful in disaster-recovery and business-continuity situations.*
- *Support for the IBM TotalStorage Enterprise Storage Server (ESS) FlashCopy V2 providing increased flexibility for improved capacity management and utilization*
- *Support for the IBM ESS Peer-to-Peer Remote Copy Extended Distance (PPRC-XD) function, extending the distance, well beyond the 103 km supported with PPRC synchronous mode. PPRC-XD is suitable for data migration, backup, and disaster recovery procedures. PPRC Version 2 (V2) is also supported for guest operating systems, offering an asynchronous cascading solution providing a complete, consistent, and coherent copy of data at a remote site.*

Systems Management

Improvements in systems management some of which help to provide self-configuring, self-managing, and self-optimization facilities in z/VM 4.4 include:

- *Functions that may be called by client applications to allocate and manage resources for guests running in z/VM virtual machines (virtual images). Use of the application programming interfaces (APIs) through an application provided by a customer or solution provider are designed so that such applications can allow administrators who lack in-depth VM knowledge to manage a large number of virtual images, running in a single z/VM system.*

- *Hardware Configuration Manager (HCM) and Hardware Configuration Definition (HCD) components to create and manage your I/O configuration. This new support provides a comprehensive, easy-to-use I/O-configuration-management environment similar to that available with the z/OS operating system.*
- *Performance Toolkit for VM that provides enhanced capabilities for a z/VM systems programmer, operator, or performance analyst to monitor and report performance data. The toolkit is an optional, per-engine-priced feature derived from the FCON/ESA program (5788-LGA), providing:*
 - *Full-screen mode system console operation and management of multiple z/VM systems*
 - *Post-processing of Performance Toolkit for VM history files and of VM monitor data captured by the MON-WRITE utility*
 - *Viewing of performance monitor data using either Web browsers or PC-based 3270 emulator graphics*

The toolkit also provides the capability to monitor TCP/IP for z/VM, as well as to process Linux performance data.

Application Enablement

CMS will host a new C/C++ for z/VM compiler (5654-A22). This environment allows C/C++ programs to be compiled and executed on CMS and creates portability between z/VM and z/OS C/C++ programs. C/C++ source files can be read from a CMS minidisk, the SFS, or the Byte File System (BFS) and output can be written to any of these file systems. C/C++ will only execute on z/VM 4.4 and can only be licensed to operate on standard processor engines. In order to support the C/C++ for z/VM compiler, the C/C++, the Language Environment has been updated to the level shipped with z/OS 1.4 and is integrated into the base of z/VM 4.4.

Networking with z/VM

TCP/IP for z/VM, formerly a priced, optional feature of VM/ESA and z/VM V3, is packaged at no additional charge and shipped enabled for use with z/VM V4. The former priced, optional features of TCP/IP — the Network File System (NFS) server and TCP/IP source — are also packaged with TCP/IP for z/VM at no additional charge.

TCP/IP for z/VM delivers expanded Internet/intranet access, improved e-business performance and extended function. Performance of the TCP/IP stack was enhanced by redesigning algorithms to reduce path lengths, recoding procedures to optimize high-use paths, identifying and implementing performance improvement items, and adding virtual multiprocessing capabilities.

TCP/IP is designed to support the z/Architecture HiperSockets function for high-speed communication among virtual machines and logical partitions within the same zSeries server. The HiperSockets function allows virtual machines and logical partitions to communicate internally over the memory bus using the internal-queued-direct (IQD) channel type in the z990, z900, and z800. TCP/IP broadcast support is now available for the HiperSockets environment when utilizing Internet Protocol version 4 (IPv4) with z/VM 4.4. Applications that use the broadcast function can now propagate frames to all TCP/IP applications.

The z990 servers include an important performance enhancement that virtualizes adapter interruptions and can be used with V=V guests (pageable guests) on z/VM 4.4. With the enhancement of the TCP/IP stack in z/VM 4.4 to use adapter interruptions for OSA-Express, TCP/IP for VM can benefit from this performance assist for both HiperSockets and OSA-Express adapters.

z/VM 4.4 exploits the Virtual Local Area Network (VLAN) technology. VLANs ease the administration of logical groups of users so that they can communicate as if they were on the same physical LAN. VLANs help increase traffic flow and may help reduce overhead to allow the organization of networks by traffic patterns rather than by physical location. To support VLAN, z/VM 4.4 provides:

- *Enhancements to TCP/IP for z/VM to enable membership in a VLAN for QDIO and HiperSockets adapters*
- *Enhancements to z/VM Guest-LAN simulation to allow virtual QDIO and HiperSockets adapters to participate in a VLAN*
- *Management and control of VLAN topology by the z/VM virtual switch*

The Guest LAN support provided in z/VM 4.2 simulates the HiperSockets function for communication among virtual machines without the need for real IQD channels, much as VM simulates channel-to-channel adapters for communication among virtual machines without the need for ESCON, FICON, or other real channel-to-channel connections. With the Guest LAN capability, customers with S/390 servers can gain the benefits of HiperSockets communication among the virtual machines within a VM image, since no real IQD channels are required.

z/VM 4.4 further enhances its virtualization technology by providing the capability to deploy virtual IP switches in the Guest LAN environment. The z/VM virtual switch replaces the need for virtual machines acting as routers to provide IPv4 connectivity to a physical LAN through an OSA-Express adapter. Routers consume valuable processor cycles and require additional copying of data being transported. The virtual-switch function alleviates this problem and also provides centralized network configuration and control. These controls allow the LAN administrator to more easily grant and revoke access to the network and to manage the configuration of VLAN segments.

TCP/IP for z/VM provides numerous self-protection functions. A Secure Sockets Layer (SSL) server is available to facilitate secure and private conversations between z/VM servers and external clients. The upgraded SSL server in z/VM 4.4 provides appropriate Red Hat Package Manager (RPM) packages for the SuSE Linux 2.4.7 SLES 7 and 2.4.19 SLES 8 distributions. Security of the TCP/IP stack has been improved to prevent additional types of Denial of Service (DoS) attacks including: Smurf, Fraggle, Ping-of-Death, Kiss of Death (KOD), KOX, Blat, SynFlood, Stream, and R4P3D. The overall security and auditability of the TCP/IP for z/VM stack and the integrity of the z/VM system have been improved by providing better controls, monitoring, and defaults. An IMAP user authentication exit has been added that removes prior user ID and password length restrictions and eliminates the need for every IMAP client to have a VM user ID and password.

In addition to the new function provided by the Performance Toolkit for VM, RealTime Monitor (RTM), and Performance Reporting Facility (PRF) are still available to support new and changed monitor records in z/VM. RTM simplifies performance analysis and the installation management of VM environments. PRF uses system monitor data to analyze system performance and to detect and diagnose performance problems. RACF for z/VM is available as an priced, optional feature of z/VM V4 and provides improved data security for an installation. RTM, PRF, and the Performance Toolkit are also priced, optional features of z/VM V4 as is the Directory Maintenance Facility (DirMaint™).

For further information see the *z/VM Reference Guide*, GM13-0137.



To learn more

Visit the zSeries World Wide Web site at ibm.com/eserver/zseries or call IBM DIRECT at 1 800 IBM-CALL in the U.S. and Canada.

Australia	132 426
Austria	0660.5109
Belgium	02-225.33.33
Brazil	0800-111426
China	(20) 8755 3828
France	0800-03-03-03
Germany	01803-313233
Hong Kong	(20) 2825 6222
Hungary	165-4422
India	(80) 526 9050
Indonesia	(21) 252 1222
Ireland	1-850-205-205
Israel	03-6978111
Italy	167-017001
Japan	0120 300 426
Korea	(02) 781 7800
Malaysia	(03) 717 7890
Mexico	91-800-00316
Netherlands	020-513.5151
New Zealand	0800-801-800
Philippines	(02) 819 2426
Poland	(022) 878-6777
Singapore	1800 320 1975
South Africa	0800-130130
Spain	900-100400
Sweden	020-220222
Switzerland	0800 55 12 25
Taiwan	(06) 2725 9300
Thailand	(02) 273 4444
Vietnam Hanoi	(04) 843 6675
Vietnam HCM	(08) 829 8342
United Kingdom	0990-390390

Copyright IBM Corporation 2003

Integrated Marketing Communications, Server Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America

05/03

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, IBM eServer, the IBM logo, the e-business logo, AFP, AIX, AnyNet, APPN, BookManager, CICS, DB2, DB2 Connect, DB2 Universal Database, DFSMSdftp, DFSMSdss, DFSMSshm, DFSMSrmm, DFSMS/MVS, DFSORT, DirMaint, e-business on demand, Encina, Enterprise Storage Server, Enterprise Tape Controller, ESCON, FICON, FICON Express, FlashCopy, FFST, GDDM, GDPS, Geographically Dispersed Parallel Sysplex, HiperSockets, HiperSpace, HyperSwap, IMS, IMS/ESA, InfoPrint, Intelligent Miner, IP PrintWay, iSeries, Language Environment, MQSeries, Multiprise, MVS, Netfinity, NetSpool, NetView, Open Class, OS/2, OS/390, Parallel Sysplex, Performance Toolkit for VM, PR/SM, Processor Resource/Systems Manager, pSeries, RACF, Resource Link, RMF, RS/6000, S/390, S/390 Parallel Enterprise Server, SecureWay, Sysplex Timer, SystemPac, Tivoli, Tivoli Storage Manager, TotalStorage, Virtual Image Facility, VM/ESA, VSE/ESA, VTAM, WebSphere, xSeries, z/Architecture, z/OS, z/VM, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Intel is a trademark of the Intel Corporation in the United States and other countries.

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.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

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.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

GM13-0229-00