



© Copyright IBM Corporation 2001

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

Produced in the United States of America, 10-01  
All Rights Reserved

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

IBM, the IBM logo, the e-business logo, AnyNet, APPN, CICS, DB2, DB2 Connect, DB2 Universal Database, DFSMSdtp, DFSMSdss, DFSMSHsm, DFSMSrmm, DFSORT, DFSMS/MVS, ENCINA, Enterprise Storage Server, ESCON, FlashCopy, FFST, FICON, GDDM, GDPS, Geographically Dispersed Parallel Sysplex, HyperSockets, Hiperspace, IMS, InfoPrint, Intelligent Miner, Language Environment, Magstar, MQSeries, Multiprise, MVS, Netfinity, NetSpool, NetView, OS/2, OS/390, Parallel Sysplex, PrintWay, PR/SM, Processor Resource/Systems Manager, pSeries, RACF, RMF, RS/6000, S/390, S/390 Parallel Enterprise Server, SecureWay, Sysplex Timer, Systempac, Tivoli, Virtual Image Facility, VM/ESA, VSE/ESA, VTAM, WebSphere, xSeries, z/Architecture, z/OS, z/VM, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation.

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

Tux appears courtesy of Larry Eung.

UNIX is a registered trademark of The Open Group.

Windows NT is a registered trademarks of Microsoft Corporation.

Linux is a registered trademark of Linus Torvalds.

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

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

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

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

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

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

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.

# IBM server

## zSeries 900 and z/OS

## Reference Guide

*October 2001*



## **zSeries Overview**

The IBM® <sup>^</sup> zSeries™ 900 family of servers offers 26 air-cooled models, from one-way to 16-way, utilizing from three to twenty processors. These servers can be configured in numerous ways to offer unparalleled flexibility to speed deployment of e-business solutions. Each z900 can operate independently or as part of a Parallel Sysplex® cluster of servers with as many as 640 processors. In addition to supporting z/OS™, the z900 can host hundreds of Linux® images running open source applications using Linux or z/VM™.

The z900 family is based on the z/Architecture™, which supports a new standard of performance and integration by expanding on the balanced system approach of the S/390® architecture.

The z900 can eliminate bottlenecks associated with the lack of addressable memory through its virtually unlimited 64-bit addressing capability, providing plenty of “headroom” for unpredictable workloads and growing enterprise applications.

A robust I/O subsystem complements the greater number of processors and larger main memory. High-speed interconnects for TCP/IP communication, known as HiperSockets™, let TCP/IP traffic travel between partitions at memory speed, rather than network speed. A high performance Gigabit Ethernet feature is one of the first in the industry capable of achieving line speed:

one gigabit per second. Furthermore, the availability of native FICON™ devices and FICON CTC can increase I/O performance, simplify/consolidate channel configuration, and help reduce total cost of ownership. The total result is ultra high-speed communications within the server, between servers, to devices, and out to users, thereby allowing greater integration between traditional and Web applications to maximize e-business effectiveness.

The z900 family of servers also automatically directs resources to priority work through Intelligent Resource Director (IRD). The z900 IRD combines the strengths of three key technologies: z/OS Workload Manager, Logical Partitioning and Parallel Sysplex clustering.

## **New Tools for Managing e-business**

The IBM <sup>^</sup> 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.

\* The S/390 Parallel Enterprise Server and OS/390 Reference Guide (G326-3070-12) describes all functions and features available on the prior S/390 CMOS servers which are also available on the z900.

### **Ease of Use and Self-Management**

To help organizations deal effectively with complexity, IBM has announced Project eLiza, a blueprint for self-managing systems. The goal is to use technology to manage technology, creating an intelligent, self-managing IT infrastructure that minimizes complexity and gives customers the ability to manage environments that are hundreds of times more complex and more broadly distributed than exist today. This enables increased utilization of technology without the spiraling pressure on critical skills, software and service/support costs.

Project eLiza represents a major shift in the way the industry approaches reliability, availability and serviceability (RAS). It harnesses the strengths of IBM and its partners to deliver open, standards-based servers and operating systems that are self-configuring, self-protecting, self-healing and self-optimizing. Project eLiza technology helps ensure that critical operations continue without interruption and with minimal need for operator intervention.

Project eLiza will help customers dramatically reduce the cost and complexity of their e-business infrastructures, and overcome the challenges of

systems management. zSeries plays a major role in Project eLiza, since the self-management capabilities available for the zSeries will function as a model for other IBM servers.

The latest z900 and z/OS V1R2 announcement gives the ability to configure, connect, extend, operate and optimize the computing resources to efficiently meet the "always-on" demands of e-business.

One of the key functions of z/OS is Intelligent Resource Director (IRD), an exclusive IBM technology that makes the z900 the only server capable of automatically reallocating processing power to a given application on the fly, based on the workload demands being experienced by the system at that exact moment. This advanced technology, often described as the "living, breathing server," allows the z900 and z/OS to provide nearly unlimited capacity and nondisruptive scalability, according to priorities determined by the customer.

## z/Architecture

The zSeries is based on the z/Architecture, which is designed to eliminate 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.

This architecture has been implemented on the z900 to allow full 64-bit real and virtual storage support. A maximum 64 GB of real storage is available on z900 servers. z900 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 2GB 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 VM/ESA®.
- Use of 64-bit operands and general registers for all Cryptographic Coprocessors instructions and Peripheral Component Interconnect Cryptographic Coprocessors (PCICC) instructions is added.

The implementation of 64-bit z/Architecture can eliminate any bottlenecks associated with lack of addressable memory by making the addressing capability virtually unlimited (16 Exabytes from the current capability of 2 GB).

## 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 will run unmodified on the z900.

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

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

Some of the exploiters of z/Architecture for z/OS and OS/390® Release 10 include:

- DB2® Universal Database™ Server for OS/390
- IMS™
- Hierarchical File System (HFS)
- Virtual Storage Access Method (VSAM)
- Remote Dual Copy (XRC)
- Tape and DASD access methods

Operating System	ESA/390 (31-bit)	z/Arch. (64-bit)
z/OS Version 1 Release 1, 2 and 3	No	Yes
OS/390 Version 2 Release 10	Yes	Yes
OS/390 Version 2 Release 6, 7, 8, and 9	Yes	No
Linux for zSeries	No	Yes
Linux for S/390	Yes	No
S/390 Virtual Image Facility for Linux	Yes	No
z/VM Version 4 Release 1 and 2	Yes	Yes
z/VM Version 3 Release 1	Yes	Yes
VM/ESA Version 2 Release 3 and 4	Yes	No
VSE/ESA Version 2 Release 3, 4 and 5	Yes	No
TPF Version 4 Release 1 (ESA mode only)	Yes	No

\* SOD z/OS 1.4 supports ESA/390 or G5/G6 and MP 3000

## 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 z900 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 z900 server. 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 e-business workloads unrivaled in the industry.

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

### LPAR CPU Management


LPAR CPU Management allows WLM working in goal mode to manage the processor weighting and logical processors across an LPAR cluster. With z/OS Version 1 Release 2, WLM can even direct CPU resources outside a z/OS LPAR cluster, to an LPAR running either z/VM or Linux. CPU resources are automatically moved toward LPARs with the greatest need by adjusting the partition's weight. WLM also manages the available processors by adjusting the number of logical CPs in each LPAR. This optimizes the processor speed and multiprogramming level for each workload, reduces MP overhead, and gives z/OS more control over how CP resources are distributed to meet your business goals.

z/OS V1R2 enhances the LPAR CPU management capabilities and will allow the dynamic assignment of CPU resources to non-z/OS partitions such as Linux.

### 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 z900, paths may be dynamically assigned to control units to reflect the I/O

## zSeries IRD Scope



load. For example, in an environment where an installation normally requires four channels to several control units, but occasionally needs as many as six, system programmers must currently define all six channels to each control unit that may require them. With Dynamic Channel Path Management, the system programmer need only define the four channels to the control units, and indicate that DCM may add an additional two. As the control unit becomes more heavily used, DCM may assign channels from a pool of managed channels, identified by the system programmer, to the control unit. If the work shifts to other control units, DCM will unassign them from lesser utilized control units and assign them to what are now the more heavily used ones. This helps reduce the requirement for greater than 256 channels. DCM can also reduce the cost of the fibre infrastructure required for connectivity between multiple data centers.

### 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 and FICON Express channels.

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 the z900 working in synergy extends the industry leading workload management tradition of S/390 and OS/390 to ensure that the most important work on a server meets its goals, to increase the efficiency of existing hardware, and to reduce the amount of intervention in a constantly changing environment.

### HiperSockets

HiperSockets, a new feature unique to the zSeries, provides a “TCP/IP network in the server” that allows high speed any-to-any connectivity among virtual servers (TCP/IP images) within a z900 without any physical cabling. HiperSockets minimizes network latency and maximizes bandwidth between combinations of Linux, z/OS and z/VM virtual servers. These OS images can be first level (directly under an LPAR), or second level images (under z/VM).

With up to four HiperSockets per LPAR connection, one could separate traffic to different HiperSockets for security (separation of LAN traffic, no external wire-tapping, monitoring) and performance and management reasons (separate sysplex traffic Linux or non-sysplex LPAR traffic).

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



## **z900 Support for Linux**

HiperSockets acts like any other TCP/IP network interface, so TCP/IP features like IP Security (IPSec) in Virtual Private Networks (VPN) and SSL can be used to provide heightened security for flows within the same CHPID. HiperSockets supports multiple frame sizes, which is configured on a per HiperSockets CHPID basis. This support gives the user the flexibility to optimize and tune each HiperSockets to the predominate traffic profile, for example 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, bounded only by the combinations of operating systems and their respective applications. HiperSockets will provide the fastest z900 connection between e-business and 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 LPAR or z/VM guest machine and will be able to use HiperSockets for very fast TCP/IP traffic transfer to a DB2 database server running in a z/OS V1.2 LPAR. System performance is optimized because this allows you to keep your Web and transaction application environments in close proximity to your data and eliminates any exposure to network related outages, thus improving availability.

The new 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.

### **Linux for zSeries**

Linux and zSeries make a great team. The flexibility and openness of Linux bring with it access to a very large portfolio of applications. zSeries incorporates the qualities of service that deliver an industrial strength environment for these Linux applications. In addition zSeries enables massive scalability within a single server. Hundreds of Linux images can run simultaneously, providing unique server consolidation capabilities and reducing both cost and complexity.

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

IBM has developed a series of patches that enable Linux to run on the S/390 platform. These were released to the Open Source community in December, 1999.

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

#### *IBM Software*

- IBM WebSphere Application Server Advanced Edition 3.5
- WebSphere Commerce Suite, Pro Edition for Linux for zSeries Version 5.1
- WebSphere Personalization Version 3.5
- MQSeries® Client for Linux for S/390 Version 5.2
- WebSphere Host-On-Demand Version 5.0.3
- CICS® Transaction Gateway Version 4.0
- Tivoli® Storage Manager Linux for zSeries and S/390 Client Version 4.2
- IMS Connect
- DB2 Universal Database Enterprise Edition Version 7.2
- DB2 Connect™ Unlimited Edition Version 2
- DB2 Connect Web Starter Kit Version 7.2
- DB2 Intelligent Miner™ Scoring Version 7.2
- DB2 Net Search Extender Version 7.2

### Linux Distribution Partners

- SuSE Linux Enterprise Server for S/390  
Product Information available at [suse.de/en/produkte/susesoft/S390/](http://suse.de/en/produkte/susesoft/S390/). There are many different options for ordering. Customers should contact SuSE by sending an e-mail to [s390-sales@suse.de](mailto:s390-sales@suse.de).
- TurboLinux Server 6 for zSeries and S/390  
Product Information at [turbolinux.com/products/s390](http://turbolinux.com/products/s390). Customers should contact the TurboLinux sales department by sending an e-mail to [sales@turbolinux.com](mailto:sales@turbolinux.com).
- Red Hat Server for zSeries  
Product Information at [redhat.com](http://redhat.com).

### z/VM Version 4

z/VM enables large scale horizontal growth of Linux images when using Linux for zSeries. Only z/VM gives the capability of running hundreds of Linux for zSeries or Linux for S/390 images. This version of z/VM is priced on a per-engine basis (onetime charge) and will support IBM Integrated Facility for Linux (IFL) processor features for Linux-based workloads and standard engines for all other zSeries and S/390 workloads.

### IBM S/390 Integrated Facility for Linux

This optional feature provides a way to add processing capacity, exclusively for Linux workload, with no effect on the model designation. No traditional zSeries workload will be able to run in this area. Consequently, these engines will not affect the IBM S/390 and zSeries software charges for workload running on the other engines in the system.

### OSA-Express Gigabit Ethernet for Linux

Driver support is provided for the functions of the OSA-Express Gigabit Ethernet feature. This driver supports the IPv4 protocol, delivering the advantages of more rapid communication across a network. This improvement may be between virtual Linux instances on a single machine (either in LPAR or virtual mode) communicating across a network, or a Linux for zSeries or Linux for S/390 instance communicating with another physical system across a network.

### HiperSockets

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

### Crypto

Linux for zSeries running on standard z900 engines is capable of exploiting the hardware cryptographic feature provided by the PCI card for SSL acceleration. This enables customers implementing e-business applications on Linux for zSeries to utilize this enhanced security of the hardware.

### Linux Support

Environment

- zSeries or S/390 single image
- zSeries or S/390 LPAR
- Virtual Image Facility image
- 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, Ethernet, Fast Ethernet, Token-Ring)
- HiperSockets
- 3172
- IUCV
- Character devices
- 3215 console
- Integrated console

Additional information is available at [ibm.com/linux/](http://ibm.com/linux/) and at [ibm.com/zseries/linux/](http://ibm.com/zseries/linux/).



## 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 ensure time synchronization.

Parallel Sysplex Resource Sharing enables multiple system resources to be managed as a single logical resource shared among all of the images. Some examples of resource sharing include Automatic Tape Switching, 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 – GF22-5115-00 on the Parallel Sysplex home page at [ibm.com/servers/eservers/zseries/pso](http://ibm.com/servers/eservers/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, provide 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/eservers/zseries/pso](http://ibm.com/servers/eservers/zseries/pso).

The IBM 9037 Sysplex Timer provides a common time reference to all images which assists in managing the

cluster of multiple footprints as a single operational image. The common time source also enables proper sequencing and time stamping of updates to shared databases, a feature critical to recoverability of the shared data.

### Coupling Facility Configuration Alternatives

IBM offers different options for configuring a functioning Coupling Facility:

- *Standalone Coupling Facility:* z900 Model 100, and 9672-R06 models provide a physically isolated, totally independent CF environment. There are no software charges associated with standalone CF models. An ICF or CF partition sharing a server with z/OS or OS/390 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. An ICF feature is a processor that can only run Parallel Sysplex coupling code (CFCC) in a partition. Since CF LPARs on ICFs are restricted to running only Parallel Sysplex coupling code, there are no software charges associated with ICFs. ICFs are ideal for Intelligent Resource Director and resource sharing environments.
- *Coupling Facility partition on a z900 or 9672 server using standard LPAR:* A CF can be configured to run in either a dedicated or shared CP partition. 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 backup 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, and “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 failure such as loss of a single structure or CF or loss of connectivity to a single CF, through rapid failover to the other structure instance of the duplex pair.

Benefits of System-Managed CF Structure Duplexing include:

- **Availability**  
Faster recovery of structures by having the data already there in the second CF. Furthermore, if a potential IBM, vendor or customer CF exploitation implementation were being prevented by the effort of providing alternative recovery mechanism such as structure rebuild, log recovery, etc., that constraint might be removed by the much simpler exploitation requirements for system-managed duplexing.
- **Manageability and Usability**  
A consistent procedure to set up 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**

Enables the use of non-standalone CFs (e.g. ICFs) for all resource sharing and data sharing environments.

- **Flexibility**

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



*Note:* System-Managed CF Duplexing requires CFCC Level 10 and z/OS 1.2.

To understand which of the options and capabilities discussed above are suitable for your environment, please review GF22-5042, *Coupling Facility Configuration Options: A Positioning Paper* at [ibm.com/servers/eservers/zseries/library/whitepapers/gf225042.html](http://ibm.com/servers/eservers/zseries/library/whitepapers/gf225042.html).

### 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 z900 server to a z900 Model 100 CF, 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 Gb/s capacity, the ICB-3 link with 1 GB/s peak capacity, and the IC-3 link with 1.25 GB/s capacity. Additional Peer Mode benefits are obtained by enabling the link to be MIFed between z/OS (or OS/390) and CF LPARs. The link acts simultaneously as both a CF Sender and CF Receiver link, reducing the number of links required. Larger data buffers and improved protocols also improve long distance performance. For connectivity to 9672s, z900 CF links can be configured to run in Compatibility mode with the same characteristics as links on the 9672 (single-mode CF links with 100 Gb/s or ICB-3 links with 333 MB/s). All of these above coupling link speeds are theoretical maximums.

#### Theoretical Maximum Coupling Link Speed

Connectivity Options	z900 ISC-3	z900 ICB	z900 ICB-3
G2-G6 ISC	1 Gb/s	n/a	n/a
z900 ISC-3	2 Gb/s Peer Mode	n/a	n/a
G5-G6 ICB	n/a	333 MB/s	n/a
z900 ICB-3	n/a	n/a	1 GB/s Peer Mode

- **ISC-3.** InterSystem Coupling Facility-3 channels provide the connectivity required for 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 z900 general purpose models and z900 coupling facilities. ISC-3 channels operating in Compatibility Mode provide connection between z900 systems and HiPerLink (ISC) channels on 9672 Models. A four port ISC-3 card struc-

ture is provided on the z900 family of 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 Gb/s in Compatibility Mode or 2 Gb/s in native mode up to a distance of 10 km. From 10 to 20 km, an RPQ card is available which runs at 1 Gb/s 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.

- **HiPerLinks.** HiPerLinks, based on single-mode CF links, are available on 9672s and 9674s only. ISC-3s replace HiPerLinks on zSeries 900 models.
- **ICB.** The Integrated Cluster Bus is used to provide high speed coupling communication between 9672 G5/G6 and/or z900 servers over short distances (~7 meters). For longer distances, ISC links must be used. Up to 8 ICB links (16 possible via RPQ) are available on the general purpose z900 models and up to 16 ICB links are available on the z900 Coupling Facility Model 100.
- **ICB-3.** The Integrated Cluster Bus-3 is used to provide high speed coupling communication between two z900 systems over short distances (~7 meters). For longer distances, ISC-3 links must be used. Up to 16 ICB-3 links are available on both the general purpose z900 models and the z900 Coupling Facility Model 100. The performance of the ICB-3 link has been improved by higher data rates and new buffering capabilities.

**IC-3.** The Internal Coupling-3 channel emulates the Coupling links between images within a single server. No hardware is required, however a minimum of 2 CHPID numbers must be defined in the IOCDs. IC-3 links provide the fastest Parallel Sysplex connectivity.

#### 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 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/psa/services.html](http://ibm.com/servers/eserver/zseries/psa/services.html)

### **Geographically Dispersed Parallel Sysplex**

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

GDPS requires Tivoli NetView® for OS/390, System Automation for OS/390, and remote copy technologies. GDPS supports both the synchronous Peer-to-Peer Remote Copy (PPRC) as well as the asynchronous Extended Remote Copy (XRC) forms of remote copy. GDPS/PPRC is a continuous availability solution and near transparent business continuity/disaster recovery solution that allows a customer to meet a Recovery Time Objective (RTO) of less than an hour, a Recovery Point Objective (RPO) of no data loss, and protects against metropolitan area disasters (up to 40 km between sites). GDPS/XRC is a business continuity/disaster recovery solution that allows a customer to meet a RTO of one to two hours, a RPO of less than a minute, and protects against metropolitan as well as regional disasters, since the distance between sites is unlimited. XRC can use either common communication links and channel extender technology between sites or dark fiber.

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

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

A new I/O VTS selection option is provided especially for use with GDPS, so that all virtual volumes are

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

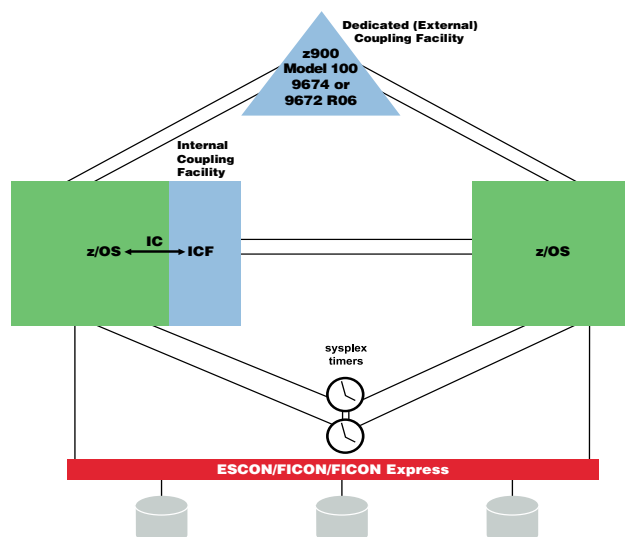
Operational data, data that is used directly by applications supporting end users, is normally found on disk. For the past several years, GDPS has provided continuous availability and near transparent business continuity for disk resident data. However, there is another category of data that “supports”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

### **Automatic Enablement of CBU for Geographically Dispersed Parallel Sysplex**

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

GDPS is discussed in detail in two white papers which are available at [ibm.com/server/eserver/zseries/psolibrary.html](http://ibm.com/server/eserver/zseries/psolibrary.html). GDPS is a service offering of IBM Global Services. For IBM Installation Services for GDPS refer to the IBM Web site.

## Continuous Availability Recommended Configuration



### Components and assumptions

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

### Key attributes can include

- No single points 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

For additional information, see the *Five Nines/Five Minutes: Achieving Near Continuous Availability* white paper at [ibm.com/s390/psa/library](http://ibm.com/s390/psa/library).

## IBM ^ zSeries 900

IBM's zSeries is the enterprise class e-business server optimized for integration, transactions and data of the next generation e-business world. In implementing the new z/Architecture with new technology solutions, the z900 models are designed to facilitate the IT business transformation and reduce the stress of business-to-business and business-to-customer growth pressure. The z900 represents a new generation of servers that feature enhanced performance, support for S/390 Parallel Sysplex clustering, improved hardware management controls and innovative functions to address e-business processing.

The z900 has 26 models available as new build systems and as upgrades from the S/390 Parallel Enterprise Server™— Generation 5/6. The S/390 R06 Coupling Facility is upgradable to the z900 as well.

The z900 processor enhances performance by exploiting the new architecture and technology and through many design enhancements. The z900 has sixteen central processors in a symmetrical processing complex, up from twelve in the S/390 G6. This provides customers with up to 60 percent performance improvement over the G6 model ZZ7. At the uniprocessor level performance improvements are in the 20-30 percent range when compared to the G6 Model Z17, depending on workload and environment. The Large System Performance Reference (LSPR) should be referenced when considering performance on the z900. The Web site for LSPR is [ibm.com/servers/eserver/zseries/lspr/](http://ibm.com/servers/eserver/zseries/lspr/).

The I/O subsystem includes Dynamic Channel Path Management (DCM) and Channel CHPID Assignment. These two functions effectively increase the number of CHPIDS that can be used for I/O connectivity. DCM allows channel paths to be dynamically and automatically moved from less utilized devices to constrained devices under the supervision of the Workload Manager. Channel CHPID Assignment permits the assignment of a CHPID to any physical port. This allows the assignment of all 256 CHPIDs to usable channel paths. Combined, the use of these two functions allows the full exploitation of the I/O bandwidth inherent in the 256 channels available on the z900.

Design and technology advances also include introducing an improved Fiber CONnectivity (FICON) channel card called FICON Express and FICON CTC.

These are some of the significant enhancements in the zSeries 900 processor that bring improved

performance, availability and function to the platform. The following sections highlight the functions and features of the hardware platform.

### **z900 Design and Technology**

The z900 is designed to provide balanced system performance. From processor storage to the system's I/O and network channels, end-to-end bandwidth is provided to deliver data where and when it is needed.

The base of packaging technology for the z900 processor is the MultiChip Module (MCM). The MCM contains the processor unit (PU) chips, the cache structure chips and the processor storage and I/O interface chips. The modules measure approximately 127 by 127 millimeters. There is one MCM per system and there are two MCM types. One z900 MCM type provides support for 20 PUs and 32 MB of level 2 cache making it the world's densest chip module. This module contains 101 layers of glass ceramic and 6 layers of thin film to provide interconnection between the chips and the off-module environment. In total, there is approximately 1 kilometer of wiring on this module. This MCM type is positioned for the highest performance models. The second MCM type supports 12 PUs and 16 MB of level 2 cache. This module contains 45 ceramic and 6 thin film layers and is price/performance situated for the lower performance models.

The design and MCM technology on the z900 provides the flexibility to configure the PUs for different uses. PUs can be used as Central Processors (CPs), System Assist Processors (SAPs), Internal Coupling Facility processors (ICFs) and Integrated Facility for Linux processors (IFLs). For the general purpose processor models, the number of CPs and SAPs (provide I/O processing) are specified by the model number. Optionally, as features on a system, the remaining inactive PUs on the MCM are available as ICF processors for Coupling Facility applications, IFLs for Linux applications or as optional SAPs. These capabilities are shown, by model, in the Processor Options table on page 54 and provide the customer with tremendous flexibility in establishing the best system for running applications. One PU is always reserved as a spare.

The Processor Unit (PU) uses CMOS 8S with copper interconnect and 47 million transistors with a cycle time of 1.3 nanoseconds for all models. This 179 x 9.9 millimeter chip is from the latest family of CMOS chips

coming from IBM's semiconductor laboratories. Implemented on this chip is the new z/Architecture with its 64-bit capabilities including new instructions, 64-bit General Purpose Registers and translation facilities.

The z900 system has a processor storage capability of up to 64 GB, depending on model. The processor storage is implemented in 4, 8 and 16 GB cards. Some z900 models have 2 memory cards, other z900 models have 4 memory cards. The recovery design of the memory cards has been enhanced with additional spare memory chips.

The MultiChip Module is the technology cornerstone for flexible PU deployment in the z900 models. For most models the ability of the MCM to have inactive PUs allows such features as Capacity Upgrade on Demand (CUoD), the ability to add CPs, ICFs and IFLs dynamically providing nondisruptive upgrade of processing capability. Also, the ability to add CPs lets a small z900 system be a backup for another system in the enterprise; expanding the z900 system to meet an emergency outage situation. This is called Capacity BackUp (CBU). From a capacity point of view, the MCM can package up to a full 16-way processor.

The I/O subsystem provides significant bandwidth to support the number of processors and the continuing advances in high speed I/O. The z900 I/O subsystem supports up to 96 FICON and/or FICON Express channels. This represents up to 768 equivalent ESCON channels. (Note that this configuration could also support 160 ESCON channels {256 CHPIDS total} which would represent a combined equivalence of up to 928 ESCON channels.) Also, a single zSeries I/O Cage can package the maximum of 256 ESCON channels, a feat that would have required three I/O cages previously. 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 a new high speed capability to communicate between Logical Partitions; this is based on high speed TCP/IP memory speed transfers and will provide value in allowing applications running in one partition to communicate with applications running in another. The z900 Gigabit Ethernet feature is one of the first adapters in the industry capable of achieving line speed: one Gigabit per second. Industry standard and openness are design objectives for I/O in z900. The improved I/O subsystem is delivering new horizons in I/O capability.

## Relative Performance of z900 to G5/G6 Servers

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 and the use of FICON, FICON Express and associated CTC, tape, printer and disk devices, thus providing for larger and quicker data movement into 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. New coupling links in the architecture allow customers to continue to exploit the advantage of the leading clustering solution in the industry, Parallel Sysplex technology. The combined balanced system effect allows for increases in performance across a broad spectrum of work.



When comparing a 16-way zSeries Model 116 to a G6 turbo ZZ7, the following levels of performance can be obtained:

- Savings in batch elapsed time of up to 17% for the CPU component of an individual batch job. A batch stream can see up to a 60% throughput improvement with a potential for further savings via implementation of FICON, compression, and ESS DASD.
  - Due to the increase in hardware capacity, additional processing power can be used to reduce the CPU component of the batch job. Combined with the support of larger data movement, this allows for improved turnaround time in CPU constrained batch jobs.
- Ability to support 55% additional Linux images
  - The z/Architecture not only supports traditional work but also allows Linux work to execute within the

same server. The number of Linux images that can be supported for productive work is significantly more than was previously available due to the balanced system approach. This allows for industry leading levels of server consolidation for Linux within the z/Architecture while providing the performance required to support all work types.

- Support of up to 50% more SAP throughput
  - The balanced system approach enables additional capacity to support a larger number of SAP users through additional server processing power as well as through the ability to accommodate more users within a single operating system image through the new storage hierarchy.
- Support of up to 50% more Lotus® users
  - The balanced system approach enables additional capacity to support a larger number of Lotus users through additional server processing power as well as the ability to accommodate more users within a single operating system image through the new storage hierarchy.
- Support of up to 45% more throughput in Web-initiated e-business transactions
  - The unique performance characteristic of the z/Architecture is that the balanced system approach not only enables additional transactions to be executed but it is also a design point that is well suited to contain and manage spikes in work based upon an influx of Web-initiated transactions. This capability is achieved through a combination of the operating system as well as increased throughput delivered with the z900 server.

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



## z900 Family Models

The z900 has a total of 26 models to offer flexibility in selecting a system to meet the customer's needs. Twenty-five of the models are general purpose systems. The remaining model is the Coupling Facility Model 100. There are a wide range of upgrade options available which are described below and shown on the following pages. Capacity Upgrade on Demand and Capacity Backup (CBU) are available. The z900 has also been designed to offer a high performance and efficient I/O structure to meet the demands of e-business as well as high demand transaction processing applications. Up to 256 ESCON channels will now fit into a single I/O cage; or a total of 96 FICON and/or FICON Express channels and 160 ESCON channels can be accommodated in a fully configured system. To provide the best choice of processor for the application, two designs have been developed. Although similar in structure, one design has a 12-processor unit (PU) Multichip Module (MCM) and two memory cards. The other design has a 20 PU MCM and 4 memory cards. Both have equivalent I/O capability. The processor models and a discussion of configurations follow.

### z900 Models 101-109

These nine models are general purpose systems and range from a 1-way to 9-way symmetrical multiprocessor (SMP). The systems have a 12-PU MCM, two memory buses, and can support up to 32 GB processor storage (entry storage is 5 GB). The PU has a cycle time of 1.3 nanoseconds (ns). These models can easily upgrade from one model to the next through Capacity Upgrade on Demand. Also, there are upgrade paths to the Models 110 through 116; however, this upgrade will require a system outage. The Models 101-109 have 2 System Assist Processors (SAPs) as standard and up to 24 STI links for I/O attachment. Any of the spare PUs on the MCM can be assigned as a CP, SAP, Integrated Coupling Facility (ICF) or Integrated Facility for Linux (IFL).

### z900 Models 110-116

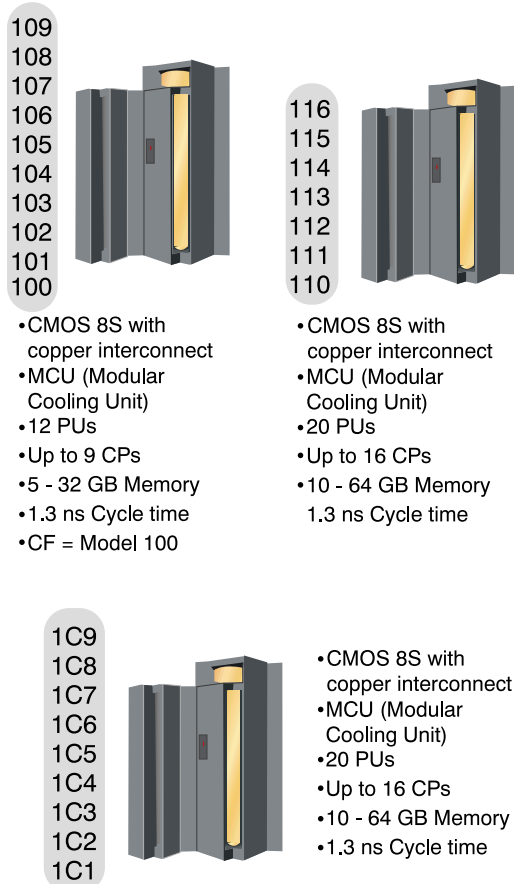
These seven models are general purpose systems and range from a 10-way to a 16-way symmetrical multiprocessor (SMP). The systems have the 20-PU MCM, four memory buses and can support up to 64 GB processor storage (entry storage is 10 GB). The PU, as

with the Models 101-109, has a cycle time of 1.3 nanoseconds. The Models 110-116 can easily upgrade from one model to the next through Capacity Upgrade on Demand and support Capacity Backup (CBU). The Models 110-116 have 3 SAPs as standard and up to 24 STI links for I/O attachment. Any of the spare PUs on the MCM can be assigned as a CP, SAP, ICF, or IFL.

### z900 Models 1C1-1C9

These nine models are general purpose systems and range from a 1-way to a 9-way symmetrical multiprocessor. The systems have a design and cycle

### z900 Models





time identical to the models 110-116 including the 20 PU MCM, four memory buses and supporting up to 64 GB processor storage (entry storage is 10 GB). The Models 1C1-1C9 are available as an option for Capacity Upgrade on Demand and Capacity Backup (CBU) requirements and can be upgraded to a 16-way z900 without a system outage. Customers whose capacity requirements are likely to exceed the Model 101-109 range should consider the 1C1-1C9 as an alternative.

**z900 Coupling Facility Model 100**

The Model 100 is the standalone Coupling Facility in the z900 family. This Model can have up to 9 ICF engines. It is recommended that the z900 CF Model 100 be used in production data sharing configurations for its improved coupling efficiency.

Customers can upgrade current 9672 R06 Models to the z900 Coupling Facility Model 100, maximizing the coupling efficiency. The z900 CF Model 100 can be upgraded to the z900 general purpose models.

**I/O Connectivity**

The z900 contains an I/O subsystem infrastructure which uses an I/O cage that provides 28 I/O slots compared to the G5/G6 style cage with 22 slots. ESCON, FICON, FICON Express and OSA-Express cards plug into the zSeries I/O (nI/O) cage. Parallel channel, OSA-2 Token-Ring and FDDI and ESCON 4-port cards are accommodated in a Compatibility I/O (cI/O) cage (22 slots). ESCON 4-port cards are used only in upgrading from a G5/G6 model. All I/O cards and their support cards can be hot-plugged in the nI/O cage. Installation of an I/O cage remains a disruptive MES so the Plan Ahead feature remains an important consideration when ordering a z900 system.

The nI/O cage takes advantage of an exclusive IBM packaging technology that provides a subsystem with approximately seven times higher bandwidth than the previous G5/G6 I/O cage. Each general purpose z900 model comes with one nI/O cage standard in the A-Frame (the A-Frame also contains the processor CEC cage). The nI/O cage, using new 16 port ESCON cards, can hold 256 ESCON channels; previous packaging required three I/O cages to package the

same number of channels. For FICON and FICON Express, the nI/O cage can accommodate up to 16 cards or 32 channels per cage; with the previous technology, up to 36 channels would require three I/O cages. Thus, much denser packaging and higher bandwidth has been achieved.


The Z-Frame is an optional I/O Frame and attaches to the A-Frame. The Z-Frame can contain up to two of the new zSeries I/O (nI/O) cages, up to two compatibility I/O (cI/O) cages or a mixture of both. Following is the layout of the A- and Z-Frames and both types of I/O cages.

**I/O Cards:** The next topic has a chart that shows cards supported by the two types of I/O cages.

**Up to 256 ESCON Channels**

- Two different ESCON channel cards are available on the z900 system:
- A new 16 port channel card which plugs into the new I/O cage will be used for all new ESCON channel orders and any overflow from Compatibility I/O cages on G5/G6 upgrades. Up to 15 ports will be used for ESCON connectivity; one port is reserved as a spare port.
  - Four port channel cards which plug into the Compatibility I/O cage are used only in upgrades from G5/G6 to z900.
  - Channels are available in four port increments for both card types.

**Cage Layout and Options**



### Up to 96 FICON or FICON Express Channels

The z900 supports up to 96 FICON and/or FICON Express channels. Both FICON and FICON Express are available in long wave (LX) and short wave (SX) operation. Each FICON and FICON Express card has two channels per card. The LX and SX cannot be intermixed on a single card. The maximum number of FICON cards is 48 installed in 3 I/O cages.

### Up to 88 Parallel Channels

The four port parallel channel card is the same card used on G5/G6 models and is orderable on the z900. However, it must plug into a Compatibility I/O cage. The three port parallel card, if present during an upgrade from a G5/G6, will be carried forward.

I/O Type	zSeries I/O Cage	Compatibility I/O Cage
ESCON	Yes (16 port)	Yes (4 port/MES)
Parallel Channel	No	Yes
ISC-3 (1 & 2 Gigabit)	Yes	No
FICON FICON Express	Yes	No
Fast Ethernet	Yes	No
ATM 155	Yes	No
PCICC	Yes	No
PCICA	Yes	No
Token-Ring	Yes	No
GbEthernet	Yes	No
FDDI	No	Yes
OSA-2 Token-Ring	No	Yes

Eighty-eight (88) parallel channels can be ordered on a new-build z900 and up to 96 via RPQ 8P2198. This RPQ provides an additional Compatibility I/O cage to enable installation of the extra parallel channel cards. This RPQ is not required if a G5/G6 Model with greater than 88 parallel channels is upgraded to a z900.

### ISC-3

A four port ISC-3 card is provided on the z900 family of processors. It consists of a mother card with two daughter cards which have two ports each. Each daughter card port is capable of operation at 1 Gb/s in Compatibility Mode or 2 Gb/s in native mode up to a distance of 10 km. The mode is selected for each port via CHPID type in the IOCDs. The ports are orderable in one port increments.

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

### Channel CHPID Assignment

The z900 provides customers with the option of mapping CHPID numbers to physical channel ports. This is called Channel CHPID Assignment. Channel CHPID Assignment helps customers maintain G5/G6 I/O definitions during system upgrades and in establishing new I/O definitions. It also allows for a logical sequence of CHPID numbers to be assigned to a control unit. A CHPID Mapping tool is available on Resource Link, [ibm.com/servers/resourcelink](http://ibm.com/servers/resourcelink), to assist with mapping.

## Cryptography

IBM leads the industry in bringing greater security to e-business with its high availability CMOS Cryptographic Coprocessors. This feature has earned Federal Information Processing Standard (FIPS) 140-1 level 4, the highest certification for commercial security ever awarded by the U.S. Government. For the z900, the two Cryptographic Coprocessors Single Chip Modules (SCMs) have been moved from the MCM to the CPC Cage. The SCMs are plugged directly into the rear of the CPC backplane. The SCMs are individually serviceable, minimizing system outage.

The z900 servers can also support a combination of up to 8 optional Peripheral Component Interconnect Cryptographic Coprocessor (PCICC) or the PCI Cryptographic Accelerator (PCICA) features. Each PCICC/PCICA feature contains 2 cryptographic coprocessor accelerators, for a maximum of 16 coprocessor accelerators. This provides the capability to support up to 3800 SSL transactions/second.

The combination of the two coprocessor types enables applications to invoke industry-standard cryptographic capabilities — such as DES, Triple DES, or RSA — for scalable e-transaction security and the flexibility to adopt new standards quickly.

The new PCICA feature is designed to address the high performance SSL needs of e-business applications, and has a design point different from the existing zSeries CMOS Cryptographic Coprocessor and zSeries PCICC feature.

### **z900's Capacity Upgrade on Demand (CUoD)**

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

The CUoD functions are:

- z900 continues to support the dynamic CUoD function introduced on G5/G6.
  - Nondisruptive CP, ICF, and IFL upgrades are available within minutes.
- Dynamic upgrade of all I/O cards in the nl/O Cage.
- Dynamic upgrade of spare installed memory.

### **Plan Ahead and Concurrent Conditioning**

Concurrent Conditioning configures a system for hot plugging of I/O based on a future specified target configuration. Concurrent Conditioning of the zSeries I/O is minimized by the fact that all I/O cards, plugging into the zSeries I/O (nl/O) cage, are hot pluggable. This means that the only I/O to be conditioned is the nl/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, FICON Express or additional OSA-Express networking is made exceptionally easy and nondisruptive with the appropriate microcode load and if the cage space is available. Planning ahead is an important activity when purchasing a typical new build system (which will probably have one nl/O cage, one cl/O cage and space for one additional nl/O cage in the Z-Frame). This is even more important when considering a G5/G6 system upgrade where the Z-Frame space is occupied with cl/O cages. FICON, FICON Express, OSA-Express, Crypto PCICC and ISC-3 cards plug only into nl/O cages.

The z900 now 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, 16 or 32 GB increments. A customer with a Model 106, for example, with 5 GB of storage will be able to concurrently upgrade to 6, 7 and 8 GB but will not be able to get to the next increment of 10 GB without a disruption.

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 z900 capacity. This creates an attractive new option for businesses to quickly respond to changing environments, bringing new applications online or growing existing applications without disrupting users.

### **z900 Server Capacity BackUp (CBU)**

Capacity BackUp (CBU) is offered with the z900 processors to provide reserved emergency backup CPU capacity for situations where customers have lost capacity in another part of their establishment and want to recover by adding reserved capacity on a designated z900 system. A CBU system normally operates with a "base" CPU configuration and with a preconfigured number of additional Processor Units (PUs) reserved for activation in case of an emergency.

The z900 technology is ideally suited for providing capacity backup since the reserved CBU processing units are on the same technology building block, the MCM, as the regular CPs. Therefore, a single processor can support two diverse configurations with the same MCM. For CBU purposes, the Models 101-109 can scale from a 1-way to a 9-way; the Models 110-116 from a 10-way through a 16-way; and the Models 110-116 in conjunction with the Capacity models (1C1-1C9) can scale from 1-way through 16-way with the purpose of providing capacity backup.

The "base" CBU configuration must have sufficient memory and channels to accommodate the potential

needs of the larger CBU target machine. When capacity is needed in an emergency, the primary operation performed is activating the emergency CBU configuration with the reserved PUs added into the configuration as CPs.

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

The z900 supports concurrent CBU downgrade. This function enables a Capacity Backup Server to be returned to its normal configuration without an outage (i.e. Power-on reset).

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

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

## **S/390 Fibre (Channel) CONnection (FICON)**

### **FICON**

The e-business world is demanding. With the requirements for fast data access, continuous data availability, improved flexibility — all with lower cost of ownership — zSeries with native FICON-attached devices can help achieve these goals.

Native FICON delivers leading edge high speed data transfer technology to servers, switches, control units, and storage devices. FICON and FICON Express channels allow a more efficient and faster data transfer while, at the same time, allowing customers to use their currently installed single mode and multimode fiber optic cables. FICON and FICON Express channels are available with both longwave and shortwave options.

### ***FICON and FICON Express***

FICON was first introduced on the G5 servers which support a maximum of 24 FICON channels. The G6 servers support a maximum of 36 FICON channels, and the z900 servers support a maximum of 96 FICON channels in any mix of long wavelength and short wavelength features. A single FICON channel is capable of supporting up to 3600 I/O operations per second on G5/G6 and up to 5000 I/O operations/sec on z900. A FICON Express channel can support up to 7200 I/O operations/sec in native mode on z900.

With its new internal bus and LC connectors, FICON Express is the z900's latest implementation of the Fibre Channel Architecture. Available only on the zSeries, this card has two channels and can achieve improved performance over the previous generation FICON channel card. The z900 servers can support a maximum of 96 FICON Express channels in any mix of longwave and shortwave features. A single FICON Express channel is capable of supporting up to 7200 I/O operations/sec on z900. Customers can leverage this additional capacity to consolidate channels and reduce configuration complexity, infrastructure costs, and the number of channels that must be managed.

A FICON channel can coexist with ESCON and parallel channels in any model S/390 G5/G6 or z900 server and is supported by Multiple Image Facility (MIF). A FICON Express channel can coexist with ESCON and FICON channels in any model z900 server and is supported by MIF as well. The G5/G6 FICON channels are

available in increments of one and the z900 FICON and FICON Express channels are available in increments of two. On the z900 server, a maximum of 32 channels (16 cards) can be plugged into a single new z900 I/O cage, with a total of 96 channels (48 cards) in three cages.

To be consistent with the requirement to provide an "OPEN" platform for connectivity, FICON and FICON Express are compatible with the Fibre Channel Physical and Signaling Standard (FC-PH) and are in the approval process to NCITS for formal adoption into the FC-4 layer of the Fibre Channel standard under the name FC-SB-2.

FICON and FICON Express support a maximum unrepeated distance of up to 10 km (6.2 miles) (20 km via RPQ, up to 100 km with repeaters) for nine micron single mode fiber and up to 550 meters (1,804 feet) for 50 or 62.5 micron multimode fiber through the use of long wavelength FICON channel cards and Mode Conditioning Jumper (MCP) cables. FICON and FICON Express use different jumper cables which are supplied as required during configuration/ordering process. The maximum unrepeated distances for the short wavelength FICON and FICON Express channel cards are up to 500 meters for 50 micron multimode fiber and 300 meters (984 feet) for 62.5 micron multimode fiber. FICON and FICON Express are designed to reduce the data droop effect that made long distances not viable for ESCON. This distance capability is becoming increasingly important as customers are moving toward remote I/O, vaulting for disaster recovery and Geographically Dispersed Parallel Sysplexes for availability.

### ***FICON Bridge***

The 9032 Model 005 ESCON Director supports FICON and FICON Express longwave features through the use of a bridge card and is intended to help provide investment protection for customers' currently installed ESCON control units. Up to 16 bridge cards are supportable on a single 9032 Model 005 with each card capable of sustaining up to eight concurrent ESCON data transfers. Current 9032 Model 005 ESCON Directors can be field upgradable at no charge to support the FICON bridge cards, and both bridge cards and ESCON cards can coexist in the same box.

### Native FICON

Native FICON support is available on G5/G6 and zSeries servers.

Currently, the IBM ESS models F10 and F20 have two new host adapters to support native FICON. These host adapters each have one port per card and can either be FC 3021 Fibre Channel/FICON (long wave) Host Adapter or FC 3032 Fibre Channel/FICON (short wave) Host Adapter. The F10 and F20 models can support up to 16 FICON ports per ESS.

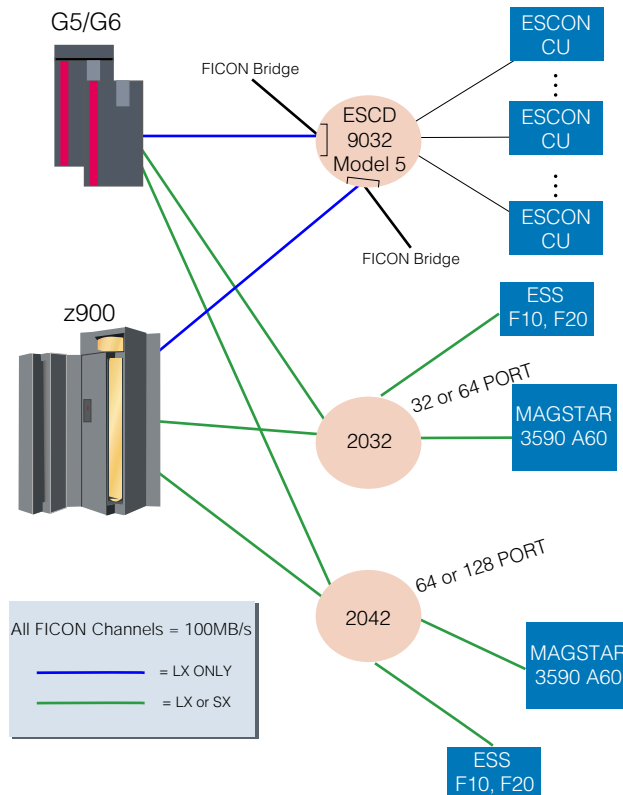
The IBM Magstar® 3590 A60 Tape controller provides up to two FICON interfaces which can coexist with ESCON on the same box. Customers can utilize IBM's highest capacity, highest performance tape drive to support their new business models.

Four Fibre Channel directors are now available to provide dynamic connectivity to native FICON control units. The IBM 2032 models 001 and 064 (resells of the McDATA ED-5000 and ED-6064) are 32 and 64 port high availability directors. The IBM 2042 models 001 and 128 (a resell of the Inrange FC/9000) are 64 and 128 port high availability directors. All have features that provide interface support to allow the unit to be managed by System Automation for OS/390.

OS/390 Resource Measurement Facility (RMF™) reports on the bus utilization, as well as port bandwidth for performance management and capacity planning. OS/390 R3 or higher with PTFs is required for FICON and FICON Express and associated RMF support in the 9032 Model 005 FICON bridge environment and OS/390 R6 or higher with PTFs is required for non-bridge environments.

Native FICON channels will 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.

### FICON Connectivity



### FICON CTC

FICON and FICON Express on z900 now support CTC. G5 and G6 servers can connect to z900 FICON CTC. This channel to channel connectivity will increase bandwidth between G5, G6 and z900 systems.

In addition, FICON CTC on the z900 is not limited to intersystem connectivity (as is the case with ESCON), but will also support multiple device definitions. ESCON channels that are dedicated as CTC cannot communicate with any other device. z900 FICON and FICON Express channels are not dedicated to CTC only. They can support both device and CTC mode definition concurrently, allowing for greater connectivity flexibility.

## Open Systems Adapter-Express Features (OSA-Express)

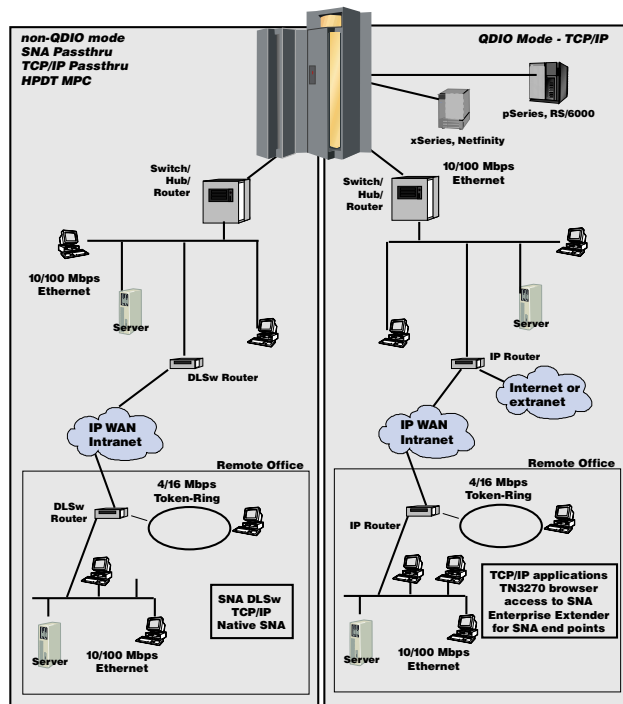
The Open Systems Adapter-Express (OSA-Express) features (Gigabit Ethernet, Fast Ethernet, and 155 ATM, and the new Token-Ring) have been redesigned to support the new I/O infrastructure of zSeries. These features combined with z/OS or OS/390, z/VM or VM/ESA, Linux and VSE/ESA in the future, deliver a balanced system solution to maximize throughput and minimize host interrupts to continue to satisfy your business goals. The redesigned zSeries OSA-Express feature requires the new I/O Cage.

The OSA-Express features currently available on the S/390 Parallel Enterprise Servers G5 and G6 are not available on z900. A Feature Conversion is available to help G5/G6 Gigabit Ethernet, Fast Ethernet, ATM and Token-Ring users to move z900 OSA-Express. Users of OSA-2 FDDI will require a compatibility I/O cage on the z900.

A higher performing Gigabit Ethernet feature with a 66 MHz, 64-bit PCI bus capable of achieving line speeds of 1 Gb/s is available on the z900. This new design incorporates two ports of connectivity in a single I/O slot. Each port uses one CHPID. These features attach to the STI bus at 333 MB/s.

The infrastructure for Fast Ethernet and 155 ATM and the new Token-Ring available on the z900 offers two ports of connectivity in a single I/O slot. The Fast Ethernet runs at either 10 or 100 Mbps. The ATM runs at 155 Mbps. The new Token-Ring runs at either 4 or 16 Mbps permitting connection to higher speed Token-Ring switch backbones and servers. Each OSA-Express port uses one CHPID. These features also attach to the STI bus at 333 MB/s.

### z900 OSA-Express Fast Ethernet

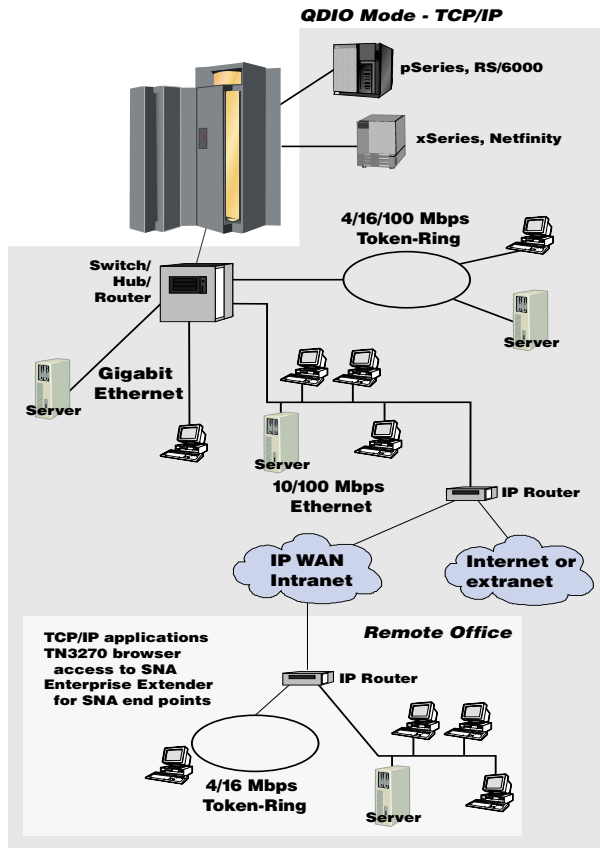


### Queued Direct Input/Output (QDIO)

The OSA-Express Gigabit Ethernet, Fast Ethernet, and 155 ATM (Ethernet LAN Emulation), and Token-Ring features support QDIO, which was first introduced in Communication Server for OS/390 Version 2 Release 7.

Queued Direct Input/Output (QDIO), a highly efficient data transfer architecture, breaks the barriers associated with the Channel Control Word (CCW/ESCON) architecture increasing data rates and reducing CPU cycle consumption. QDIO is designed to dramatically reduce system overhead and improves throughput using S/390 memory queues and a signaling protocol to directly exchange data between the OSA-Express microprocessor and CS for OS/390 R10 and z/OS.

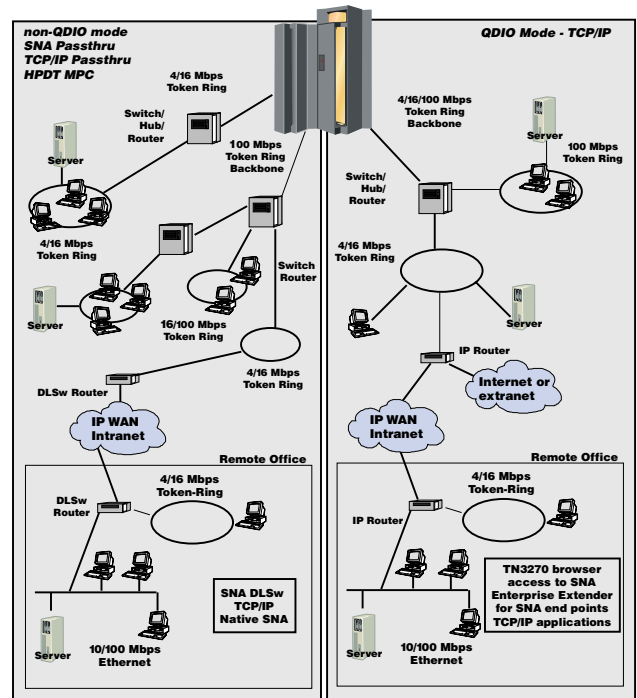
## z900 OSA-Express Gigabit Ethernet



### NON-QDIO operational mode

The OSA-Express Fast Ethernet, the new Token-Ring, and the 155 ATM (native, Ethernet and Token-Ring LAN emulation) features also support the non-QDIO mode of operation. The adapter can only be set (via the CHPID type parameter) to one mode at a time. The non-QDIO, mode does not provide the benefits of QDIO. However it does provide for protocol support similar to the OSA-2 adapter but at higher levels of performance. This support includes native SNA, APPN®, High Performance Routing, TCP/IP passthru, and HPDT MPC.

## z900 OSA-Express Token-Ring



### Server to User connections

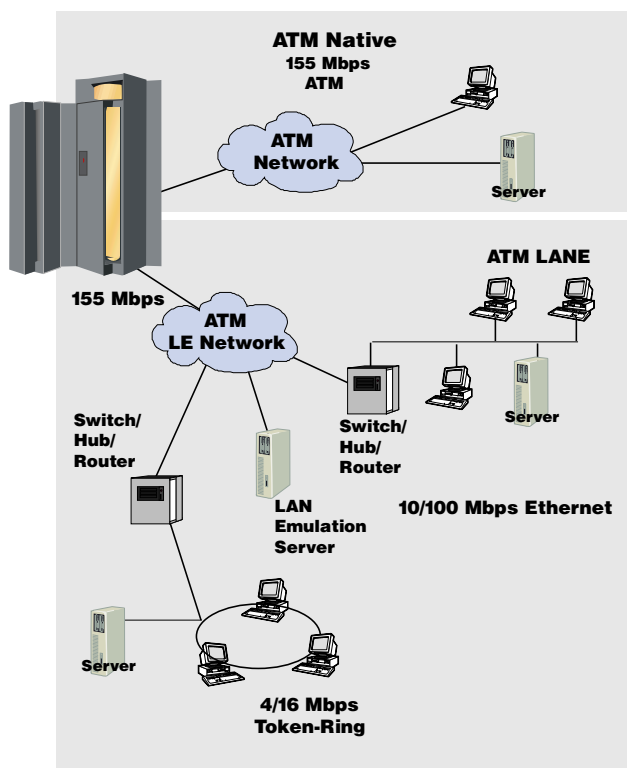
A key strength of OSA-Express and associated Communications Server protocol support is the ability to accommodate the customer's attachment requirements, spanning combinations of TCP/IP and SNA applications, and devices. Customers requiring TCP/IP connections from the remote site to the TCP/IP or SNA applications on z900 and S/390 can use OSA-Express with QDIO and either direct TCP/IP access or appropriate SNA to IP integration technologies such as TN3270 Server and Enterprise Extender. Customers who require the use of SNA-based connections from the remote site can use a TCP/IP or SNA transport to the data center and connection into z900 and S/390 using OSA-Express in non-QDIO mode.



## LPAR Support of OSA-Express

For customers who use the Processor Resource/ Systems Manager™ (PR/SM™) capabilities of the z900 and the S/390 servers, IBM offers the G5 and G6 Multiple Image Facility (MIF), allowing the sharing of physical channels by any number of logical partitions (LPARs). Since a port on an OSA-Express feature is like a channel, sharing of an OSA-Express port is done using MIF. The LPARs are defined in the Hardware Configuration Definition (HCD), and each LPAR is given a specific partition ID. Depending upon the feature, and how it is defined, SNA/APPN/HPR and TCP/IP traffic can flow simultaneously through any given port.

## z900 OSA-Express Asynchronous Transfer Mode



## Open Systems Adapter 2

The OSA-2 Ethernet-Token-Ring adapter feature #5201 (sometimes called the ENTR or OSA-2 Token-Ring) will not be available on the z900 with the availability of the new OSA-Express Token-Ring adapter.

The Open Systems Adapter 2 (OSA-2) Fiber Distributed Data Interface (FDDI) feature continues to be supported on the z900, and is available for new builds and for upgrades.

For planning purposes, the OSA-2 Token-Ring and FDDI feature can only be used in the compatibility I/O cage, while the new OSA-Express Token-Ring uses the new I/O cage.

The OSA-2 feature require the S/390 Open Systems Adapter Support Facility (OSA/SF) when configuring and customizing the feature, and updating the supported software.

*Note:* OSA-Express Token-Ring can now be used instead of OSA-2 Token-Ring, removing the requirement for the CI/O cage slot.

## **Advanced Availability Functions**

### **Transparent Sparing**

z900 offers two MCMs, a 20-processor version (max 16 Central Processors) and a 12-processor version (max 9 CPs). The added capacity is availability for nondisruptive customer growth or, in the case of processor failure, for transparent sparing.

### **Enhanced Dynamic Memory Sparing**

The z900 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:* z900 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**

*Upgrade for Coupling Links:* z900 has concurrent maintenance for the ISC-3 adapter card. Also, Coupling Links can be added concurrently. This eliminates a need for scheduled downtime in the demanding sysplex environment.

*Cryptographic Coprocessors:* The Cryptographic Coprocessors are designed as Single Chip Modules mounted on the Processor Board and are individually serviceable. This eliminates the need to change the MCM (where they previously resided) in the event of a crypto chip failure, reducing downtime.

*Concurrent Service for I/O Cards:* All the cards which plug into the new I/O Cage are able to be added and replaced concurrent with system operation. This eliminates any need to schedule outage to service or upgrade the I/O subsystem on this cage.

*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 eliminates customer scheduled outage.

*Auto-Switchover for Service Element:* The z900 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**

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

### **Concurrent Memory Upgrade**

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

## zSeries 900 Exclusive Functions


- z/Architecture
  - 64-Bit Architecture
- Intelligent Resource Director
  - LPAR CPU Management
  - Dynamic Channel Path Management
  - Channel Subsystem Priority Queuing
- HiperSockets
- Internal System Channel-3
- Integrated Cluster Bus-3
- Internal Coupling Channel-3
- Channel CHPID Assignment
- Nondisruptive replacement of I/O
- Two port OSA-Express adapters at line speed
- ESCON Port Sparing
- FICON Express
- FICON CTC
- Concurrent Maintenance for ISC-3 adapter card
- Concurrent upgrade for Coupling Links
- Concurrent Service for all I/O cards
- Total 512K subchannel addresses
- Redundant Cage Controllers
- Auto-Switchover for Service Element
- More Dense Logic Modules (MCMs)
- Enhanced Storage Protect Keys
- Enhanced Dynamic Memory Sparing
- Enhanced Hardware Compression
- Cryptographic Coprocessors (Single Chip Module)
- Integrated Facilities for Linux (IFLs)
- Nondisruptive CBU downgrade
- Concurrent memory upgrade

### -Note:

- G5/G6 supports IFL
- z900 supports all G5/G6 functions, except
  - > Asynchronous Data Mover Facility (ADMF)
  - > Integrated Coupling Migration Facility (ICMF)

## zSeries 900 Model Upgrades

- Comprehensive upgrade paths available
- G5 to z900 Models
- G6 to z900 Models
- G5 R06 Coupling Facility to z900 Model 100 CF
- G5 R06 CF to z900 Models
- Vertical upgrades within z900
- Vertical upgrades within z900 Model 100 CF
- z900 Model 100 CF to z900 Models



## Processor Options z900 Models 101-116, 1C1-1C9

	PUs	CPs	SAPs		IFLs*/	ICB-3*	ICB***	ISC-3*	STIs
			Std	Opt*	ICFs*				
101	12	1	2	3	8	16	8**	32	24
102	12	2	2	3	7	16	8**	32	24
103	12	3	2	3	6	16	8**	32	24
104	12	4	2	3	5	16	8**	32	24
105	12	5	2	3	4	16	8**	32	24
106	12	6	2	3	3	16	8**	32	24
107	12	7	2	2	2	16	8**	32	24
108	12	8	2	1	1	16	8**	32	24
109	12	9	2	0	0	16	8**	32	24
110	20	10	3	5	6	16	8**	32	24
111	20	11	3	5	5	16	8**	32	24
112	20	12	3	4	4	16	8**	32	24
113	20	13	3	3	3	16	8**	32	24
114	20	14	3	2	2	16	8**	32	24
115	20	15	3	1	1	16	8**	32	24
116	20	16	3	0	0	16	8**	32	24
100	12	0	2	0	9 (ICF)	16	16	32****	24
1C1	20	1	3	5	15	16	8**	32	24
1C2	20	2	3	5	14	16	8**	32	24
1C3	20	3	3	5	13	16	8**	32	24
1C4	20	4	3	5	12	16	8**	32	24
1C5	20	5	3	5	11	16	8**	32	24
1C6	20	6	3	5	10	16	8**	32	24
1C7	20	7	3	5	9	16	8**	32	24
1C8	20	8	3	5	8	16	8**	32	24
1C9	20	9	3	5	7	16	8**	32	24
* up to ** up to 16 ICB links available via an RPQ if Compatibility I/O Cage is not required. ***ICB Compatibility Links **** Up to 42 with RPQ 8P2233									

## zSeries 900 Family Configuration Detail

I/O Channel / ports or STIs/Feature <sup>1</sup>	I/O Cage min 1 - max 3 <sup>2</sup>	Compatibility I/O Cage min 0 - max 2 <sup>3</sup>	Total Quantity of Channels <sup>4</sup>	Channel Increments	Comments Total I/O cages min 1 - max 3 <sup>5</sup>
ESCON - 4 port	No	Yes	176 / 44	4	G5/G6 upgrades only
ESCON - 15/16 port	Yes	No	256 / 18	4 (LIC CC)	1 spare port per card
Parallel - 4 port	No	Yes	88 / 22	4	96 via RPQ or G5/G6 u/g
Parallel - 3 port	No	Yes	96 / 32	3	MES upgrade only
FICON/FICON Express - 2 port <sup>7</sup>	Yes	No	96 / 48	2	—
PCICC - 2 processors	Yes	No	16 / 8	2	Not defined in IOCP
PCICA - 2 processors	Yes	No	12 / 6	2	—
Fast Ethernet - 2 port	Yes	No	24 / 12	2	G5/G6 upgrades only
Gbit Ethernet <sup>7</sup> - 2 port	Yes	No	24 / 12	2	Different features
155 ATM <sup>7</sup> - 2 port	Yes	No	24 / 12	2	Different features
FE, Gbit and ATM	—	—	—	—	24 ports (12 card)
Token Ring	Yes	No	24 / 12	2	G5/G6 upgrades only
FICON, FICON Express, PCI-CC/CA, FE, GbE, ATM, TR	—	—	—	—	48 card/system 16 card/ I/O cage
OSA-2 - TR - FDDI - 1 port	No	Yes	12 / 12	1	FDDI + TR = 12 cards max
OSA-2 - TR - 2 ports	No	Yes	24 / 12	1	FDDI + TR = 12 cards max
ISC-3 (1 & 2 Gbit) -2 + 2 ports	Yes	No	32 / 8	1 (LIC CC)	20 km via RPQ (H/W 1Gb)
ICB-3 - 1 STI	No	No	16 / 0	1	1 GB STI
ICB - 1 STI	No	No	8 / 0	1	333 MB STI <sup>6</sup>
ISC-3, ICB-3, ICB	—	—	—	—	32 channel maximum
IC Channel <sup>8</sup>	No	No	32 / 0	2	u-code support
ISC-3, ICB-3, ICB, IC Channels	—	—	—	—	64 channel maximum
HiperSockets	—	—	4/0	1	microcode support -1 CHPID for each defined HiperSockets

1. maximum 256 CHPIDS; 2. 28 I/O slots per nI/O Cage; 3. 22 I/O slots per cI/O Cage; 4. Supported and I/O Slots required; 5. nI/O Cage plus cI/O Cage; 6. RPQ to 12: if no cI/O Cage present--RPQ to 16; 7. MM & SM; 8. 2 defined channel


nI/O=new zSeries I/O Cage; cI/O=Compatibility I/O Cage

## Summary of z900 Features

### Processor Unit Assignments

Models 101-109

Models 110-116, 1C1-1C9



### Processor Storage

z900 Models	100-109	110-116	1C1-1C9
Minimum	5 GB	10 GB	10 GB
Maximum	32 GB	64 GB	64 GB
Storage sizes: 5, 6, 7, 8, 10, 12, 14, 16, 18, 20, 24, 28, 32, 40, 48, 56, 64 GB			

### Channels

	Parallel	ESCON	FICON/FICON Express
Minimum	0	0	0
Maximum	88/96*	256	96
Increments	3**/4	4	2

\*Available with G5/G6 upgrade or via RPQ for new builds.  
 \*\*Available with G5/G6 upgrade only (existing 3-port parallel card).  
 Notes:  
 - Min. 1 Feature of Parallel, ESCON or FICON or FICON Express required  
 - Total CHPIDS all types (parallel, ESCON, FICON, FICON Express, OSA, PCICC/CA, ISC-3, ICB, ICB-3, IC-3): Maximum 256

### Coupling Links

Links	IC	ICB-3	ICB	ISC-3	Max # Links
z900 - 100 CF	32	16	16	32 42 w/RPQ	64
z900 Server	32	16	8 12/16 w/RPQ	32	32

### PCI Cryptographic Coprocessor

PCICC	Dual PCICC Features*	Cryptographic Engines**
Minimum	0	0
Maximum	8	16***
Increments	1	2
* requires nI/O Cage; 1 card slot/feature ** 2 Crypto engines/feature; 2 CHPIDs per feature *** Combined total of PCICC and PCICA engines cannot exceed 16		

### PCI Cryptographic Accelerator

PCICA	Dual PCICA Features*	Cryptographic Engines**
Minimum	0	0
Maximum	6	12***
Increments	1	2
* requires nI/O Cage; 1 card slot/feature ** 2 Crypto engines/feature; 2 CHPIDs per feature *** Combined total of PCICC and PCICA engines cannot exceed 16		

### OSA Networking Features

	OSA-Express*	OSA-2**
Minimum	0	0
Maximum	12	12
Increments	1	1
* GbE, FEN, ATM, TR; 2-ports/feature; 2 CHPIDs/feature; requires nI/O Cage; one card slot/feature ** Requires cI/O Cage: TR: 2-ports/feature, FDDI: 1 port/feature; 1 CHPID/feature; 1 card slot/feature		

## z900 Frame and I/O Configuration Content: Planning for I/O


The following frame configurations and I/O tables show the capability and flexibility built into the I/O subsystem. The single frame system table (new build systems) shows that up to 256 ESCON channels can be contained in one new zSeries I/O (nl/O) Cage. 256 ESCON channels, with the 16 port card (1 port is a spare), will fill 18 of the 28 card slots in the nl/O Cage demonstrating the packaging efficiency of the current design. A more typical configuration of 200 ESCON (14 cards) and 20 FICON (10 cards) channels, 4 OSA-Express ports (2 cards) and 8 ISC links (2 cards) still fits within the single nl/O Cage (using all 28 card slots).

A customer planning to migrate to more FICON or FICON Express channels or to a larger configuration, at a later date, might consider using the Plan Ahead feature to replace the single frame, single nl/O Cage configuration with a two frame and two nl/O Cage configuration. This action will allow nondisruptive changes up to the limit of the two nl/O Cages. The advantage of doing this is that: 1) the channel upgrade will be completely nondisruptive and 2) the original I/O will be distributed across the two nl/O Cages giving better balance and availability characteristics to the upgraded system. Frame and cage configurations and their I/O capabilities are shown on the following pages.

For a new build system with parallel channel and OSA-2 FDDI requirements, a Compatibility I/O (cl/O) Cage is required. The single cl/O Cage has 22 card slots and can house up to 88 parallel channels (4 channels per card) and up to 12 OSA-2 features. If requirements for greater than 22 card slots are required, a second cl/O Cage can be obtained through an RPQ; however, this will reduce the number of nl/O Cages to a single cage in the z900 A-Frame.

Most upgrades from a G5/G6 system into a z900 system will have two cl/O Cages in the Z-frame for packaging

parallel channels, OSA-2 Token-Ring/FDDI ports and ESCON 4 port cards from the previous G5/G6 system and one nl/O Cage to support the new I/O cards. Each cl/O Cage has 22 card slots (2 cages, 44 slots). The ESCON 4 port cards are available on the z900 only through an upgrade from a G5/G6. Any ESCON overflow out of the cl/O Cages to accomplish the upgrade and any ESCON new channel requirements are provided by the new ESCON 16-port card in the nl/O Cage. A typical upgrade consisting of 200 ESCON channels (50 cards), 24 parallel channels (6 four port cards) and 2 FDDI ports (2 cards) would have an ESCON overflow of 56 ESCON channels into the nl/O Cage. Fifty-six ESCON (new) channels would require four of the twenty-eight slots in the nl/O Cage. Customers upgrading from G5/G6 to z900 should consider future channel requirements, FICON migration, and move to newer OSA-Express Gigabit Ethernet, Fast Ethernet, ATM and Token Ring adapters as part of the upgrade planning. The IBM Account team will help with this planning.




### New Build 1 nl/O Cage

Channel Type	Features <sup>1</sup>	Maximum
ESCON (new) <sup>3</sup>	18 cards max	256 channels
FICON/FICON Express	16*	32 channels
OSA-Express <sup>4</sup>	12*	24 ports
PCICC	8*	16 engines
PCICA	6*	12 engines
Parallel, OSA-2 <sup>2</sup>	none	none

\* Maximum combined FICON, FICON Express, OSA-Express, PCICC/CA features is 16

See notes on page 57.


## z900 Frame and I/O Configuration Content: Planning for I/O



**New Build  
2 nI/O Cages**

Channel Type	Features <sup>1</sup>	Maximum
ESCON (new) <sup>3</sup>	18 cards max	256 channels
FICON/FICON Express	32*	64 channels
OSA-Express <sup>4</sup>	12*	24 ports
PCICC	8*	16 engines
PCICA	6*	12 engines
Parallel, OSA-2 <sup>2</sup>	none	none


\* Maximum combined FICON, FICON Express, OSA-Express, PCICC/CA features is 32



**New Build  
3 nI/O Cages**

Channel Type	Features <sup>1</sup>	Maximum
ESCON (new) <sup>3</sup>	18 cards max	256 channels
FICON/FICON Express	48*	64 channels
OSA-Express <sup>4</sup>	12*	24 ports
PCICC	8*	16 engines
PCICA	6*	12 engines
Parallel, OSA-2 <sup>2</sup>	none	none


\* Maximum combined FICON, FICON Express, OSA-Express, PCICC/CA features is 48



**New Build  
1 nI/O Cage  
1 cI/O Cage**

Channel Type	Features <sup>1</sup>	Maximum
ESCON (new) <sup>3</sup>	18 cards max	256 channels
FICON/FICON Express	16*	32 channels
OSA-Express <sup>4</sup>	12*	24 ports
PCICC	8*	16 engines
PCICA	6*	12 engines
Parallel <sup>2</sup>	22	88 channels
FDDI <sup>2</sup>	12	12 ports


\* Maximum combined FICON, FICON Express, OSA-Express, PCICC/CA features is 16



**New Build  
2 nI/O Cages  
1 cI/O Cage**

Channel Type	Features <sup>1</sup>	Maximum
ESCON (new) <sup>3</sup>	18 cards max	256 channels
FICON/FICON Express	32*	32 channels
OSA-Express <sup>4</sup>	12*	24 ports
PCICC	8*	16 engines
PCICA	6*	12 engines
Parallel <sup>2</sup>	22	88 channels
FDDI <sup>2</sup>	12	12 ports

\* Maximum combined FICON, FICON Express, OSA-Express, PCICC/CA features is 32



**G5/G6 Upgrade to z900  
1 nI/O Cage  
2 cI/O Cages  
(typical upgrade)**

Channel Type	Features <sup>1</sup>	Maximum
ESCON (new) <sup>3</sup>	<18 cards**	<256 channels
FICON/FICON Express	16*	32 channels
OSA-Express <sup>4</sup>	12*	24 ports
PCICC	8*	16 engines
PCICA	6*	12 engines
ESCON (4-port)**	44	176 channels
Parallel <sup>2</sup>	24 4-port 32 3-port	96 channels
Token-Ring <sup>2</sup>	12	24 ports
FDDI <sup>2</sup>	12	12 ports

\* Maximum combined FICON, FICON Express, OSA-Express, PCICC/CA features is 16

\*\* For upgrades, the ESCON 16 port card is used for ESCON overflow out of the Compatibility I/O cage and for fulfilling new ESCON requirements.

### Notes:

- 1) Up to this number. nI/O Cage 28 features (cards) maximum, cI/O Cage 22 features maximum. (see notes 3 & 6)
- 2) cI/O Cage only.
- 3) ESCON Configured in 4 port increments. Up to 28 channels in 2 cards, 60 channels in 4 cards, 88 in 6 cards, 120 in 8 cards, 148 in 10 cards, 180 in 12 cards, 208 in 14 cards, 240 in 16 cards, 256 in 18 cards.
- 4) OSA-Express can be Gbit Ethernet/ Fast Ethernet/Token Ring/ATM
- 5) Up to 256 channels for the system including ISC-3, ICB and IC links, and Crypto PCICC/CA.
- 6) ISC links, although not in the tables, are packaged in the nI/O Cage. To 8 ports in 2 cards, to 12 ports in 3 cards, 16 ports in 4 cards, 20 in 5 cards, 24 in 6 cards, 28 in 7 cards, 32 in 8 cards.
- 7) One nI/O Cage is standard in the A-Frame on all general purpose z900 processors.
- 8) Total number of PCICC and PCICA cards cannot exceed 8 with a maximum number of 6 for PCICA.

## Physical Characteristics

### Models 101-116 and 1C1-1C9

	Minimum 1 Frame System without B-Frame	Maximum 2 Frame System with B-Frame
Power Requirements 50/60 Hz, kVA	5.3	13.7
Heat output KBTU/hr	18.1	46.8
Air flow CFM Air flow m <sup>3</sup> /m in	800 22.2	1990 55.2
Floor space – Sq meters – Sq feet	1.32 14.2	2.81 30.3
Including service clearance – Sq meters – Sq feet	3.04 32.7	6.18 66.5
Approximate weight – kg – lbs	917 2021	1866 4113
Height – cm – inches	200.4 79.8 <sup>1</sup>	200.4 79.8 <sup>1</sup>
<sup>1</sup> Reduced height feature = 70.5 inches with packaging Without covers = 73.3 inches		

### Model 100 Coupling Facility

	Minimum 1 Frame System without B-Frame and I/O Cage	Maximum 1 Frame System with B-Frame and ISC Links
Power Requirements 50/60 Hz, kVA	4.0	5.0
Heat output KBTU/hr	13.7	17.1
Air flow CFM Air flow m <sup>3</sup> /m in	600 16.6	750 20.8
Floor space – Sq meters – Sq feet	1.32 14.2	1.56 16.8
Including service clearance – Sq meters – Sq feet	3.04 32.7	3.28 35.3
Approximate weight – kg – lbs	644 1450	1189 2621
Height – cm – inches	200.4 79.8 <sup>1</sup>	200.4 79.8 <sup>1</sup>
<sup>1</sup> Reduced height feature = 70.5 inches with packaging Without covers = 73.3 inches		

## Coupling Facility—CF Level of Support

CF Level	Function
10	System-Managed CF Structure Duplexing
9	MQSeries shared message queues WLM Multi-System Enclave Support Intelligent Resource Director ** IC3/ISC3/ICB3 peer mode**
8	Dynamic ICF expansion into shared ICF pool Systems-Managed Rebuild
7	Shared ICF partitions on server models DB2 Delete Name optimization
6	ICB & IC TPF support
5	DB2 cache structure duplexing DB2 castout performance improvement Dynamic ICF expansion into shared CP pool*
4	Performance optimization for IMS & VSAM RLS Dynamic CF Dispatching Internal Coupling Facility* IMS shared message queue extensions
3	IMS shared message queue base
2	DB2 performance VSAM RLS 255 Connectors/1023 structures for IMS Batch DL1 (non-BMP)
1	Dynamic Alter support CICS temporary storage queues System logger

#### Notes:

- G1 base can be upgraded to CF Level 4.
- G2 base can be upgraded to CF Level 5.
- G3 and G4 can be upgraded to CF Level 8.
- G5 base CF level code is CF Level 6.
- G6 base CF level code is CF level 8.
- Detailed information regarding CF Levels can be found in "Coupling Facility Level (CFLEVEL) Considerations" at url [ibm.com/s390/psa/cftable.html](http://ibm.com/s390/psa/cftable.html)

\* G3, G4, G5, and G6 only

\*\* - z900 required

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



### **Fiber-Optic Cabling and System Connectivity**

In designing the z900 I/O subsystem, primary goals included openness, industry standards (Fibre Channel, SANs, connectors, etc.), significant bandwidth and I/O packaging density to handle the new e-business workloads as well as the traditional transaction processing environment. The connectivity alternatives available in this high performance environment have significantly increased the number of cabling alternatives available to connect systems to I/O and networks. It is important that the customer's infrastructure planning consider the technology alternatives. In addition, fiber-optic trunking via Fiber Quick Connect is available for ESCON on the z900. In planning for z900 systems, the customer should reference *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 on Resource Link ([ibm.com/servers/resourcelink](http://ibm.com/servers/resourcelink)).

IBM offers fiber-optic cable features orderable with the z900 and/or IBM Global Services to provide the cable connectivity necessary for installing the z900. A limited number of cable features are orderable with a z900 system which adapt the system to an existing cable infrastructure. Included are two 2 meter conversion cables which provide connector integration to the new ESCON, ETR and ISC-3 I/O cards and five 2 meter Mode Conditioning Jumper (MCP) kit features that

handle single mode to multi-mode infrastructure conversions. In addition, there are ESCON trunk (Fiber Quick Connect; harnesses and brackets) features available for customers desiring the efficiency and convenience of trunking solutions.

IBM Global Services can supply the full complement of cabling solutions including the conversion cables and MCP kit cables plus jumper cables of various lengths, connector arrangements, cable types, and ratings for new connection requirements. IBM Global Services also offers Fiber Transport Configuration Service; a specific cable service contract for the planning, supply and installation of up to 50 cables or a full service contract for implementation of a Fiber Transport structured cabling system.

Bottom line: IBM offers a wide range of cables and cabling solutions to meet a customer's requirements for connecting to I/O and networks.

While z900 servers are supported by a multitude of 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 a new, robust operating system that is based on the new 64-bit z/Architecture. It delivers the highest qualities of service for enterprise transactions and data, and extends these qualities to new applications using the latest software technologies. It provides a highly secure, scalable, high-performance base on which to build and deploy Internet and Java®-enabled applications, providing a comprehensive and diverse application execution environment. z/OS takes advantage of the latest software technologies: new object-oriented programming models that permit the rapid design, development and deployment of applications essential to enterprise e-business. It protects your investment in your present S/390 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 Java Beans™, XML, HTML, 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.

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 instantly to changing conditions and prioritize critical business workloads.
- Support for 64-bit real memory and initial support of 64-bit virtual storage.

- A new installation and configuration infrastructure that reduces the skills required to install and configure z/OS and related products.
- A new software pricing model designed to support e-business reality.
- System-managed CF Structure duplexing.

### **Intelligent Resource Director**

Intelligent Resource Director (IRD) is a new feature of the z/Architecture which extends the Workload Manager to work with PR/SM on z900 servers to dynamically manage resources across an LPAR cluster. An LPAR cluster is the subset of the systems that are running as LPARs on the same CEC. 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 which work together to respond to the demands of e-business:

- LPAR CPU Management
- Dynamic Channel Path Management
- Channel Subsystem Priority Queuing

(see page 8)

### **64-bit Real Storage Support**

z/OS provides 64-bit real storage support for the new z900 servers. z/OS will continue to provide 31-bit real and expanded storage support for G5/G6, and Multiprise 3000 servers. 64-bit real support will eliminate expanded storage and help eliminate paging. The 64-bit real support may allow you to consolidate your current systems into fewer LPARs or to a single native image.

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

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

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

- DB2 Version 6 (with PTF)
- IMS Version 7

IBM is working with many software vendors to enhance 64-bit exploitation. IBM and ISV software products that run in 31-bit mode under OS/390 V2R6 through OS/390 V2R10 will run unchanged on z900 servers under OS/390 V2R6 through OS/390 V2R10.

In addition, z/OS V1R2 delivers basic 64-bit virtual storage management support. Assembler programs can obtain virtual storage above 2 GB for storing and manipulating data.

### **Managed System Infrastructure for Setup (msys for Setup)**

z/OS starts a major ease-of-manageability initiative with the introduction of Managed System Infrastructure for Setup (msys for Setup) as a new base element. msys for Setup offers a new approach for installing and configuring z/OS and products running on z/OS that result in major productivity improvements.

msys for Setup allows for the usage of consistent interfaces with wizard-like configuration dialogs. These new dialogs reduce the skill requirements for setting up products, freeing up scarce and valuable personnel for other tasks.

msys for Setup builds upon the Web-based wizard technology that has been introduced in a number of areas. msys for Setup employs the same easy, interview style as the wizards for defining the customization parameters.

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. Furthermore, msys for Setup introduces an LDAP-based z/OS management directory that will become the central repository for all configuration data.

msys for Setup makes setting up a Parallel Sysplex resource sharing environment much quicker and easier by defining and implementing the required:

- policies
- parmlib specifications
- security

z/OS V1R1 focuses on resource sharing configurations, including support for Intelligent Resource Director, XCF signaling, and Global Resource Serialization star.

z/OS V1R2 extended this support to cover additional tasks, such as setting up a system logger environment as well as customizing enhanced catalog sharing. Also you can define basic IP settings, define network devices and links or set up FTP or Telnet 3270 servers. Finally, ISPF definitions can be generated through msys for setup.

### **Workload License Charges — Pay for what you need**

Workload License Charges is IBM's newest software pricing strategy designed to support e-business reality. With Workload License Charges, customers pay only for the software capacity that they need. Because the **new pricing structure** is designed around the concept of paying for what you need, customers can grow a single workload at a time, with ease and granularity.

#### *Highlights*

- New software pricing strategy designed to support e-business reality
- Changes the metric of software pricing from machine-capacity to sub-machine capacity
- Allows customers to grow one workload at a time, in a simple and granular fashion
- Provides enhanced price performance

- Enables customers to add excess hardware capacity without incurring additional software charges
- Permits customers to define and pay for software on their (four hour rolling) average requirement instead of the peak requirement

### System-Managed CF Structure Duplexing

z/OS V1R2 introduces System-Managed Coupling Facility (CF) Structure Duplexing to significantly enhance Parallel Sysplex availability. It provides a robust failure recovery capability via CF structure redundancy, and it enhances Parallel Sysplex ease of use by reducing the complexity of CF structure recovery. These benefits are achieved by creating a duplicate (or duplexed) copy of a CF structure and then maintaining the two structure instances in a synchronized state during normal CF operation. In the event of a CF related failure (or even a planned outage of a CF), failover to the remaining copy of the duplexed structures is initiated and quickly completed transparent to the CF structure user and without manual intervention.

### TCP/IP Networking Enhancements

z/OS can provide near continuous availability for TCP/IP applications and their users with two key features in z/OS: Sysplex Distributor and VIPA nondisruptive Takeover. 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 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 is that if an adapter fails, the IP protocol can find an alternate path to the same software, be it the TCP/IP services on z900 or an application.

VIPA Takeover introduced in OS/390 R8 supports movement to a backup IP stack on a different server in a Parallel Sysplex cluster in case of a failure of the primary IP stack. VIPA Nondisruptive Takeover enhances the initial R8 functions, providing VIPA takeback support. This allows the movement of workload back from the alternate to the primary IP stack.

#### Sysplex Distributor

Introduced in OS/390 R10, Sysplex Distributor is a software-only means of distributing IP workload across a Parallel Sysplex cluster. Client connections appear to be connected to a single IP address yet the connections are routed to servers on different z900 or S/390 servers. In addition to load balancing, Sysplex Distributor simplifies the task of moving applications within a Parallel Sysplex environment.

### System Services

#### z/OS Version 1 Release 2 base elements

Base Control Program (BCP)  
JES2  
ESCON Director Support  
MICR/OCR Support  
Bulk Data Transfer base  
DFSMSdfp™  
EREP/MVS™ V3R5  
High Level Assembler V1R4  
ICKDSF R16  
ISPF  
TSO/E  
3270 PC File Transfer Program V1R11  
FFST™/ESA  
TIOC  
IBM License Manager

#### z/OS Version 1 Release 2 optional priced features

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

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.

#### Highlights:

**Workload Manager** (WLM) addresses the need for managing mixed workload distribution, load balancing and the distribution of computing resources to competing workloads. It does this while providing fewer, simpler system externals, performance management goals expressed in Service Level Agreement terms, automatic work and resource management. All this is

done with a single policy that can be used across the sysplex to provide a single control point, eliminating the need to manage each individual image.

**IBM License Manager**, a new base element of z/OS, is not yet enabled for use.

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

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

### **Parallel Sysplex**

#### *Dynamic Workload Balancing*

Data sharing and workload balancing enable work to be directed to available processors “on the fly” and also permit servers to be dynamically added to the cluster without requiring costly downtime. This can be accomplished without splitting applications or databases across multiple servers, a time-consuming and expensive process, or requiring significant new investments in staff or systems management.

#### *Auto Alter*

Dynamically tunes many of the CF structure sizes and object ratios.

#### *Automatic Restart Manager (ARM)*

To minimize the impact of a failed database manager or a failed system, this database manager with its associated transaction manager regions must be restarted as soon as possible. Since ARM is part of z/OS, it quickly knows of address space, JOB, and system terminations. ARM knows when all resource

managers have cleaned up for the failing system and can therefore restart registered JOBS/STCs as quickly as possible. In system failure cases, ARM uses policy and Workload Manager (WLM) to balance the restart of failed elements across surviving systems.

### **Base Control Program**

- The BCP has support for the IEEE floating point hardware which is standard on all servers supported by z/OS.
- Allocation improvements allow DB2 to support customers with a requirement for more than 10,000 DB2 dynamically-allocated table spaces.
- Open Systems Adapter 2 (OSA-2) — support provides Fast Ethernet LAN attach communication.
- Dump command parameters can be supplied via a PARMLIB member. Parallel Sysplex problem determination provides convenient specification of multiaddress space and multisystem dumps.
- Display Logger command improves Parallel Sysplex-wide management of the System Logger. Logger serviceability is improved with additional diagnostic and data capture capability. DASD Logger Offload processing is improved by operating asynchronously to overlap operations and more efficiently utilize the I/O subsystem.
- IOS Channel Path Availability with asynchronous WTOR processing and enhanced recovery actions.
- Recoverable Resource Management Services (RRMS) recognizes Resource Managers that manage work contexts as Work Managers. RRMS systems management supports Work Manager names, wild card filtering, sorted displays, and exit duration reports with the result that it is easier to locate and manage RRMS coordinated transactions. Connections from any RRMS managed application-enabled environment to access IMS DB databases is supported by IMS Version 6 Open Database Access (ODBA) using the IMS Database Resource Adapter (DRA).

#### *JES2 and JES3*

- In z/OS V1R2, 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.

## System Management Services

### ***z/OS Version 1 Release 2 base elements***

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

### ***z/OS Version 1 Release 2 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.

z/OS introduces Managed System Infrastructure for Setup (msys for setup), a new base element, offering a new approach for installing and configuring z/OS and products running on z/OS. msys for Setup allows usage of consistent interfaces with wizard-like configuration dialogs. These new dialogs reduce the skill requirements for setting up products, freeing up scarce and valuable personnel for other tasks.

z/OS Managed System Infrastructure for Operations (msys for Operations) in z/OS V1R2 makes well-proven automation technology part of z/OS. It provides self-healing for some critical system and sysplex resources and simplifies the day-to-day operation of a z/OS Parallel Sysplex environment. It monitors sysplex specific events to avoid single points of failure (such as no alternate Couple Data Set) or sysplex-wide outages due to operator mistakes or system load.

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.

The unique and rich functions of System Automation for OS/390 (SA OS/390) Version 2.1 (separately orderable) can ease z/OS and OS/390 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.

#### *Highlights:*

- In z/OS V1R1, msys for Setup makes setting up a Parallel Sysplex resource sharing environment much quicker and easier by defining and implementing required policies, parmlib specifications and security settings
- z/OS V1R2 makes it easier to setup ISPF and IP Services, and provides increased benefits for setup of a Base or Parallel Sysplex environment
- RMF can show response time distributions and work manager delay data, as well as reporting information according to periods, even for report classes
- RMF can also show the contention for Cryptographic Coprocessors, including a description of which workloads are using or are delayed in access to the cryptographic coprocessors

#### *Benefits can include:*

- Improve the productivity of system programmers
- Reduce the skill level for configuring z/OS components or products

## Security Services

### **z/OS Version 1 Release 2 base elements**

Cryptographic Services (Integrated Cryptographic Service Facility, Open Cryptographic Services Facility, System SSL)

### **z/OS Version 1 Release 2 optional priced features**

SecureWay® Security Server, which includes:

- RACF®
- Network Authentication and Privacy Service<sup>2</sup> (Kerberos services)
- LDAP Server<sup>2</sup>
- Firewall Technologies
- DCE Security Server (OSF DCE level 1.2.2)
- Open Cryptographic Enhanced Plug-ins

### **z/OS Version 1 Release 2 optional no charge features**

Network Authentication Service Level 3<sup>1</sup>

System SSL Security Level 3<sup>1</sup>

Open Cryptographic Services Facility Security Level 3<sup>1</sup>

<sup>1</sup>Export considerations

<sup>2</sup>These functions are part of the base z/OS and do not require SecureWay Security Server feature license

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

#### *Highlights:*

- Integrated package of security and directory functions, with advanced security capabilities
- Tivoli administration interface

#### *Benefits can include:*

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

#### *RACF*

- Digital Certificates can be automatically authenticated without administrator action.
- Administrative enhancements enable definition of profiles granting partial authority. Handling of new passwords and removal of class authority are simplified.
- Network Qualified LU Names allow non-unique LU Names within interconnected networks.
- A protected user ID can be created to ensure user IDs assigned to UNIX, UNIX daemons, and other important started tasks and subsystems cannot be used for other purposes. It will also protect the IDs from being revoked, either accidentally or intentionally, with invalid password attempts.
- Identification and audit trail for a UNIX System user.
- Support for Tivoli management of all user segments.
- Customers designing e-business applications need a way to associate more users under a RACF Group definition, so RACF allows you to create a new kind of Group that can contain an unlimited number of users.

#### *SSL*

- Support for the Transport Layer Security (TLS) Version 1 protocol (RFC 2246).
- Support for PKIX-compliant Certificate Revocations Lists (CRLs) created by the Tivoli SecureWay Public Key Infrastructure.
- 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.

#### *LDAP*

- Client access to information in multiple directories is supported with the LDAP protocol.
- LDAP Server now at V3 level. Bulk load utility for populating the LDAP Directory.
- The LDAP server supports thousands of concurrent clients, increasing the maximum number of concurrently connected clients by an order of magnitude.
- The LDAP Server supports authentication using Kerberos credentials.
- An LDAP Configuration Utility will easily automate a basic setup.
- An LDAP client can find an LDAP server with information in a Domain Name System (DNS) server without knowing the LDAP server's host name or IP address in advance.

### ICSF

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

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

z/OS V1R2 provides leading edge SSL performance through support of the z900 IBM PCI Cryptographic Accelerator (PCICA), a fast cryptographic processor designed to provide leading edge performance of the complex RSA cryptographic operations used in the Secure Sockets Layer (SSL) protocol. This will benefit functions that use System SSL, such as the z/OS HTTP Server (and WebSphere), TN3270 server, LDAP server and CICS Transaction Gateway server.

#### *Network Authentication and Privacy Service*

- A component of SecureWay Security Server, Network Authentication and Privacy 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 and Privacy 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.

## **Application Enablement Services**

### **z/OS Version 1 Release 2 base elements**

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

### **z/OS Version 1 Release 2 optional priced features**

C/C++ with Debug Tool  
C/C++ without Debug Tool  
DFSORT™  
GDDM®-PGF V2R1.3¹  
GDDM-REXX V3R2¹  
HLASM Toolkit V1R3¹

1. z/OS V1R2 non-exclusive elements can be purchased standalone

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

#### *Highlights:*

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

Today's applications require much more than a set of traditional development tools. To provide true business value they must also:

- Enable rapid application development to deliver timely business function
- Be portable and interoperable
- Enable function to be delivered with performance across the network
- Allow access to data in a heterogeneous environment
- Be manageable and highly secure



*Benefits can include:*

- Improved development productivity, cycle time and quality
- Reuse of existing code as objects through “wrapping”
- Support for multiple language environments
- Improved end user productivity through provision of GUIs for existing applications
- Choice of many high level languages enables use of existing skills
- Improved interoperation and portability of code and objects between heterogeneous platforms
- Distributed applications development

*Language Environment (LE)*

- Performance improved by HEAPPOOLS, message filtering, and pathlength reductions.
- Serviceability changes adding CEEDUMP and IPCS support.

*C/C++*

- OS V1R2 provides two levels of the C/C++ compiler as migration aid. One is identical to the compiler provided with OS/390 V2 R10. 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 as on ASCII platforms, thus eliminating the need to find all such places in the code and converting them to EBCDIC when porting UNIX applications to z/OS.
- The C/C++ Compiler includes support for IEEE Floating Point, and 64-bit long integer format. Support is also added to the C/C++ Runtime Library.
- Performance and serviceability are improved, with a single optimization level, BCD format enhancements, and allocation for stack storage.
- Class library header in HFS.
- Extra Performance Linkage (XPLINK)

When you run a C or C++ application, there is overhead associated with each function call. The more highly functionalized a program, the more overhead. XPLINK cuts down on the overhead associated with

these function calls and improves the performance of these applications. In order to exploit the bulk of “high performance linkage” customers must recompile their C and C++ programs under the new XPLINK environment.

## **Network Communication Services**

### **z/OS Version 1 Release 2 base elements**

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

### **z/OS Version 1 Release 2 optional no-charge features**

z/OS Communications Server Security Level 3<sup>1</sup>  
z/OS Communications Server Network Print Facility

*1. Export considerations*

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

*Highlights:*

- HiperSockets is a new function introduced in z/OS V1R2 for ^ zSeries which provides very high speed, low latency TCP/IP data communications across LPARs resident within the same z900 server. HiperSockets acts like a TCP/IP network within the server.
- HiperSockets Accelerator, provides an “accelerated routing path” which concentrates traffic between OSA-Express external network connections and HiperSockets connected LPARs. This function can improve performance, simplify configuration, and increase scalability while lowering cost by reducing the number of networking adapters and associated I/O cage slots required for large numbers of virtual servers.
- The Sysplex Distributor has been extended to control workload balancing, working in conjunction with Cisco MultiNode Load Balancing (MNLB).
- Host-based Intrusion Detection Services (IDS) complements network-based IDS sensors and scanners by providing defense mechanisms that discard attacking packets before they cause damage, discard packets exceeding established thresholds, and limit the number of connections from “greedy” users.
- OSA-Express Queued Direct Input/Output support
- LPAR-to-LPAR communication sharing a single OSA-Express ethernet or Token Ring adapter

- Integration of SNA/APPN, TCP/IP and AnyNet®
- High performance TCP/IP stack for all applications
- High Performance Native Sockets (HPNS) for TCP/IP applications
- High Performance Data Transfer for SNA applications
- Multinode Persistent Sessions for SNA applications running in a Parallel Sysplex environment
- SNA 3270, Sockets, and APPC application support
- High Performance Routing enhanced to additional S/390 network configurations
- Native ATM support for high speed networking
- Easy access to host applications from Java-enabled Web browser
- Greater performance and mobility for CICS Sockets

*Benefits can include:*

- Simplified deployment of client/server applications
- Function for new e-business Internet and intranet applications
- Multivendor, multiplatform connectivity
- High performance, high availability, network choice
- SNA class of service over IP networks
- World class TCP/IP services
- Dramatic improvements in TCP/IP performance include optimization of the TCP/IP stack, and inclusion of a number of performance related capabilities such as:
  - UNIX Sendmail is supported including the POP3 server.
  - TN3270 sessions use Secure Sockets Layer, with increased number of supported ports (255).
  - Multi-Node Persistent Sessions (MNPS) includes recovery for Dependent LU Requester (DLUR)-owned LUs. This leads to full Parallel Sysplex exploitation with native IP networks, including nondisruptive session-switching.
  - Client/Server Affinity improvements to allow recovery over sync points in multiphase commit giving higher availability for these applications.
  - Simple Network Management Protocol (SNMP) Distributed Protocol Interface (DPI) instance level registration.
  - X-Motif uses Data Link Libraries (DLL) to reduce disk space.
  - Virtual Private Network (VPN) easier to establish.

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

- Enhanced Multipath Load Balancing.
- Support for dynamic updates to Service Policy.
- Higher availability in Parallel Sysplex environments with Virtual Addressing Takeover (VIPA).
- Support of Internet Key Exchange (IKE) protocol to automatically create and distribute encryption keys for dynamic IP clients.
- Addition of Secure Sockets Layer (SLL) authentication to the TN3270e server, to protect against unauthorized access to SNA applications from TCP/IP clients.
- Support of Triple DES session encryption for SNA applications
- Service Policy Agent for IBM Communications Server Enhancements

## **UNIX System Services**

### **z/OS Version 1 Release 2 base elements**

---

z/OS UNIX System Services Application Services

---

z/OS contains the UNIX Application Services (shell, utilities and debugger) and the UNIX System Services (kernel and runtime environment). The Application Services Shell and Utilities provides 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. UNIX System Services provide the world of open UNIX-based computing to the

z/OS operating system. With Language Environment, they support industry standards for C programming, shell and utilities, client/server applications, and the majority of the standards for thread management and the X/Open Single UNIX Specification. The combination of open computing and z/OS allows the transparent exchange of data, easy portability of applications, cross-network management of data and applications, and the exploitation of traditional MVS system strengths in an open environment.

*Highlights:*

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

*Benefits can include:*

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

UNIX

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

*UNIX Parallel Environment*

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

*USS*

- Superuser controls allow for selective assignment of security by specific responsibility.
- More control for message queues.
- Make porting from other platforms easier with support of the Magic Number (#!), double-square-bracket conditional testing, and other functions.
- Parallel Sysplex support to share UNIX file systems benefits 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.

## Distributed Computing Services

### ***z/OS Version 1 Release 2 Base elements***

Network File System (NFS)  
DCE Base Services (OSF DCE Level 1.1)  
Distributed File Service (OSF DCE level 1.2.2)

- 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 to appear as local directories and files on the client system. NFS also provides access to the Hierarchical File System (HFS).
- DCE enables data encryption using the data encryption standard (DES) algorithm and the commercial data masking facility (CDMF).

*Highlights:*

- z/OS R2 introduces the zSeries File System (zFS). zFS is a UNIX file system that can be accessed with the z/OS hierarchical file system file Application Programming Interfaces (APIs). It will provide significant performance gains in most customer environments requiring files 8K in size or greater.
- Remote Procedure Call (RPC) lets calls between programs running on different platforms appear as local procedure calls.

- Directory Services allows resources to be found anywhere in an enterprise without the need to know local names.
- Security Services helps solve security problems common in a distributed environment by handling identification and certification of users, clients, servers and systems.
- Distributed Time Services synchronizes clocks running on different nodes
- DFS allows access to both z/OS UNIX Services Hierarchical File System and the OSF's DCE Local File System.
- DCE DFS brings the strengths of z/OS systems management software and high-speed unique file naming across enterprises.
- Client support for record-level access to PS, PDS, PDS/E and VSAM data sets.
- Security based on Kerberos authentication with access to RACF data as well.

*Benefits can include:*

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

#### *DFS*

- Windows networking compatible file and print serving is available with Native SMB File and Print Serving for Windows Clients. SMB file serving enables z/OS to share HFS files with Windows workstations. This allows development of applications on Windows, for deployment on z/OS. z/OS also supports printing of SMB files without requiring that code be installed on the clients and without requiring unique printer setup on the workstations. OS/2® client file serving support is also available.

## **e-business Services**

### **z/OS Version 1 Release 2 Base elements**

IBM HTTP Server  
Text Search

### **z/OS Version 1 Release 2 Optional no charge features**

IBM HTTP Server North America Secure

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

Text Search is a powerful, full-text indexing and search server. It supports high-speed searching of Web sites, as well as 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 allows the server to respond to requests with the version of a Web page or document appropriate for that browser.
- Page counter and date/time information can be displayed on a page as graphical images.
- Dramatic improvements in single and multiple processor environments coupled with significant CPU requirement reductions provide higher throughput and shorter response time.
- Platform for Internet Content Selection (PICS) provides a way for users to filter material they encounter on the Internet and accept or reject the material based on its rating.
- Web usage statistics.

*Benefits can include:*

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

## **Print Services**

### **z/OS Version 1 Release 2 optional priced**

#### **Infoprint Server**

IP Printway™  
NetSpool™  
z/OS Print Interface

Today, businesses are looking for better ways to get information to knowledge 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, VTAM® application output capture program, and print drivers and management tools that let you manage any print job to any printer defined to Infoprint Server. These capabilities give you the flexibility to deliver the output of e-business anywhere you need it:

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

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

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

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

## **Test**

To further reduce migration time, the z/OS product is system integration tested using a production-like environment. This environment includes subsystems, such as CICS, IMS and DB2. 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 function.

## **Publications**

For a list of the publications available for z/OS, visit the z/OS library Web site.

## Installation Considerations

OS/390 V2 R10 provides a strong foundation for z/OS. You can install OS/390 V2R10 now and upgrade to z/OS V1R1 with the installation of the special z/OS V1R1 Upgrade Package for OS/390 V2R10. This product upgrade package (PUP) will provide an easy path to z/OS V1R1, and will install like a monthly RSU service upgrade. Non-OS/390 V2R10 customers, or customers who would like to install z/OS V1R2, can order z/OS via Customized Offerings (ServerPac, Systempac® and CBPDO) for a complete system.

Customers are strongly encouraged to implement WLM goal mode. WLM in goal mode continues to expand its role in providing enhanced system efficiencies and workload balancing on the S/390 and zSeries platforms. Goal mode is required for much of the Intelligent Resource Director functions. z/OS V1R2 will be the last release to support WLM compatibility mode.

### z/OS V1R2 is supported on the following IBM servers:

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

*SOD: z/OS V1R3 and z/OS V1R4, expected to be available in 2H2002, will run on G5, G6, MP3000 and z900 Servers (or compatible servers) (31-bit mode).*

## Coexistence

z/OS continues to give customers optimum compatibility and flexibility as they migrate systems in a multisystem configuration by allowing up to four consecutive releases to coexist. Coexistence considerations apply to multisystem configurations in which there is resource sharing. This includes non-Parallel Sysplex and Parallel Sysplex multisystem configurations.

The following table describes the effect of this special provision. The release shown in column 1 is the highest release running in a multisystem configuration.

**Table: Coexistence Supported Releases (OS/390 & z/OS)**

Release	Coexistence Release
R6	R6, R5, R4, R3, R2 <sup>(1)</sup>
R7	R7, R6, R5, R4
R8	R8, R7, R6, R5
R9	R9, R8, R7, R6
R10	R10, R9, R8, R7, R6 <sup>(2)</sup>
z/OS R1	z/OS R1 OS/390 R10 <sup>(3)</sup> , R9, R8, R7, R6 <sup>(2)</sup>
z/OS R2	z/OS R2, z/OS R1, OS/390 R10 <sup>(3)</sup> , R9, R8

### Notes:

1. Coexistence of OS/390 R2 with OS/390 R6 supported as a special provision
2. Coexistence of OS/390 R6 with OS/390 R10 supported as a special provision
3. 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 ensure appropriate positioning for future software migrations. Since each release can normally be ordered for only a six month window, it is very important that you order the required releases while they are available.

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

For additional information on the general z/OS policy, in addition to the specific JES releases that can coexist with z/OS Release 2 JES in a multisystem configuration, refer to *z/OS Planning for Installation*.

*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.

z/VM is the follow-on product for VM/ESA customers. It opens new and exciting opportunities for customers who have built enterprise-wide automation and infrastructure enhancements on the VM platform in support of their applications, database systems, and e-business solutions.

Using virtualization technology as a foundation, z/VM offers new function and technology for customers who want to exploit IBM Virtual Machine (VM) capabilities on the mainframe. Virtualization technology allows customers to virtualize processor, communications, storage and I/O resources thus avoiding the overhead of planning, purchasing and installing new hardware to support new workloads.

z/VM supports a wide variety of existing IBM zSeries and S/390 servers. This offers today's VM/ESA customer the ability to maintain their existing VM environments on the latest technology while positioning themselves to grow into the zSeries.

z/VM offers an ideal platform for consolidating Linux workloads onto a single zSeries 900 or S/390 server. Running Linux as a guest of z/VM enables you to run hundreds of Linux images while benefiting from the reliability, scalability, availability, and serviceability characteristics of zSeries 900 or S/390 servers.

#### *z/VM Version 3*

z/VM offers a new technology base for customers looking to use IBM Virtual Machine technology on the industry's best of breed server platform, zSeries 900. When z/VM is running on a zSeries server, it's possible to run 64-bit capable OS/390 Release 10, z/OS, and Linux for zSeries as guest systems of z/VM, in addition to ESA/390 guest operating systems such as OS/390, VSE/ESA, TPF, and Linux for S/390. To operate z/OS as a guest of z/VM on a zSeries server, z/VM and z/OS 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 reduces the need to invest in separate standalone configurations.

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

Additional enhancements in z/VM Version 3 include:

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

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

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

#### *z/VM Version 4*

Licensed using the IBM International Program License Agreement (IPLA) which includes a one-time charge (OTC) pricing structure, z/VM Version 4 provides a great opportunity for customers to exploit zSeries 900 or S/390 servers at a substantially lower cost. z/VM Version 4 is priced on a per-engine basis and is designed to support the IBM Integrated Facility for Linux (IFL) processor features for Linux-based workloads and standard engines for all other zSeries 900 and S/390

workloads (an engine is also referred to as a central processor). z/VM Version 4 supports the IBM zSeries 900, S/390 Parallel Enterprise Server™ Generation 5 and 6, and the S/390 Multiprise 3000, or equivalent.

To help improve the performance of I/O to network devices, the z/VM Control Program (CP) provides a fast CCW translator in z/VM V4R1 to translate a wide range of channel programs that perform I/O to network adapters. This fast translator for network devices includes support for 64-bit indirect data address words (IDAWs), which allows guest machines to read and write data above the 2 GB limit using 64-bit I/O addressing. In z/VM V4R2, 64-bit IDAW support was further extended to the existing DASD fast CCW translator. This enhancement is also designed to include 64-bit IDAW support for DASD channel programs that are simulated via the minidisk cache (MDC). Additional guest-support improvements in z/VM V4R2 that may benefit Linux virtual machine performance include DASD fast CCW translation using 64-bit IDAWs and (when corresponding function is available from Linux for zSeries) enhanced page-fault handling which allows a Linux guest to continue processing while the page faults are handled by z/VM.

When corresponding function is available from Linux for zSeries, z/VM V4R2 is also designed to provide the capability for Linux guest virtual machines to use clear-key RSA functions of the IBM PCI Cryptographic Coprocessor (PCICC) or the IBM PCI Cryptographic Accelerator (PCICA) for Secure Socket Layer (SSL) acceleration on the zSeries 900 S/390 servers. z/VM V4R2 offers some ease-of-use functions to manage Linux images and a migration facility to help existing customers of the Virtual Image Facility move their configuration files and data to a z/VM V4R2 system.

z/VM V4R2 extends its hardware support by exploiting technological enhancements for the IBM ^ zSeries 900 (z900) server, such as:

- faster communications between and among virtual machines and logical partitions with HiperSockets
- high-speed OSA-Express Token-Ring
- improved z/OS Parallel Sysplex support with virtualized Coupling Facility (CF) duplexing
- FICON channel-to-channel adapter (CTCA) communications for guests

Connectivity enhancements to TCP/IP for z/VM, Level 420 are designed to include a new Internet Message Access Protocol (IMAP) server, improved security of the TCP/IP stack by preventing some Denial of Service (DOS) attacks, and support for HiperSockets. A “Guest LAN” capability is designed to allow a VM guest to install a virtual HiperSockets adapter and connect it with other virtual network adapters on the same VM host system to form an emulated LAN segment.

To aid application development within the Language Environment (LE), z/VM V4R2 consolidates existing VM C-language sockets APIs into a single consistent socket library.

RealTime Monitor (RTM) and Performance Reporting Facility (PRF) have been enhanced to support z/VM, including new and changed monitor records in z/VM Version 3. 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. RTM and PRF are priced, optional features of z/VM Version 4 as is the Directory Maintenance Facility (DirMaint).

For further information see the z/VM and VM/ESA Reference Guide, G320-9997-17.



## Notes

## To learn more

Visit the S/390 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