

IBM @server zSeries 890 and z/OS Reference Guide



Table of Contents

<i>z/Architecture</i>	<i>page 5</i>
<i>IBM @server zSeries 990</i>	<i>page 6</i>
<i>z890 Family Models</i>	<i>page 10</i>
<i>z890 and z900 Performance Comparison</i>	<i>page 12</i>
<i>z890 I/O SubSystem</i>	<i>page 14</i>
<i>z890 Channels and I/O Connectivity</i>	<i>page 16</i>
<i>Fibre Channel Connectivity</i>	<i>page 17</i>
<i>Open Systems Adapter-Express Features</i> <i>(OSA-Express)</i>	<i>page 22</i>
<i>HiperSockets</i>	<i>page 26</i>
<i>Cryptography</i>	<i>page 28</i>
<i>Availability</i>	<i>page 29</i>
<i>Advanced Availability Functions</i>	<i>page 31</i>
<i>Parallel Sysplex Cluster Technology</i>	<i>page 33</i>
<i>z890 Support for Linux</i>	<i>page 47</i>
<i>zSeries 890 Family Configuration Detail</i>	<i>page 51</i>
<i>Physical Characteristics</i>	<i>page 53</i>
<i>Coupling Facility - CF Level of Support</i>	<i>page 54</i>
<i>z/OS</i>	<i>page 56</i>
<i>z/VM</i>	<i>page 79</i>
<i>VSE/ESA</i>	<i>page 87</i>
<i>To Learn More</i>	<i>page 88</i>

zSeries Overview

Technology has always accelerated the pace of change. New technologies enable new ways of doing business, shifting markets, changing customer expectations, 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, 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 an e-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 must be capable of employing 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 change 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 890 (z890), brings enriched functions that are required for the on demand data center to a larger audience.

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, helping to ensure that upgrading the servers is timely and meets the needs of your business. It's not just the capacity of the servers that can be changed on demand, but the mix of workload and the allocation of resources can be changed 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® technology 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, as the resources required by different applications, middleware 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 designed to be resilient. The zSeries has been renowned for reliability and availability. The zSeries platform will help protect from both scheduled and unscheduled outages, and GDPS® automates site reconfiguration for both planned and unplanned outages.

The New zSeries from IBM

IBM's ongoing investment in zSeries technology produced a re-invention of the zSeries server — the z990 which was developed at a cost in excess of \$1 Billion. The new technology introduced a host of new benefits to meet today's on demand business.

The z890 takes advantage of the new z990 technology in a package that help deliver excellent price/performance for those customers requiring a lower capacity entry point and more granularity of upgrades than offered with the z990 but still delivering the ability to add incremental capacity as it is required in an on demand world.

And IBM's investment in zSeries doesn't stop here. To solidify the commitment to zSeries, IBM introduced the "Mainframe Charter" that provides a framework for future investment and a statement of IBM's dedication to deliver ongoing value to zSeries customers in their transformation to on demand business.

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 (IRD). The z/Architecture is a 64-bit superset of ESA/390.

z/Architecture is implemented on the z890 to allow full 64-bit real and virtual storage support. A maximum 32 GB of real storage is available on z890 servers. z890 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 help 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 and z/OS.e do 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 890

Operating System	ESA/390 (31-bit)	z/Arch (64-bit)	Compatibility	Exploitation
OS/390® V2R10	Yes	Yes	Yes	No
z/OS V1R2	No*	Yes	Yes	No
z/OS and z/OS.e V1R3	No*	Yes	Yes	No
z/OS and z/OS.e V1R4	No*	Yes	Yes	Yes
z/OS and z/OS.e V1R5, 6	No	Yes	Included	Included
Linux on S/390®	Yes	No	Yes	Yes
Linux on zSeries	No	Yes	Yes	Yes
z/VM V3R1	Yes	Yes	Yes	No
z/VM V4R2, 3	Yes	Yes	Yes	No
z/VM V4R4	Yes	Yes	Included	Yes
z/VM V5R1 (3Q04)	No	Yes	Included	Yes
VSE/ESA™ V2R6, 7	Yes	No	Yes	Yes
z/VSE V3R1	Yes	No	Yes	Yes
TPF V4R1 (ESA mode only)	Yes	No	Yes	No

* Customer with z/OS Bimodal Migration Accommodation Offering may run in 31-bit support per the terms and conditions of the Offering. Bimodal Offering available for z/OS ONLY.

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 z890 server, the latest addition to the zSeries family, takes advantage of the design enhancements and new technology of the z990 in a package that helps deliver excellent price/performance for those customers requiring a lower capacity entry point than that offered by the z990. With a new superscalar microprocessor and the CMOS 9S-SOI technology, the z890 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 traditional zSeries Quality of Service to emerging e-business applications such as WebSphere®, Java and Linux.

The new IBM @server zSeries Application Assist Processor (zAAP), planned to be available on the IBM @server zSeries 990 (z990) and zSeries 890 (z890) servers, is an attractively priced specialized processing unit that provides an economical Java execution environment for customers who desire the traditional Qualities of Service and the integration advantages of the zSeries platform.

When configured with general purpose processors within logical partitions running z/OS, zAAPs may help increase general purpose processor productivity and may contribute to lowering the overall cost of computing for z/OS Java technology-based applications. zAAPs are designed to operate asynchronously with the general processors to execute Java programming under control of the IBM Java Virtual Machine (JVM). This can help reduce the demands and capacity requirements on general purpose processors which may then be available for reallocation to other zSeries workloads.

The IBM JVM processing cycles can be executed on the configured zAAPs with minor or no anticipated modifications to the Java application(s). Execution of the JVM processing cycles on a zAAP is a function of the Software Developer's Kit (SDK) 1.4.1 for zSeries, z/OS 1.6, and the Processor Resource/Systems Manager™ (PR/SM™).

The amount of general purpose processor savings will vary based on the amount of Java application code executed by zAAP(s). This is dependent upon the amount of Java cycles used by the relevant application(s) and on the zAAP execution mode selected by the customer.

Execution of the Java applications on zAAPs, within the same z/OS SMP LPAR as their associated database subsystems, can also help simplify the server infrastructures and improve operational efficiencies. For example, use of zAAPs could reduce the number of TCP/IP programming stacks, firewalls, and physical interconnections (and their associated processing) that might otherwise be required when the application servers and their database servers are deployed on separate physical server platforms.

IBM does not impose software charges on zAAP capacity. Additional IBM software charges will apply when additional general purpose CP capacity is used.

Customers are encouraged to contact their specific ISVs/USVs directly to determine if their charges will be affected.

z890 has a single model, the A04, and will support a wide performance range from sub uni models through to a 4-way multiprocessor. There will be a total of 28 capacity settings which will offer outstanding granularity for customers' workload and cost management. The A04 model will provide up to four processor units (PUs) that can be characterized as either Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facilities (ICFs) or zAAPs. This model will be offered as a new build and as upgrades from designated models of the z800.

With the introduction of the z890 customers can expect to see the following performance improvements:

Number of CPs	Base	Estimated Ratio
1	z800	2.0
2	z800	2.1
3	z800	2.1
4	z800	2.1

The Large System Performance Reference (LSPR) should be referenced when considering performance on the z890. Visit: ibm.com/servers/eserver/zseries/lspr/ for more information on LSPR.

As with the z990, an improvement to the I/O subsystem has been introduced on the z890 to “break the barrier” of the 256 CHPIDs per Channel Subsystem. Horizontal growth is provided by allowing the definition of two Logical Channel SubSystems (LCSSs) each capable of supporting up to 256 CHPIDs, giving a total of 512 CHPIDs per z890 server. The z890 can support up to 30 Logical Partitions (LPARs). There is still a 256 CHPID limit per operating system. (Note: The lowest Capacity Setting 110 can support up to 240 CHPIDs, dependent on channel type, and up to a maximum of 15 LPARs.)

These are some of the significant enhancements in the zSeries 890 server that bring improved performance, availability and function to the platform. The following sections highlight the functions and features of the server.

z890 Design and Technology

The z890 is designed to provide balanced system performance. From server to the storage to the system's I/O and network channels, end-to-end bandwidth is provided and designed to deliver data where and when it is needed. The z890 is a scaled down version of the z990 using one of the most advanced chip and packaging technology in the industry. The z890 has more technology similarities to the z990 than the z800 does to the z900. Both the z890 and z990 use the latest MCM packaging and book package.

The z890 compared to the z800 provides a significant increase in system scalability and opportunity for server consolidation by providing improved model granularity while scaling to 2X the performance of the z800. The z890's processor cycle time has been improved to 1.0 ns

from the 1.6 ns of the z800. All z890's have a single MultiChip Module (MCM) that is nearly identical to the z990's MultiChip module. The differences are minor. With fewer functional processors and a slower cycle time less power is required for the z890 MCM enabling air cooling. The air cooled z890 MCM uses approximately 450 Watts compared to the liquid cooled z990 MCM which requires approximately 650 Watts. The z890 MCM can deliver 1 to 4-way configurations with a variety of engine sizes. The z890 MCM, which measures approximately 95 x 95 millimeters, is about 44% smaller than the z900's 127 mm x 127 mm MCM. The z890 MCM contains 5 Processor Unit (PU) chips, four Storage Data (Level 2 cache) chips, a Level 2 Cache Storage Controller chip and two main Memory Storage Controller (MSC) chips which control the main memory or Level 3 storage. The MCM contains 101 glass ceramic layers of which 23 layers are to provide interconnection between the chips and the off-module environment. In total, there are approximately 378 meters of internal copper wiring on both the z890 and z990 MCMs compared to the 997 meters of internal wiring needed for the MCMs in the z900 190 and 290 models. The new smaller MCM packaging delivers an MCM with many shorter paths. The z890 MCM is 40% smaller than the z900 MCM. It has over 20% more I/O connections and over a 130% I/O density improvement compared to the z900 MCM.

The z890's MCM provides support for 5 PUs and 32 MB level 2 cache which can be shared by all PUs. Each PU chip contains approximately 122 million transistors and measures 14.1 mm x 18.9 mm. The design of the MCM technology on the z890 provides the flexibility to configure

the PUs for different uses. One of the PUs is reserved for use as dedicated I/O Processor (IOP), commonly referred to as a System Assist Processor (SAP) enabling the Central Processor to avoid the burden of I/O set ups and operations. Any unused PUs can be utilized as spares. The inactive PUs on the MCM are available to be characterized as either Central Processors (CPs), Internal Coupling Facility (ICF) processors for Coupling Facility applications, Integrated Facility Linux (IFL) for Linux applications, or as IBM **@server** zSeries Application Assist Processor (zAAP) engines, providing enterprises with tremendous flexibility in configuring the best system for running applications. Each model of the z890 must always be ordered with at least one CP, IFL or ICF.

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 1.0 nanoseconds and is identical to the z990 PU except for the cycle time. Implemented on this chip are leading edge functions as the z/Architecture with its 64-bit capabilities including instructions, 64-bit General Purpose Registers, SuperScalar function, translation facilities, ability to use long displacement operations and on board Cryptographic Assist functions.

The z890 is a single book design. The book can support up to 32 GB of memory on a single memory card and 8 Self-Timed Interconnects (STIs) for I/O Connectivity. The memory which can be purchased in 8 GB increments, is delivered on 8 GB, 16 GB or 32 GB cards. The minimum memory is 8 GB. The book has a 2 Memory Bus Adapters (MBAs) and each MBA supports 4 STIs.

The z890 supports LPAR mode only (basic mode is no longer supported).

The MultiChip Module (MCM) is the technology cornerstone for flexible PU deployment in the z890 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, IFLs and zAAPs dynamically providing nondisruptive upgrade of processing capability. Also, the ability to add CPs lets a z890 with spare PU capacity become a backup for other systems in the enterprise; expanding the z890 system to meet an emergency outage situation. This is called Capacity BackUp (CBU).

The I/O SubSystem has been enhanced to support two Logical Channel SubSystems (LCSSs) which provide the capability to define up to 512 CHPIDs (256 CHPIDs 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™ provides high-speed capability to communicate among virtual servers and 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 without dependency on an external network. Industry standard and openness are design objectives for I/O in z890. 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.

z890 Family Models

The z890 offers a single model, the A04 and supports a wide performance range from sub uni models up to a 4-way multiprocessor. The z890 offers 7 general purpose CP sizes (1 full performance plus 6 sub uni's). Each engine size is available in 2-way, 3-way and 4-way multiprocessors. Seven engine sizes across 4 engines gives outstanding capacity granularity over the 28 capacity settings.


The new model provides 1-4 configurable processor units which can be characterized as either CPs, IFLs, ICFs, or zAAPs.

A new easy-to-enable ability to “turn off” capacity is available on z890. The objective is to allow customers to purchase capacity for future use with virtually no impact on software billing. An MES feature will enable the capacity for use where the customer requires the increased capacity.

There are a wide range of upgrade options available which are indicated in the z890 Models chart. The introduction of the capacity setting allows capacity upgrades and downgrades between capacity settings independent of the CP quantity.

Unlike other zSeries server offerings, it is no longer possible to tell by the hardware model A04, the number of PUs or the capacity of the server that is being used as CPs. The capacity setting is in the form xy0, where the X represents the number of CPs and Y represents the sub-capacity, from 7 (most capacity) to 1 (least capacity). All IFL and ICF configurations are software model 070. The capacity setting will be returned by the Store System Information (STSI) instruction for software billing purposes.

z890 Performance Comparison



z890 and IBM @server On/Off Capacity on Demand

IBM @server On/Off Capacity on Demand (On/Off CoD) is offered with z890 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 turn on temporary capacity for CPs, IFLs, ICFs or zAAPs. This is achieved by either increasing CP capacity through changing the Capacity Setting of the CPs to a new Model Capacity Indicator or by using unassigned/unowned PUs available for use as CPs, IFLs, ICFs, or zAAPs. The On/Off CoD is supported for PUs only, temporary use of memory and channels is not supported.

Although the hardware upgrade does not need a POR/IML if there is a change in engine size, a z/OS IPL will be required.

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.

Although Capacity BackUp (CBU) and On/Off Capacity on Demand can both reside on the server, the activation of On/Off Capacity on Demand is mutually exclusive with 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. On/Off CoD extends zSeries capacity on demand offerings to the next level of flexibility. It is designed to help customers match cost with capacity utilization and manage periodic business spikes. On/Off Capacity on Demand is designed to provide a low risk way to deploy new pilot applications, and it is designed to enable a customer to grow capacity smartly and proportionately with market demand.

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.

The z890 is designed to offer an efficient I/O structure using the same I/O cages as the z990. The single I/O cage z890, has the capability of plugging up to 28 I/O cards which would support a maximum of 420 ESCON® channels. Seven of the eight available STIs on the z890 are required to support the 28 channel slots in the I/O cage.

The following chart shows the upgrade paths from z800 to z890 and from z890 to z990. The z800 Models 002 and 004 can be upgraded to any capacity z890. Any z890 with more than 220 MIPS can be upgraded to a z990 Model A08. There are no upgrades from 9672 G5/G6 or z900 server.

z800 to z890 and z890 Model Upgrades



z890 Performance Comparisons

The IBM z890 is the latest member of the zSeries family. The design of the z890 is a continuation of the major change in the direction of the zSeries platform started with the IBM **@server** z990. The implementation of the superscalar microprocessor provides for improvements in the performance of new workloads as well as maintaining excellent performance for traditional workloads. The z890 provides a lower capacity entry point as compared to the z800, an increase in total system capacity almost 2.1 times that of the z800 Model 004, doubling the number of CHPIDs, doubling the number of LPARs, quadrupling the number of HiperSockets, and increasing the number of ESCON channels by 75%, the number of FICON™ channels by 25%, the number of ISC-3s by 100% and the number of OSA-Express Network connections by 67%. This server is designed to provide you with the ability to improve application performance, increase the number of users supported, support more transactions, increase scalability, and consolidate workloads beyond what is available on a z800.

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 great 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 wide range of performance levels, z890 offers from 1-4 Central Processors each with 7 different “dial points” (28 choices in all) and resource management within the system, it is expected that there will be a large performance variability than has been previously seen by our traditional customer set. The z890 supports an estimated performance range of 0.14 to 7.40 compared to a z800 Model 001. 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 available on the z890. 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 z890 has been optimized to run many workloads at high utilization rates.

It is also important to notice that the LSPR workloads for z890 and 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 z890/ z990 will be the legacy online workload, IMS, and two legacy batch workloads CB84 and CBW2. The z890/ 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 z890/ 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 z800 Model 001 to a one-way z890 Model A04 170, it is estimated that the z890 model has 2.0 times the capacity of z800 Model 001.
- Comparing a two-way z800 Model 002 to a two-way z890 Model A04 270, it is estimated that the z890 model has 2.1 times the capacity of z800 Model 002.
- Comparing a three-way z800 Model 003 to a three-way z890 Model A04 370, it is estimated that the z890 model has 2.1 times the capacity of z800 Model 003.

- Comparing a four-way z800 Model 004 to a four-way z890 Model A04 470, it is estimated that the z890 model has 2.1 times the capacity of z800 Model 004.

z890's Positioning in the zSeries Family



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

z890 I/O Subsystem


The z890 contains an I/O subsystem infrastructure which uses an I/O cage that provides 28 I/O slots. Each I/O cage has seven I/O domains with four features in each I/O domain, and each domain has one STI for connectivity to support each I/O cage. ESCON, FICON Express™ and OSA-Express features plug into the z890 I/O cage along with any ISC-3s and STI-3 distribution cards, and PCICA and PCIXCC cryptographic features. All I/O cards and their support cards can be hot-plugged in the I/O cage. The z890 provides approximately a 78 percent increase in I/O bandwidth provided by the STIs when compared to the z800.

z890 Cage Layout



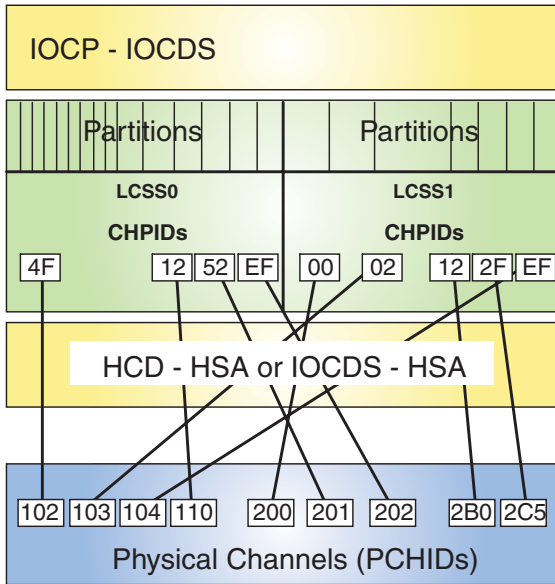
z890 Logical Channel SubSystems (LCSSs) and support for greater than 15 Logical Partitions (LPARs)

The z890 I/O subsystem is a breakthrough in connectivity by providing up to 2 LCSS per server, and each LCSS can support up to 256 CHPIDs when exploitation software is installed. As the z890 is a single I/O cage configuration it will support up to a maximum of 512 CHPIDs of which up to 420 CHPIDs can be ESCON channels. This support is provided in such a way that is transparent to the programs operating in the Logical Partition (LP). Each Logical Channel SubSystem may have from 1 to 256 CHPIDs 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. The z890 supports up to 30 Logical Partitions. (Note: The lowest Capacity Setting 110 can support up to 240 CHPIDs and up to a maximum of 15 LPARs.)



Physical Channel IDs (PCHIDs) SubSystem

In order to accommodate the support for up to a maximum of 512 CHPIDs on z890 with the Logical Channel SubSystem (LCSS) a new Physical Channel ID (PCHID) is available. 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 2 LCSSs, there is still a single IOCDs to define the I/O subsystem, each LCSS can support CHPID 00-FF, and the CHPID must be unique within the LCSS. 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

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. The z890 supports the spanning of the channels types: IC, HiperSockets, FICON Express, OSA-Express, ISC-3s and ICBs (ICB-3 and ICB-4). Note: Spanning of ESCON channels is not supported.

Channel Spanning



z890 Channels and I/O Connectivity

A z890 has a total of 28 I/O slots in one I/O cage. These slots can be plugged with a mixture of cards providing the I/O connectivity, networking connectivity, coupling and cryptographic capability of the server.

Up to 420 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, Configuration Control (LIC CC). There is a 75% increase in ESCON connectivity over z800.

Up to 40 FICON Express Channels

The z890 supports up to 40 FICON Express channels. It offers a 25% increase in FICON Express channels over the z800. 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 which can be installed in the I/O cage is 20.

The z890 supports up to 40 FCP channels for attachment to Small Computer System Interface (SCSI) disks in a Linux environment. The same two-port FICON Express feature card used for FICON Express channels can also be used for Fibre Channel Protocol (FCP) channels. FCP channels are enabled on these existing features via a

Licensed Internal Code (LIC) with a unique mode of operation and CHPID definition (FCP). 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 quantity of FICON Express, OSA-Express, PCICA, and PCIXCC features in combination cannot exceed 20 features per server.

InterSystem Channel-3 (ISC-3)

A four link ISC-3 feature is provided on the z890 family of servers. It consists of a mother card with two daughter cards each with two links. Each link is capable of operating 1 gigabits per second (Gbps) in Compatibility Mode or 2 Gbps in peer mode up to an unrepeated distance of 10 km (6.2 miles). The mode is selected for each link via the CHPID type in the IOCDs. ISC-3 links are ordered in increments of one. It is recommended that an initial order for ISC-3 include two links.

ISC-3 links in Compatibility Mode must be used for connectivity to 9672s (ICB-2 links are not supported on the z890).

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. A unique daughter card with two links per card is shipped.

Integrated Cluster Bus-3 (ICB-3)

The ICB-3 feature is a coupling link used to provide high-speed communication between z990 and z890 servers and z900 and z800 General Purpose Servers and Coupling Facilities 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 2.0 GBps input into two 1 GBps ICB-3s.

Fibre Channel Connectivity

Integrated Cluster Bus-4 (ICB-4)

The ICB-4 feature is a coupling link used to provide high-speed communication between a z990 and/or z890 servers over a short distance (less than 7 meters). The ICB-4 consists of one link that attaches directly to a 2.0 GBps STI port on the server and does not require connectivity to an I/O cage.

Internal Coupling Channel (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.

The on demand operating environment requires fast data access, continuous data availability, and improved flexibility all with lower cost of ownership. The zSeries Fibre Channel Connectivity can help you achieve these goals.

FICON Express Channel Card Features

Performance

With its 2 Gigabit per second link data rate 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 channels and can achieve improved performance over the previous generation FICON features. 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 7200 I/O operations/sec (channel is 100% utilized) and an aggregate total throughput of 120 MBps on z890.

With 2 Gbps links, customers may expect up to 170 MBps of total throughput. The 2 Gbps link data rates 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 Licensed Internal Code is needed in order to obtain 2 Gbps link data rates. The functionality was incorporated in all zSeries with March 2002 LIC. The link data rate 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 disks and may access these devices from Linux on zSeries, and new with z/VM Version 5 Release 1, installation and operation of z/VM on a SCSI disk.

Distance

All channels defined on FICON Express LX channel features at 1 Gbps link data rates 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 data rates, FICON Express LX 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 data rates 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 FICON Express SX features are up to 500 meters (1640 feet) and 250 meters (820 feet) for 50 and 62.5 micron multimode fiber, respectively. The maximum unrepeated distances for 2 Gbps links defined on the FICON Express SX features are up to 300 meters (984 feet) and 120 meters (394 feet) for 50 and 62.5 micron multimode

fiber, respectively. The FICON Express features are designed to reduce the data droop effect that made long distances not viable for ESCON. This distance capability is becoming increasingly important as enterprises are moving toward remote I/O, vaulting for disaster recovery and Geographically Dispersed Parallel Sysplex™ for availability.

Shared infrastructure

FICON (FC-SB-2 Fibre Channel Single-Byte Command Code Set-2) has been adopted by INCITS (International Committee for Information Technology Standards) as a standard 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 TotalStorage® Enterprise Storage Server® (ESS) Models 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. The IBM TotalStorage Enterprise Tape Controller 3592-J70 provides up to four FICON interfaces which can exist with ESCON on the same box. The 3592-J70 is designed to provide up to 1.5 times the throughput of the Model A60. 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, and Intrepid 6000 Series Directors) are 32-, 64- and 140-port high availability directors. The IBM 2042 Models 001, 128 and 256 (resell of the CNT FC/9000 Directors) 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 Intrepid 6000 Series Directors and CNT (FC/9000 Directors) FC/9000-001/-128/-256 support 2 Gbps links as well.

The FICON Express features now support attachment to the IBM M12 Director (2109-M12). The IBM M12 Director supports attachment of FICON Express channels on the z890 via native FICON (FC CHPID type) and Fibre Channel Protocol (FCP CHPID type) supporting attachment to SCSI devices in Linux environments.

Wave Division Multiplexor and Optical Amplifiers that support 2 Gbps FICON Express links are: Cisco Systems ONS 15530 and 15540 ESP (LX, SX) and optical amplifier (LX, SX), Nortel Networks Optera Metro 5100, 5200 and 5300E and optical amplifier, ADVA Fiber Service Platform (FSP) 2000 system, and the IBM 2029 Fiber Saver.

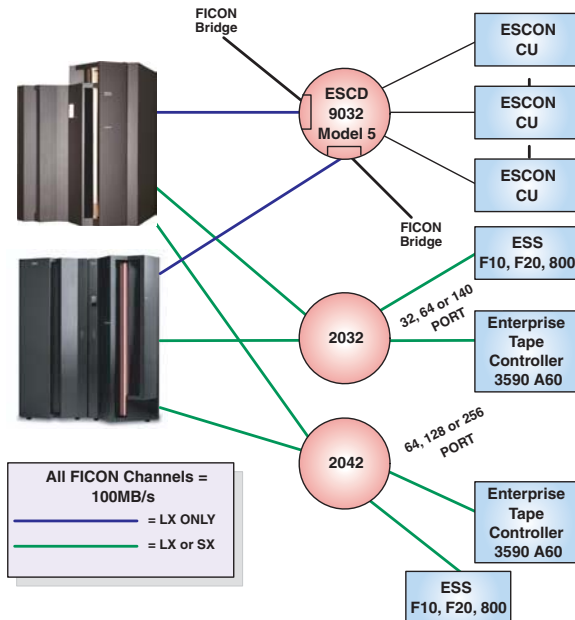
The raw bandwidth and distance capabilities that native FICON end-to-end connectivity has to offer makes it of interest for anyone with a need for high performance, large data transfers or enhanced multi-site solutions.

FICON CTC function

Native FICON channels support CTC on the z990, z900, z890 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, z890 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

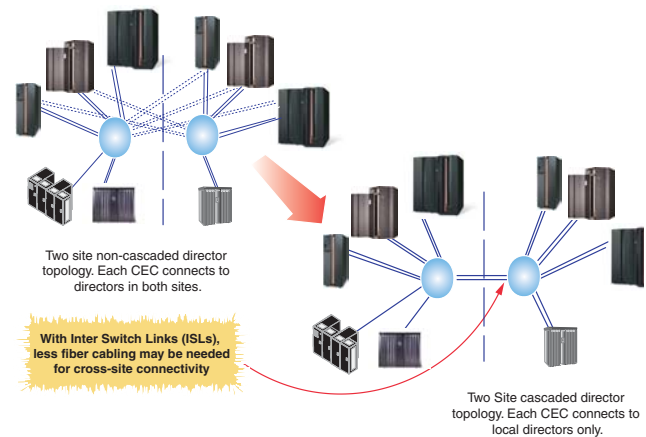


FICON Support for Cascaded Directors

Native FICON (FC) channels now support cascaded directors. This support is for a 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 help 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 high integrity connectivity. New integrity features introduced within the FICON Express channel and the FICON cascaded switch fabric to aid in the detection and reporting of any miscabling actions occurring within the fabric helps prevent data from being delivered to the wrong end point.

FICON cascaded directors are offered in conjunction with IBM, CNT, and McDATA Directors.



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 Bridge 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 on zSeries and z/VM Version 4 Release 3 and later. 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 on 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 helping to leverage existing investments and lower total cost of ownership for their Linux implementation.

For details of supported FICON and FCP attachments access Resource Link at ibm.com/servers/resourcelink and in the Planning section go to z890/z990 I/O Connection information.

The support for FCP channels is for Linux and z/VM 4.3 and later. Linux may be the native operating system on the zSeries server (note z890 runs LPAR mode only), or it can be in LPAR mode or, operating as a guest under z/VM 4.3 or later. The z890 provides support for IPL of Linux guest images from appropriate FCP attached devices.


Now, z/VM V5.1, support for SCSI FCP disks enables the deployment of a Linux server farm running under VM configured only with SCSI disks. With this support you can install, IPL, and operate z/VM from SCSI disks.

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

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 and directors from IBM, CNT, 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.

When configured as FCP CHPID type, the z890 FICON Express features support the industry standard interface for Storage Area Network (SAN) management tools.



Open Systems Adapter-Express Features (OSA-Express)

With the introduction of the z890, the OSA-Express Adapter family of Local Area Network (LAN) features is offering a maximum of 20 features per system, versus the maximum of 12 features per system on prior generations, expanding the z890 balanced solution to increase throughput and responsiveness in an on demand operating environment. These features combined with z/OS and z/OS.e, or OS/390, z/VM, Linux on zSeries, TPF, and VSE/ESA can help deliver a balanced system solution to increase throughput and decrease host interrupts to continue to help 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 40 ports per z890. Each port uses a single CHPID and can be separately configured. For a new z890 build, 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 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 z890

z890 OSA-Express 1000BASE-T Ethernet

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 (1 Gbps) using the same copper cabling infrastructure as Fast Ethernet, making the transition to this higher speed Ethernet feature a straightforward process. It is designed to support Auto-negotiation, QDIO and non-QDIO environments on each port allowing you to make the most of your TCP/IP and SNA/APPN® and HPR environments at up to gigabit speeds.

When this adapter is operating at gigabit Ethernet speed it runs full duplex only. It also can support standard (1492 or 1500 byte) and jumbo (8992 byte) frames.

The new Checksum offload support on the 1000BASE-T Ethernet feature, when operating in QDIO mode at gigabit speed, is designed to offload z/OS 1.5 and Linux TCP/IP stack processing of Checksum packet headers for TCP/IP and UDP.

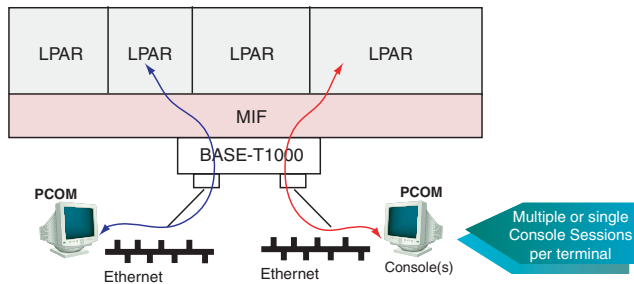


OSA-Integrated Console Controller

The new Open Systems Adapter Integrated Console Controller function (OSA-ICC) which is exclusive to IBM and to the IBM z890 and z990 servers since it is based on the OSA-Express feature, supports the attachment of non-SNA 3270 terminals for operator console operations. Now, 3270 emulation for console session connections is integrated as part of the OSA-Express 1000BASE-T Ethernet feature of the zSeries and can help eliminate the requirement for external console controllers (2074, 3174) helping to reducing cost and complexity. The OSA-ICC can use one or both ports on an OSA-Express 1000BASE-T Ethernet feature with the appropriate Licensed Internal Code (LIC). The OSA-ICC is enabled using CHPID type OSC.

The OSA-ICC is designed to support up to 120 client console LAN attached sessions either locally or remotely.

Support for this new feature will be available for z/VM Version 4 Release 4 with a PTF, z/OS Version 1 Release 3, VSE/ESA Version 2 Release 6 onwards and TPF.



Queued Direct Input/Output (QDIO)

The OSA-Express Gigabit Ethernet, 1000BASE-T Ethernet and Token-Ring features support QDIO, which is unique to IBM. QDIO was first introduced to the world on the z900, in Communication Server for OS/390 2.7.

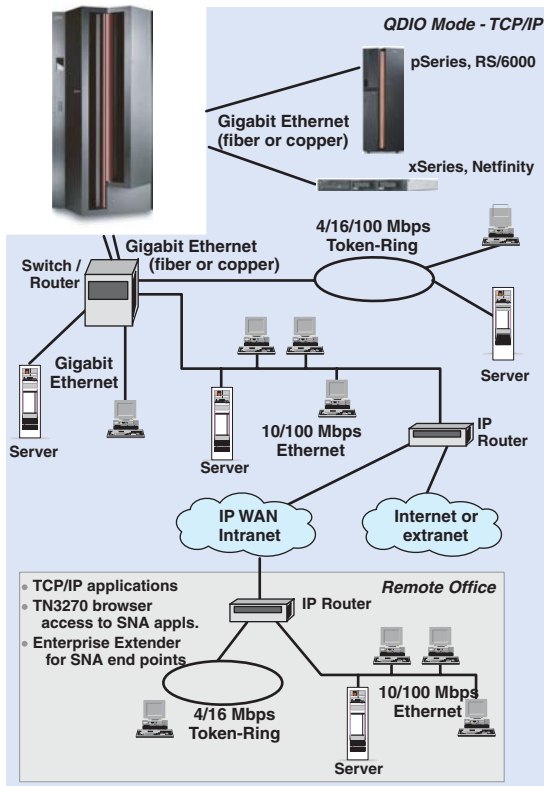
Queued Direct Input/Output (QDIO), a highly efficient data transfer architecture, allows an OSA-Express feature to directly communicate with the server's communications program through the use of data queues in memory. QDIO eliminates the use of channel programs and channel control words (CCWs), helping to reduce host interrupts and accelerate TCP/IP packet transmission.

TCP/IP connectivity is increased with the capability to allow up to a maximum of 160 IP stacks per OSA-Express port and 480 devices. This support is applicable to all the OSA-Express features available on the z890 and is provided through the Licensed Code (LIC).

Full Virtual Local Area Network (VLAN) support is available on z890 in z/OS and z/OS.e 1.5 Communications Server (CS) for the OSA-Express 1000BASE-T Ethernet, Fast Ethernet and Gigabit Ethernet features 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.

z890 OSA-Express Gigabit Ethernet

The new OSA-Express Gigabit Ethernet LX and Gigabit Ethernet SX features introduced on the z990 replace the z900/z800 Gigabit Ethernet features for new build z890. The new OSA-Express GbE features have a new connector type, LC Duplex, replacing the current SC Duplex connectors used on the prior z900/z800 Gigabit Ethernet features. The new Checksum offload support on these z890 features is designed to offload z/OS 1.5 and Linux TCP/IP stack processing of Checksum packet headers for TCP, IP and UDP.

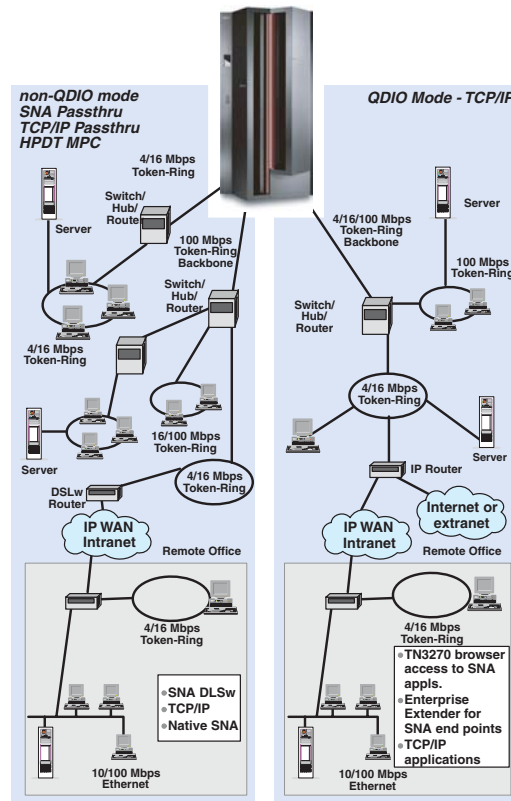


NON-QDIO operational mode

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. This support includes native SNA/APPN, High Performance Routing, TCP/IP passthrough, 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.

z890 OSA-Express Token-Ring

The same OSA-Express Token-Ring feature is supported on z890, z800, z990 and z900. This Token-Ring supports a range of speed including 4, 16 and 100 Mbps, and can operate in both QDIO and non-QDIO modes.



Note: OSA-Express 155 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.

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 can use TCP/IP connections from the remote site to either their TCP/IP or SNA applications on zSeries and S/390 by configuring OSA-Express with QDIO and using either direct TCP/IP access or appropriate SNA to IP integration technologies, such as TN3270 Server and Enterprise Extender for access to SNA applications. Customers who require the use of native SNA-based connections from the remote site can use a TCP/IP or SNA transport to the data center and then connect into zSeries and S/390 using appropriate SNA support on OSA-Express features configured in non-QDIO mode.

LPAR Support of OSA-Express

For z890 customers or customers who use the Processor Resource/Systems Manager (PR/SM) capabilities of the z800, 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 operates as 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.

IPv6 Support

IPv6 requires the use of an OSA-Express adapter running in QDIO mode and is supported only on OSA-Express features on zSeries at driver level 3G or above. IPv6 is supported on OSA-Express for zSeries Fast Ethernet, 1000BASE-T Ethernet and Gigabit Ethernet when running with Linux on zSeries, z/VM 5.1, and z/OS 1.4 and later.

Performance enhancements for virtual servers

Two important networking technology advancements are announced in z/VM 4.4 and Linux on z990

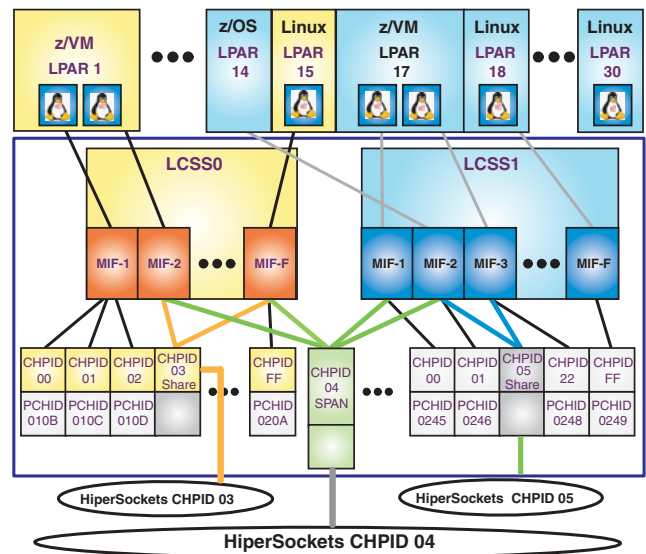
- *The high performance adapter interrupt handling first introduced with HiperSockets is now available for both OSA-Express in QDIO mode (CHPID=OSD) and FICON Express (CHPID=FCP). This advancement provides a more efficient technique for I/O interruptions designed to reduce path lengths and overhead in both the host operating system and in the adapter. This benefits OSA-Express TCP/IP support in both Linux for zSeries and z/VM.*
- *z990/z890 support of virtual machine technology has been enhanced to include a new performance assist for virtualization of adapter interruptions. This new z890/z990 performance assist is available to V=V guests (pageable guests) that support QDIO on z/VM 4.4 and later. The deployment of adapter interruptions improves efficiency and performance by reducing z/VM Control Program overhead.*

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) and LPARs within a zSeries system without any physical cabling. HiperSockets decreases network latency and increases 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 (virtual servers under z/VM).

With new support for up to sixteen HiperSockets the z890 and z990 provides four times more HiperSockets, and up to 4,096 TCP/IP images (stack) connections, which is also four times more capacity than the z800 and z900. The increased HiperSockets capacity and expanded connectivity provides additional flexibility in designing the networking to accommodate consolidated and multiple partitioned systems. HiperSockets can be divided among Logical Channel SubSystems for separation between various LPARs, while at the same time a single LPAR could have access to all sixteen HiperSockets if the HiperSockets all are assigned to the same LCSS.

A HiperSockets channel also supports spanned channels in order to communicate between LPARs connected to different LCSSs. HiperSockets (IQD CHPID type) can be configured to Channel SubSystems and transparently shared by any or all configured LPARs without regard to the LCSS to which the LPAR is configured. This means one HiperSockets could be common to all 30 z890 or z990 LPARs. This support is exclusive to z890 and z990. Different HiperSockets can be used for security (separation of 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, reducing 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 LPARs within the same z890. Furthermore, HiperSockets can 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 Secure Sockets layer (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 helps eliminate 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.

HiperSockets VLAN support in a Linux environment: Virtual Local Area Networks (VLANs), IEEE standard 802.1q, is now being offered for HiperSockets in a Linux for zSeries environment. VLANs can help 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 between applications over HiperSockets and between applications on HiperSockets connecting to an OSA-Express Gigabit Ethernet, 1000BASE-T Ethernet, or Fast Ethernet feature.

HiperSockets broadcast support for IPv4 packets – Linux, z/OS, z/VM: Internet Protocol Version 4 (IPv4) broadcast packets are now supported over HiperSockets internal LANs. TCP/IP applications that support IPv4 broadcast, such as z/OS OMPROUTE when running Routing Information Protocol Version 1 (RIPv1), can send and receive broadcast packets over HiperSockets interfaces. This support is exclusive to z890 and z990. Broadcast for IPv4 packets is supported by Linux for zSeries. Support is available in z/OS 1.5. Support is also offered in z/VM 4.4 and later.

HiperSockets Network Concentrator

HiperSockets Network Concentrator support, exclusive to z890 and z990, can simplify network addressing between HiperSockets and OSA-Express. You can now integrate HiperSockets-connected operating systems into external networks, without requiring intervening network routing overhead, thus helping to increase performance and simplify configuration. With the HiperSockets Network Concentrator support, you can configure a special purpose Linux operating system instance, which can transparently bridge traffic between a HiperSockets internal LAN and an external OSA-Express network attachment, similar to a real Layer 2 switch which bridges between different network segments. This support can make the internal HiperSockets network address connection appear as if it were directly connected to the external network.

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 security-rich, symmetric performance for financial and banking applications using PIN/POS type encryption. As in the past zSeries will be designed to deliver seamless integration of the cryptography facilities through use of ICSF. Use of ICSF will enable applications to be designed to work without change regardless of how and where the cryptographic functions are implemented, and also assure that the cryptography work to be 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 z890 and z990 – for SSL acceleration / clear key operations. To support the increased number of LPARs available on z890 and 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) is 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, and also as an execution environment for customer written programs and a high perfor-

mance path for Public Key / SSL operations. The PCIXCC hardware and microcode design will support almost all of the past Cryptographic functions that were provided on the zSeries 800 and 900 via the CMOS Cryptographic Coprocessor (CCF) and 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 z890 and z990 processor (CPU). With enhanced scalability and data rates the z890 and z990 processor is designed to provide 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 z890 or z990 processor speed, an order of magnitude faster than the CMOS Crypto Coprocessor in the zSeries 800 or 900. As these crypto functions are implemented in each and every CPU the affinity problem of pre-z990/z890 systems (which had only two CMOS Crypto Coprocessors) is virtually 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 help reduce programming overhead. To confirm with US Export and Import Regulations of other countries, an SE panel is provided for proper enable/disable of 'strong' cryptographic functions.

Availability

The Trusted Key Entry (TKE) 4.1 code level workstation is an optional feature that can provide a basic key management system and Operational Key Entry support. The key management system allows an authorized person a method for key identification, exchange, separation, update, backup, and management. The TKE workstation and 4.0 code level are designed to provide a secure, remote, and flexible method of providing Master Key Entry and to remotely manage PCIX Cryptographic Coprocessors.

zSeries Security Certification

Cryptography

- *z890/z990 PCIXCC:*
 - *Designed for FIPS 140-2 level 4 certification*
- *Logical Partitions*
 - *z900 and z800 servers are the first and only to receive Common Criteria certification at EAL5*
- *Operating Systems Common Criteria Certification*
 - *SUSE LINUX on zSeries*
 - *SUSE SLES 8 has been certified at Controlled Access Protection Profile (CAPP) EAL3+*
- *z/OS 1.6*
 - *z/OS 1.6 is under evaluation for Controlled Access Protection Profile (CAPP) EAL3+ and Labeled Security Protection Profile (LSPP) EAL3+*
- *z/VM*
 - *IBM has applied for Common Criteria Controlled (ISO/IEC 15408) certification of z/VM V5.1 with the RACF[®] for z/VM optional feature against the Controlled Access Protection Profile (CAPP) and the Labeled Security Protection Profile (LSPP), both at the EAL3+ assurance level*

z890 Capacity Upgrade on Demand (CUoD)

Capacity Upgrade on Demand allows for the nondisruptive addition of additional Central Processor (CP) Capacity, Internal Coupling Facilities (ICFs), Integrated Facility for Linux (IFLs) and/or IBM **@server** zSeries Application Assist Processor (zAAP). 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, can enable virtually unlimited capacity upgrade capability.

The CUoD functions are:

- *Nondisruptive CP, ICF, IFL, and zAAP upgrades within minutes*
- *Dynamic upgrade of all I/O cards in the I/O Cage*
- *Dynamic upgrade of spare installed memory (24 - 32 GB only)*
- *Plan Ahead and Concurrent Conditioning*

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 non-disruptive with the appropriate microcode load and if the cage space is available.

The z890 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 8 or 16 GB increments. A customer with a Model A04 with 24 GB of storage will be able to concurrently upgrade to 32 GB as there is spare memory capacity on the card. However, for a customer with 8 or 16 GB of memory installed getting to the next increment would be disruptive.

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 z890 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.

z890 Server Capacity BackUp (CBU)

Capacity BackUp (CBU) is offered with the z890 servers 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 z890 server. 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 z890 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.

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.

Customers who have an active Remote Support Facility connection can perform a CBU upgrade automatically and within a matter of minutes, right from their CBU machine's Hardware Management Console. For more information on how a CBU upgrade can be activated automatically, please refer to the *z890 Capacity Backup Users Guide* found on IBM Resource Link.

The z890 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).

Advanced Availability Functions

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 or site failure, 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. GDPS-CBU management helps to minimize manual customer intervention and the potential for errors, thereby reducing the outage time for critical workloads from hours to minutes. Similarly, GDPS-CBU management can also automate the process of dynamically returning the reserved CPs when the temporary period has expired.

z890 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 will be able to 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. CIU is a low-risk, well tested-and-tried facility for ordering capacity on demand.

Transparent Sparing

z890 offers a 5 PU MCM in the case of processor failure. If there is a spare PU available it will be used for transparent sparing. Note this can only be achieved on 1-way, 2-way and 3-way models.

Enhanced Dynamic Memory Sparing

The z890 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: z890 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 virtually eliminates any need to schedule outage to service or upgrade the I/O subsystem on this cage.*

- *Upgrade for Coupling Links: z890 has concurrent maintenance for the ISC-3 adapter card. Also, Coupling links can be added concurrently. This virtually 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 nondisruptive service to the controllers and virtually eliminates customer scheduled outage.*
- *Auto-Switchover for Service Element: The z890 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).

Fault Tolerant Interconnect Design: The memory design of the z890 provides a fault tolerant dual interconnect fabric. The z890 is designed to prevent a single failure within the fabric from disrupting full memory access.

Concurrent Capacity Upgrade: A z890 can be upgraded via the concurrent addition of additional Processor Units (PUs) and I/O and limited memory upgrades (24 Gb to 32 Gb increment) with no disruption to current operations.

However, while capacity upgrades to the processor itself are concurrent, your software may not be able to take advantage of the increased capacity without performing an Initial Programming Load (IPL).

Sparing for Storage Protect Keys: The robust configuration of the Storage Protect Keys has been enhanced with chip sparing. Already a triple-array with parity protection and voting, the Key structure for z890 has added sparing, similar to Main Memory chip sparing, to further enhance the availability of this critical function.

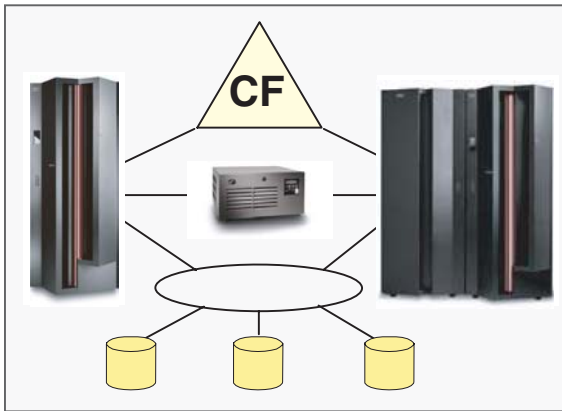
Failure Containment for the Memory Bus Adapter (MBA): Memory Bus Adapters are designed to provide the critical link between the z890 Processor Units and the I/O subsystem. The Model A04 has two MBAs. In the unlikely event of a catastrophic failure of an MBA chip, the z890 is designed to isolate the failure of that chip such that the remaining MBA chips continue to operate. This helps minimize the impact of a failure and allows for scheduling maintenance.

Enhanced Firmware Simulation: The z890 process for design, development, and test of Licensed Internal Code (LIC) has been significantly enhanced with the use of simulation to improve quality and early availability. Virtually every action/reaction of the code can be tested with the simulated hardware/code of the rest of the server. The result is to discover and correct design errors much earlier in the process.

These new features, together with legacy features such as high levels of recovery, concurrent processor upgrade, concurrent memory upgrade, concurrent I/O upgrade, and concurrent maintenance for hardware and LIC give the z890 a very impressive RAS structure.

Parallel Sysplex Cluster Technology

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 address 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, increased performance and/or scalability. For more detail, please see *S/390 Value of Resource Sharing White Paper – GF22-5115* 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/pso.

Coupling Facility Configuration Alternatives

IBM offers different options for configuring a functioning Coupling Facility:

- *Standalone Coupling Facility: z900 Model 100, z800 Model OCF and 9672-R06 models provide a physically isolated, totally independent CF environment. Like the z990, there is no unique standalone coupling facility model offered with the z890. Customers can achieve the same physically isolated environment as on prior mainframe families by ordering a z890 or 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, z890 and 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 dramatically reducing 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 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.



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, please 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

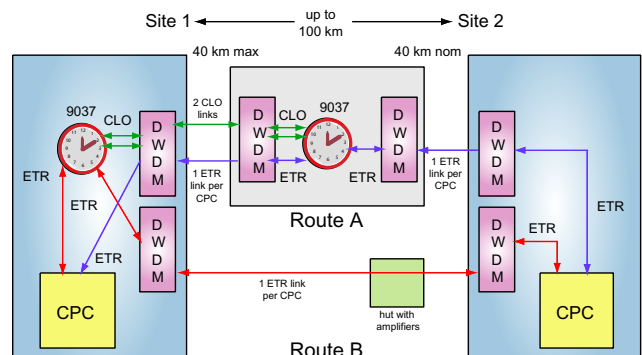
The Coupling Facilities communicate with z/OS and OS/390 images in the Parallel Sysplex environment over specialized high-speed links. For availability purposes, it is recommended that there be at least two links connecting each z/OS or OS/390 image to each CF in a Parallel Sysplex cluster. As processor 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/z890/z800) to a z800 Model 0CF, 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 12 km, 100 Gbps up to 20 km, the ICB-3 link with 1 GB/s peak capacity, the ICB-4 for z890/z990 to z890/z990 connection at 2.0 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 ICB coupling link speeds described above are theoretical maximums. Note: The z890 does not support ICB-2 links for connectivity to 9672s. ISC-3 links in compatibility mode must be used.

GDPS/PPRC Cross Site Parallel Sysplex distance Extended to 100 km

When using a Dense Wave Division Multiplexor (DWDM), it will be possible via an RPQ to configure GDPS/PPRC or a multi-site Parallel Sysplex with up to 100 km between the two sites. The immediate advantage of this extended distance is to potentially decrease the risk that the same

disaster will affect both sites, thus providing the ability for customers to recover their production applications at another site. Support for the External Timer Reference (ETR) links and the Inter System Channel (ISC-3) links has been increased from the current capability of 50 km to an extended capability of up to 100 km. The extended distance support for ETR and ISC-3 links is now consistent with other cross site technologies that already support 100 km (such as FICON, Peer-to-Peer Remote Copy (PPRC), and Peer-to-Peer Virtual Tape Server (PtP VTS)). It should be noted that the maximum distance between a pair of 9037 Sysplex Timers in an Expanded Availability configuration remains at 40 km. Therefore, to achieve the extended distance of 100 km between the two sites, one of the options to be considered is locating one of the Sysplex Timers in an intermediary site that is less than 40 km from one of the two sites (as can be seen in the diagram below). Other potential options can be evaluated when the RPQ is ordered.



z890 and z990 Theoretical Maximum Coupling Link Speed


Connectivity Options	z890 ISC-3	z890 ICB-3	z890 ICB-4
G5/G6 ISC	1 Gbps Compatibility	n/a	n/a
z800/z900 ISC-3	2 Gbps Peer Mode	n/a	n/a
z890/z990 ISC-3	2 Gbps Peer Mode	n/a	n/a
G5/G6 ICB	n/a	n/a	n/a
z900 ICB-2	n/a	n/a	n/a
z990 ICB-2	n/a	n/a	n/a
z900 ICB-3	n/a	1 GBps Peer Mode	n/a
z890/z990 ICB-3	n/a	1 GBps Peer Mode Recommendation use ICB-4s	n/a
z890/z990 ICB-4	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/z890/z900/z990) general purpose models and zSeries Coupling Facilities. ISC-3 channels operating in Compatibility Mode provide connection between zSeries 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 z890 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 2 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. The z890 supports a maximum of 48 links (12 features, 4 links per feature). Note: A maximum of 32 ISC-3s can be defined in Compatibility Mode.

- **ISC-2. HiPerLinks.** HiPerLinks, based on single-mode CF links, are available on 9672s (G3 - G6) and 9674s only. ISC-3s replace HiPerLinks on zSeries 800, 900, 890 and 990 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 MBps. The ports are activated in one port increments. Up to 4 STI-2 cards, 8 ICB-2 links are available on the z990. ICB-2 is not offered on the z890.
- **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 z890/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 GBps. The ports are activated in one port increments. Up to 8 STI-3 cards, 16 ICB-3 links are available on the z890. ICB-3 links operate in "peer mode."

- *ICB-4. The Integrated Coupling Bus-4 is a “native” coupling connection available for connecting a z890/z990 server or CF to another z890/ z990 server or CF over short distances. Capable of up to 2.0 GBps, the ICB-4 is the fast external coupling connection available for the z890/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 8 ICB-4 features can be configured on a z890.*
- *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. The z890 supports up to a maximum of 32 ICs.*




Intelligent Resource Director

Exclusive to IBM's 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 z890/ 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 Sub-system Priority Queuing.

zSeries 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, helps reduce MP overhead, and helps give z/OS more control over how CP resources are distributed to help 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 (DCM), 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 help reduce the cost of the fiber infrastructure required for connectivity between multiple data centers. On a z890/ z990 with Logical Channel SubSystems (LCSSs), the scope of DCM management is within a Logical Channel SubSystem. Although an LPAR cluster can span LCSSs, 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 z890/ z990, the scope of Channel Subsystem I/O Priority Queuing is a Logical Channel SubSystem.

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 world-class workload management tradition of S/390 and OS/390 so 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/psos/services.html.

Geographically Dispersed Parallel Sysplex

The GDPS solution, based on Peer-to-Peer Remote Copy (PPRC), referred to as GDPS/PPRC, is designed with the attributes of a continuously availability solution. PPRC is a hardware solution that is designed to synchronously mirror

data residing on a set of disk volumes, called the primary volumes in site 1, to secondary disk volumes on a second system in site 2. Only when the primary storage subsystem receives "write complete" from the secondary storage subsystem is the application I/O signaled completed. GDPS/PPRC 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/PPRC is application independent and therefore can cover the customer's complete application environment. GDPS supports both the synchronous 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 is designed to allow a customer to meet a Recovery Time Objective (RTO) of less than an hour, a Recovery Point Objective (RPO) of no data loss, and helps protect against metropolitan area disasters (up to 40 km between sites). GDPS/XRC is a business continuity/disaster recovery solution that is designed to allow a customer to meet a RTO of one to two hours, an RPO of less than a minute, and helps protect 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.

On the other hand, the GDPS solution based on Extended Remote Copy (XRC), referred to as GDPS/XRC, has the attributes of a Disaster Recovery solution. XRC is a combined hardware and software asynchronous remote copy solution. The application I/O is signaled completed when the data update to the primary storage is completed. Subsequently, a DFSMSdfp™ component called System Data Mover (SDM), typically running in site 2, is designed to asynchronously offload data from the primary storage subsystem's cache and update the secondary disk volumes.

In GDPS/XRC, the production system(s) located in site 1 can be a single system, multiple systems sharing disk, or a base or Parallel Sysplex cluster1. GDPS/XRC provides a single, automated solution designed to dynamically manage storage subsystem mirroring (disk and tape) to allow a business to attain "near transparent" disaster recovery with minimal data loss. GDPS/XRC is designed to provide the ability to perform a controlled site switch for an unplanned site outage, maintaining data integrity across multiple volumes and storage subsystems. GDPS/XRC is application independent and therefore is capable of covering the customer's complete application environment.

GDPS requires Tivoli® NetView® for z/OS or Tivoli NetView for OS/390, Tivoli System Automation for OS/390, and remote copy technologies.

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 mask planned and unplanned disk and site reconfigurations by transparently switching to use the secondary PPRC volumes. 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 planned provides the ability to transparently switch all primary PPRC disk subsystems with the secondary PPRC disk subsystems for a planned switch configuration. It 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.

The unplanned HyperSwap function contains additional function to transparently switch to use secondary PPRC disk subsystems, in the event of unplanned outages of the primary PPRC disk subsystems or a failure of the site containing the primary PPRC disk subsystems. With unplanned HyperSwap function, disk subsystem failures no longer constitute a single point of failure for an entire sysplex. If applications are cloned and exploiting data

sharing across the two sites, the GDPS/PPRC unplanned HyperSwap capability lays 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. An improvement in the Recovery Time Objective (RTO) can be accomplished.

Geographically Dispersed Parallel Sysplex (GDPS) enhancements

Geographically Dispersed Parallel Sysplex, (GDPS), a world-class e-business continuity solution, is a multisite solution that is designed to provide the capability to manage the remote copy configuration and storage subsystems, automate Parallel Sysplex operational tasks, and perform failure recovery from a single point of control, thereby helping to improve application availability. GDPS supports both the synchronous Peer-to-Peer Remote Copy (PPRC), as well as the asynchronous Extended Remote Copy (XRC) forms of remote copy. Depending on the form of remote copy, the solution is referred to as GDPS/PPRC or GDPS/XRC.

GDPS/PPRC and GDPS/XRC have been enhanced to include new functions.

GDPS/PPRC HyperSwap function: 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.

Planned HyperSwap function provides the ability to:

- *Transparently switch all primary PPRC disk subsystems with the secondary PPRC disk subsystems for a planned reconfiguration*
- *Perform disk configuration maintenance and planned site maintenance without requiring any applications to be quiesced.*

Planned HyperSwap function became generally available December 2002.

Unplanned HyperSwap function contains additional function to transparently switch to use secondary PPRC disk subsystems in the event of unplanned outages of the primary PPRC disk subsystems or a failure of the site containing the primary PPRC disk subsystems. Unplanned HyperSwap support allows:

- *Production systems to remain active during a disk subsystem failure. Disk subsystem failures will no longer constitute a single point of failure for an entire Parallel Sysplex.*
- *Production servers to remain active during a failure of the site containing the primary PPRC disk subsystems if applications are cloned and exploiting data sharing across the two sites. Even though the workload in the second site will need to be restarted, an improvement in the Recovery Time Objective (RTO) will be accomplished.*

Unplanned HyperSwap function became generally available February 2004.

GDPS/PPRC management for open systems LUNs (Logical Unit Numbers): GDPS/PPRC technology has been extended to manage a heterogeneous environment of z/OS and open systems data. If installations share their disk subsystems between the z/OS and open systems platforms, GDPS/PPRC, running in a z/OS system, can manage the PPRC status of devices that belong to the other platforms and are not even defined to the z/OS platform. GDPS/PPRC will also provide data consistency across both z/OS and open systems data.

GDPS/PPRC management of open systems LUNs became generally available February 2004.

GDPS supports PPRC over Fiber Channel links: In 2003, IBM TotalStorage Enterprise Storage Server (ESS) announced support of PPRC over Fiber Channel for the ESS Model 800. Refer to Hardware Announcement 103-298, (RFA38991) dated October 14, 2003. This support is designed to provide improved throughput (compared to ESCON), and a reduction in cross-site connectivity (two PPRC Fiber Channel links per ESS are considered sufficient for most customer workloads). One of the potential benefits of this support is the ability for customers to increase the distance between sites while maintaining acceptable performance.

GDPS/PPRC support for PPRC over Fiber Channel became generally available February 2004.

GDPS supports FlashCopy® V2 elimination of the LSS constraint: In 2003, IBM TotalStorage Enterprise Storage Server (ESS) announced support of FlashCopy V2. Refer to Hardware Announcement 103-141, dated May 13,

2003. Prior to this announcement, both source and target volumes had to reside on the same Logical SubSystem (LSS) within the disk subsystem. Since this constraint has been removed with FlashCopy V2, GDPS will now allow a FlashCopy from a source in one LSS to a target in a different LSS within the same disk subsystem. This new flexibility can help simplify administration and capacity planning for FlashCopy.

GDPS/PPRC support for FlashCopy V2 became generally available February 2004.

GDPS/PPRC and Cross-site Parallel Sysplex distance extended to 100 km: On October 31, 2003, IBM delivered, via a Request for Price Quote (RPQ), the capability to configure GDPS/PPRC or a multi-site Parallel Sysplex up to a distance of 100 kilometers (62 miles) between two sites. This extended distance can potentially decrease the risk that the same disaster will affect both sites, thus permitting enterprises to recover production applications at another site. Support has been extended up to a distance of 100 km from the current capability of 50 km (31 miles) for:

- *External Time Reference (ETR) links*
- *An ETR link on a zSeries or S/390 server provides attachment to the Sysplex Timer.*
- *InterSystem Channel-3 (ISC-3) links operating in Peer mode ISC-3 links, supported on all zSeries servers, connect z/OS and OS/390 systems to Coupling Facilities in a Parallel Sysplex environment.*

The extended distance support for ETR and ISC-3 links is now consistent with other cross-site link technologies that currently support 100 km between two sites, such as

FICON, Peer-to-Peer Remote Copy (PPRC), Peer-to-Peer Virtual Tape Server (PtP VTS). It should be noted that the maximum fiber optic cable distance between a pair of Sysplex Timers in an Expanded Availability configuration remains at 40 km (25 miles). Therefore, to achieve the extended distance of 100 km between sites, one of the options to be considered is locating one of the Sysplex Timers in an intermediary site that is less than 40 km from one of the two sites. Other potential options can be evaluated when the RPQ request is submitted to IBM for review.

Coordinated continuous availability and disaster recovery for Linux guests: z/VM is providing a new HyperSwap function so that the virtual device associated with one real disk can be swapped transparently to another. HyperSwap can be used to switch to secondary disk storage subsystems mirrored by Peer-to-Peer Remote Copy (PPRC).

HyperSwap can also be helpful in data migration scenarios to allow applications to use new disk volumes.

GDPS plans to exploit the new z/VM HyperSwap function to provide a coordinated continuous availability and disaster recovery solution for z/OS and Linux guests running under z/VM. This innovative disaster recovery solution requires GDPS, IBM Tivoli System Automation for Linux, Linux on zSeries, and z/VM V5.1 to help anticipate and rapidly respond to business objectives and technical requirements while maintaining unsurpassed system availability. This solution is especially valuable for customers who share data and storage subsystems between z/OS and Linux on zSeries.

To support planned and unplanned outages, GDPS is designed to provide the following recovery actions:

- *Re-IPL in place of failing operating system images*
- *Site takeover/failover of a complete production site*
- *Coordinated planned and unplanned HyperSwap of storage subsystems transparently to the operating system images and applications using the storage*

Performance enhancements for GDPS/PPRC and GDPS/XRC configurations

- *Concurrent activation of Capacity Backup Upgrade (CBU) can now be performed in parallel across multiple servers, which results in an improved RTO. This improvement applies to both the GDPS/PPRC and GDPS/XRC configurations.*
- *In a GDPS/XRC configuration, it is often necessary to have multiple System Data Movers (SDMs). The number of SDMs is based on many factors, such as the number of volumes being copied and the I/O rate. Functions are now capable of being executed in parallel across multiple SDMs, thus providing improved scalability for a coupled SDM configuration.*
- *Analysis has shown that PPRC commands issued by GDPS will generate a large number of Write to Operator messages (WTOs) that may cause WTO buffer shortages and temporarily adversely impact system performance. The Message Flooding Automation function will substantially reduce the WTO message traffic and improve system performance by suppressing redundant WTOs.*

Performance enhancements for GDPS/PPRC and GDPS/XRC became generally available March 24, 2003.

These GDPS enhancements are applicable to z800, z900, z890, and z990. For a complete list of other supported hardware platforms and software prerequisites, refer to the GDPS executive summary white paper, available at: ibm.com/server/eserver/zseries/psa.

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, GDPS will dynamically add PUs to the configuration in the takeover site to restore processing power for mission-critical production workloads. GDPS-CBU management helps to minimize manual customer intervention and the potential for errors, thereby helping to reduce the outage time for critical workloads from hours to minutes. Similarly, GDPS-CBU management can also automate the process of dynamically returning the reserved CPs when the temporary period has expired.

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.

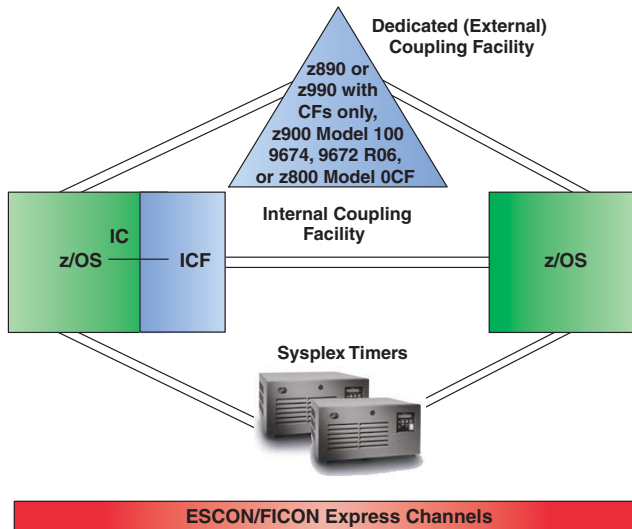
Message Time Ordering (Sysplex Timer Connectivity to Coupling Facilities)

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 z890, 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 z890, 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.



Continuous Availability Recommended Configuration for Parallel Sysplex



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

Note: z890/z990 will attach to 9037-001 or 9037-002. Service for 9037-001 discontinued end of 2003.

Components and assumptions

- Two Coupling Facilities; at least one external or else using System-Managed CF Structure Duplexing
- 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

z890 Support for Linux

Linux on 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 aiding in reducing both cost and complexity.

Of course, no matter which Linux applications are brought to the zSeries platform, they can all benefit from high-speed access to the applications and corporate data that reside 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 help eliminate the need for an external 2074 Console controller and associated consoles, an administrator may utilize the Hardware Management 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 was made available in z/VM Version 4 Release 4. An update for Linux will be made available to IBM Linux Distribution Partners.

Linux on zSeries supports the 64-bit architecture available on zSeries processors. z/Architecture eliminates the existing main storage limitation of 2 GB. Linux on zSeries provides full exploitation of the architecture in both real and virtual modes. Linux on zSeries is based on the Linux 2.4 kernel. Linux on 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 Versions 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™ Versions 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 Versions 1.4 and 1.3.1
 - IBM Directory Integrator Version 5.1
 - IBM Directory Server Versions 5.1 and 4.1
 - IBM HTTP Server Version 1.3.19.3
 - IBM Object RE90 Version 2.2
 - IBM Screen Customizer Versions 2.0.7 and 2.0.6

Linux Distribution Partners

- *SUSE LINUX*
Product Information at suse.de/en/produkte/susesoft/S390/
- *Turbolinux*
Product Information at turbolinux.com/products/s390
- *Red Hat Linux*
Product Information at redhat.com/software/S390

z/VM Version 4 and Version 5

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 servers, both z/VM V4 and V5 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. Engine-based Value Unit pricing for z/VM V5.1 is replacing the pricing model available with z/VM V4. Engine-based Value Unit pricing is designed to provide a lower entry point and a decreasing price curve which will help provide improved price/performance as hardware capacities and workload grow. Value Unit pricing helps you to add capacity and workload with an incremental and improved price and the ability to aggregate licenses acquired across machines that are part of your enterprise.

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 workloads, with minimal or 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 necessarily 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 on 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.

Two important networking technology advancements are announced in z/VM 4.4 and Linux on z990:

- *The high performance adapter interrupt handling first introduced with HiperSockets is now available for both OSA-Express in QDIO mode (CHPID=IQD) and FICON Express (CHPID=FCP). This advancement provides a more efficient technique for I/O interruptions designed to reduce path lengths and overhead in both the host operating system and in the adapter. This benefits OSA-Express TCP/IP support in both Linux for zSeries and z/VM.*
- *The z990's support of virtual machine technology has been enhanced to include a new performance assist for virtualization of adapter interruptions. This new z990 performance assist is available to V=V guests (pageable guests) that support QDIO on z/VM 4.4. The deployment of adapter interruptions improves efficiency and performance by reducing z/VM Control Program overhead when handling Linux guest virtual servers.*

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 on 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 OMROUTE 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 help 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. 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 on zSeries. This expanded attachability provides a larger selection of storage solutions for Linux implementations.

zSeries 890 Family Configuration Detail

Cryptographic Support for Linux

Linux on 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 on zSeries to utilize this enhanced hardware security.

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 16 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; One I/O cages: 28 I/O slots

Feature	Per System			Ports/ Channels/ Increments per Feature	Purchase Increments
	Minimum Features	Maximum I/O Slots used by Features	Maximum Connections		
ESCON, 16 port	0 ¹	28	420 channels	16 channels ²	4 channels ³
FICON Express	0 ¹	20 ⁴	40 channels	2 channels	1 feature
STI-3 ⁵ ICB-3 link	0 0 ¹	8 N/A	N/A 16 links ⁶	2 outputs N/A	N/A 1 link
ICB-4 link	0 ¹	N/A ⁷	8 links ⁶	N/A	1 link
ISC-3	0 ¹	12	48 links ^{6,8}	4 links	1 link ⁹
OSA- Express	0	20 ⁴	40 ports	2 ports	1 feature
PCICA	0	2 ⁴	4 accelerator cards	2 accelerator cards	1 feature
PCIXCC	0 ¹⁰	4 ⁴	4 coprocessor	1 coprocessor	1 feature ¹⁰

1) A minimum of one I/O feature (ESCON, FICON Express) or one Coupling Link (ICB, ISC-3) is required.

2) Each ESCON feature has 16 channels of which 15 channels may be activated. One channel is always reserved as a spare.

3) ESCON channels are purchased in increments of four and are activated via Licensed Internal Code, Configuration Control (LIC CC). Channels are activated equally across all installed 16-port ESCON features for high availability.

4) The maximum quantity of FICON Express, OSA-Express, PCICA, and PCIXCC in combination cannot exceed 20 features per server.

5) Each STI-3 distribution card, which supports the ICB-3s, resides in the I/O cage, occupying one I/O slot.

6) The maximum number of Coupling Links combined (ICs, ICB-3s, ICB-4s, and active ISC-3 links) cannot exceed 64 per server.

7) ICB-4s do not require connectivity to a card in the I/O cage. ICB-4s are not included in the maximum feature count for I/O slots.

8) A maximum of 32 ISC-3s can be defined in compatibility mode (operating at 1 Gbps, instead of 2 Gbps).

9) It is recommended that an initial order for ISC-3 include two links. When two links are purchased, two ISC-3 features are shipped and activated links are balanced across ISC-3 features for high availability.

10) PCIXCC feature increments are 0, 2, 3, or 4.

Note: The Capacity Model 110 has an I/O cage with 16 I/O slots maximum which will support up to 240 channels which can be any mix of the following as long as it does not exceed 240 channels — maximum of 240 ESCON channels, 32 FICON Express channels, 24 OSA-Express ports.

Processor Unit Assignments

Model A04***	Total PUs	Characterized PUs*	SAP Standard	Spares
1-way**	5	1	1	3
2-way**	5	2	1	2
3-way**	5	3	1	1
4-way**	5	4	1	0

*PU can be characterized as CP, IFL, ICF, or zAAP

**These values are for all machines including the sub-uni, sub 2-way models, etc.

***Customer will be required to purchase at least one CP, IFL, ICF or zAAP

Processor Memory

z890 Model	Minimum	Maximum
A04	8 GB	32 GB

Single memory card. Memory cards 8 GB, 16 GB or 32 GB. Memory increments 8 GB.

Channels

Model	ESCON Min	**ESCON Max	*FICON Min	FICON Max
A04	0	420	0	40

*FICON Express and FCP configured on the same FICON Express features. Max channels total 40.

**ESCON increments of 4 channels

Coupling Links

Links	IC	ICB-3*	ICB-4	ISC-3	Max Links
	0-32	0-16	0-8	0-48	Max combined IC, ICB-3, ICB-4, active ISC-3 = 64

*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	2 ³	4 ⁴

1. Max two PCICA features per system

2. Max eight PCICA and PCIXCC features per system

3. Two accelerator cards per PCICA feature

4. One coprocessor per PCIXCC feature (minimum orderable 2 features)


OSA-Express Features

	OSA-Express* Features
Minimum	0
Maximum	20

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

z890 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 a single frame, the A-Frame has a single I/O cage with 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 ports
PCICA	2	4 ports

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

Physical Characteristics

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 420 channels on 28 cards
- OSA-Express can be Gigabit Ethernet, 1000BASE-T Ethernet or Token-Ring
- Total number of PCIXCC / PCICA is 8 per system
- If ICB-3 are required on the system, these will use up a single I/O slot for each ICB-3 to accommodate the STI-3 cards. The first ICB-3 required a slot. The second to the fourth required another slot. The fifth to the sixth required another slot. (STI-3 cards each provide two ICBs)

Channels and channel adapters no longer supported on z890

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 migrate to Ethernet as soon as possible
- 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
- PCICC - replaced with PCIXCC for most functions
- ICB-2 - replace with ISC-3 links in compatibility mode.

z890 Power/Heating/Cooling

System Power Consumption (50/60 Hz, KVA)

Model / Config	1 I/O Cage
A04	1.5 - 4.7

Note: Assumes 60 amp cords

System Cooling (Air Flow Rate - CFM)

Model / Config	1 I/O Cage
A04	640

Note: Assumes chilled underfloor temperature of 24°C

Heat Output (kBTU/hr)

Model / Config	1 I/O Cage
A04	512 - 1605

z890 Dimensions

	z890
# of Frames	1 Frame IBF contained within the single frame
Height (w/ covers) Width (w/ covers) Depth (w/ covers)	194.1 cm / 76.4 in (40 EIA) 78.5 cm / 30.9 in 157.7 cm / 62.1 in
Height Reduction Width Reduction	178.5 cm / 70.3 in (38 EIA) None
Machine Area Service Clearance	1.24 sq. meters / 13.31 sq. feet 3.03 sq. meters / 32.61 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 Toleration for >15 LPAR ID on z990, z890 Enhanced Storage Protect DB2 Performance Message Time Ordering
11	9672 G5/G6 CF Duplexing Toleration for >15 LPAR ID on z990, z890
10	z900 GA2 support
9	MQSeries® shared message queues WLM Multi-system Enclave Support Intelligent Resource Director** IC-3/ISC-3/ICB-3 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, z890 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

**zSeries 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 world class 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 address a solution for your fiber cable installation. Enterprise fiber cabling offers two additional options to help meet your structured (trunking) environments requirements.

zSeries fiber cabling options:

- *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-route 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 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, helps reduce the cable bulk associated with the installation of potentially 240 (z800) to 256 (z900) to 420 (z890/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 is designed to deliver 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 helps protect 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, TCP/IP networking capability and complies with industry standards.

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.

This integrated set of system services can help a customer to 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 designed to be minimized and problem determination is facilitated with the sophisticated recovery, reporting and debug facilities of z/OS. 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 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 and z890 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 exploit the fast growing world of next generation e-business by making the deployment of new applications on the z800 and z890 very attractively priced.

z/OS.e uses the same code base as z/OS with custom parameters and invokes an operating environment that is comparable to z/OS in service, management, reporting, and reliability. In addition, z/OS.e invokes zSeries hardware functionality just as z/OS does. No new z/OS skills and service procedures are required for z/OS.e.

z/OS.e is specifically for Java, Enterprise Java, C/C++, and Web-based data transaction processing applications, giving these workloads a price/performance customers expect.

z/OS.e disables traditional workloads. z/OS.e will not execute CICS, IMS, COBOL, or FORTRAN applications. However, pre-compiled COBOL DB2 stored procedures and other pre-compiled COBOL applications using the Language Environment® preinitialization interface (CEEPIPI) are supported. You cannot use the following compilers: COBOL, PL/I, VisualAge® PL/I, and FORTRAN. However, z/OS.e supports execution of pre-compiled PL/I and VisualAge PL/I applications.

At only a fraction of the cost of z/OS, z/OS.e makes the decision to run new workloads on the mainframe easy due to its reduced total cost of ownership and exceptional robustness and functionality. z/OS.e and z800 or z890 server can help to reduce the total cost of ownership of hardware, software, people, and environments – making the combination very cost-effective for new applications.

With very few exceptions, z/OS.e has the same features and functionality as z/OS. As a result, unless otherwise specified, z/OS features described herein are applicable to z/OS.e running on a z800 or z890 server.

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 z/OS systems that are running as LPARs on the same CEC in the same Parallel Sysplex cluster. 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 an evolutionary step. It enables the system to continuously allocate resources for different applications, and this helps to reduce the total cost of ownership of the system. 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*

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 V1.2 and higher.

zSeries Application Assist Processor

The IBM @server zSeries Application Assist Processor (zAAP), planned to be available on the z990 and z890 servers, is an attractively priced specialized processing unit that provides an economical Java execution environment for z/OS for customers who desire the traditional qualities of service and the integration advantages of the zSeries platform.

When configured with general purpose processors within logical partitions running z/OS, zAAPs may help increase general purpose processor productivity and may contribute to lowering the overall cost of computing for z/OS Java technology-based applications. zAAPs are designed to

operate asynchronously with the general processors to execute Java programming under control of the IBM Java Virtual Machine (JVM). This can help reduce the demands and capacity requirements on general purpose processors which may then be available for reallocation to other zSeries workloads.

The IBM JVM processing cycles can be executed on the configured zAAPs with no anticipated modifications to the Java application(s). Execution of the JVM processing cycles on a zAAP is a function of the Software Developer's Kit (SDK) 1.4.1 for zSeries, z/OS V1.6, and the Processor Resource/Systems Manager (PR/SM).

z/OS V1.6 is planned to be available September of 2004.

64-bit Support

z/OS provides 64-bit real storage and virtual storage support for the zSeries servers, while continuing to support 24-bit and 31-bit applications.

The 64-bit real support is intended to eliminate expanded storage, helps eliminate paging and may allow you to consolidate your current systems into fewer LPARs or to a single native image. z/OS V1.5 delivers 64-bit shared memory support to allow middleware to share large amounts of 64-bit virtual storage among multiple address spaces. This is expected to provide a significant capacity enhancement for relieving shared virtual storage constraints.

Applications that can be written to 64-bit virtual storage have significantly larger addressability to data. With z/OS 1.2, assembler programs can obtain virtual storage above 2 GB for storing and manipulating data. This 64-bit support is used by DB2 V8 and other middleware. Planned for z/OS 1.6 is C/C++ support for the development of 64-bit applications, including debug and runtime support.

Single image scale

z/OS V1.6 plans to support up to 24 CPs in a single LPAR on the z990. In conjunction with z/OS V1.6, the maximum number of combined zAAPs and z/OS V1.6 CPs supported in a single LPAR is planned to be 24.

Statement of Direction: IBM intends to support greater than 24 CPs, or combined CPs and zAAPs, in a single LPAR in the future on appropriate releases of z/OS in combination with designated zSeries server(s).

Automation Support

z/OS Managed System Infrastructure for Operations (msys for Operations) provides automation for single system and sysplex operations to help simplify operations and improve availability. msys for Operations plays an important role in outage avoidance.

Functions control and manage both hardware and software resources making fully automated solutions possible. 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.

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 V1.4, 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 has 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.

Also, with z/OS V1.5 msys for Setup can now use the IBM Directory Server, OpenLDAP, on any IBM platform. This includes OpenLDAP on z/OS UNIX System Services, which can make the Management Directory virtually transparent to the user.

The following functions can be configured using msys for Setup: Parallel Sysplex clustering, TCP/IP, UNIX System Service, Language Environment, LDAP. z/OS V1.5 has added support for configuring RMF™, ISPF, FTP, and DB2 for z/OS V8.

System Services

z/OS Version 1 Release 5 base elements

Base Control Program (BCP)
UNIX System Services
JES2⁴
ESCON Director Support
MICR/OCR Support²
Bulk Data Transfer base
DFSMSdftp
EREP/zSeries V3R5
High Level Assembler V1R4
ICKDSF R16
ISPF
TSO/E³
3270 PC File Transfer Program V1R1¹
FFST™/ESA
TIOC

z/OS Version 1 Release 5 optional priced features

DFSMSdss™
DFSMSHsm™
DFSMSrmm™
JES3⁴
Bulk Data Transfer (BDT) File to File¹
BDT, SNA, NJE

1. Disabled on z/OS.e

2. Not licensed on z/OS.e

3. TSO usage is limited to 8 concurrent users on z/OS.e

4. z/OS.e only supports the latest JES that ships with the product

The backbone of the z/OS system is the Base Control Program (BCP) with JES2 or JES3. These provide the essential services that make z/OS the system of choice when workloads must be processed reliably, securely, with complete data integrity and without interruption. The BCP includes the I/O configuration program (IOCP), the Workload Manager (WLM), systems management facilities (SMF), the z/OS UNIX Systems Services kernel, and support for Unicode.

Sense and Respond with Workload Manager

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. Performance management goals are expressed in Service Level Agreement terms. 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 a sysplex has been enhanced in z/OS V1.4 to provide better utilization of processor resources. WLM is designed to check 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 helping to ensure that processors are more evenly utilized.

WLM Improvements for WebSphere

z/OS V1.5 can simplify WLM control for WebSphere. Customers now have the choice to manually define WebSphere application environments for WLM or have WebSphere define them as and when required.

Performance block reporting for enclaves and multi-period classes are designed to provide 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.

z/OS UNIX System Services

z/OS UNIX System Services are an integral element of z/OS and are a key element of the zSeries' open and distributed computing strategy. Many middleware and application products that run on z/OS use z/OS UNIX System Services.

z/OS UNIX System Services has a hierarchical file system familiar to UNIX users. Applications can work with data in both UNIX hierarchical file systems and traditional zSeries data sets. The SMB File and Print Server enables a distributed file sharing infrastructure for z/OS HFS files and Windows workstations. Users can choose which interface they want to use: the standard shell, 3270, and ISPF interfaces.

Data Management with DFSMS

DFSMS comprises a suite of related data and storage management functions for the z/OS system. DFSMS is an operating environment that helps automate and centralize the management of storage based on the policies that your installation defines for availability, performance, space, and security. The heart of DFSMS is the Storage Management Subsystem (SMS). Using SMS, the storage administrator defines policies that can automate the management of storage and hardware devices. These policies

describe data allocation characteristics, performance and availability goals, backup and retention requirements, and storage requirements for the system.

DFSMSdftp is automatically included with z/OS. DFSMSdftp performs the essential data, storage, and device management functions of the system. DFSMSdftp and DFSMSshsm can provide disaster recovery functions such as Advanced Copy Services and aggregate backup and recovery support (ABARS).

The other elements of DFSMS – DFSMSdss, DFSMSshsm, DFSMSrmm, and DFSMStvs – complement DFSMSdftp to provide a comprehensive approach to data and storage management. In a system-managed storage environment, DFSMS can automate and centralize storage management based on the policies that your installation defines for availability, performance, space, and security. With these optional features enabled, you can take full advantage of all the functions that DFSMS offers.

z/OS V1.5 can help significantly enhance application backup with enhancements to DFSMSshsm to utilize volume level fast replication. The fast backup is designed to exploit the FlashCopy and virtual concurrent copy capability of IBM TotalStorage Enterprise Storage Server and IBM RAMAC® Virtual Array (RVA) respectively. DFSMSshsm Fast Replication in z/OS V1.5 is also intended to provide a fast, easy to use point-in-time backup and recovery solution specifically designed for DB2 Universal Database (UDB) for z/OS V8. It is designed to allow fast, nondisruptive backups to be taken at appropriate events when there is minimum activity at the application level or when a fast point-in-time backup is desired.

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 33, Parallel Sysplex Clustering Technology).

CICS/VSAM enabled for 24x7 availability

DFSMS Transactional VSAM Services (DFSMSStvs) allows you to share VSAM data sets across CICS, batch, and object-oriented applications on z/OS or distributed systems. DFSMSStvs can enable concurrent shared updates of recoverable VSAM data sets by CICS transactions and multiple batch applications. DFSMSStvs can enable 24-hour availability of CICS and batch applications. DFSMSStvs is built on top of VSAM record-level sharing (RLS), which permits sharing of recoverable VSAM data sets at the record level.

System-Managed CF Structure Duplexing

System-Managed Coupling Facility (CF) Structure Duplexing is designed to significantly enhance Parallel Sysplex availability. It can provide a robust failure recovery capability via CF structure redundancy, and it can enhance Parallel Sysplex ease of use by helping to reduce the complexity of CF structure recovery. These benefits can be 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 can be initiated and quickly completed transparent to the CF structure user and without manual intervention.

JES2 and JES3

JES2 and JES3 handle the submission of work for the Base Control Program. JES2 exercises independent control over its job processing functions, whereas JES3 exercises centralized control.

- In z/OS V1.2, 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 V1.4, 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 V1.4, JES3 provides additional tolerance for initialization errors and the MAINPROC refresh function which provides dynamic addition of systems to the sysplex

System Management Services

z/OS Version 1 Release 5 base elements

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

z/OS Version 1 Release 5 optional priced 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

- *z/OS V1.4 Consoles Enhancements Feature is an optional no-charge feature which is designed to improve system availability by enhancing the capacity and reliability of message delivery. Major changes to the message production and consumption flow can help reduce the possibility of bottlenecks which can cause of backlog of undelivered messages. This console enhancement will be integrated into z/OS V1.5*
- *msys for Setup has been enhanced in z/OS V1.4 to allow multiple users to log 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. 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.*

The following functions can be configured using msys for Setup: Parallel Sysplex clustering, TCP/IP, UNIX System Service, Language Environment, LDAP. z/OS V1.5 has added support for configuring RMF, ISPF, FTP, and DB2 for z/OS V8.

The msys Management Directory in z/OS 1.5 includes new options to use the IBM Directory Server, OpenLDAP, on any IBM platform.

z/OS msys for Operations is a base element in z/OS V1R2 that incorporates automation technology into z/OS. It provides self-healing attributes for some critical system and

sysplex resources and can simplify the day-to-day operation of a single z/OS image or of a Parallel Sysplex cluster.

- *msys for Operations enhancements in z/OS V1.3 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 deactivating 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*
- *Application State Recording, a new feature of z/OS V1.4 provides more granular performance reporting for middleware such as WebSphere*
- *In z/OS V1.5, RMF Monitor II and Monitor III performance data is now RACF protected.*

Enhancements in z/OS V1.5 to the RMF Spreadsheet Reporter include a new and intuitive resource oriented flow. Working set generation is simplified to one single step and a batch facility is now available which can run without user interaction. Benefits can include:

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

System Automation

The unique and rich functions of System Automation for OS/390 (SA OS/390) Version 2.2 (separately orderable) can ease z/OS, and z/OS.e management, reduce costs, and increase application availability. SA OS/390 automates I/O, processor, and system operations, and includes “canned” automation for IMS, CICS, Tivoli OPC, and DB2. Its focus is on Parallel Sysplex automation, including multi- and single-system configurations, and on integration with end-to-end Tivoli enterprise solutions. With the new patented manager/agent design, it is now possible to automate applications distributed over a sysplex by virtually removing system boundaries for automation.

Security Services

z/OS Version 1 Release 5 base elements

Integrated Security Services include:

- Public Key Infrastructure Services
- DCE Security Server
- Open Cryptographic Enhanced Plug-ins
- Firewall Technologies
- LDAP Server
- Network Authentication Service
- Enterprise Identity Manager

Cryptographic Services:

- Integrated Cryptographic Service Facility
 - Open Cryptographic Services Facility
 - System SSL
 - Data Confidentiality Services
-

z/OS Version 1 Release 5 optional priced features

Security Server includes:

- RACF
-

z/OS Version 1 Release 5 optional no-charge features

- Network Authentication Service Level 3
 - System SSL Security Level 3
 - Open Cryptographic Services Facility Security Level 3
-

z/OS extends its robust mainframe security features to address the demands of e-business enterprises. Technologies such as LDAP, Secure Sockets Layer (SSL), Kerberos V5, Public Key Infrastructure, and exploitation of zSeries cryptographic features are available in z/OS.

z/OS V1.5 together with DB2 Universal Database for z/OS Version 8, provide a solution for multilevel security on the zSeries mainframe. This support will provide row-level security labeling in DB2 and protection in z/OS designed to meet the stringent security requirements of multi-agency access to data. This solution leverages zSeries leadership in scale, high availability, and self-managing capabilities for a secure single-database hosting infrastructure.

- *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) services supporting thousands of concurrent clients.*

RACF

Resource Access Control Facility (RACF) provides the functions of authentication and access control for z/OS resources and data, including the ability to control access to DB2 objects using RACF profiles. Using an entity known as the RACF user ID, RACF can identify users requesting access to the system. The RACF user password (or valid substitute, such as a RACF PassTicket or a digital certificate) authenticates the RACF user ID.

Once a user is authenticated, RACF and the resource managers control the interaction between that user and the objects it tries to gain access to. These objects include: commands, datasets, programs, tape volumes, terminals and objects that you define. RACF supports flexibility in auditing access attempts and changes to security controls. To audit security-relevant events, you can use the RACF system management facility (SMF) unload utility and a variety of reporting tools.

With one command, a security administrator can update remote RACF databases without logging on to remote systems. Throughout the enterprise, RACF commands can be sent automatically to synchronize multiple databases. In addition, RACF can automatically propagate RACF database updates made by applications. With RACF, users can keep passwords synchronized for specific user IDs. When you change one password, RACF can change passwords for your user ID on different systems and for several user IDs on the same system. Also, passwords can be changed automatically for the same user ID on different systems. This way, several RACF databases can be kept synchronized with the same password information.

Recent RACF enhancements:

- *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 the creation of a new kind of Group that can contain an unlimited number of users.*

Multilevel Security

z/OS V1.5 is the first and only IBM operating system to provide Multilevel Security. This technology can help improve the way government agencies and other organizations share critical classified information. Combined with IBM's DB2 UDB for z/OS Version 8, z/OS provides multilevel security on the zSeries mainframe to help meet the stringent security requirements of government agencies and financial institutions, and can help open up new options for e-hosting facilities. Multilevel security technology allows IT administrators to give users access to information based on their need to know, or clearance level. It is designed to prevent individuals from accessing unauthorized information and to prevent individuals from declassifying information.

With multilevel security support in IBM's z/OS V1.5 and DB2 V8, customers can enable a single repository of data to be managed at the row level and accessed by individuals based on their need to know.

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 the AES Symmetric Cipher for SSL V3 and TLS Connections: System SSL supports the Advanced Encryption Standard (AES), which provides data encryption using 128-bit or 256-bit keys for SSL V3.0 and TLS V1.0 connections.*
- *Support for DSS (Digital Signature Standard) Certificates: System SSL has been enhanced to support Digital Signature Standard certificates defined by the FIPS (Federal Information Processing Standard) 186-1 Standard.*
- *System SSL of RSA Private Keys Stored in ICSF: With z/OS V1R4, support is introduced that is designed to 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.*
- *Failover 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.*
- *TLS: z/OS LDAP now provides support for TLS (Transport Layer Security) as defined in IETF RFC 2830 as an alternative to SSL support. It also provides support, via an LDAP extended operation, that allows applications to selectively activate TLS for certain LDAP operations at the application's discretion.*

ICSF

Integrated Cryptographic Service Facility (ICSF) is a part of z/OS which provides cryptographic functions for data security, data integrity, personal identification, digital signatures, and the management of cryptographic keys. These functions are provided via APIs intended to deliver the highly scalable and available security features of z/OS and the zSeries servers. Together with cryptography features of the IBM zSeries servers, z/OS is designed to provide high performance SSL, which can benefit applications that use System SSL, such as the z/OS HTTP Server and WebSphere, TN3270, and CICS Transaction Gateway server.

ICSF provides support for the z990 and z890 PCIX Cryptographic Coprocessor (PCIXCC), a replacement for the PCICC and the CMOS Cryptographic Coprocessor Facility that were found on the z900 and z800. All of the equivalent PCICC functions offered on the PCIXCC are expected to be implemented with higher performance. In addition, PCIXCC implements the functions on the CMOS Cryptographic Coprocessor Facility used by known applications. PCIXCC supports secure cryptographic functions, use of secure encrypted key values and user-defined extensions.

LDAP

- *Client access to information in multiple directories is supported with the LDAP protocol.*
- *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 V1.4: 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. The LDAP Server, the C/C++ APIs, and the utilities are updated with this support. Interoperability is improved for any applications that make use of these methods.*

PKI Services

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

Firewall

- Firewall Technologies provide sysplex-wide Security Association Support: This function is designed to enable VPN (virtual private network) security associations to be dynamically reestablished 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 is designed to 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

- 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 V1.4 for improved network security scalability.

- Kerberos in z/OS V1.4 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 5 base elements

Language Environment³
C/C++ IBM Open Class[®] Library
DCE Application Support¹
Encina[®] Toolkit Executive²

z/OS Version 1 Release 5 optional priced features

C/C++ with Debug Tool
C/C++ without Debug Tool
DFSORT
GDDM[®]-PGF V2R1.3¹
GDDM-RE90 V3R2¹
HLASM Toolkit V1R4

1. Disabled on z/OS.e

2. Not licensed on z/OS.e

3. z/OS.e only supports the latest Language Environment that ships with this product.
See LE limitations for z/OS.e.

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

Highlights

Language Environment

Language Environment is a base element of z/OS and provides the run-time environment for programs generated with C, C++, COBOL, FORTRAN, and PL/1.

C/C++

Extra Performance Linkage (XPLINK) is provided in z/OS V1.2.

A C or C++ application has overhead associated with each function call. The more highly functionalized a program, the more overhead. XPLINK helps cut down on the overhead associated with these function calls and can improve 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 rewritten to take advantage of this unique z/OS function, which can result in performance improvements.

z/OS V1.2 and V1.3 provide two levels of the C/C++ compiler as a migration aid. One is identical to the compiler provided with OS/390 V2.10. 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.

z/OS V1.5 includes several C/C++ enhancements:

- *Support for 64-bit compiles: The compiler has been enhanced to generate z/Architecture instructions which include utilizing the 64-bit general purpose registers. This new 64-bit support will enable C and C++ developers to recompile existing 32-bit C/C++ applications into 64-bit code and to compile new 64-bit C/C++ code. Object code is not generated in this release.*

- *Performance enhancements: A new higher optimization level, OPTIMIZE(3), provides the compiler's highest and most aggressive level of optimization. OPTIMIZE(3) is suggested when the desire for run-time improvement outweighs the concern for minimizing compilation resources.*
- *DB2 preprocessor integration: The C/C++ compiler has been enhanced to integrate the functionality of the DB2 precompiler. A new SQL compiler option enables the compiler to process embedded SQL statements.*

C/C++ IBM Open Class Library: As previously announced in Software Announcement 203-131, dated May 13, 2003, the application development support (that is, the headers, source, sidedecks, objects, and samples from the Application Support Class and Collection Class libraries) is withdrawn from the C/C++ IBM Open Class Library (IOC) in z/OS V1.5. Applications that use these IOC libraries cannot be compiled nor linked using z/OS V1.5. Run-time support for the execution of existing applications which use IOC libraries is provided with z/OS V1.5, but is planned to be removed in a future release.

z/OS 64-bit C/C++ environment: z/OS V1.6 is planned to deliver the capability to exploit 64-bit virtual in developing and deploying emerging e-business applications that require a significantly larger addressability of data. This capability is planned to be provided with enhanced UNIX System Services, 64-bit Language Environment (LE) runtime developed with the C/C++ compiler 64-bit support, and the Program Management Binder 64-bit support. The availability of this support will complete the major steps of the z/OS 64-bit virtual roadmap.

z/OS V1.6 C/C++ is planned to support the generation and debugging of 64-bit applications. This capability is expected to remove the current limitations on virtual

memory and address space sizing, and aid in porting of applications from other platforms. The C/C++ compiler is planned to support the industry-standard LP64 programming model to allow C/C++ code to exploit virtual memory above the bar. In addition, the architecture level of ARCH(6) will be supported to allow for exploitation of the new z990 hardware features for performance enhancements.

Debugger support: With z/OS V1.6, dbx is planned to support 64-bit virtual C/C++ code. In addition, Unicode support for “wide-character ASCII,” which is a subset of UCS-2 within dbx, is planned to enhance the ability of dbx to be recognized for ASCII support by providing full debugging support of wide characters.

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 and Web services*
- *Continue to take advantage of:*
 - *Common cross platform programming Security APIs within Java framework*
 - *Java Record Input/Output (JRIO) APIs to provide record-oriented access to VSAM datasets, System catalogs, and PDS directory*
 - *Persistent reusable JVM technology for CICS, IMS, and DB2*

- *Leverage traditional zSeries software and server benefits: scalability, reliability, availability, performance and serviceability*

The Java SDK for z/OS is available via download from the IBM **@server** zSeries Java Web site and by tape from IBM Software Delivery and Fulfillment (SDF) in SMP/E format.

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

Unicode

z/OS provides 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.*

RE90 Functions

z/OS V1.4 extends the RE90 language on z/OS when used in a UNIX System Services zSeries RE90 environment. It includes functions for standard RE90 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 RE90 program and enter an immediate RE90 command such as TS to start interactive tracing.

Network Communication Services

z/OS Version 1 Release 5 base elements

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

OSA Support Facility

z/OS Version 1 Release 5 optional no-charge features

z/OS Communications Server Security Level 3

The z/OS base includes z/OS Communication 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.

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 z800/z900 and S/390 networking environment.

Virtual IP Address Nondisruptive Takeover

VIPA represents an IP address that is not tied to a specific hardware adapter address. The benefit can be that if an adapter fails, the IP protocol can find an alternate path to the same software, be it the TCP/IP services on a zSeries server z800/z900 or an application.

In case of a failure of the primary IP stack, VIPA Takeover introduced in OS/390 V2.8 can support movement to a backup IP stack on a different server in a Parallel Sysplex

cluster. VIPA Nondisruptive Takeover can enhance the initial V2.8 functions, providing VIPA takeback support. This can allow the movement of workload back from the alternate to the primary IP stack.

With Sysplex-Wide Security Associations (SWSA) in z/OS V1.4, IPsec protected workloads are expected to now realize all the benefits derived from workload balancing, such as optimal routing of new work to the target system and server application based on QoS and WLM advice, increased availability by routing around failed components and increased flexibility in adding additional workload in a nondisruptive manner.

Sysplex Distributor

Introduced in OS/390 V2.10, 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 z/OS images on servers on different zSeries z800/z900 or S/390 servers. In addition to load balancing, Sysplex Distributor simplifies the task of moving applications within a Parallel Sysplex environment.

In z/OS, we have taken the functions provided by the Cisco MNLB Workload Agent and Systems Manager, and integrated them into Enhanced Sysplex Distributor. This can eliminate the need for separate Cisco LocalDirector machines in the network and the need for MNLB workload agents to be run on the zSeries servers. It can also improve performance, while allowing the Sysplex Distributor to decide, based on priority supplied by WLM, the Service Policy Agent and the TCP/IP stack status, which application instance to forward the packet to.

Security Server (RACF) in z/OS V1.4 supports Enterprise Identity Mapping (EIM). EIM defines a user's security context that is consistent throughout an enterprise, regardless of the User ID used and regardless of which platform the user is accessing. RACF commands are enhanced to allow a security administrator to define EIM information for EIM applications to use. The EIM information consists of the LDAP host name where the EIM domain resides, the EIM domain name, and the bind distinguished name and password an application may use to establish a connection with the domain.

Intrusion Detection Services (IDS) Introduced in z/OS V1.2 and enhanced in V1.5, enables the detection of attacks on the TCP/IP stack and the application of defensive mechanisms on the z/OS server. The focus of IDS is self-protection. IDS can be used alone or in combination with an external network-based Intrusion Detection System. IDS is integrated into the z/OS Communications Server stack.

Highlights

- *IPv6 support provides increased scalability and simplified configuration.*
- *HiperSockets, introduced in z/OS V1.2, 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.*

Significant enhancements in z/OS V1.5. OSA-Express Queued Direct Input/Output support (now includes IP Broadcast support, used by DHCP, for IPv4)

- *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*
- *Multinode Persistent Sessions for SNA applications running in a Parallel Sysplex environment*
- *Simple Network Time Protocol Support (SNTP) for client/server synchronization*

Benefits can include:

- *Function for new e-business Internet and intranet applications*
- *Multivendor, multiplatform connectivity*
- *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*

Further improvement have been made 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 may 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.

UNIX System Services

z/OS Version 1 Release 5 base elements

z/OS UNIX System Services

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 zSeries system strengths in an open environment

Highlights

- *X/Open UNIX 95 Branded*
- *Permanent z/OS UNIX Kernel (restart capability provided in z/OS V1.3)*
- *Integrated/converged z/OS UNIX Sockets*
- *Web application and UNIX C program performance improvements*
- *Improved z/OS UNIX setup*
- *Multiprocess/Multiuser Kernel Support*
- *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 can allow 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) can reduce real storage when sharing large amounts of virtual storage.*
- *UNIX System Services and UNIX debugger add support for IEEE-floating point.*

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*
- *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 V1.4. 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. The availability of this function will improve performance and workload balancing.*

Distributed Computing Services

z/OS Version 1 Release 5 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).

zSeries File System (zFS)

zFile System (zFS) is a UNIX file system that complements the z/OS Hierarchical File System (HFS). It uses the same APIs as HFS.

zFS can provide the following benefits over HFS:

- *Improved performance*
- *Additional function*
 - *Disk space can be shared between file systems in the same data set*
 - *File system quota (maximum file system size). With zFS the file system quota can be increased with a simple administrative command*
 - *Can improve failure recovery. zFS performs asynchronous writes to disk and does not wait for a synch interval to begin writes*

zFS or zFS-related administration, system management, performance, configuration support, and scalability have been further enhanced in z/OS V1.4:

- *Dynamic reconfiguration for file system configuration options.*
- *Dynamic use of secondary allocation for a zFS aggregate (data set) or file system.*
- *Improvements in the UNIX System Services automount support for zFS.*
- *Addition of ISHELL support for zFS.*
- *Ability to perform I/O operations in parallel for a zFS aggregate that spans multiple DASD volumes. This is designed to provide improved performance when using multi-DASD volume aggregates.*
- *Support for 64-bit user virtual buffer addresses.*

zSeries File System (zFS) is the strategic UNIX Systems Services file system for z/OS. The Hierarchical File System (HFS) functionality has been stabilized. HFS is expected to continue shipping as part of the operating system and will be supported in accordance with the terms of a customer's applicable support agreement. IBM intends to continue enhancing zFS functionality, including RAS and performance capabilities, in future z/OS releases. All requirements for UNIX file services are expected to be addressed in the context of zFS only.

Distributed File Services (DFS) Server Message Block (SMB)

Microsoft Windows networking compatible file and print serving is available in z/OS 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.

SMB Support has been further enhanced in z/OS V1.4 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 5 base elements

IBM HTTP Server
Text Search

z/OS Version 1 Release 5 optional no-charge features

IBM HTTP Server North America Secure

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:

- *Utilization of large storage capacity*
- *Single point of entry and control*
- *Consolidation of multiple Web sites*
- *Exploitation of z/OS WLM*

Print Services

z/OS Version 1 Release 5 optional priced features

Infoprint Services:
- IP Printway™
- NetSpool™
- z/OS Print Interfaces

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, high availability, 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, application output capture program, and print drivers and management tools that let you manage any print job to any printer defined to Infoprint Server.

Enhancements in z/OS V1.5 include:

Infoprint Central: Continuing the IBM commitment to help lower customers' overall cost of distributed print operations, Infoprint Server includes a new component called Infoprint Central. Infoprint Central is a Web-based, GUI for managing print jobs and printers throughout the enterprise from anywhere in the enterprise using a Web browser. Intended primarily for help desk operators, it lets users query the status of jobs and printers, see job and printer messages, stop and start printers, move jobs from one printer to another, cancel or hold jobs, and many other functions. Infoprint Central can use integrated z/OS security services so that users can be authorized to perform only certain tasks, or to perform tasks only on designated devices.

IP PrintWay extended mode: Infoprint Central is backed by an architecture in the component that delivers print or e-mail output to printers, servers or users over TCP/IP or Internet Printing Protocol (IPP). IP PrintWay extended

mode uses the Sysout Application Programming Interface (SAPI) to access print jobs and job information from the JES spool. The advantage of this change can be higher availability and throughput, more flexibility for handling print-related tasks, and scalability of Infoprint Server for very large distributed print environments.

Common message log: A new common message log helps to improve productivity of help desk operators for print problem diagnosis and resolution, thus helping to increase system availability and user satisfaction. Messages can easily be accessed from Infoprint Central for a particular job or printer.

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 5 base elements

BookManager® READ V3¹
 Library Server
 GDDM (including PLCK and OS/2® Link and RE90 code)¹

z/OS Version 1 Release 5 optional no-charge features

BookManager Build¹

1. Disabled on z/OS.e

BookManager READ is used to display, search, and manage online documents and bookshelves. BookManager BUILD is an optional feature that allows the creation of softcopy documents that can be used by any of the BookManager products.

Library Server converts BookManager documents to HTML for display through a Web browser.

Library Center

IBM is providing an alternative way to navigate our z/OS library on the Internet. Beginning with z/OS V1.5, the Library Center for z/OS provides a Microsoft Windows Explorer-like view of the contents of the entire z/OS and Software Products DVD Collection. The Library Center uses the new IBM Library Server with new advanced search functions to help users find information “on demand.”

The Library Center offers:

- *Easier navigation: When opened, the Library Center presents a left-hand navigation frame that lets the user visually navigate the entire repository (over 2000 manuals). The user can expand each shelf to see all the books in that shelf and expand each book to see all the topics.*
- *New advanced search features: A search scope pull-down on the Library Center lets the user search the entire repository, a particular shelf, or a specific book.*
- *Available single frame navigation: A single frame mode makes the new Library Center more accessible to people using screen readers, such as the IBM Homepage Reader.*
- *A portal to other information: An IBM Redbooks bookshelf lets the user perform a BookManager search and locate a corresponding Redbook in PDF format. The search scope pull-down lets the user launch searches in other repositories such as the WebSphere Application Server for z/OS or Google.*
- *Built-in handheld support: The Library Center also provides a handheld mode to support both connected and disconnected handhelds*

Integrated Testing

To further reduce migration time, the z/OS and z/OS.e products are 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.

Publications

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

Installation Considerations

Customers are able to upgrade to z/OS V1.5 from z/OS V1.2, V1.3 or V1.4 in a single step.

CustomPac is a suite of services designed to help you efficiently install, migrate and maintain a z/OS system. It can also help with migrating and maintaining z/OS system-related products and/or third parties' software vendor products. Options include:

- *RefreshPac®*, which includes preventative software services
- *ProductPac®* for custom-built products
- *SystemPac®* for installation or system replacement

Highlights

- *Customized services and products to help meet customer specific needs*
- *Efficient and effective exploitation of new z/OS functions*
- *A solution for businesses that lack skilled system programmers*

- *Installation of IBM and selected third-party software vendor products in one package*

z/OS.e can be ordered by ServerPac and SystemPac only.

For more information visit: ibm.com/ca/custompac

Note: Workload Manager Goal Mode is required for z/OS V1.3 and later.

z/OS V1.4 and V1.5 will be Supported on the Following IBM Servers:

- *S/390 Parallel Enterprise Server™ G5/G6 (or compatible server) (ESA/390 mode)*
- *Multiprise® 3000 Enterprise Server (or compatible server) (ESA/390 mode)*
- *IBM @server zSeries 800, 890, 900 and 990 servers (or compatible servers) (z/Architecture mode)*

Migration/Coexistence

General coexistence, release migrations, and fallback: z/OS continues to deliver compatibility and flexibility as the user migrates systems in a multisystem configuration by allowing several releases of z/OS, OS/390, and z/OS.e to coexist. This includes Parallel Sysplex and non-Parallel Sysplex multisystem configurations.

Coexistence allows systems within a multisystem configuration to be upgraded to a new release level of the operating system one system at a time. This is contingent on the fact that the release the user is migrating to can coexist with the lowest release running in the user's multisystem configuration.

IBM has a consistent coexistence, migration, and fallback policy. Migration forward as well as backward should be made within the same releases supported by the coexistence policy.

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.*

For additional information on 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
z/OS R2 z/OS R3	z/OS R2, {z/OS R1, OS/390 R10} ¹ , OS/390 R9, R8 {z/OS.e R3, z/OS R3}, z/OS R2, {z/OS R1, OS/390 R10} ¹ , OS/390 R9
z/OS.e R3	{z/OS.e R3, z/OS R3}, z/OS R2, {z/OS R1, OS/390 R10}, OS/390 R9
z/OS R4	{z/OS.e R4, z/OS R4}, {z/OS.e R3, z/OS R3}, z/OS R2, {z/OS R1, OS/390 R10}
z/OS.e R4	{z/OS.e R4, z/OS R4}, {z/OS.e R3, z/OS R3}, z/OS R2, {z/OS R1, OS/390 R10}
z/OS R5	{z/OS.e R5, z/OS R5}, {z/OS.e R4, z/OS R4}, {z/OS.e R3, z/OS R3}, z/OS R2
z/OS.e R5	{z/OS.e R5, z/OS R5}, {z/OS.e R4, z/OS R4}, {z/OS.e R3, z/OS R3}, z/OS R2
z/OS R6	{z/OS.e R6, z/OS R6}, {z/OS.e R5, z/OS R5}, {z/OS.e R4, z/OS R4}, {z/OS.e R3, z/OS R3}
z/OS.e R6	{z/OS.e R6, z/OS R6}, {z/OS.e R5, z/OS R5}, {z/OS.e R4, z/OS R4}, {z/OS.e R3, z/OS R3}

¹. OS/390 R10 and z/OS R1 treated as one coexistence level

This information is being provided to you to facilitate release planning and to help ensure appropriate positioning for future software migrations.

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 and z/OS.e planning, refer to *z/OS and z/OS.e Planning for Installation*, available on the z/OS and z/OS.e Web sites in the z/OS Library at: ibm.com/servers/eserver/zseries/zos and ibm.com/servers/eserver/zseries/zose.

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.

Migration, installation and customization enhancements

For information on migration to the latest z/OS releases, see the z/OS Migration Web page ibm.com/zseries/zos/migration. The publications, *z/OS V1R5.0 Migration and z/OS V1R4 Migration*, are available from this site for helping to plan the migration to z/OS V1.4 or V1.5.

zSeries Bimodal Support for z/OS

IBM provides is announcing the z/OS Bimodal Migration Accommodation Offering to assist customers in migrating from OS/390 to z/OS. This addresses many customer requests to have a “fallback” option to 31-bit mode when first migrating to z/OS in 64-bit mode on a z/Architecture server. IBM has reconsidered the investment in 31-bit support on a zSeries server to encourage our customers to move forward. This offering is available for 6 months for each z/OS license (5694-A01) starting from the registration of a z/OS license to a z/Architecture server. This offering is available for It only applies to z/OS Version 1 Release 2, 3, and 4, and will not be provided for z/OS V1.5 or for any z/OS.e releases.

Wizards

z/OS wizards are available on the Internet to help to set up some of the z/OS functions. These wizards are interactive assistants that ask a series of questions about the task to be performed (for example, setting up a Parallel Sysplex environment). The wizards can simplify planning and configuration by exploiting recommended values and by building customized checklists. For configuration tasks, these wizards also generate outputs like jobs, policies, or parmlib members that can be uploaded to z/OS.

For more information about wizards visit the Web site at: www-1.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) engines 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 V5 offers new levels of price/performance, functional capabilities, and hardware exploitation that increase the attractiveness of deploying Linux solutions on the mainframe. You can add capacity to existing zSeries systems for hosting Linux-on-z/VM workloads by configuring their server with IFL engines. z/VM V5 is the follow-on product to z/VM V4.

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 on zSeries as guest systems of z/VM, in addition to ESA/390 guest operating systems such as OS/390, VSE/ESA, TPF, and Linux on 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 can also reduce 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 on 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 Sockets 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, z890, z900, z800, S/390 Parallel Enterprise Server Generation 5 and 6, S/390 Multiprise 3000, or equivalent servers. Support for IFL processor features by z/VM V4 or later 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, z890, z900, z800, and S/390 servers in other logical partitions.

Engine-based pricing for z/VM V4 and its optional features allows customers the opportunity 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 engines. Traditional operating systems such as z/OS, z/OS.e, OS/390, TPF, VSE/ESA, z/VM V3.1, or VM/ESA are not supported nor can they operate on IFL engines. Only Linux workloads in an LPAR or Linux guests of z/VM V4 or later can operate on the IFL engine.

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 V4.4 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 (HMC) 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 standalone 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 similar to the way VM data is viewed and presented.*

With corresponding function available in Linux on zSeries and S/390, z/VM V4.4 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, z890, z900, or z800 server equipped with the SCSI IPL Feature Enabler*
- *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 on 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 V4.4 provide:

- *Exploitation of the zSeries 890 and 990 server*
 - *Extending Dynamic-I/O configuration support allows channel paths, control units, and devices to be dynamically added, changed, and deleted in a Logical Channel SubSystem (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, 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 that provides a complete, consistent, and coherent copy of data at a remote site.
- Support for IBM TotalStorage Enterprise Tape Controller 3592 Model J70 and Tape Drive 3592 Model J1A

Systems Management

Improvements in systems management, some of which help to provide self-configuring, self-managing, and self-optimization facilities in z/VM V4.4, 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 the 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 V4.4 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 V4.4.

Networking with z/VM

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, z890, z900, and z800. TCP/IP broadcast support is now available for the HiperSockets environment when utilizing Internet Protocol version 4 (IPv4) with z/VM V4.4. Applications that use the broadcast function can now propagate frames to all TCP/IP applications.

The z890 and z990 servers include an important performance enhancement that virtualizes adapter interruptions and can be used with V=V guests (pageable guests) on z/VM V4.4. With the enhancement of the TCP/IP stack in z/VM V4.4 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 V4.4 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 V4.4 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 V4.2 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 V4.4 provides appropriate RPM format packages for the SUSE LINUX Enterprise Server 7 (SLES 7) at the 2.4.7 kernel level, SUSE LINUX Enterprise Server 8 (SLES 8) powered by UnitedLinux at the 2.4.19 kernel level, and Turbolinux Enterprise Server 8 (TLES 8) powered by UnitedLinux at the 2.4.19 kernel level. Security of the TCP/IP stack has been improved to help prevent additional

types of Denial of Service (DoS) attacks including: Smurf, Fraggle, Ping-o-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.

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 and V5. 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.

In addition to the new function provided by the Performance Toolkit for VM, RealTime Monitor (RTM), and Performance Reporting Facility (PRF) are still available in z/VM V4.4 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™).

z/VM Version 5 (V5)

z/VM Version 5 Release 1 (V5.1) continues the evolution of its premier and world-class zSeries virtualization technology with a new version that can offer traditional capabilities to manage zSeries operating systems, including Linux

on a single mainframe as guests of z/VM. z/VM V5.1 is designed to operate only on zSeries servers that support the z/Architecture (64-bit) including the z990, z890, z900, and z800 or equivalent.

Engine-based Value Unit Pricing

z/VM V5 introduces engine-based Value Unit pricing which replaces the per-engine pricing model that is available with z/VM V4 as well as providing a lower entry price. Engine-based Value Unit pricing is designed to provide a decreasing price curve which may help provide improved price/performance as hardware capacities and workload grow. Value Unit pricing for z/VM V5 can provide for a lower price per processor engine as more processor engines are licensed with z/VM V5.1 across the enterprise.

Value Unit pricing helps you to:

- *Add capacity and workload with an incremental and improved price*
- *Manage software costs better*
- *Aggregate licenses acquired across machines that are part of your enterprise.*

Engine-base Value Unit pricing of z/VM V5 should not be tied, or associated with, MSU-based Value Unit pricing.

Enhancements in z/VM V5.1 include:

Virtualization Technology and Linux Enablement

- *Support for SCSI FCP disks enable the deployment of a Linux server farm on z/VM using only SCSI disks. SCSI disks can be used as such by guests through dedicated FCP subchannels, and are also supported as emulated 9336 Fixed-Block Architecture (FBA) devices for use by guests, CMS, and CP. With this support, you can install, IPL, and operate z/VM from SCSI disks.*

- z/VM V5.1 includes the capability to install z/VM both from a DVD to an ESS SCSI disk emulated as a FBA device and a 3390 DASD. Installing from a DVD can significantly reduce the required installation media and allows you to install to a zSeries server using only SCSI disks. This is expected to be most beneficial in a z/VM environment with Linux guests and without traditional installation devices such as IBM TotalStorage tape drives attached to the IBM zSeries server.
- Coordinated near-continuous availability and disaster recovery for Linux guests by providing a new HyperSwap function so that the virtual devices associated with one real disk can be swapped transparently to another. HyperSwap can be used to switch to secondary disk storage subsystems mirrored by Peer-to-Peer Remote Copy (PPRC). HyperSwap is planned to be exploited by Geographically Dispersed Parallel Sysplex (GDPS) 3.1 to provide a coordinated near-continuous availability and disaster recovery solution for distributed applications, such as WebSphere, that span z/OS images running natively and Linux guests running under z/VM.
- PCIX Cryptographic Coprocessor (PCIXCC) support provides z/OS and Linux guest support for the PCIXCC Feature available with the z990 and z890 servers. Delivery of the z/VM PCIXCC support satisfies the Statement of Direction made on May 13, 2003.
- The Systems Management APIs, introduced in z/VM V4.4, provided a basic set of functions that may be called by applications to allocate and manage resources for guests running in z/VM virtual machines (virtual images). Although these APIs are primarily intended for managing Linux virtual images, they can be used to manage many types of z/VM virtual machine. All enhancements to the APIs in z/VM V5.1 have been implemented using Version 2 (V2) of the RPC server.

In addition to usability enhancements, new functions include:

- DASD volume management for virtual images
- VMRM configuration file management
- Query status of active images
- Query VMRM measurement data
- Removal of userid entries in an authorization file with a single request
- Query all shared storage segments instead of one at a time
- A new programming service is provided by an emulated DIAGNOSE instruction that helps enable a guest virtual machine to specify an action to be taken by CP when the guest becomes unresponsive. A time interval and action are specified by the guest. If the guest fails to reissue the DIAGNOSE instruction within the specified time interval, CP performs the action.
- A new publication, *Getting Started with Linux on zSeries*, describes z/VM basics and how to configure and use z/VM functions and facilities to create and manage Linux servers running on zSeries processors. The publication is designed to help systems personnel (system programmers, administrators, and operators) with limited knowledge of z/VM deploy Linux servers on z/VM more quickly and more easily.

Network Virtualization and Security

- The virtual IP switch, introduced in z/VM V4.4, was designed to improve connectivity to a physical LAN for hosts coupled to a guest LAN. The virtual switch has been enhanced to provide enhanced failover support for less disruptive recovery for some common network failures helping to provide business continuity as well as infrastructure reliability and availability.

- *Authorization capabilities have been enhanced for z/VM guest LANs and virtual switches by using Resource Access Control Facility (RACF) or any equivalent External Security Manager (ESM) that supports this function. It is designed to provide ESM centralized control of authorizations and Virtual LAN (VLAN) assignment.*

Technology Exploitation

- *z/VM V5.1 supports the new z890 as well as the new enhancements to the z990 including:

 - *Four Logical Channel SubSystems (LCSSs) on the z990 and two on the z890*
 - *Transparent sharing of internal and external channel types across LCSSs such as ICB-3, ICB-4, ISC-3, FICON Express, and OSA-Express*
 - *Open Systems Adapter-Express Integrated Console Controller (OSA-ICC) function**
- *Up to 24 real processor engines in a single z/VM image on a z990 satisfies the Statement of Direction made on May 13, 2003.*
- *IPv6 support for guest LANs has been enhanced to allow the z/VM TCP/IP stack to be configured for IPv6 networks connected through OSA-Express operating in QDIO mode. The stack can be configured to provide static routing of IPv6 packets and to send IPv6 Router Advertisements. In addition, support is being provided to help application developers to develop socket applications for IPv6 communications.*

Systems Management Improvements

The Performance Toolkit for VM has been enhanced in z/VM V5.1 to provide functional equivalence to the Performance Reporting Facility (PRF) priced, optional feature, thereby virtually eliminating the need for separate products (PRF and RealTime Monitor (RTM)) to help manage your performance more efficiently.

Other new function includes:

- *New high-level Linux reports based on Application Monitor records from Linux*
- *A new report for SCSI disks*

Delivery of equivalent function to PRF in the Performance Toolkit for VM satisfies the Statement of Direction made on May 13, 2003 to remove the RTM and PRF features in a future release of z/VM. The RTM and PRF features have been withdrawn from z/VM V5.1. These features are still available with z/VM V4.4 but cannot be licensed with z/VM V5.1

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

VSE/ESA

VSE/ESA

VSE clients with growing core VSE workloads, growing Linux workloads, or both will find the z890 to be a robust, cost-effective server solution and a unique opportunity to reduce IT cost through infrastructure simplification.

VSE/ESA is designed to protect extensive customer investments in existing core VSE applications, data, IT skills, equipment, business processes, end user training.

Standards-based, no-charge connectors provide flexible interoperability between VSE and Java technology-capable platforms such as Linux on zSeries. Connectors enable Linux-based applications to access core VSE applications and data. In addition, optional priced IBM middleware such as CICS Transaction Gateway, DB2 UDB or DB2 Connect, and WebSphere MQ complement VSE connectors.

VSE, Linux (exploiting IFLs), z/VM, and z890 are a compelling combination for VSE clients seeking a competitive IT advantage. New, state-of-the-art Linux-based Web applications and existing VSE applications and data can form comprehensive, integrated solutions that combine the best of e-business on demand with traditional transaction and batch processing.

z/OS remains an attractive migration target for some VSE clients. Affinity between VSE and z/OS has been an enduring part of IBM's strategy for VSE. One objective of affinity is to help ensure that clients can adapt their core VSE assets and skills to z/OS.

z890 is supported by VSE/ESA Version 2 Release 6 or later.

VSE/ESA V2.6 support includes zSeries features such as FICON Express, OSA-Express, OSA-Integrated Console Controller, up to 30 LPARs, and up to two Logical Channel SubSystems (4 on z990). VSE/ESA V2.7 adds HiperSockets

(to make the combination of VSE and Linux even more effective), PCICA encryption assist, and OSA-Express adapter interrupts. VSE/ESA V2.6 and V2.7 operate in 31-bit mode in LPARs, or as a 31-bit guest under z/VM. Linux for zSeries is available for new applications requiring 64-bit capability.

z/VSE

z/VSE will be the replacement product for VSE/ESA and is anticipated to be the basis for future VSE development. z/VSE will continue the VSE tradition of providing robust, cost-effective mainframe solutions. z/VSE is designed to bring the value of innovative zSeries technologies to VSE clients. It will continue to emphasize investment protection, network integration, and infrastructure simplification. z/VSE will focus on exploiting synergies with Linux on zSeries.

When available, z/VSE V3.1 plans to add support for FCP-attached SCSI disks to the z890 items already included in VSE/ESA V2.7. z/VSE can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. z/VSE is designed to exploit selected features of IBM zSeries hardware. Linux for zSeries is available for new e-business on demand applications requiring 64-bit capability.

Note: All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notification. For more information on VSE, see the VSE Web site at ibm.com/servers/eserver/zseries/os/vseStream, 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.



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 2004

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

Produced in the United States of America
04/04
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, IBM @server, the IBM logo, the e-business logo, AFP, APPN, BookManager, CICS, DB2, DB2 Connect, DB2 Universal Database, DFSMSdftp, DFSMSdss, DFSMSshsm, DFSMSsrm, DFSMS/MVS, DFSORT, DirMaint, e-business on demand, EKCD, Encina, Enterprise Storage Server, ESCON, FICON, FICON Express, FlashCopy, FFST, GDDM, GDPS, Geographically Dispersed Parallel Sysplex, HiperSockets, HiperSpace, HyperSwap, IMS, InfoPrint, Intelligent Miner, IP PrintWay, Language Environment, MQSeries, Multiprise, MVS, Net.Data, Netfinity, NetSpool, NetView, Open Class, OS/2, OS/390, Parallel Sysplex, Performance Toolkit for VM, PR/SM, Processor Resource/Systems Manager, RACF, RAMAC, RefreshPac, Resource Link, RMF, RS/6000, S/390, S/390 Parallel Enterprise Server, ProductPac, pSeries, Sysplex Timer, SystemPac, Tivoli, Tivoli Storage Manager, TotalStorage, VisualAge, 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 or both.

Java and all Java-based trademarks are trademarks or 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.

Intel is a trademark of Intel Corporation in the United States, other countries or both.

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.

GM13-0522-00