IBM

IBM @server zSeries 890 Technical Introduction

- Hardware description and software support
- Planning and migration considerations
- Frequently asked questions

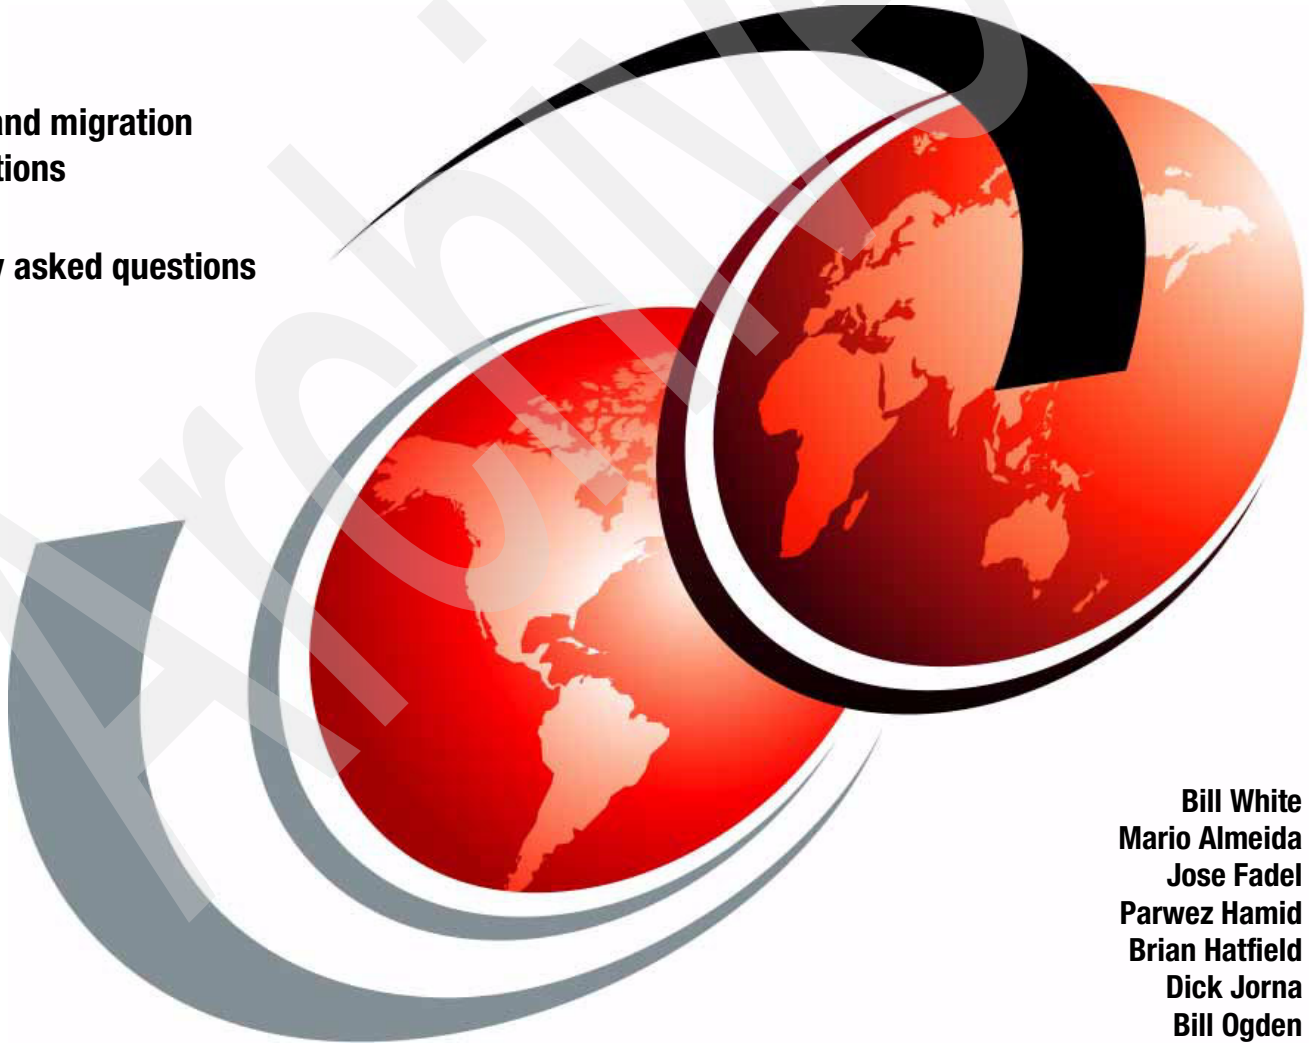

Bill White
Mario Almeida
Jose Fadel
Parwez Hamid
Brian Hatfield
Dick Jorna
Bill Ogden

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM @server zSeries 890 Technical Introduction

May 2004

SG24-6310-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (May 2004)

This edition applies to IBM eServer zSeries 890 at hardware Driver Level 55.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

CICS®
DB2®
Enterprise Storage Server®
ESCON®
IBM @server
IBM @server
FICON®
GDPS®
HiperSockets™
IBM®
ibm.com®
IMS™
Multiprise®
MVS™
OS/2®
OS/390®
Parallel Sysplex®
PR/SM™
Redbooks (logo) ™
Redbooks™
Resource Link™
RMF™
S/360™
S/370™
S/390®
Sysplex Timer®
ThinkPad®
Tivoli®
TotalStorage®
VM/ESA®
VSE/ESA™
VTAM®
Wave®
WebSphere®
z/Architecture™
z/OS®
z/VM®
zSeries®

The following terms are trademarks of other companies:

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

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

Other company, product, and service names may be trademarks or service marks of others.

Preface

This IBM® Redbook introduces the IBM @server zSeries® 890, which represents the continuation of the scalable servers featured with the IBM eServer zSeries 990. The z890 is based on z/Architecture™, the zSeries building blocks of the z990, and the virtualization technology of passed sever families. It is designed to be resilient in the unpredictable on demand world.

The z890 is a single model server with a wide range of capacity settings, delivering significantly improved granularity and enriched functions over its predecessor. At the same time, the z890 is also introducing the new eServer zSeries Application Assist Processor (zAAP), which provides a Java™ execution environment.

This publication provides information on the hardware and software features available with the z890. It also includes planning and migration considerations.

This technical introduction is intended for hardware planners, system engineers, and consultants that need to understand the capabilities of the z890.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Bill White is a Project Leader and Senior Networking Specialist at the International Technical Support Organization, Poughkeepsie Center.

Mario Almeida is a Certified Consulting IT Specialist in Brazil. He has 29 years of experience in IBM Large Systems. His areas of expertise include zSeries and S/390® servers technical support, large systems design, data center and backup site design and configuration, and FICON™ channels.

Jose Fadel is a zSeries System Architect from Brazil. He has 27 years of experience with IBM Large Systems. Most of the 14 years with IBM, he has worked as a pre-sales technical support specialist for zSeries. His areas of expertise include zSeries hardware and software.

Parwez Hamid is a Consulting IT Specialist from the United Kingdom. He has 30 years of experience in IBM Large Systems. During this time he has worked in various roles within IBM, and currently provides pre-sales technical support for the IBM zSeries product portfolio. In addition he has written and produced technical material for all IBM S/390 and zSeries severs since G3. Parwez has also written and contributed to a number of IBM Redbooks™.

Brian Hatfield is a Senior Education Specialist in the US. He has 26 years of experience in IBM Large Systems, ranging from customer engineer to instructor/course developer. He is currently a course designer, developer and instructor, responsible for several technical education offerings in the z/OS® curriculum, including Parallel Sysplex®, zSeries Architecture and FICON environment courses addressing implementation, operations, recovery and availability concerns. He has presented at several IBM conferences and contributed to other zSeries Redbooks.

Dick Jorna is a Certified Senior Consulting IT Specialist in the Netherlands. He has 35 years of experience in IBM Large Systems. During this time he has worked in various roles within IBM, and currently provides pre-sales technical support for the IBM zSeries product portfolio. In addition he is zSeries product manager, and is responsible for all zSeries activities in his country.

Bill Ogden is a retired Senior Technical Staff member at the International Technical Support Center, Poughkeepsie, who has worked with S/360™, S/370™, S/390, and zSeries machines since 1966. He is currently working remotely from the Navajo reservation in Arizona.

Thanks to the following people for their invaluable contributions to this project:

Franck Injey and Julie Czubik
International Technical Support Organization, Poughkeepsie Center

Darelle Gent
zSeries Hardware Product Planning, IBM Poughkeepsie

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an Internet note to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYJ Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

Introduction

The IBM eServer zSeries 890 represents the continuation of the new generation of scalable servers introduced with the IBM eServer zSeries 990. Using z/Architecture and zSeries latest building blocks and virtualization technology, the z890 extends zSeries key platform characteristics of reliability, availability, scalability, clustering, and quality of service to respond to the ever changing business climate with a processor that delivers extensive growth options and excellent price/performance for those customers requiring a lower capacity entry point than offered with the z990.

The z890 is designed to help enable your businesses to be resilient in the unpredictable on demand world. With a single model and a wide range of capacity settings, the newest member of the zSeries family delivers significantly improved granularity and enriched functions over its predecessor, the z800; it can provide up to twice the processing power and significantly increase I/O capacity. The z890 can support double the number of logical partitions (LPs), and the Logical Channel SubSystem (LCSS) feature introduced on z990 facilitates horizontal growth. The z890 is also introducing the new eServer zSeries Application Assist Processor (zAAP), which reduces overall capacity that is needed for Java workloads.

1.1 Evolution

The z890 server is, in many ways, a smaller version of the z990 server—both are members of the z/Architecture family that extended the S/390 series of servers. Both machines, the z990 and the z890, continue the evolution of the original S/390 design. Two fundamental characteristics of this evolution have been:

- *Compatible evolution*. IBM recognizes the total customer investment in data processing solutions and recognizes that application software (and the non-DP processes that are intertwined with these applications) holds the bulk of this investment. All of the evolution of the system architecture has been planned with application compatibility in mind.
- *Resolving constraints*. Evolving hardware capabilities and evolving application designs produce various stress points in system designs. Over the years these stress points have included CPU performance, usable memory size, addressable memory size, I/O speed, I/O management overhead, I/O attachment capacity, and so forth. System architecture has evolved to relieve constraints in many areas.

Examples of this evolution include:

- Single processor systems evolved to multiple processor systems.
- Real memory systems evolved to virtual memory systems.
- Memory addressing (real and virtual) evolved from 24 bit to 31 bit to 64 bit addressing.
- I/O channels evolved from 16 channels to 256 channels.
- A single operating system in the machine has evolved to Logical Partitions providing up to 15 system images in one physical system.
- I/O management evolved from complex operating system functionality to System Assistance Processors (SAPs) that handle much of the work.
- I/O performance has evolved from parallel channels to ESCON® channels to FICON channels.
- Communication between similar systems has evolved from simple shared DASD to interconnected channels (CTCs) to Coupling Facilities and Parallel Sysplex.
- Advanced I/O interfaces have evolved from intricately programmed control units (such as the IBM 2701, 2702, 2703 products) to front-end processors (such as the IBM 3745) to OSA Express designs.
- I/O programming has evolved from the basic SIO-type instructions of the S/360 to the SSCH-related instructions and subsystem of the S/390, and to the QDIO model available with the z/Series.
- System monitoring and tracing has grown from simple sampling routines (with high overhead) to hardware-assisted tracing instructions and to hardware-assisted channel and device performance measurements.
- Instructions for small program models (with strict 12-bit displacements) have been expanded through the relative and immediate instructions added to the architecture a few years ago. These instructions, intended primarily for compilers, provide more efficient implementations of larger programs.
- Processor workloads have been partly offloaded into segregated processors, such as IFLs and ICFs.
- Certain processing has been made more efficient by the use of “microcoded” assists; the best-known of these is the SIE instruction for VM assistance.

z890 and z990 enhancements

The z890 and z990 servers provide additional major evolutionary steps. Key architectural limits of older systems are addressed without significant impact to the application

compatibility that is so important to IBM's customers. Among other elements, the z890 and z990 servers address:

- Total system processing speed. The largest z890 has roughly twice the processing capacity of the largest z800 server. (The z990 has several times the processing capacity of the largest z900 server.)
- Total system memory. A z890 can have 8 GB to 32 GB memory. (A z990 can have 16 GB to 256 GB memory.)
- Total number of I/O channels. A z890 may have up to 420 channels.[1] (A z990 may have up to 1024 channels.)
- Total number of CHPIDs. A z890 can have up to 512 CHPIDs. (A z990 can have up to 1024 CHPIDs.)
- Number of system images. A z890 may have up to 30 logical partitions.
- Processing offload. A z890 may have segregated engines for Java processing. (This new feature will also be available for z990 servers.) This new engine type is the eServer zSeries Application Assist Processor, commonly known as zAAP.
- Enhanced concurrent maintenance and upgrade capabilities are available with both the z890 and z990.
- Additional channels and device measurements are available with both servers.
- More efficient large program models are possible on both servers, using long-displacement instructions.
- More efficient LAN operation is possible on both servers, using larger data blocks and offloading more processing from the main processor to the I/O features.
- I/O connections for operating system consoles (and other local 3270 terminals) have been simplified by new OSA Express 3270 features. (This feature will also be available for z990 servers.)
- More performance assists are available for z/VM® and Linux.

z800 comparison

Compared to the z800, the new z890 has the capability of providing the following when properly configured:

- Increased scalability—up to 123 percent more total capacity of the largest z800
- A lower capacity entry point into the family with 32 percent less capacity than the smallest z800
- Twenty-eight levels of capacity offering enhanced incremental growth options
- Double the number of supported CHPIDs
- Double the number of logical partitions (except on the smallest sub uniprocessor equipped capacity setting)
- Quadruple the number of HiperSockets (internal LANs)
- Performance assists for Linux and z/VM
- OSA-Express enhancements including the new Integrated Console Controller
- FICON Express, Integrated Cluster Bus (ICB), InterSystem Channel-3 (ISC-3), and OSA-Express spanned channels
- Increased channel maximums for ESCON, FICON Express, ISC-3, and OSA-Express (except on the smallest sub uniprocessor equipped capacity setting)
- On/Off Capacity Upgrade on Demand (On/Off CoD)

[1] The architecture, with two channel subsystems, can address up to 512 channels. The limitation of 420 is related to the number of I/O card slots and the maximum number of channel ports on an I/O card. The 420 limit is reached with 28 ESCON cards having 15 usable ports each. If we include an ICB-4 connection (which is possible with the free STI port in the system), we could say the maximum is 421 active channels.

- Fiber optic cabling services from IBM Networking Services
- zSeries industry-leading SSL performance
- zSeries Application Assist Processor (zAAP)
- Internal Battery Feature (IBF)
- GDPS/PPRC cross site Parallel Sysplex distance now up to 100 km
- Up to 48 ISC-3 links in peer mode, CFCC level 13, CFCC enhanced patch apply

Multiprise® and 9672 comparison

Compared to the Multiprise 3000 and the 9672 servers, the z890 differs in a number of ways. Below is a list of the key differences in function and features. These differences are in addition to the list for the z800 comparison:

- z890 has considerably more capacity and I/O bandwidth compared to the Multiprise 3000 and the pre-G5/G6 servers.
- FICON is not available on Multiprise 3000 or pre-G5/G6 servers.
- Multiprise 3000 supports Cryptographic Coprocessors only. PCICA/PCICC support on recent servers and PCIXCC are only on the z890 and z990.
- Multiprise 3000 has integrated DASD adapter(s), integrated LAN, integrated tape, and an integrated TN3270 server. z890 uses external disks and OSA-Express. External devices allow for larger capacity disks and OSA-Express provides more bandwidth for the LAN environment. With internal disks, there is no shared DASD with other systems.
- z890 has up to eight times more ESCON connectivity compared to a maximum of 56 channels on the Multiprise 3000.
- Practical LAN bandwidth of Multiprise 3000 is much less than that of the OSA-Express features.
- Multiprise 3000 has no external CF links (ICB, ISC) for external CF connection.
- Very limited concurrent maintenance on Multiprise 3000.
- Multiprise 3000 cannot switch an application from a failing CP to another CP via hardware scanouts (for example, no Application Protection Function is available).

1.2 z890 highlights

As mentioned, the z890 is a member of the zSeries family and uses the z/Architecture and instruction set (with some extensions) of the z900 and z990 machines. (This architecture, formerly known as ESAME Architecture, is commonly known as *64-bit* architecture, although it provides much more than 64-bit capability.)

The z890 is offered as a single hardware model—z890 Model A04 and has a machine type of 2086. The model naming is representative of the single-book design and indicates the number of processor units (PUs). Unlike prior processor model names, which indicate the number of purchased PUs, the z890 model number indicates the maximum number of processor units potentially orderable, and not the actual number that has been ordered.

A z890 server has one frame that is the same size as one z990 frame. (All z990 servers have two frames.) The frame is air cooled, with a new air management design. The frame normally uses three-phase power, but can optionally be used with single phase power. An Internal Battery Feature (IBF) is also provided.

In addition, the frame contains two ThinkPads that are used as Support Elements, and one I/O cage with 28 I/O slots (see Figure 1-1 on page 5). Some I/O cards from a z800 or z900 can be moved forward to the z890.

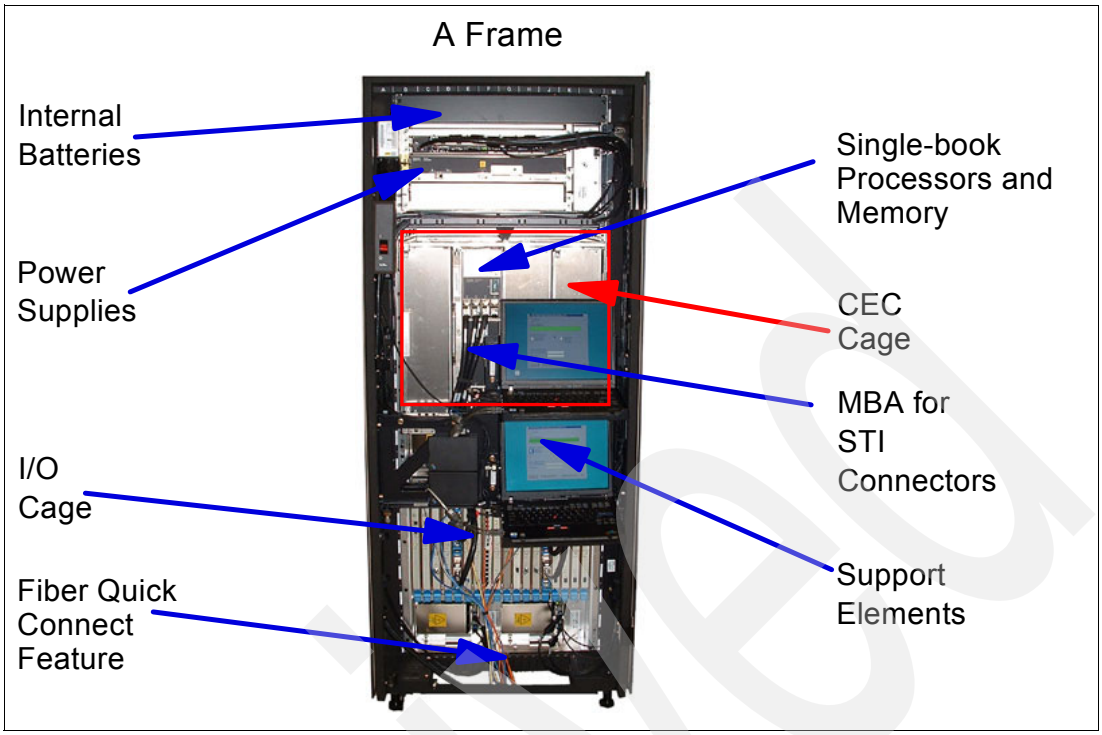

Figure 1-1 z890 - Under the covers - Front view

A more extensive description of the frame appears in Chapter 2, “Technical details” on page 13.

The z990 introduced the concept of a *book* within a zSeries machine. A z890 has a single book, while a z990 can have 1 to 4 books. A book contains processor units, memory, and STI connectors providing connectivity to the I/O cage and for ICB coupling channels.

The z990 has additional attributes associated with multiple books. These are described in the *IBM eServer zSeries 990 Technical Guide*, SG24-6947.

The z890 book has from 8 GB to 32 GB of memory and five Processor Units (PUs) with 28 unique capacity settings, increasing granularity and providing enhanced incremental growth options. Up to four PUs can be characterized or spare PUs, depending on the z890 features ordered. Note that the z890 does not have any pre-assigned PUs as spares and one PU must be configured as a System Assist Processor (SAP). The z890 also introduces the zSeries Application Assist Processor (zAAP), enabling movement of Java processing cycles to a lower cost zSeries operating environment.

A PU is the generic term for the z/Architecture processor on the MultiChip Module (MCM) that can be characterized as:

- Central Processor (CP) to be used by the operating system
- Internal Coupling Facility (ICF) to be used by the Coupling Facility Control Code (CFCC)
- Integrated Facility for Linux (IFL) to be used by z/VM and/or Linux on zSeries
- zSeries Application Assist Processor (zAAP) for Java processing
- A standard System Assist Processor (SAP) to be used by the Channel SubSystem (CSS)

Note: The zAAP requires z/OS V1.6 or z/OS.e V1.6.

The z890 book supports up to 16 GB/sec of bandwidth for data communication between I/O and memory through up to eight Self-Timed Interconnect (STI) host buses.

A z890 server runs only in *LPAR mode*; *Basic mode* is not supported. Up to 30 logical partitions (LPs) may be defined to the z890.[2] There are new rules associated with LP definitions, and a new type of definition for one or two Logical Channel SubSystems (LCSSs) is required. This is covered in more detail in “Channel subsystem” on page 42. Briefly, the key rules include:

- Every LP must be associated with a LCSS. Multiple LPs (up to 15) can be associated with a single LCSS.
- A LCSS can define and use a maximum of 256 channels (CHPIDs).
- Some channel types, such as ESCON, can be associated with only a single channel subsystem. Other channel types, such as FICON and OSA-Express, can be associated with multiple channel subsystems.
- LCSS definitions are part of the overall IOCDS definition. These definitions may be done through HCD or through direct IOCP statements in the Support Element.

The z890 offers an entry-level system (see “Smallest sub-uniprocessor - Capacity Setting Model 110” on page 7) that has reduced limits for some functions. These reduced functions include the following:

- The I/O cage is limited to 16 slots (as opposed to 28 slots in other z890 systems).
- There is a maximum of 15 Logical Partitions (as opposed to 30 LPs in other z890 systems).
- There is a maximum of 240 ESCON channels (as opposed to 420 ESCON channels in other z890 servers.)
- There is a maximum of 32 FICON channels (as opposed to 40 FICON channels in other z890 severs.)
- There is a maximum of 24 OSA-Express ports (as opposed to 40 OSA-Express ports in other z890 servers.)

Important: These reduced functions for the smallest sub-uniprocessor may not be generally noted in all tables and discussions in this publication. However, you should keep these reduced capabilities in mind while using this book.

The z890 also introduces OSA-Express enhancements including the new OSA-Integrated Console Controller (OSA-ICC). The OSA-ICC feature uses the OSA-Express 1000BASE-T to provide TN3270E and non-SNA DFT 3270 emulation. This enhancement can eliminate the need for console controllers such as 2074s and 3174s.

1.3 Capacity Settings and MSUs

With the z890, you can start at the capacity setting you need to meet your IT infrastructure requirements, then easily add additional capacity to meet your changing needs as you grow. Each of the four processors on the z890 can be divided into seven sub-units. This creates a 7 by 4 matrix of settings. As long as upgrades are positive MIPS growth, you can move around anywhere within the matrix when adding capacity. The two key things to remember about the matrix is that (1) the multiple engine, general purpose engines cannot 'mix' subcapacity settings in one machine and (2) upgrades can only be done disruptively (either hardware or software) where there is no change in engine size, that is, staying on the same 'horizontal'

[2] The smallest sub-uniprocessor z890 is limited to 15 logical partitions.

row. Upgrades where engine size changes are non disruptive from hardware, but may require a software IPL.

The z890 Model A04 offers 28 capacity settings. Each capacity settings will have a 4-digit associated CP feature code, as follows:

- Position one = fixed value of 6 (can be ignored)
- Position two = number of CPs (1 to 4)
- Position three = capacity setting (1 to 7, one being the smallest subcapacity setting and seven being a full processor)
- Position four = engine identifier (0=CP, 1=On/Off CoD Use Day, 2=downgrade record)

Each capacity setting is based on its related z890 CP feature code ('6xxx'), using the CP engine identifier (last position = 0) and removing the first position ('6').

Using the above information, Table 1-1 shows the capacity settings for the z890.

Table 1-1 z890 capacity settings

1-WAY	MSUs	2-WAY	MSUs	3-WAY	MSUs	4-WAY	MSUs
110	4	210	8	310	11	410	15
120	7	220	13	320	20	420	26
130	13	230	26	330	38	430	49
140	17	240	32	340	47	440	62
150	26	250	50	350	74	450	97
160	32	260	62	360	91	460	119
170 Full 1-way	56	270 Full 2-way	107	370 Full 3-way	158	470 Full 4-way	208

See more information in "z890 capacity settings" on page 28.

Smallest sub-uniprocessor - Capacity Setting Model 110

The entry model z890 has capacity setting 110. The 110 has most of the same features as the other 27 capacity settings; however, it has smaller I/O configuration/expansion capabilities. The 110 will only be allowed to run 15 logical partitions; it can have a maximum of 240 ESCON channels, 32 FICON Express, and 24 OSA-Express. The 110 is also different from the other models because it offers zELC pricing. This zELC price will enable entry level customers to get many of the technical advantages of the largest zSeries customers at a very attractive price.

Model upgrade paths

The z890 have comprehensive options for upgrades within the range providing full flexibility and choice of four full capacity CPs, each with seven capacity settings giving up to 28 granular models. As shown in Figure 1-2 on page 8, upgrades can be horizontal, vertical, diagonal or whichever way best fits your capacity needs as long as the upgraded systems result in an increased capacity.

One key point to remember is that in multi-CP systems all engines have to be of the same size capacity. For example, full capacity CPs cannot be mixed with CPs of lesser capacity.

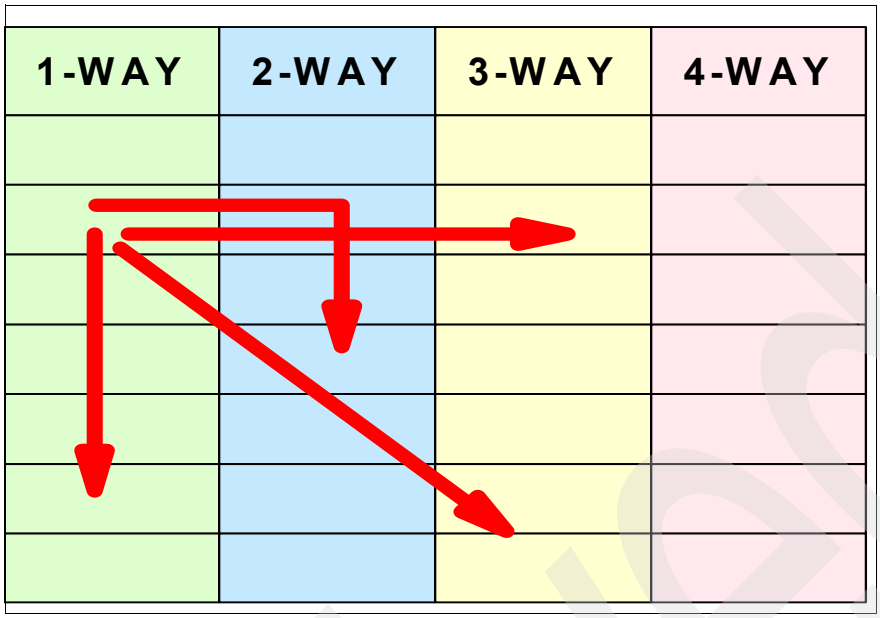

Figure 1-2 z890 model granularity

Existing z800 can be upgraded to the z890. Figure 1-3 shows the models that can be upgraded to a z890, as well as from the various z890 capacity settings to a z990 Model A08.

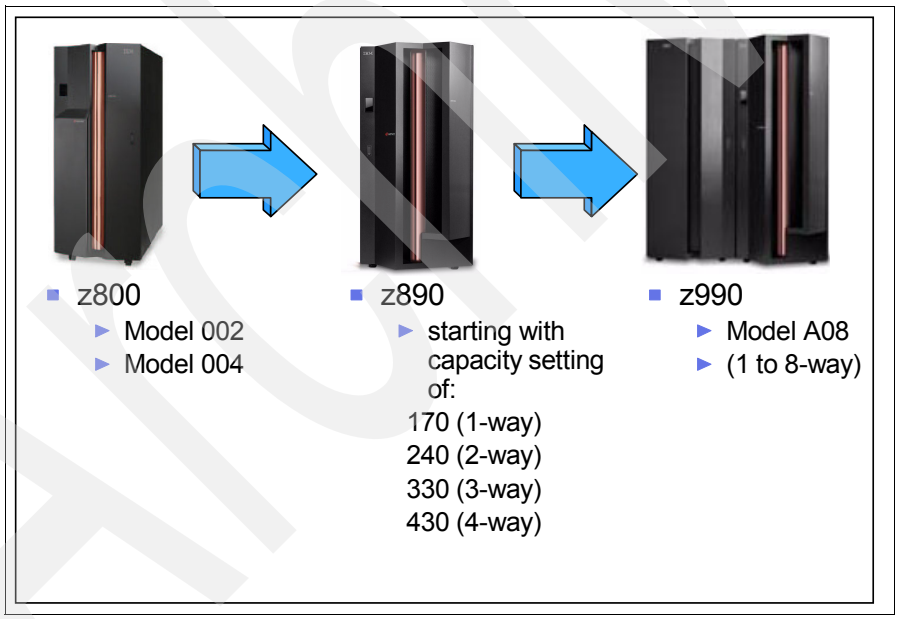

Figure 1-3 z890 model upgrade paths from z800

Model downgrade paths

You are allowed to downgrade your machine via CIU, CUoD, or MES. The primary benefit to downgrading is a reduction in software charges based on a lower reported machine capacity. The downgrade can be vertical (capacity setting), horizontal (number of CPs), or diagonal (capacity setting and the number of CPs).

1.4 Performance comparison

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

The performance design of the z/Architecture enables the entire server to support a new standard of performance for all applications through expanding upon a balanced system approach. As CMOS technology has been enhanced to support not only additional processing power, but also more engines. The entire server is modified to support the increase in processing power. The I/O subsystem supports a great amount of bandwidth through internal changes, thus providing for larger and quicker data movement into and out of the server. Support of larger amounts of data within the server required improved management of storage configurations made available through integration of the software operating system and hardware support of 64-bit addressing. The combined balanced system effect allows for increases in performance across a broad spectrum of work. However, due to the wide range of performance levels, z890 offers from 1–4 Central Processors, each with 7 different "dial points" (28 choices in all) and resource management within the system. It is expected that there will be a larger performance variability than has been previously seen by our traditional customer set. The z890 supports an estimated performance range of 0.14 to 7.40 compared to a z800 Model 001. This variability may be observed in several ways. The range of performance ratings across the individual LSPR workloads is likely to have a larger spread than past processors. There will also be more performance variation of individual logical partitions as the impact of fluctuating resource requirements of other partitions can be more pronounced with the increased number of partitions available on the z890. The customer impact of this increased variability will be seen as increased deviations of workloads from single-number-metric based factors such as MIPS, MSUs and CPU time chargeback algorithms. It is important to realize the z890 has been optimized to run many workloads at high utilization rates.

It is also important to notice that the LSPR workloads for z890 and z990 have been updated to reflect more closely our customers' current and growth workloads. The traditional TSO LSPR workload is replaced by a new, heavy Java technology-based online workload referred to as Trade2-EJB (a stock trading application). The traditional CICS®/DB2® LSPR online workload has been updated to have a Web-frontend, which then connects to CICS. This updated workload is referred to as WEB/CICS/DB2 and is representative of customers who Web-enable access to their legacy applications. Continuing in the LSPR for z890/ z990 will be the legacy online workload, IMS™, and two legacy batch workloads CB84 and CBW2. The z890/ z990 LSPR will provide performance ratios for individual workloads as well as a "default mixed workload", which is used to establish single-number-metrics such as MIPs, MSUs and SRM constants. The z990 default mixed workload will be composed of equal amounts of five workloads: Trade2-EJB, WEB/CICS/DB2, IMS, CB84 and CBW2. Additionally, the z890/ z990 LSPR will rate all z/Architecture processors running in LPAR mode and 64-bit mode. The existing z900 processors have all been re-measured using the new workloads—all running in LPAR mode and 64-bit mode.

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

- Comparing a one-way z800 Model 001 to a one-way z890 Model A04 170, it is estimated that the z890 model has 2.0 times the capacity of z800 Model 001.
- Comparing a two-way z800 Model 002 to a two-way z890 Model A04 270, it is estimated that the z890 model has 2.1 times the capacity of z800 Model 002.
- Comparing a three-way z800 Model 003 to a three-way z890 Model A04 370, it is estimated that the z890 model has 2.1 times the capacity of z800 Model 003.
- Comparing a four-way z800 Model 004 to a four-way z890 Model A04 470, it is estimated that the z890 model has 2.1 times the capacity of z800 Model 004 (see Figure 1-4).

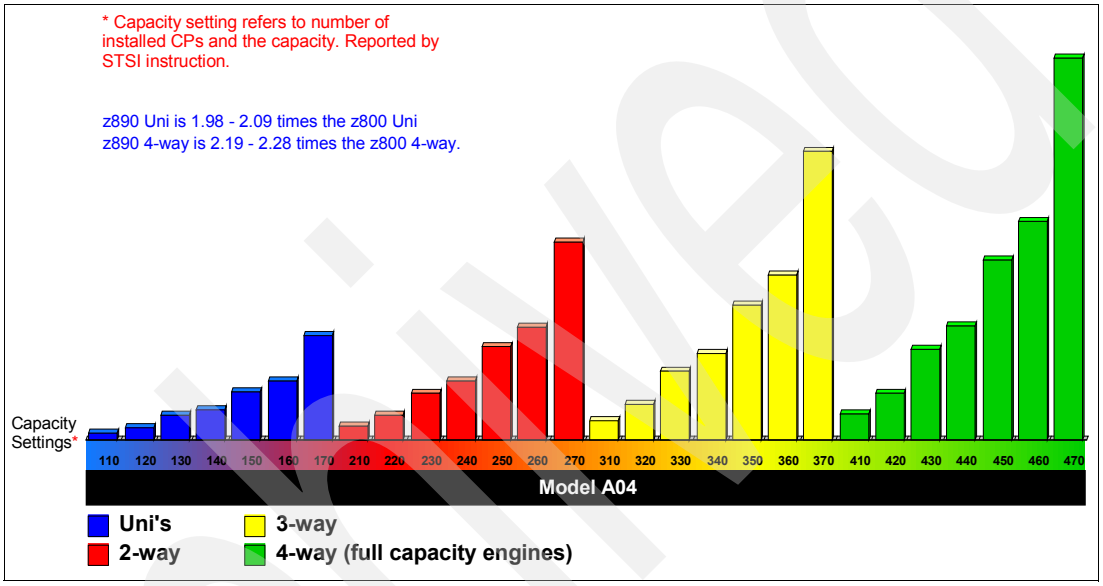

Figure 1-4 z890 capacity settings comparison

Further LSPR documentations can be obtained at:

http://www.ibm.com/servers/eserver/zseries/lspr

1.5 Migration

The z890 contains significant architectural changes from earlier systems. Software exploitation of these changes is available (see Chapter 3, “Software support” on page 63).

Migration from an earlier system can be somewhat more complex. This complexity is primarily associated with migrating an existing I/O environment (possibly for several earlier S/390 machines) into the Logical Channel Subsystem environment of the z890.

The most significant migration activity at this level might be an IOCDS (or HCP equivalent) move. A z890 contains more information in an IOCDS. A CHPID mapping tool is available and should ease this migration. The new PCHID addressing level can be used to make the CHPID numbers for the z890 generally match the CHPID numbers in an existing system. This can considerably reduce the effort in migrating a complex IOCDS (see “I/O configuration definition and management” on page 109 and “LPAR planning” on page 113 for details).

1.6 Considerations

There are a few z890 considerations when migrating from earlier systems. These include:

- Parallel (“bus and tag”) channels are not supported. Converters (which convert an ESCON channel to a parallel channel) are supported. The IBM 9034 converters may be used, as well as converters from Optica Technologies, Incorporated.
- Hardware cryptographic functions are different:
  - The CMOS cryptographic coprocessors are not available.
  - The PCICC cryptographic features are not available; replaced by PCIXCC features.
  - The PCICA cryptographic feature is available.
  - New cryptographic-assist instructions are included in *every* processor unit.

  For details on zSeries cryptography setup, refer to *zSeries 990 Cryptographic Coprocessor Configuration*, Redp3747, at:

  `http://www.redbooks.ibm.com/redpapers/pdfs/redp3747.pdf`
- Older channel and communications features cannot be used. These include:
  - FICON features; FICON Express features are used instead.
  - OSA-2 features; OSA-Express features are used instead.
  - FDDI and ATM features; an outboard solution must be used.
- Existing OSA-Express Fast Ethernet features on z800 or z900 servers can be moved to a z890, but new orders must be for the OSA-Express 1000BASE-T Ethernet feature.
- ISC connectivity to G3/G4 Servers are not supported.
- ICB-2 are not supported.
- A z890 IOCDS must contain PCHID numbers that assign CHPIDs to specific hardware channels. These assignments must be added to older HCD (and IOCP) definitions.

For more information on these and other considerations, refer to Chapter 4, “Planning and migration considerations” on page 91.

2

Technical details

With a single model and a wide range of capacity settings, the newest member of the zSeries family delivers significantly improved granularity and enriched functions over its predecessor, the z800. With almost twice the processing power of the z800, and significantly increased I/O capacity, the z890 supports double the number of Logical Partitions (LPs) with multiple Logical Channel SubSystem (LCSS), facilitating horizontal growth. New with the z890 is the zSeries Application Assist Processor (zAAP), which enables movement of JAVA processing cycles to a lower cost zSeries operating environment.

In this chapter we look at the inside of the z890 server in detail.

2.1 Frame and CEC cage

The single z890 frame, called the ‘A’ frame, is built to Electronic Industry Association (EIA) standards. It contains the Central Electronic Complex (CEC) cage at the top and I/O cage at the bottom. The basic frame is 40 EIU units high, with a breakpoint at 38 EIUs for the IBF area. The single I/O cage is always present. The approximate dimensions (including covers) are:

- 78.5 cm (30.9 inches) wide
- 157.7 cm (62.1 inches) deep
- 194.1 cm (76.4 inches) high
- 699 kg (1542 lbs) weight, maximum configuration without the IBF feature
- 785 kg (1730 lbs) weight, maximum configuration with the IBF feature

The top section (used for the IBF feature) can be removed to move the system through a doorway.

2.1.1 A-frame

As shown in Figure 2-1 on page 15, the main components in the A-frame are:

- The Internal Battery Features (IBFs).
  - The IBF further enhances the robustness of the power design, increasing Power Line Disturbance immunity. It provides battery power to preserve processor data in case of a loss of power on both of the AC feeders from the utility company. The IBF can hold power briefly over a “brownout”, or for orderly shutdown in case of a longer outage. The IBF provides up to 20 minutes of full power, depending on I/O configuration.
  - The optional Internal Battery Feature provides the function of a local uninterrupted power source.
- The CEC cage, contains the single book.
- A single I/O cage can house all supported types of channel cards. An I/O cage accommodates up to 420 ESCON channels or up to 40 FICON channels in the absence of any other card.
- Air-moving devices (AMD) providing N+1 cooling for the MBAs, memory, and DCAs.
- The Support Element (SE) tray, which is located in front of the I/O cage, contains the two Support Elements.

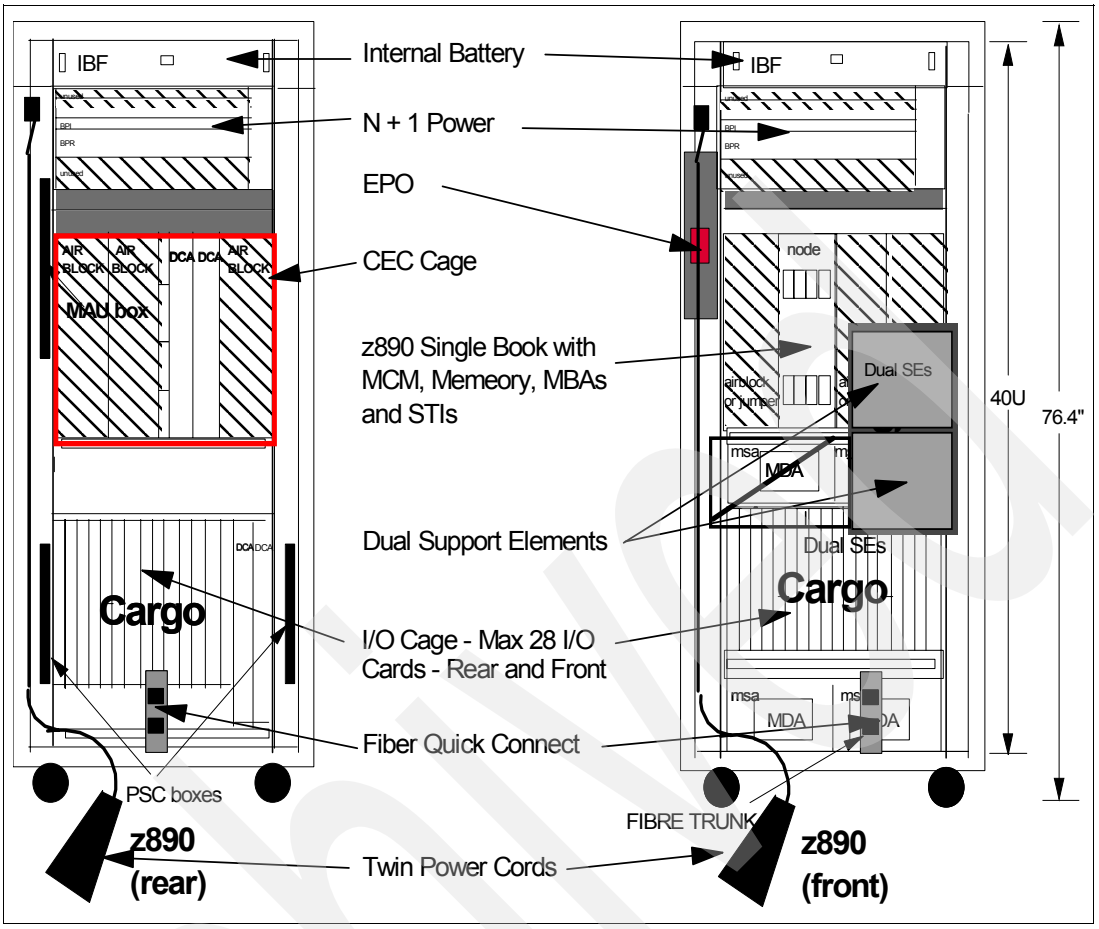

Figure 2-1 CEC cage and I/O cage location

2.1.2 CEC cage

The CEC cage contains the processor package known as a *book*. A z890 has one book; a z990 has up to four. Physically it appears to be a box that is plugged into the area indicated in Figure 2-1. The book contains the MCM, system memory, MBA chips, and the STI connectors. (Self Timed Interface or STI is the interface between the processor complex and an I/O cage.)

The book (Figure 2-2 on page 16) is physically large, about 50.8 cm (20 inches) high, 47 cm (18.5 inches) deep, and 15.25 (6 inches) wide. The rear side contains connectors with 36 power pins and 160 signal pins for the controllers. The front contains 8 STI connectors, of which a maximum of 7 can be connected to the I/O cage. The large physical size of the book makes cooling easier and provides a rugged base for the connectors and their associated cables.

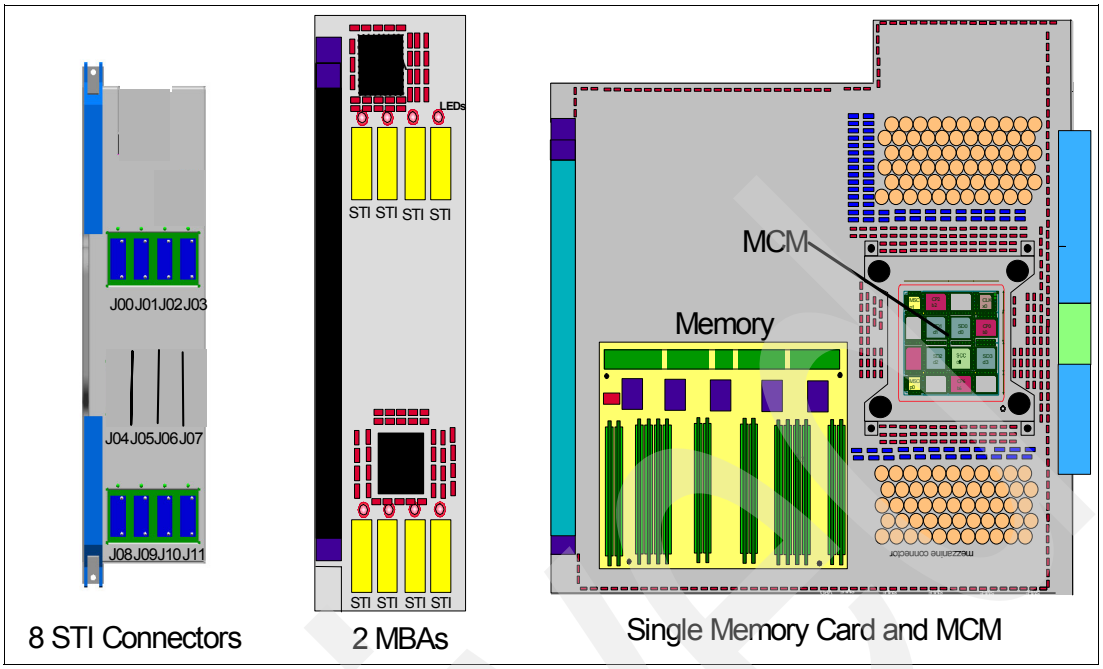

Figure 2-2 z890 book (showing STI connectors, MBA riser, memory, and MCM)

The PUs in the z890 are the same core processors as used with the z990, although the clocking speed is different. There is a single processor available on a chip, and *each* processor has duplicate I and E units that are used for internal checking in all zSeries machines. Processor cycle time is 1.0 ns. Every processor provides superscalar operation, meaning that multiple instructions may be executed in a cycle and some elements of an instruction may be executed out of sequential order.

z900 and z800 machines can have multiple instructions in process during a single cycle, while z990 and z890 machines can have multiple instructions complete during a single cycle. Instruction completion is always in the order of the original instructions, but the z990 and z890 can perform memory accesses (required by an instruction) out of order. The net effect is improved performance. The degree of improved performance depends somewhat on the nature of the workload.

The logical book structure is show in the Figure 2-3 on page 17. There are up to 8 STI buses to transfer data, and each STI has a bidirectional bandwidth of 2.0 GB/sec.

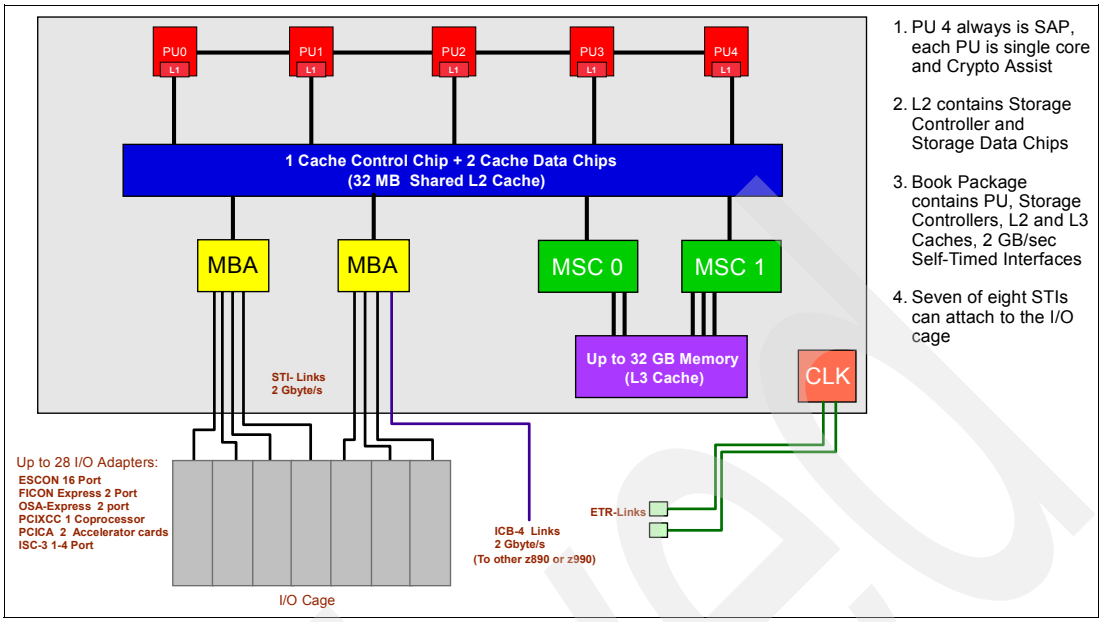

Figure 2-3 z890 logical structure

An STI is an interface from the Memory Bus Adapter (MBA) to:

- ESCON channels (16 port cards) in an I/O cage.
- FICON-Express channels (FICON or FCP modes, 2 port cards) in an I/O cage.
- OSA-Express channels (all on 2 port cards) in an I/O cage.
  - OSA-Express Gb Ethernet
  - OSA-Express Fast Ethernet
  - OSA-Express 1000BASE-T Ethernet
  - Token Ring
- ISC-3 links (up to 4 coupling links per mother card), via an ISC Mother (ISC-M) card in an I/O cage.
- Integrated Cluster Bus-3 (ICB-3) links. The ICB-3 (1 GB/sec) is in the I/O cage and requires a STI-3 extender card.
- ICB-4 links, directly attached to the 2.0 GB/sec STI interface for the communication to other z890s or z990s.
- PCIX Cryptographic Coprocessors (PCIXCC) in an I/O cage. Each PCIX Cryptographic Coprocessor feature contains one cryptographic coprocessor (daughter) card.
- PCI Cryptographic Accelerator (PCICA) in an I/O cage. Each PCI Cryptographic Accelerator feature contains two cryptographic accelerator (daughter) cards.

Data transfer between the CEC memory and attached I/O devices or CPCs is done through the Memory Bus Adapter. The physical path includes the Channel card (except for STI connected CPCs), the Self-Timed Interconnect bus, and possibly a STI extender card, the Storage Control, and the Storage Data chips.

Processor Unit design

Each PU is optimized to meet the demands of new e-business workloads without compromising the performance characteristics of traditional workloads. The PUs in the z890 have a superscalar design.

Superscalar

A scalar processor is a processor that is based on a single issue architecture, which means that only a single instruction is executed at a time. A superscalar processor allows concurrent execution of instruction by adding additional resources onto the microprocessor to achieve more parallelism by creating multiple pipelines, each working on their own set of instructions.

A superscalar processor is based on a multi-issue architecture. In such a processor where multiple instructions can be executed at each cycle, a higher level of complexity is reached because an operation in one pipeline may depend on data in another pipeline. A superscalar design therefore demands careful consideration of which instruction sequences can successfully operate in a multi-pipeline environment.

As an example the following deserves some thought. If the branch prediction logic of the microprocessor makes the wrong prediction, it might be necessary to remove all instructions in the parallel pipelines also.

There are challenges in creating an efficient superscalar processor. The superscalar design of the z890 PU has made big strides in avoiding address generation interlock situations. Instructions requiring getting information from memory locations may suffer multi cycle delays to get the memory content. The superscalar design of the z890 PU tries to overcome these delays by continuing to execute (single cycle) instructions that do not cause delays. The technique used is called “out of order operand fetching”. This means that some instructions in the instruction stream are already underway, while earlier instructions in the instruction stream that cause delays due to storage references take longer. Eventually the delayed instructions catch up with the already fetched instructions and all are executed in the designated order. The z890 PU gets much of its superscalar performance benefits from avoiding address generation interlocks.

It is not only the processor that contributes to the capability of successful execution of instructions in parallel. Given a superscalar design, compilers and interpreters must create code that benefit optimally from the particular superscalar processor implementation. Work is under way to update the C++ compiler and Java Virtual Machine for z/OS to better exploit the z890 microprocessor superscalar implementation. The intent is improve the performance advantage for e-business workloads such as WebSphere and Java applications.

By the time the Java Virtual Machine (JVM) and compilers are available, more improvement in the throughput of the superscalar processor is expected. In order to create instruction sequences that are least affected by interlock situations, instruction grouping rules are enforced to create instruction streams that benefit most from the superscalar processor. It is expected that e-business workloads will primarily benefit from this design since they tend to use more computational instructions.

A Websphere Application Server workload environment that runs a mix of Java and DB2 code will greatly benefit from the superscalar processor design of the z890. The superscalar design of the z890 microprocessor means that some instructions are processed immediately and that processing steps of other instructions may occur out of the normal sequential order, called pipelining. The superscalar design of the z890 offers:

- Decoding of two instructions per cycle
- Execution of three instructions per cycle (given that the oldest instruction is a branch)
- In-order execution
- Out-of-order operand fetching

Other features of the microprocessor, aimed at improving the performance of the emerging e-business application environment are:

- Floating point performance for IEEE Binary Floating Point arithmetic is improved to assist further exploitation of Java application environments.
- A secondary cache for Dynamic Address Translation, called the Secondary level Translation Look aside Buffer (TLB), is provided for both L2 instruction and data caches, increasing the number of buffer entries by a factor of eight.
- The CP Assist for Cryptographic Function (CPACF) accelerates the encryption and decryption of SSL transactions and VPN encrypted data transfers. The assist function uses five new instructions for symmetrical clear key cryptographic encryption and encryption operations.

Compression Unit

Each z990 PU has a Compression Unit on the chip, providing excellent hardware compression performance. The Compression Unit is integrated with the CP Assist for Cryptographic Function benefiting from combining the use of buffers and interfaces.

CP Assist for Cryptographic Function

Each z990 PU has a CP Assist for Cryptographic Function on the chip. The assist provides high performance hardware encryption and decryption support for clear key operations. To that end five new instructions are introduced with the cryptographic assist function.

The CP Assist for Cryptographic Function offers a set of symmetric cryptographic functions that enhance the encryption and decryption performance of clear key operations for SSL, VPN and data storing applications that do not require FIPS 140-2 level 4 security. The cryptographic architecture includes DES, T-DES data encryption and decryption, MAC message authorization and SHA-1 hashing.

The CP Assist for Cryptographic Function complements public key (RSA) functions and the secure cryptographic operations provided by the PCIXCC cryptographic coprocessor card.

The Multi-Chip Module (MCM)

The z890 is designed to provide balanced system performance. From processor to the 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 z890 is a scaled down version of the z990 using the most advanced chip and packaging technology in the industry. The z890 has more technology similarities to the z990 than the z800 does to the z900. Both the z890 and z990 use the latest MCM packaging and the processor book package. In fact, the chip layout of the z890 and z990 MCMs are identical.

The z890 MCM uses single core processor chips and blank "heat flow" chips which are non functional PU chips to assist drawing off heat from the MCM. Of the 16 chip sites 13 are used for functional chips. The glass ceramic has 101 layers and approximately 0.4 km of internal wire to connect the chips and provide paths to the Land Grid Array connectors on the bottom of the module. The MCM technology is denser than either the z800 or z900 technology because of a more advanced CMOS technology (9S versus. 8S) as well as a denser "pin out" from the bottom of the module.

Of distinction is the z890's MultiChip Module's (MCM's) Uni Processor Performance / Watt (UPP/W) ratio. High performance with low power assists the attainment of scalable reliable systems. The z890 achieves the highest UPP/W ratio to date of any large mainframe:

Year	Generation	Approx 1-Way Perf.	Power (W)	Ratio (Higher is better)
1998	G5	152	800	.190
1999	G6	205	900	.228
2000	z900	250	1100	.227
2002	z900 T	303	1300	.233
2003	z990	450	650	.692
2004	z890	370	450	.822

The z890 compared to the z800 provides a significant increase in system scalability and opportunity for server consolidation by providing improved model granularity while scaling to over two times the performance of the z800. The z890's processor cycle time has been improved to 1.0 ns from the 1.6 ns of the z800. This is achieved through shorter paths combined with cooler technology. All z890's have a single MultiChip Module (MCM) that is nearly identical to the z990's MultiChip module. The differences are minor. In the z890 MCM has 3 of the Processor Unit (PU) chip sites populated with blank chips that have the same thermal and physical characteristics as the z990's, but are not functional.

The z890 make use of single core Processor Units rather than utilization of dual core Processor Units, common on the z990. With fewer functional processors and a slower cycle time less power is required for the z890 MCM enabling air cooling. The air cooled z890 MCM uses approximately 450 Watts compared to the liquid cooled z990 MCM which requires approximately 650 Watts. The same z890 MCM can deliver 1 to 4-way configurations with a variety of capacity settings using microcode speed controls. The z890 MCM is contained in a single System book package. The System book package is comprised of a MCM, a memory book and Self-Timed Interconnects which are attached to a scaled down Riser I/O hub card. The z890 MCM, which measures approximately 93 x 93 millimeters, is about 46 percent smaller than the z900's 127 mm x 127 mm MCM. The z890 MCM contains 5 Processor Unit (PU) chips, four Storage Data (Level 2 cache) chips, a Level 2 Cache Storage Controller chip and two main Memory Storage Controller (MSC) chips which control the main memory or Level 3 storage. The MCM contains 101 glass ceramic layers of which 23 layers are to provide interconnection between the chips and the off-module environment. In total, there is approximately 378 meters of internal copper wiring on both the z890 and z990 MCMs compared to the 997 meters of internal wiring needed for the MCM's in the z900 1xx and 2xx models.

The z890 MCM is the world's densest logic package for mid range mainframes. It has over 20 percent more I/O connections and over a 130 percent I/O density improvement compared to the z900 MCM. The 5184 I/O's from the z890 MCM using the Land Grid Array (LGA) technology provide considerably more I/O's, 960 more, than the older pin technology used by the z900 (there are 4224 pins on the bottom of the z900 MCM). The z890's MCM provides support for 5 PUs and 32 MB level 2 cache which is shared by all PUs. Each PU chip contains up to approximately 122 million transistors and measures 14.1 mm x 18.9 mm. About half the transistors are used in the z890 PU chips. The design of the MCM technology on the z890 provides the flexibility to configure the PUs for different uses. One of the PUs is reserved for use as dedicated I/O Processor (IOP), commonly referred to as a System Assist Processor (SAP) enabling the Central Processor to avoid the burden of I/O set ups and operations. Unused PUs can be utilized as spares. The inactive PUs on the MCM are available to be

characterized as either CPs, Internal Coupling Facility (ICF) processors for Coupling Facility applications, Integrated Facility Linux (IFL) for Linux applications, Capacity Backup Upgrade (CBU) engines, On/Off processors or as zSeries Application Assist Processors (zAAPs), providing customers with tremendous flexibility in configuring the best system for running applications. Every z890 must be ordered with at least one CP, IFL or ICF.

The PU, which uses the latest chip technology from IBM semiconductor laboratories, is built on CMOS 9S-SOI with copper interconnections. The 14.1 mm x 18.9 mm chip has a cycle time of 1.0 nanoseconds and is identical to the z990 processor except for the cycle time. Implemented on this chip are leading edge functions as the z/Architecture with its 64-bit capabilities including instructions, 64-bit General Purpose Registers, superscalar function, translation facilities, ability to use long displacement operations and on board Cryptographic Assist functions. There are five processor units (PUs) in every system, although all of these might not be activated (depending on the feature codes ordered for the system). The processor chips and several other “nest” chips are mounted on a Multiple Chip Module (MCM) that contains many layers of chip interconnections. The MCM, roughly illustrated in Figure 2-4, is the heart of the system.

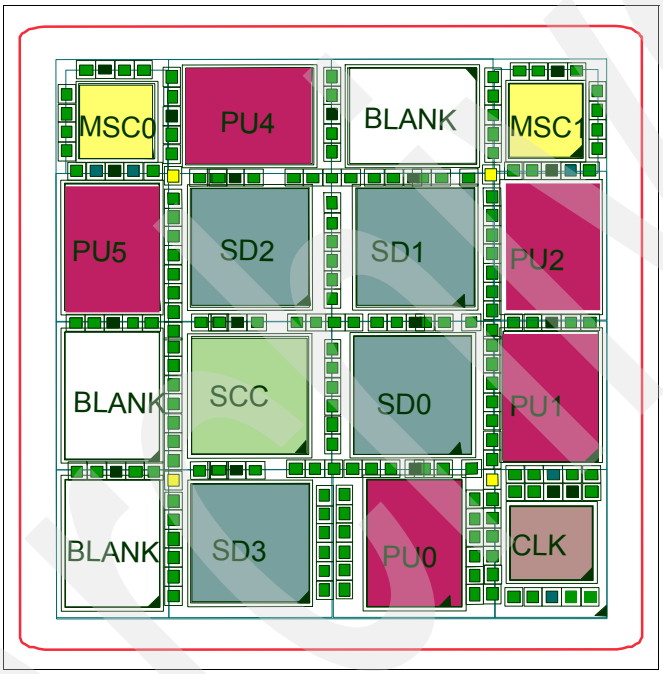

Figure 2-4 z890 MCM

Each processor chip has dual instruction and execution units (I and E units) to verify internal operations. These dual I/E units, used only to verify correct internal operation of the processor, are not further mentioned in this discussion. (They are not relevant to discussions about superscalar operation, logical partitions, threads, or any software-visible functions.)

Each processor has a superscalar design. This means that several instructions (from a single instruction stream) may be started, processed, and/or ended at the same time and that some processing steps within an instruction may be done out of sequential order.[3] However, instruction completion always appears in sequential order. A processor can decode up to 3 instructions per cycle and complete up to 3 instructions per cycle. Some instructions complete in a single cycle while others involve a pipeline of internal steps. Storage accesses may be out of sequential order.

[3] Sequential order means the order in which the instructions are issued from the program instruction stream.

Each processor has a dual L1 cache, meaning that separate caches exist for instructions and data. Both L1 caches are 256 KB. Each cache has its own Translation Lookaside Buffer (TLB) with 512 entries.[4] A second-level TLB has 4096 entries and a Branch History Table (BHT) has 8192 entries. All of the caches, TLBs and BHTs are four-way associative.

All processors share a single 32 MB L2 cache.

As with the z990, redesigned floating point hardware has enhanced Binary Floating Point (BFP) performance to bring it to at least the level of the traditional Hexadecimal Floating Point (HFP) performance.

New instructions (on both the z890 and z990) are available for certain cryptographic operations. These are new synchronous, high performance operations for symmetric, clear-key operations. Every z890 and z990 PU has these instructions.[5] The separate cryptographic coprocessors that were associated with a small number of PUs in older systems no longer exist.

The two memory interface chips (MBAs) are not included in the MCM, but are on a separate *riser card* in the processor package.

PU chip

All chips use CMOS 9SG technology, except for the clock chip (CMOS 8S). CMOS 9SG is state-of-the-art microprocessor technology based on eight-layer Copper Interconnections and Silicon-On Insulator technologies. The chip's lithography line width is 0.125 micron.

The 5 PUs are single core chips on the MCM in the single book. Four of the five PUs may be characterized for customer use. Each PU runs at a cycle time of 1.0 nanoseconds. Each PU chip measures 14.1 x 18.9 mm.

Each PU has a 512 KB on-chip Level 1 cache (L1) that is split into a 256 KB L1 cache for instructions and a 256 KB L1 cache for data, providing large bandwidth.

SC chip

The L1 caches communicate with the L2 caches (SD chips) by two bi-directional 16-byte data buses. There is a 2:1 bus/clock ratio between the L2 cache and the PU, controlled by the Storage Controller (SC chip), that also acts as an L2 cache cross-point switch for L2-to-MSC traffic, and L2-to-MBA traffic.

SD chip

The level 2 cache (L2) is implemented on the four System Data (SD) cache chips each with a capacity of 8 MB, providing a cache size of 32 MB. These chips measure 17.5 x 17.5 mm and carry 521 million transistors, making them the world's densest chips.

The single-core PU chips share the path to the SC chip (L2 control) and the clock chip (CLK)

2.1.3 Processor unit functions

One of the key components of the z890 server is the processor unit (PU). This is the microprocessor chip where instructions are executed and the related data resides. The instructions and the data are stored in the PU's high-speed buffer, called the Level 1 cache. Each PU has its own 512 KB Level 1 cache, split into 256 KB for data and 256 KB for instructions.

[4] The TLB is used to increase performance for virtual address to real address translations.
[5] A feature code must be included when ordering a system to enable these instructions.

The L1 cache is designed as a store-through cache, which means that altered data is synchronously stored into the next level, the L2 cache. Each PU has multiple processors inside and instructions are executed twice, asynchronously, on both processors.

This asymmetric mirroring of instruction execution runs one cycle behind the actual operation. This allows the circuitry on the chip to be optimized for performance and does not compromise the simplified error detection process that is inherent to a mirrored execution unit design.

One processor unit is contained on one chip. All PUs of a z890 server reside in a Multi-Chip Module. An MCM holds 5 PUs, of which 4 are available for customer use. All PUs in a z890 server are physically identical, but at initialization time PUs can be characterized to specific functions: CP, IFL, ICF, zAAP or SAP. The function assigned to a PU is set by the Licensed Internal Code loaded when the system is initialized (Power-on Reset) and the PU is "characterized". Only characterized PUs have a designated function; non-characterized PUs are considered spares.

This design brings an outstanding flexibility to the z890 server, as any PU can assume any available characterization. This also plays an essential role in z890 system availability, as these PU assignments can be done dynamically, with no server outage, allowing:

- Concurrent upgrades

  Except on a fully configured model, concurrent upgrades can be done by the Licensed Internal Code, which assigns a PU function to a previously non-characterized PU. Within the book boundary, no hardware changes are required and the upgrade can be done via Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU), On/Off Capacity on Demand (On/Off CoD), or Capacity BackUp (CBU).

- PU sparing

  In the rare case of a PU failure, the failed PU's characterization is dynamically and transparently reassigned to a spare PU.

A minimum of one PU per z890 server must be ordered as one of the following:

- A Central processor (CP)
- An Integrated Facility for Linux (IFL)
- An Internal Coupling Facility (ICF)

Central Processor

A Central Processor is a PU that has the z/Architecture and ESA/390 instruction sets. It can run z/Architecture, ESA/390, Linux, and TPF operating systems. It can also run the Coupling Facility Control Code (CFCC). The z890 can have up to 4 CPs.

The z890 can only be used in LPAR mode. In LPAR mode, CPs can be defined as dedicated or shared to a logical partition. Reserved CPs can be defined to a logical partition, to allow for non-disruptive *image* upgrades.

All CPs within a z890 configuration are grouped into a CP pool. Any z/Architecture, ESA/390, Linux, TPF operating systems, and the CFCC can run on CPs that are assigned from the CP pool.

Within the limit of all non-characterized PUs available in the installed configuration, CPs can be concurrently assigned to an existing configuration via Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU), On/Off Capacity on Demand (On/Off CoD), or Capacity BackUp (CBU).

Integrated Facility for Linux

An Integrated Facility for Linux (IFL) is a PU that can be used to run Linux on zSeries or Linux guests on z/VM operating systems. Up to 4 PUs may be characterized as IFLs. The IFL processors can be dedicated to a Linux or a z/VM logical partition, or be shared by multiple Linux guests and/or z/VM logical partitions running on the same z890 server. Only z/VM and Linux on zSeries operating systems can run on IFL processors.

All PUs characterized as IFL processors within a configuration are grouped into the ICF/IFL/zAAP processor pool. The ICF/IFL/zAAP processor pool appears on the hardware console as ICF processors. The number of ICFs shown is the sum of IFL, ICF, and zAAP processors on the server.

IFLs do not change the capacity setting model of the z890 server. Software product license charges based on the capacity setting model are not affected by the addition of IFLs.

Within the limit of all non-characterized PUs available in the installed configuration, IFLs can be concurrently added to an existing configuration via Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU), On/Off Capacity on Demand (On/Off CoD), but IFLs *cannot* be assigned via CBU.

Internal Coupling Facility

An Internal Coupling Facility (ICF) is a PU used to run the IBM Coupling Facility Control Code (CFCC) for Parallel Sysplex environments. Within the capacity of the sum of all unassigned PUs in the book, up to 4 ICFs can be characterized. ICFs can be concurrently assigned to an existing configuration via Capacity Upgrade on Demand (CUoD), On/Off Capacity on Demand (On/Off CoD), or Customer Initiated Upgrade (CIU), but ICFs *cannot* be assigned via CBU.

The ICF processors can only be used by Coupling Facility logical partitions. ICF processors can be dedicated to a CF logical partition, or shared by multiple CF logical partitions running in the same z990 server.

All ICF processors within a configuration are grouped into the ICF/IFL/zAAP processor pool. No z/Architecture, ESA/390, or TPF operating systems can run using an ICF processor from the ICF/IFL/zAAP processor pool. The ICF/IFL/zAAP processor pool appears on the hardware console as ICF processors. The number of ICFs shown is the sum of IFL, ICF, and zAAP processors on the system.

Only Coupling Facility Control Code (CFCC) can run on ICF processors; ICFs do not change the capacity setting model of the z890 server. This is important because software product license charges based on the capacity setting model are not affected by the addition of ICFs.

zSeries Application Assist Processors

The zSeries Application Assist Processor (zAAP) is a PU that is used exclusively for running Java application workloads under z/OS. One CP must be installed with or prior to any zAAP is installed. The number of zAAPs in a machine cannot exceed the number of CPs in that machine. This allows the following combinations of CPs, and zAAPs if no IFL, or ICF is present.

Table 2-1 CP and zAAP combinations in z890

z890 1-way	z890 2-way	z890 3-way	z890 4-way
CP	zAAP	-	-
CP	CP	ZAAP	-
CP	CP	zAAP	zAAP
CP	CP	CP	zAAP

Within the limit of non-characterized PUs available in the installed configuration, a zAAP can be concurrently added to an existing configuration via Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU), On/Off Capacity on Demand (On/Off CoD), as long as the number of zAAPs does not go beyond the number of CPs. zAAPs *cannot* be assigned via CBU.

With On/Off CoD you may concurrently install temporary zAAP capacity by ordering On/Off CoD Active zAAP features up to the number of current zAAPs that are permanently purchased. Also, the total number of On/Off CoD Active zAAPs plus zAAPs cannot exceed the number of On/Off Active CPs plus the number of active CPs on a z890 server.

PUs characterized as zAAPs within a configuration are grouped into the ICF/IFL/zAAP processor pool. The ICF/IFL/zAAP processor pool appears on the hardware console as ICF processors. The number of ICFs shown is the sum of IFL, ICF, and zAAP processors on the server.

zAAPs are orderable by feature code (FC 6520). Up to one zAAP can be ordered for each CP configured in the machine.

Important: The zAAP is a specific example of an assist processor that is known generically as an Integrated Facility for Applications (IFA). The generic term IFA often appears in panels, messages, and other online information relating to the zAAP.

zAAPs and logical partition definitions

zAAP processors can be defined as dedicated or shared in a logical partition and are always related to CPs. For a logical partition image both CPs and zAAPs logical processors are either dedicated or shared.

Purpose of a zAAP

zAAPs are designed for z/OS Java code execution. When Java code must be executed (i.e under control of Websphere) the z/OS Java Virtual Machine (JVM) calls the function of the zAAP. The z/OS dispatcher then suspends the JVM task on the CP it is running on and dispatches it on an available zAAP. After the Java application code execution is finished the z/OS dispatcher redispatches the JVM task on an available CP after which normal processing is resumed. This reduces the CP time needed to run WebSphere® applications, freeing capacity for other workloads.

A zAAP only executes Java Virtual Machine (JVM) code and is the only authorized user of a zAAP in association with some z/OS infrastructural code as the z/OS dispatcher, and supervisor services. A zAAP is not able to process I/O or clock comparator interruptions and does not support operator controls like IPL.

Java application code can either run on a CP or an zAAP. The user can manage the use of CPs such that Java application code runs only on a CP, only on an zAAP, or on both when zAAPs are busy.

Figure 2-5 shows at the logical level that Java code (JVM tasks) are dispatched on zAAP logical processors. PR/SM™ dispatches work for logical zAAPs on physical zAAPs as it does dispatch work for logical CPs on physical CPs.

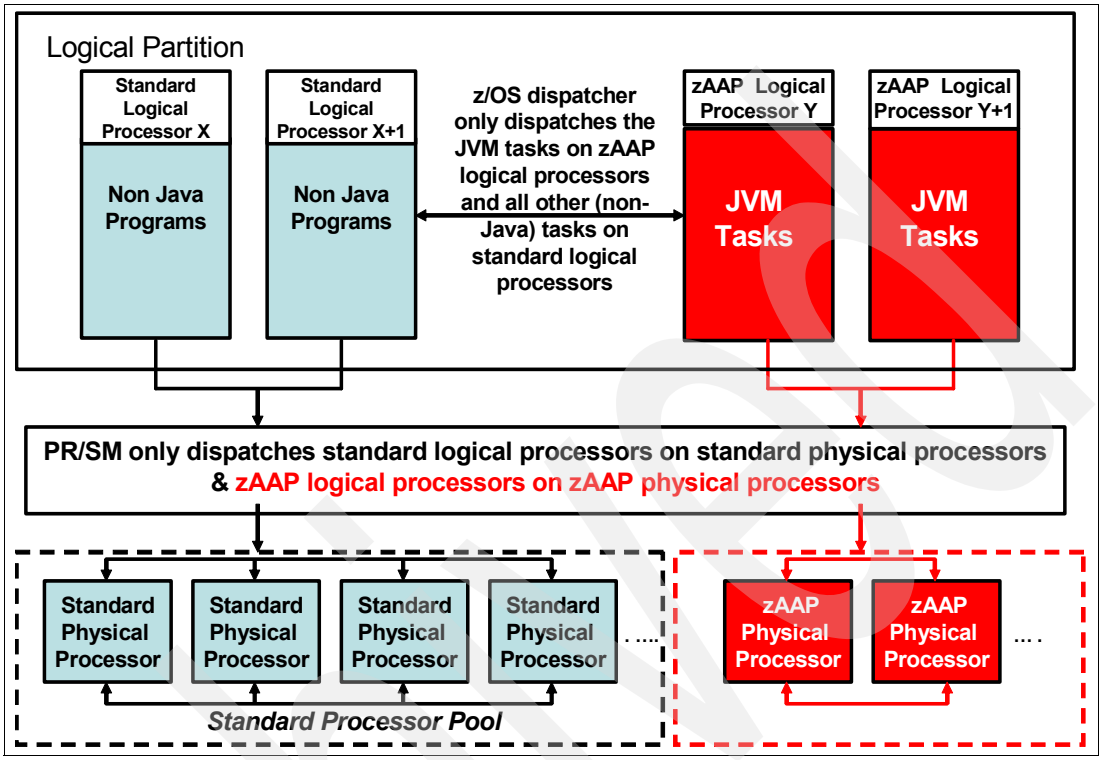

Figure 2-5 Logical and physical relationships of zAAPs and CPs

Software support

zAAPs do not change the capacity setting model of the z890 server. IBM software product license charges based on the capacity setting model are not affected by the addition of zAAPs.

z/OS Version 1.6 is a prerequisite for supporting zAAPs, together with IBM SDK for z/OS, Java 2 Technology Edition V1.4.1.

Exploiters of zAAPs include:

- WebsSphere Application Server 5.1 (WAS 5.1)
- CICS/TS 2.3
- DB2 Version 8
- IMS Version 8
- WebSphere WBI for z/OS

More on software support related to zAAPs can be found in “zSeries Application Assist Processor (zAAP)” on page 76.

System Assist Processors

A System Assist Processor (SAP) is a PU that runs the Channel Subsystem Licensed Internal Code to control I/O operations.

SAPs perform I/O operations for all logical partitions. All z890 capacity setting models have one standard SAP configured.

2.1.4 Memory

A z890 has 8 MB, 16 MB, 24 MB or 32 MB memory installed as a single memory card. IBM may install more physical memory than is activated. Memory access is interleaved two ways. Each access provides 128 bits (16 bytes) of usable data. (The storage path is actually 140 bits wide and includes a sophisticated sparing function.)

Spare memory chips are not used for this memory. Spare memory elements, used automatically as needed, are present in a distributed manner. Separate circuits and redundant memory is used for storage protection keys, and spare chips are used for the storage protection keys function.

Memory can be ordered in 8GB increments (8GB, 16GB, 24GB, 32GB) on any capacity setting model. The memory card resides on the processor board within the book. Physical card sizes are 8GB, 16GB and 32 GB. 8 and 16 GB cards come in two DRAM sizes, 256MB and 512MB. LICCC will control how much memory is to be enabled on the card which sets the total amount available for use in the system.

2.2 Internal System Control

Figure 2-6 on page 28 provides a conceptual overview of the system control design. While the details have changed, the general structure for system control is similar to other zSeries machines.

Various system elements contain *Flexible Support Processors* (FSPs).[6] A Flexible Support Processor is based on the IBM Power PC microprocessor. An FSP connects to an internal Ethernet LAN (to communicate with the Support Elements) and provides a SubSystem Interface (SSI) for controlling components.[7] The SSI includes several UARTs, digital I/O lines, and a number of unique sensing and control lines and protocols.

A typical FSP operation is to control a power supply (shown as DCAs in the figure). A Support Element might send a command to the FSP to bring up the power supply. The FSP (using the SSI connections) would cycle the various components of the power supply, monitor the success of each step, monitor the resulting voltages, and report this status to the Support Element.

A z890 has more FSPs than indicated in the figure and the interconnections are more complex, but Figure 2-6 on page 28 illustrates the general concepts involved. Most system elements are duplexed (for redundancy) and each element has an FSP. There are two internal Ethernet LANs and two Support Elements, again for redundancy. There is a crossover capability between the LANs, so that both Support Element can operate on both LANs.

The Support Elements, in turn, are connected to another (external) LAN (Ethernet or token ring) and the Hardware Management Consoles (HMCs) are connected to this external LAN. There can be one or more HMCs. In a production environment, the system hardware is normally managed from the HMCs. If necessary, the system can be managed from either Support Element. Several or all HMCs can be disconnected without affecting system operation.

---

[6] These were informally known as *cage controllers* in earlier systems.

[7] More detailed descriptions may refer to SSI-M and SSI-S notations, corresponding to master and slave functions.

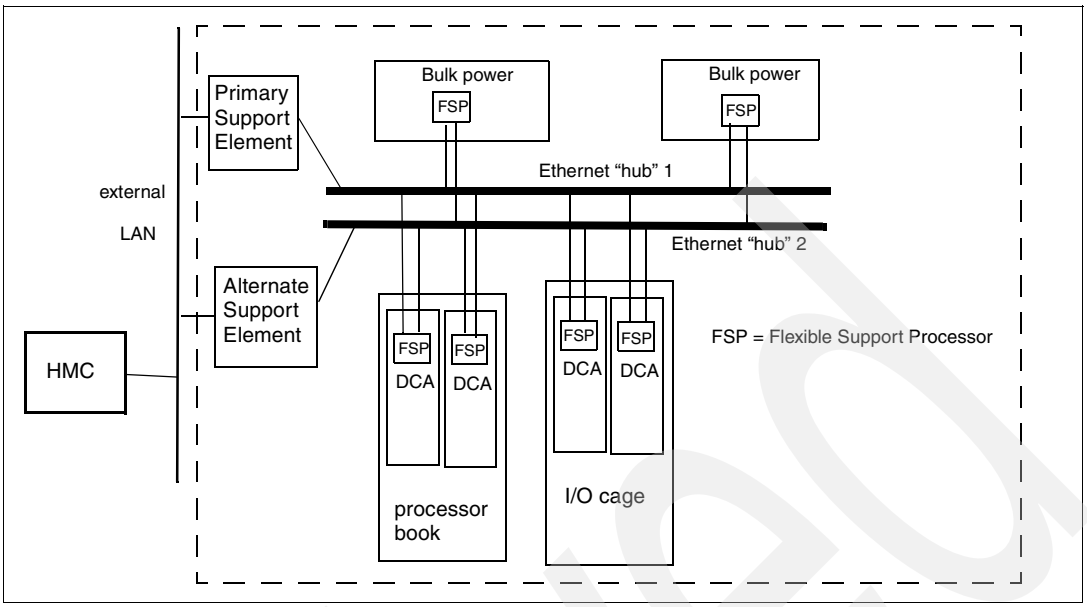

Figure 2-6 Conceptual overview of system control elements

The operation of the internal LANs and the FSPs is checked by a Network Heart Beat (NHB) function that uses complementary programs in the Support Elements and the FSPs to verify correct operation and to trigger alternative reassignments when necessary.

2.3 z890 capacity settings

There are a maximum of five PUs in a z890 server. One is always used as a SAP, leaving a maximum of four PUs for other purposes. The PUs that are to be used as CPs can be ordered at one of seven different capacity settings. (All CP engines operate at the same speed.) The CP configuration is specified by selecting *one* of the feature codes in Table 2-2 on page 29.

For example, feature code 6250 specifies two CPs operating at one of the intermediate capacity settings. It uses two of the possible four PUs. The remaining two PUs could be ordered as IFLs, ICFs, zAAPs, unassigned IFLs, or as capacity backup CPs. The feature codes for these are:

- FC 6516 - Integrated Facility for Linux (IFL)
- FC 6517 - Unassigned IFL
- FC 6518 - Internal Coupling Facility (ICF)
- FC 6520 - zSeries Application Assist Processor (zAAP)
- FC 6800 - Capacity Backup Central Processor (CBU CP)

Table 2-2 Feature codes for capacity settings

	1 CP	2 CPs	3 CPs	4 CPs
Least capacity	FC 6110	FC 6210	FC 6310	FC 6410
	FC 6120	FC 6220	FC 6320	FC 6420
	FC 6130	FC 6230	FC 6330	FC 6430
Mid-level capacity	FC 6140	FC 6240	FC 6340	FC 6440
	FC 6150	FC 6250	FC 6350	FC 6450
	FC 6160	FC 6260	FC 6360	FC 6460
Most capacity	FC 6170	FC 6270	FC 6370	FC 6470

A z890 ordered with feature code 6250 and two feature code 6516s would have two zSeries CP engines (at the appropriate capacity level), two full-speed IFL engines, and a SAP engine. A z890 ordered with feature code 6250 and no other processor feature codes would have two CP engines, one SAP engine, and two spare engines.

It is possible to order a z890 with no CP engines. (A special feature code, FC 6070, is used to indicate this.) An appropriate number of IFLs (FC 6516) and/or ICFs (FC 518) would be ordered to perform the desired function(s), as shown on Table 2-3. Note that no zAAPs can be ordered, as there are no CPs.

Table 2-3 z890 server options with no CPs (FC 6070)

FC 6070	IFLs Max [a]	ICFs Max [a]	zAAPs Max	SAPs	CBUs Max [a]
No CPs	4	4	0	1	4

a. The maximum number of combined IFLs, ICFs and CBU features is four.

zAAPs must be paired with CPs. That is, with one CP it is possible to order one zAAP. With two CPs, two zAAPs could be ordered. A zAAP cannot be ordered without a matching CP. This requirement is for system balance purposes. When operating, an zAAP is not tied to a specific CP.

Table 2-4 on page 30 shows all capacity setting models, including the CP Feature Codes, number of PUs, CPs, ICFs, IFLs, zAAPs, PUs, spare PUs, available engines for On/Off Capacity on Demand (CoD), and available engines for Capacity Backup (CBU).

Table 2-4 Capacity setting models

Capacity Setting Models	CP Feature Codes	PUs	CPs	ICFs Max [a]	IFL Max [a]	zAAPs Max [a]	SAPs	Spare PUs Max	On/Off CoD Engines Max	CBU Engines Max
110	6110	5	1	3	3	1	1	3	3	3
120	6120	5	1	3	3	1	1	3	3	3
130	6130	5	1	3	3	1	1	3	3	3
140	6140	5	1	3	3	1	1	3	3	3
150	6150	5	1	3	3	1	1	3	3	3
160	6160	5	1	3	3	1	1	3	3	3
170	6170	5	1	3	3	1	1	3	3	3
210	6210	5	2	3	3	2	1	3	3	3
220	6220	5	2	2	2	2	1	2	2	2
230	6230	5	2	2	2	2	1	2	2	2
240	6240	5	2	2	2	2	1	2	2	2
250	6250	5	2	2	2	2	1	2	2	2
260	6260	5	2	2	2	2	1	2	2	2
270	6270	5	2	2	2	2	1	2	2	2
310	6310	5	3	1	1	1	1	1	1	1
320	6320	5	3	1	1	1	1	1	1	1
330	6330	5	3	1	1	1	1	1	1	1
340	6340	5	3	1	1	1	1	1	1	1
350	6350	5	3	1	1	1	1	1	1	1
360	6360	5	3	1	1	1	1	1	1	1
370	6370	5	3	1	1	1	1	1	1	1
410	6410	5	4	0	0	0	1	0	0 [b]	0
420	6420	5	4	0	0	0	1	0	0 [b]	0
430	6430	5	4	0	0	0	1	0	0 [b]	0
440	6440	5	4	0	0	0	1	0	0 [b]	0
450	6450	5	4	0	0	0	1	0	0 [b]	0
460	6460	5	4	0	0	0	1	0	0 [b]	0
470	6470	5	4	0	0	0	1	0	0	0

a. The maximum number of combined ICFs, IFLs, and zAAPs cannot exceed the number of spare PUs.
b. Additional engines are not available, but capacity can be changed by changing capacity setting.

Some workloads run better with a smaller number of faster processors, while other workloads may run better with more, but slower processors. Preliminary information indicates that the

fastest one-way system (FC 6170) has roughly[8] the same potential capacity as a two-way (FC 6250) or a three-way (FC 6340) or a four-way (FC 6430) system. Which one you might select depends on the nature of your workload and, possibly, on the content of your software license agreements.

Additional system capacity information can be obtained from your IBM Representative.

Unassigned IFL features

A z890 server can have a total number of owned IFLs greater than the number of assigned IFLs. The owned but unassigned IFL processors on a z890 configuration are indicated by unassigned IFL features (FC 6517). One feature is used for each existent unassigned IFL on a server configuration.

On/Off CoD can temporarily turn on previously uncharacterized PUs, or any unassigned IFLs that are available within the current machine, as CPs, Integrated Coupling Facilities (ICFs), IFLs, and zSeries Application Assist Processors (zAAPs).

Downgrade features

A z890 server can also have a total number of owned CPs greater than the number of assigned CPs. However, there are no unassigned CP features for the z890. Downgrade features are used to indicate that owned but unassigned CPs exist on a z890 configuration. Downgrade features have feature codes based on their correspondent z890 CP feature codes, using the engine identifier (fourth position) as '2' (downgrade record). This results in downgrade feature codes from FC 6112 though FC 6472.

On/Off Capacity on Demand (On/Off CoD) features

On/Off CoD features are used to provide On/Off CoD upgrades and requires the On/Off CoD Enablement (FC 9896).

For Central Processors (CPs), On/Off CoD features have feature codes based on their correspondent z890 CP feature codes (except for FC 6110), using the engine identifier (fourth position) as '1' (On/Off CoD Day Use). This results in On/Off CoD feature codes from FC 6121 though FC 6471. The On/Off CoD feature code is related to the target model, after the upgrade. As an example:

- The z890 capacity setting model 220 has the CP feature code 6220 and the On/Off CoD feature code 6221.
- The z890 capacity setting model 420 has the CP feature code 6420 and the On/Off CoD feature code 6421.
- The upgrade via On/Off CoD from a model 220 server (FC 6220) to a model 420 (FC 6420) is done by the feature code 6421.

On/Off CoD capacity setting is restricted to the base capacity. For example, capacity setting 110 is eligible for On/Off CoD upgrades to only capacity settings 120 and 210. All other target capacity settings would result in more than two times the base capacity.

There are also On/Off CoD feature codes for:

- On/Off CoD Active IFLs (FC 9888)
- On/Off CoD Active ICFs (FC 9889)
- On/Off CoD Active zAAP (FC 9893)

[8] "Roughly" is the correct word here, based on very early estimates. You should consult your marketing representatives for more accurate information.

The maximum number of On/Off CoD zAAPs available for z890 cannot exceed the number of zAAPs, with the additional restriction that the sum of zAAPs and On/Off CoD zAAPs cannot exceed the number of CPs.

See more about On/Off CoD upgrades in “Capacity upgrades planning” on page 106.

Capacity Backup (CBU) features

CBU features are used to provide CBU upgrades. One CBU feature (FC 6800) is required for each additional Central Processor (CP) that is required on disaster/recovery situations.

CBU upgrades apply to whole CP engines additions and only to the largest capacity configuration (full engine). This means that a z890 capacity setting model 120 (FC 6120, one sub-capacity CP) with 2 CBU features (FC 6800) installed will result on a capacity model 370 (FC 6370, three full capacity CPs) after the CBU activation.

See more about CBU upgrades in “Capacity upgrades planning” on page 106.

2.4 Concurrent changes

A typical z890 may replace several older systems. One consequence of this is that any disruption of z890 service will probably affect more applications than on previous systems. You might view repair and configuration changes as having several levels:

- A *concurrent* repair or configuration change takes place while the system is running and the affected operating systems are running. That is, IPLs or a Power-on Reset (POR) are not required.
- A slightly more complex change, without POR or IPLs, would involve taking a CHPID offline and back online to activate a change to it.
- A *disruptive* change requires IPLing one or more operating systems, but without a POR. This implies that some operating systems (in their own logical partitions) are not affected.
- A POR affects the complete system for a relatively short time and requires all logical partitions to be IPLed again. (These IPL procedures will take varying amounts of time, depending on the complexity of the environments being started.)
- A change or repair that requires *power off* affects the complete system, of course, and implies more time than a POR.

The z890 has many internal, automatic, self-repairing functions. These are usually concurrent functions and may not be visible to any operating system. Many system upgrades and repairs are concurrent activities. Over the years, IBM has steadily increased the number of repair actions and configuration changes that can be done concurrently, and this direction is continued with the z890.

2.5 Additional hardware elements

There are new or changed hardware elements that characterize the z890.

2.5.1 Hardware System Area (HSA)

HSA is a memory area used by the system, especially by the I/O subsystem. HSA memory is taken from the total system memory. The size of HSA memory is variable and is determined by many factors. The largest factor is the number of subchannels used in each LCSS. An

IOCDS definition (via IOCP code or HCD parameters) can specify the maximum number of subchannels (MAXDEV=) per LCSS, and this number can range up to approximately 63 K. Changing these numbers (via the IOCDS) requires a Power-on Reset.

In an extreme case, where 63 K subchannels are allowed for each of two LCSS, the HSA size could be approximately 500 MB. Few installations need this number of subchannels, and a number 25 percent of this size might be considered more typical. You should work with your IBM Representative and your specific configuration requirements to obtain a valid HSA estimate.

Earlier systems allowed a *dynamic expansion factor* to be specified for HSA. This no longer exists and has been replaced by the ability to specify (in an IOCP or HCD) the maximum number of devices (MAXDEV=) that can be used in an LCSS.

The HSA Estimation Tool is available to estimate HSA size. A Hardware Management Console panel is provided for the input of configuration parameters to the tool which estimates the HSA size without requiring a Power-on Reset (POR). The following system I/O configurations require the given HSA size.

Minimum:

- No dynamic I/O
- 1 CSS
- 2 LPARs
- 2048 Devices

A total of 640 MB of memory is need for a minimum configuration.

Maximum:

- Dynamic I/O
- 2 CSSs
- 64512 devices per CSS
- 15 LPARs per CSS

A total of 1536 MB of memory is need for a maximum configuration.

2.5.2 Support Elements (SEs)

Two identical Support Elements (SEs) are standard and operate under OS/2®. The first is the primary SE, and the second is a backup. The backup can be used to preload SE code while the primary is active. The Support Elements are IBM ThinkPads. If a token ring feature is included with the Support Elements, then a Multi-station Access Unit (MAU) is included in the system frame. If the system has only Ethernet features, an Ethernet switch is included with the system.

2.5.3 Hardware Management Console (HMC)

A Hardware Management Console (HMC) must be used with the z890. An existing HMC can be used, provided it is at least at the level of feature code 0073. HMCs have a DVD drive, Token Ring and Ethernet features, and use OS/2 as the operating system.

Detailed information about available features for the HMC is found in Figure 4-2 on page 93.

2.5.4 External Time Reference (ETR)

z890 processors have 128-bit TOD clocks. This provides the capability to:

- Host dates beyond 2041 (64-bit TOD has this limit).
- Create sysplex-unique clock values.
- Provide more granular timestamps. (z/Architecture specifies that every clock query will be provided with a unique TOD value. The z890 processor speed exceeds the granularity of the 64-bit clock. If a user issues two Store Clock instructions in quick succession, the CP would have to spin, waiting for the clock to be incremented before it could return the value for the second request.)

2.5.5 Sysplex Timer®

The IBM 9037 Sysplex Timer is a mandatory hardware requirement for a Parallel Sysplex consisting of more than one zSeries or Generation 5 or 6 server.

The Sysplex Timer provides the synchronization for the time-of-day (TOD) clocks of multiple servers, and thereby allows events started by different servers to be properly sequenced in time. When multiple servers update the same database, all updates are required to be time stamped in proper sequence.

There is one model of the IBM 9037 Sysplex Timer Unit: Model 2.

2.5.6 z890 ETR (Sysplex Timer) attachment

Two optional External Time Reference (ETR) cards are features (feature code 6154) on the z890 server. These cards, located in the processor cage, provide attachment to the 9037 Sysplex Timer. Each ETR card should connect to a different 9037 Sysplex Timer in an Expanded Availability configuration.

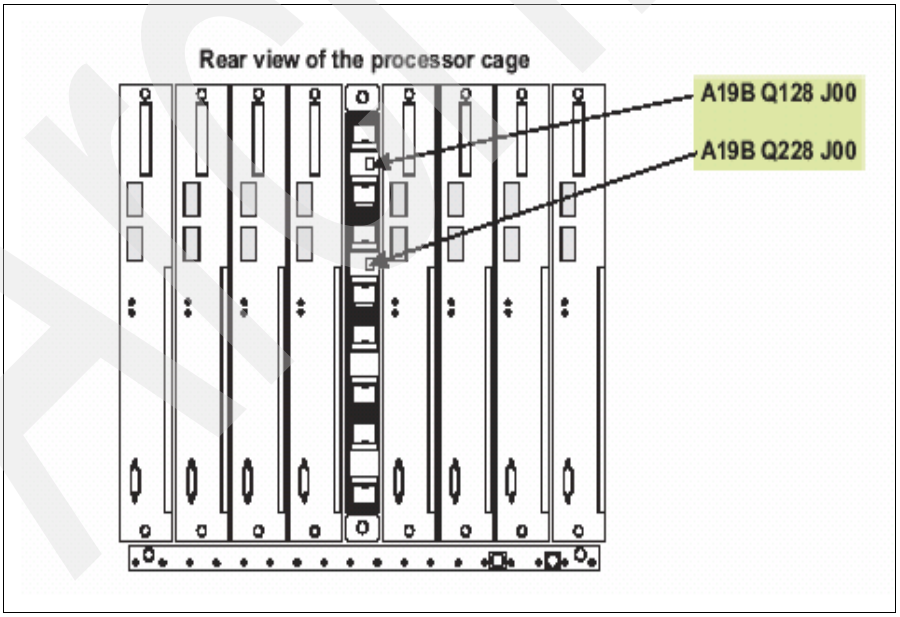

Figure 2-7 Sysplex timer connections - ETR cards

As part of the installation of a Sysplex Timer in either Basic or Expanded Availability configurations, each Sysplex Timer is assigned an ETR Network ID. More information about assigning the ETR Network ID is found in “Sysplex Timer ETR Network ID” on page 104.

Sysplex Timer synchronization

As processors and Coupling Facility link technologies have improved over the years, the synchronization tolerance between operating systems in a Parallel Sysplex has become more rigorous. In order to ensure that any exchanges of timestamped information between operating systems in a sysplex involving the Coupling Facility observe the correct time ordering, timestamps are now included in the message-transfer protocol between the server operating systems and the Coupling Facility.

Connectivity

The ETR cards are located in the CPC cage of the z890 processors as shown onFigure 2-7 on page 34 Each ETR card has a single port supporting an MT-RJ fiber optic connector to provide the capability to attach to a 9037 Sysplex Timer Unit. The MT-RJ is an industry standard connector which has a much smaller profile compared with the original ESCON Duplex connector supported on the ETR ports of G5/G6 and earlier servers. The 9037 Sysplex Timer Unit has an optical transceiver that supports an ESCON Duplex connector.

An MT-RJ/ESCON Conversion Kit supplies two 62.5 micron multimode conversion cables. The conversion cable is two meters (6.5 feet) in length and is terminated at one end with an MT-RJ connector and at the opposite end with an ESCON Duplex receptacle to attach to the under floor cabling. The same conversion kit is used on the z990 processors with the 16-port ESCON feature or ETR when reusing existing 62.5 micron multimode fiber optic cables terminated with ESCON Duplex connectors.

Note: The ETR card does not support a multimode fiber optic cable terminated with an ESCON Duplex connector. However, 62.5 micron multimode ESCON Duplex jumper cables can be reused to connect to the ETR card. This is done by installing an MT-RJ/ESCON Conversion kit between the ETR card MT-RJ port and the ESCON Duplex jumper cable. Information about ETR cable ordering is found in “Ordering ETR cables” on page 105.

Message Time Ordering

As server and Coupling Facility link technologies have improved over the years, the synchronization tolerance between servers in a Parallel Sysplex has become more rigorous. To help ensure that any exchanges of time-stamped information between servers in a Parallel Sysplex involving the Coupling Facility observe the correct time ordering, time stamps are now included in the message-transfer protocol between the servers and the Coupling Facility. Therefore, when a Coupling Facility is configured as an ICF on any z890, the Coupling Facility requires connectivity to the same Sysplex Timer (R) that the other servers in its Parallel Sysplex are using for time synchronization. If the ICF is on the same server as a member of its Parallel Sysplex, no additional Sysplex Timer connectivity is required, since the server already has connectivity to the Sysplex Timer. However, when an ICF is configured on a z890 which does not host any servers in the same Parallel Sysplex, it is necessary to attach the server to the Sysplex Timer.

This is demonstrated in Figure 2-8 on page 36.

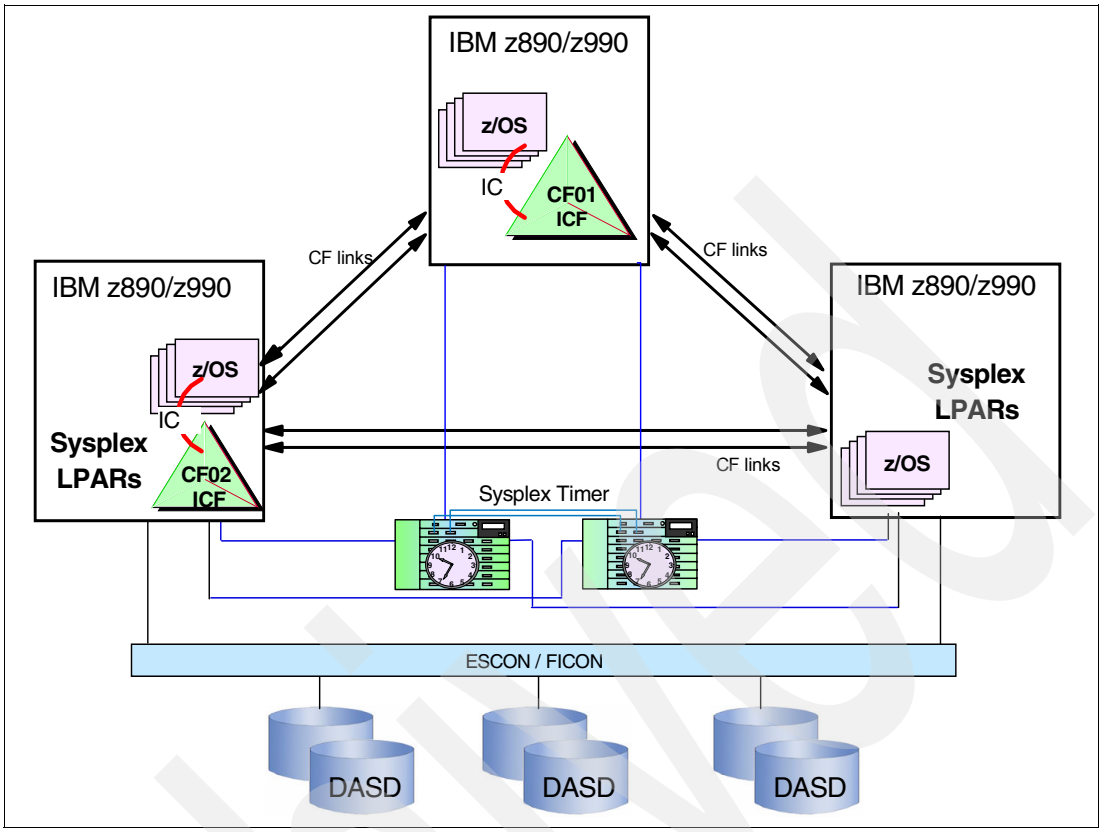

Figure 2-8 z890 and z990 Parallel Sysplex connectivity

2.5.7 Coupling Facility Control Code (CFCC) Level 13

In this section we discuss Coupling Facility Control Code (CFCC) Level 13.

Enhanced Patch Apply

The CFCC patch apply process has been enhanced to eliminate the need for a power on reset (POR) of the z890 to apply a "disruptive" CFCC patch. This enhancement provides you the ability to:

- Selectively apply the new patch to one of possibly several CFs running on a z890. For example, if you have a CF that supports a test Parallel Sysplex and a CF that supports a production Parallel Sysplex on the same z890, you now have the ability to apply a "disruptive" patch to only the test CF without affecting the production CF. After you have completed testing of the patch, it can be applied to the production CF as identified in the example. To use the updated CFCC code to a CF logical partition, simply deactivate and activated the partition. When the CF comes up, it displays its version on the OPRMSG panel for that partition.
- Allows all other logical partitions on the z890 where a "disruptive" CFCC patch will be applied to continue to run without being impacted by the application of the "disruptive" CFCC patch. This enhancement does not change the characteristics of a "concurrent" CFCC patch, but does significantly enhance the availability characteristics of a "disruptive" CFCC patch by making it much less disruptive.

Previously, small enhancements or "fixes" to the CFCC were usually distributed as a "concurrent" patch that could be applied while the CF was running. Occasionally, a CFCC patch was "disruptive". When such a "disruptive" change needed to be applied to a CF, it

required a POR of the server where the CF was running. This was especially disruptive for those enterprises that had chosen to use internal CFs, because a POR of the server affects the CF where the change was to be applied and all other logical partitions running on the same server. Further, if an enterprise was using multiple internal CFs on the same server (to support both a test and production configurations for example), there was no way to selectively apply the "disruptive" patch to just one of the CFs—once applied, all the CFs on the server had the change. Consequently, the application of the "disruptive" CFCC patch was very disruptive from an operations and availability perspective.

Additional CFCC Enhancements

CFCC Level 13 also provides changes which will affect different software environments that run within a Parallel Sysplex. For example, DB2 data sharing is expected to see a performance improvement, especially for cast-out processing against very large DB2 group buffer pool structures.

z890 fully supports System-Managed CF Structure Duplexing. This is a set of architectural extensions to Parallel Sysplex in support of duplexing of Coupling Facility structures for high availability. All three structure types—cache structures, list structures, and locking structures —can be duplexed using this architecture.

The IBM technical paper, System Managed CF Structure Duplexing (GM13-0103), includes information about:

- The cost/benefit trade-offs in duplexing
- Determining which structures should be duplexed in a specific Parallel Sysplex

This paper is available at:

http://www.ibm.com/servers/eserver/zseries/pso

2.6 I/O cage

There are 8 STI buses on the book to transfer data, with a bi-directional bandwidth of 2.0 GB/sec each. An STI is driven off an MBA. There are two MBAs in the book, each driving four STIs, providing an aggregated bandwidth of 16 GB/sec on the z890.

The STIs connect to the I/O cages that may contain a variety of channel, coupling link, OSA-Express and Cryptographic feature cards:

- ESCON channels (16 port cards).
- FICON-Express channels (FICON or FCP modes, 2 port cards).
- ISC-3 links (up to 4 coupling links, two links per daughter card (ISC-D). Two daughter cards plug into one mother card (ISC-M).
- Integrated Cluster Bus (ICB) channel - ICB-3 (1 GB/sec) require an STI extender card in the I/O cage. The STI-3 card provides two output ports to support the ICB-3 links. The STI-3 card converts the output into two 1GB/sec links.
- Integrated Cluster Bus (ICB) channel - ICB-4 (2 GB/sec) provides a direct STI to STI peer connection between z990, and z890 servers. The ICB-4 channels do not require a slot in the I/O cage and attaches directly to the STI of the communicating CEC.
- OSA-Express channels
  - OSA-E Gb Ethernet
  - Fast Ethernet (carried forward from z800 to z890)
  - 1000BASE-T Ethernet

– Token Ring (carried forward from z800 to z890)

Note: Although they reside in I/O cage slots, the cryptographic accelerator and cryptographic coprocessor are no I/O features in the proper sense but rather processors with a specific function not related to I/O.

- PCI Cryptographic Accelerator (PCICA, 2 engines per feature)
- PCIX Cryptographic Coprocessor (PCIXCC, 1 engine per feature)

See Table 2-6 on page 39 for more.

The I/O cage for the z890 is exactly the same as the standard I/O cage in a z990 system. The features are the same, with the addition of the OSA-Express Integrated Console Controller (OSA-ICC) function. (The smallest sub-uniprocessor z890 system (FC 6110) has additional limitations that are discussed later.)

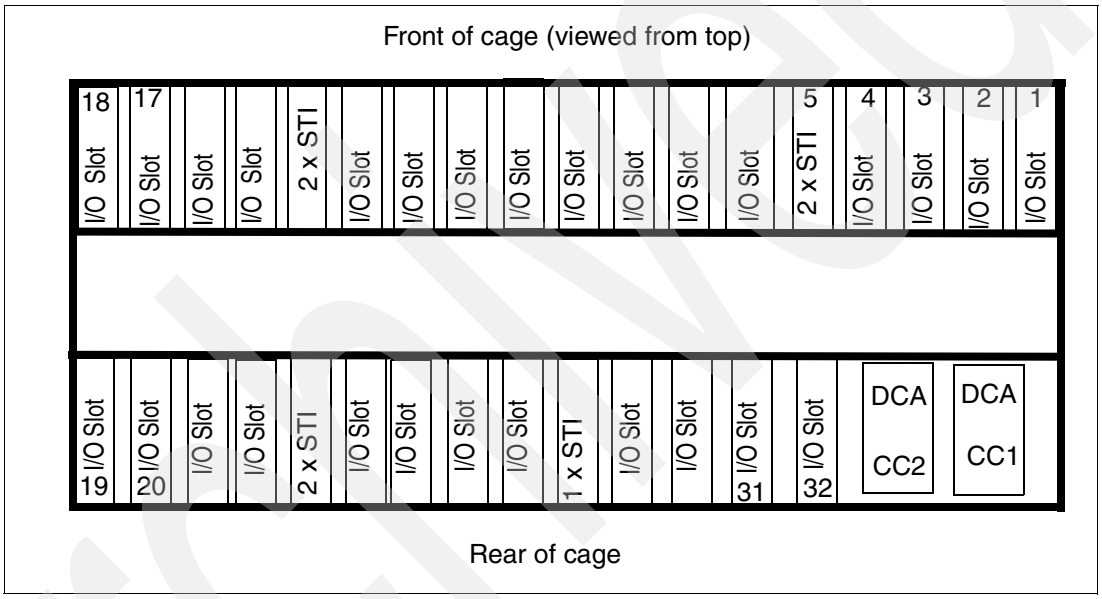

Figure 2-9 z890 I/O cage

Notice that slots 5, 14, 23, and 28 are used for STI connections. There are two such connections in slots 5, 14, and 23, while slot 28 has a single STI connection. (The “other half” of slot 28 may be used for a power sequence control (PSC) card.) IBM selects which slots are used for I/O features and STI cables. An STI cable goes directly from an I/O cage to an STI connector on the processor book. Each STI connection controls a *domain* of four I/O slots.

The assignment of slots to domains for each I/O cage is as follows.

Table 2-5 Assignment of slots to domains for each I/O cage

Domain	I/O slots in domain
0	1, 3, 6, 8
1	2, 4, 7, 9
2	10, 12, 15, 17
3	11, 13, 16, 18
4	19, 21, 24, 26

Domain	I/O slots in domain
5	20, 22, 25, 27
6	29, 30, 31, 32

2.6.1 I/O and cryptographic features

Table 2-6 provides a summary of supported I/O and cryptographic features. Note that parallel channel, FDDI, ATM, ICB-2, and OSA-2 features are not supported with the z890 server.

The PCIXCC feature was introduced on the z990 and is used with the z890. An I/O slot is required for each feature. Cryptographic features, residing in I/O cage slots, have limits regarding the number of features that can be installed.

Table 2-6 I/O and cryptographic features support

I/O feature	Feature codes	Number of		Max. number of		PCHID	CHPID definition	Config. rules notes
		Ports per card	Ports increments	Ports	I/O slots			
ESCON	2323 2324 (ports)	16 (1 spare)	4 (LIC-CC)	420	28	yes	CNC, CVC, CTC, CBY	1, 2
FICON Express LX / SX	2319 / 2320	2	2	40	20	yes	FC, FCV, FCP	3, 4
OSA-E Gbit Ethernet LX / SX	1364 / 1365	2	2	40	20	yes	OSD	4, 5, 6
OSA-E 1000BASE-T Ethernet	1366	2	2	40	20	yes	OSE, OSD, OSC	4, 5
OSA-E Fast Ethernet	2366	2	2	24	12	yes	OSE, OSD	4, 5, 6
OSA-E Token Ring	2367	2	2	40	20	yes	OSE, OSD	4, 5
ICB-3 (1 GByte/sec)	0993	2	1	16	8	yes	CBP	7
ICB-4 (2.0 GByte/sec)	3393	-	1	8	0	yes	CBP	7
ISC-3 at 10km (1 or 2 Gbit/sec)	0217 (ISC-M) 0218 (ISC-D) 0219 (ports)	4/ISC-M 2/ISC-D	1 (LIC-CC)	48	12	yes	CFP CFS, CFR	7, 8, 9
ISC-3 20km support (1 Gbit/sec)	RPQ 8P2197 (ISC-D)	4/ISC-M 2/ISC-D	2	48	12	yes	CFP CFS, CFR	7, 8, 9
HiperSockets	-	-	1	16	0	no	IQD	10
IC	-	-	2	32	0	no	ICP	7, 10
ETR	6154	1	-	2	-	no	-	11
PCIXCC	0868	1	1	4	4	yes	-	4, 12
PCICA	0862	2	2	4	2	yes	-	4, 12

Configuration rules notes:

1. The ESCON 16-port card feature code is 2323, while individual ESCON ports are ordered in increments of four using feature code 2324. The ESCON card has 1 spare port and up to 15

usable ports.

2. The maximum number of ESCON ports on the capacity setting model 110 is 240 ports (16 cards).
3. The maximum number of FICON Express ports on the capacity setting model 110 is 32 ports (16 cards).
4. The total number of FICON Express, OSA-Express, PCIXCC, and PCICA cards cannot exceed 20 in the I/O cage (maximum of 16 cards on the capacity model 110).
5. The sum of OSA-Express GbE, 1000BASE-T Ethernet, Fast Ethernet and Token Ring cards cannot exceed 20 (maximum of 16 cards on the capacity model 110).
6. OSA-Express Fast Ethernet (FC 2366) can only be brought forward on an upgrade from z800; new adapters cannot be ordered.
7. The sum of IC, ICB-3, ICB-4, active ISC-3 and RPQ 8P2197 links supported on a z890 server is limited to 64.
8. There are two feature codes for the ISC-3 card:
   – Feature code 0217 is for the ISC Mother card (ISC-M).
   – Feature code 0218 is for the ISC Daughter card (ISC-D).

   One ISC Mother card supports up to two ISC Daughter cards, and each ISC Daughter card contains two ports. Port activation must be ordered using feature code 0219. RPQ 8P2197 is available to extend the distance of ISC-3 links to 20 km at 1Gb/sec. When RPQ 8P2197 is ordered, both ports (links) in the card are activated.
9. The maximum number of ISC-3 ports in peer mode is 48, and the maximum number in compatibility mode is 32.
10. There are two types of “virtual” links that can be defined and that require CHPID numbers, but do not have PCHID numbers:
    – Internal Coupling (IC) links; each IC link pair requires two CHPID numbers.
    – HiperSockets, also called Internal Queued Direct I/O (iQDIO). Up to 16 virtual LANs can be defined, each one requiring a CHPID number.
11. Two ETR cards are automatically included in a server configuration if any coupling link I/O feature (ISC-3, ICB-3 or ICB-4) is selected.
12. The PCIXCC and PCICA features do not require CHPIDs but have PCHIDs.

The maximum number of I/O slots is 28. The capacity setting model 110 can have up to 16 I/O slots.

At least one channel I/O feature (FICON Express or ESCON) or one Coupling Facility link I/O feature (ISC-3, ICB-3 or ICB-4) must be present in a configuration.

The maximum number of configurable CHPIDs is 256 per Logical Channel Subsystem (LCSS) and per operating system image.

2.6.2 New I/O features and functions

The OSA-Express Integrated Console Controller (OSA-ICC) function is new with the z890 and will also be made available on z990 system. This provides substantial new functionality replacing external devices, for example, 3174s, and 2074s. It uses an OSA-Express 1000BASE-T Ethernet feature port. The OSA-ICC function consists of an alternate microcode load for one or both Ethernet ports on the feature.

The OSA-Express 1000BASE-T Ethernet feature was introduced on the z990 and is used with the z890. It replaces the OSA-Express Fast Ethernet feature. (As noted, existing OSA-Express Fast Ethernet features may be carried forward from earlier systems, but new ones cannot be ordered.) The new feature uses copper connections and offers 1000 Mbps, as well as 100 Mpbs and 10 Mpbs, operation. Two gigabit Ethernet ports are available:

- The OSA-Express Gigabit Ethernet feature uses fiber connections.
- The OSA-Express 1000BASE-T Ethernet feature uses copper connections.

As with other OSA-Express features, two ports (channels or CHPIDs) are installed on each feature and are used independently.

A new OSA-Express Gigabit Ethernet feature is available (introduced with the z990) and provides new fiber connectors. Both new features (1000BASE-T Ethernet, and Gb Ethernet) provide a checksum offload function that will be supported by z/OS 1.5 for use with IPv4 (but not IPv6). This function reduces CPU overhead during IP transfers.

Note: I/O feature configuration rules are found in detail in “I/O feature configuration rules” on page 102.

2.6.3 I/O interfaces and identification

The I/O design and implementation is very similar to that of the z800, z900, and z990. The primary difference is that the 1 GB/second STI links of the z800/z900 have been replaced with 2.0 GB/second STI links. The general data flow structure is conceptually illustrated in Figure 2-10, and works like this:

- The processor book has 8 STI links, regardless of the active number of PUs in the book.
- An STI link (from the book) is connected to an STI-M card in the I/O cage. Two STI-M cards fit in a single slot in the I/O cage. They use slots 5, 14, 23, and 28 in the I/O cage. Slot 28 can hold only one STI-M card. This provides for a maximum of 7 STI-M cards in the I/O cage.
- Each STI-M card (in the I/O cage) is connected to four slots. The four slots driven by an STI-M card are its *domain*.
- Directly connected ICB-4 links bypass the I/O cage and occupy a processor book STI connection.
- STI-3 features are used for ICB-3 connections.
- Non-I/O features that occupy a slot in an I/O cage (such as a cryptographic feature) are considered to be I/O features for this discussion.

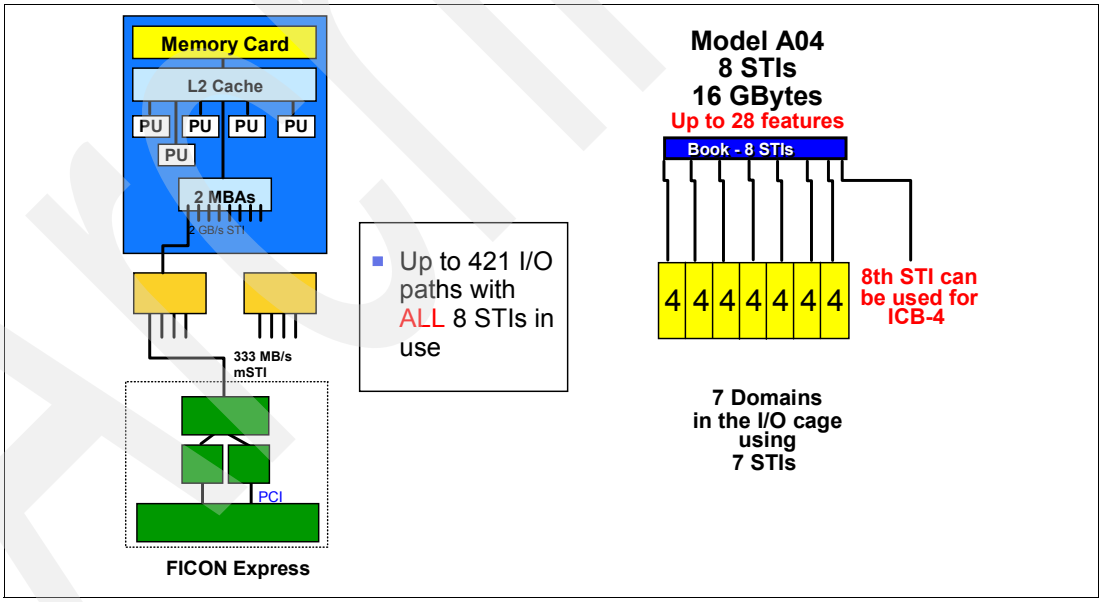

Figure 2-10 z890 logical channel configuration and STI and I/O cage structure

Each possible “physical” CHPID port/connector has a PCHID number that is fixed by the physical location of the feature and connector. An IOCDS associates a CHPID number to a PCHID number. A feature in an I/O cage slot can provide a maximum of 16 CHPIDs. (The ESCON feature does this, although only 15 can be used at one time.) The I/O cage has a maximum of 28 slots. PCHID numbers are assigned as follows:

```
I/O cage slot   PCHID numbers          I/O cage slot   PCHID numbers
```

```
1          100 - 10F          2          110 - 11F
3          120 - 12F          4          130 - 13F
6          140 - 14F          7          150 - 15F
8          160 - 16F          9          170 - 17F
10         180 - 18F          11         190 - 19F
12         1A0 - 1AF          13         1B0 - 1BF
15         1C0 - 1CF          16         1D0 - 1DF
17         1E0 - 1EF          18         1F0 - 1FF
19         200 - 20F          20         210 - 21F
21         220 - 22F          22         230 - 23F
24         240 - 24F          25         250 - 25F
26         260 - 26F          27         270 - 27F
29         280 - 28F          30         290 - 29F
31         2A0 - 2AF          32         2B0 - 2BF

PCHID numbers 100 to 1FF are in the front of the cage, 200 to 2BF in the rear.
PCHID numbers used for direct book connections: 010-013 & 018-01B
```

Note that slot numbers 5, 14, 23, and 28 are missing. These slots are used for STI connectors.

Consider an ESCON feature in slot 15 of the I/O cage. PCHID numbers 1C0 - 1CF are reserved for this slot. The first ESCON connector on the feature (whether or not that particular port is enabled or used) is PCHID 1C0, the second connector is 1C1, and so forth. The PCHID numbers are fixed. The CHPID numbers are not fixed and can be arbitrarily assigned when an IOCDS is constructed. For example, an installation might decide to make PCHID 1C0 CHPID 52 in LCSS 1.

I/O connections to books

Direct book connections (to STI ports) are used for ICB-4 channels. (These channels are used for connections to Coupling Facilities or other z890 or z990 servers.) There are 8 STI ports on the processor book, and a PCHID number is assigned to each one, using the range 010-013 & 018-01B. These PCHIDs and connections do not involve I/O cages and are used for ICB-4 connections.

2.7 Channel subsystem

One of the most striking architectural changes associated with the z890 and the z990 is the set of extensions to the channel subsystem. Perhaps the best way to start a description is to explain why these extensions were introduced. The purpose is simple. It is to permit a these systems to have more than 256 channels. The maximum number of channels has been a concern for some time, although the move to FICON channels and devices has mitigated the limitation.

The difficulty is to extend this number while maintaining compatibility with existing software. The 256 maximum CHPID number is reflected in various control blocks in operating systems, software performance measurement tools, and even some application code. Simply changing the control blocks is not a viable option since this would break too much existing code. The solution provided by the z890 and z990 is in the form of multiple Logical Channel Subsystems (LCSSs). These are implemented in a manner that has little or no impact on existing code.

It is important to note that there is no unique hardware (or feature codes) associated with LCSSs. These are not priced features. These architected functions are implemented in the firmware of the system and are, potentially, exploited by operating systems running in the system.

2.7.1 Logical Channel Subsystems (LCSSs)

Existing code works with single-byte CHPID numbers, producing the limitation of 256 CHPIDs. The architecture provides multiple sets of channel definitions, each with a maximum of 256 channels. Existing operating systems would be associated with one Logical Channel Subsystem (LCSS) and work with a maximum of 256 CHPIDs. Different logical partitions can be associated with different Logical Channel Subsystem definitions. Thus a single operating system instance (using existing code) still has a maximum of 256 CHPIDs, but the system as a whole can have more than 256 CHPIDs. It is necessary to have multiple operating images (in multiple logical partitions) to exploit more than 256 channels, but this is a common mode of operation.

In a sense, LCSSs virtualize CHPID numbers. A CHPID no longer directly corresponds to a hardware channel, and CHPID numbers may be arbitrarily assigned. A hardware channel is now identified by a PCHID, or *physical* channel identifier. A PCHID number is defined for each potential channel interface. An I/O feature card has up to 16 channel interfaces[9] and 16 PCHID numbers are reserved for each I/O feature slot in an I/O cage. Not all I/O features provide 16 channels, of course, but 16 PCHID numbers are allocated to each I/O slot. The PCHID numbers allocated to each I/O feature and port on that feature are fixed[10] and cannot be changed by the user.

The z890 can use two LCSSs. z/OS version 1.5 (and later) support the use of more than one LCSS and more than 15 logical partitions. z/VM version 4.4 (and later) provides the same support. More complete exploitation of the architecture is expected to be delivered over time, in evolutionary steps, by the operating systems and various software subsystems.

A given logical partition is associated with a single LCSS, and a single LCSS has a maximum of 256 CHPIDs. Multiple logical partitions may be associated with a given LCSS. Using two LCSSs means that up to 512 CHPIDs can be used.[11] This number is reduced by *spanned* channels and these are discussed in “Spanned channels” on page 45.

Figure 2-11 on page 44 illustrates the relationship between logical partitions (LPs), LCSSs, CHPIDs, and PCHIDs. This is an idealized illustration and ignores such complexities as spanned channels, dynamic I/O changes, and so forth. This illustration also includes the I/O devices and switches used in a later IOCP example.

[9] No existing I/O feature provides 16 usable interfaces (channels). An ESCON feature has 16 interfaces but only 15 may be used; the 16th is a spare.
[10] There is a minor exception for ESCON features when a spare port replaces a failing port.
[11] This is a theoretical limitation for the z890, since the maximum number of physical channels possible is 420.

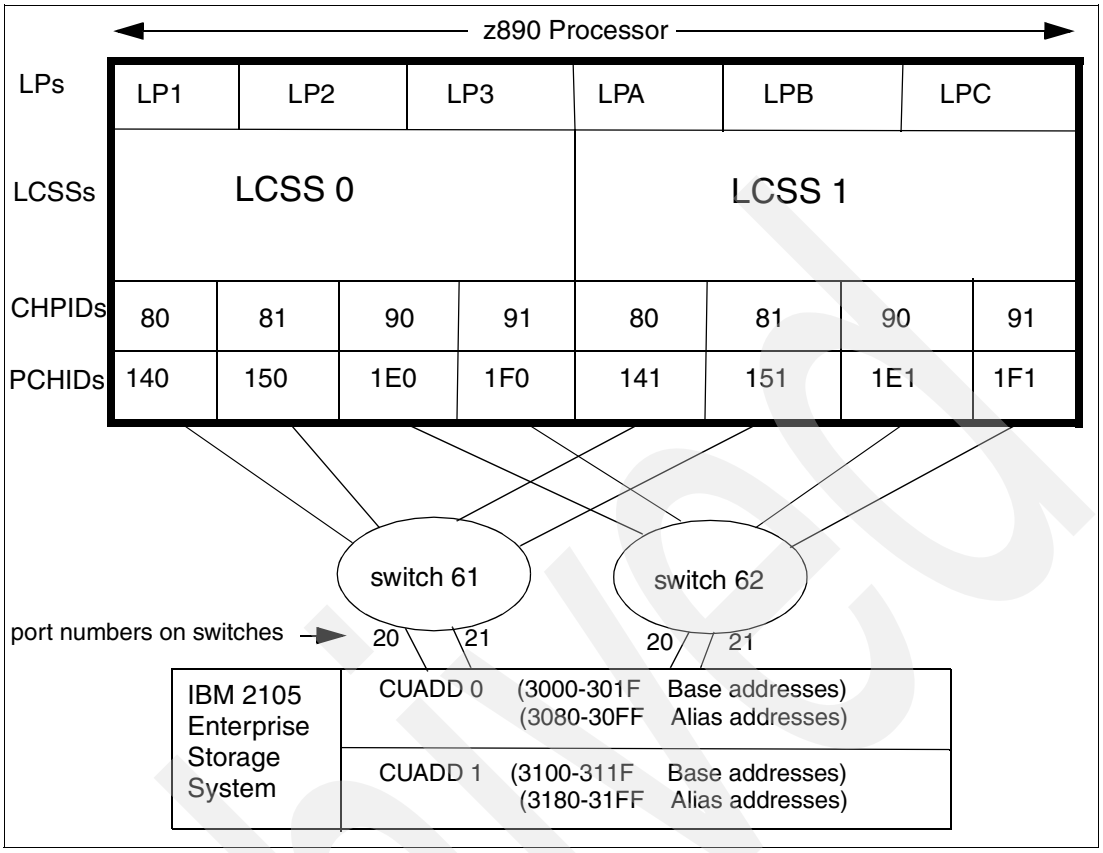

Figure 2-11 Overview of LPs, LCSSs, CHPIDs, and PCHIDs

This figure illustrates a number of important points, such as:

- Two LCSSs are shown.
- A logical partition is associated with a specific LCSS. (Also note that logical partitions have unique names across the complete system.)
- Multiple logical partitions (up to 15) may be associated with an LCSS.
- A CHPID is associated with a specific LCSS. CHPID numbers are unique within that LCSS, but may be reused in other LCSSs. (For example, there is a CHPID 80 in both LCSSs.) A CHPID number is arbitrarily selected. For example, we could change CHPID 80 (in either or both LCSSs in the illustration) to C3 simply by changing a value in the IOCDS.
- A CHPID is associated with a PCHID, and PCHID numbers are unique across the system.
- Different channels on a single I/O feature can be used by different logical partitions. In the illustration, PCHID 140 is the first channel on the feature in I/O cage 1, slot 6. PCHID 141 is the second channel on the same feature.

LCSS-related numbers

Table 2-7 lists LCSS-related information in terms of maximum values for devices, subchannels, logical partitions, and CHPIDs.

Table 2-7 z890 LCSS at a glance

	z890
Maximum number of LCSSs	2
Maximum number of CHPIDs	512
Max. number of LPs supported per LCSS	15

	z890
Max. number of LPs supported per system	30
Maximum number of HSA subchannels	1890K (63 K per partition * 30 partitions)
Maximum number of devices	126 K (2 LCSSs * 63 K devices)
Maximum number of CHPIDs per LCSS	256
Maximum number of CHPIDs per logical partition	256
Maximum number of devices/subchannels per logical partition	63 K

With two Logical Channel SubSystems come more subchannels. There has been a 63K subchannel limitation. Since there was one CSS, there was a maximum of 63K subchannels. With two Logical Channel SubSystems, each LCSS can have its own set of up to 63K subchannels. With two Logical Channel SubSystems you can have:

- Up to a maximum of 63 K devices/subchannels per LCSS
- Up to a maximum of 126 K devices for two LCSSs
  - Two LCSSs x 63k subchannels for each LCSS

Each logical partition can access the 63 K devices in its assigned LCSS. This capability helps to relieve the I/O device configuration constraints experienced by large server configurations.

2.7.2 Spanned channels

When Logical Channel SubSystems (LCSSs) were introduced, transparent sharing of internal channels was introduced—sharing of HiperSockets and Internal Coupling Channels (ICs) between LCSSs. Support on zSeries 890 includes support for sharing of internal channels (HiperSockets and ICs), as well as sharing of external channels—FICON Express (FC and FCP), ICB-3, ICB-4, ISC-3 (CFP and CFS), and OSA-Express features; they can all now be configured as Multiple Image Facility (MIF) spanning channels, allowing sharing of channel resources across Logical Partitions. Spanned channels can be shared among LPs across LCSSs. Spanned channels can be configured to multiple channel subsystems and be transparently shared by any or all of the configured logical partitions without regard to the LCSS to which the LP is configured. The following may be spanned.

Table 2-8 Spanned and shared channels

Channel type		CHPID definition	MIF shared channel	MIF spanned channel
ESCON	External	CNC, CTC	Yes	No
		CVC, CBY	No	No
FICON Express	External	FC, FCP	Yes	Yes
		FCV	Yes	No
OSA-Express	External	OSC, OSD, OSE	Yes	Yes
ICB-4	External	CBP	Yes	Yes
ICB-3	External	CBP	Yes	Yes

Channel type		CHPID definition	MIF shared channel	MIF spanned channel
ISC-3	External	CFP	Yes	Yes
		CFS	Yes	Yes
		CFR	No	No
IC	Internal	ICP	Yes	Yes
HiperSockets	Internal	IQD	Yes	Yes

For the minimum software requirements for all of the spanned channel types identified, refer to the publication, z/OS Migration (GA22-7499), for more details. An update to this publication, discussing the external spanned channel types, is planned to be available in the z/OS V1.6 timeframe.

A spanned channel occupies the same CHPID number in all LCSSs in which it is used. For example, if a HiperSocket channel is CHPID 2C in LCSS 0, it must be the same CHPID number in LCSS 1 (if LCSS 1 is used, of course, and if that HiperSocket is also defined in LCSS1). A HiperSocket that connects logical partitions in different LCSSs *must* be spanned (HiperSockets CHPID type is IQD, QDIO mode only).

Each of these channel types requires that a CHPID be defined, even if it is an internal channel and no physical hardware (channel feature) exists. Each channel, whether a “real” channel device or a virtual device (such as a HiperSocket) must be assigned a unique CHPID within the LCSS. There are no default CHPID numbers for the z890 and you can arbitrarily assign whatever number you like (within the X’00’ to X’FF’ range, of course).

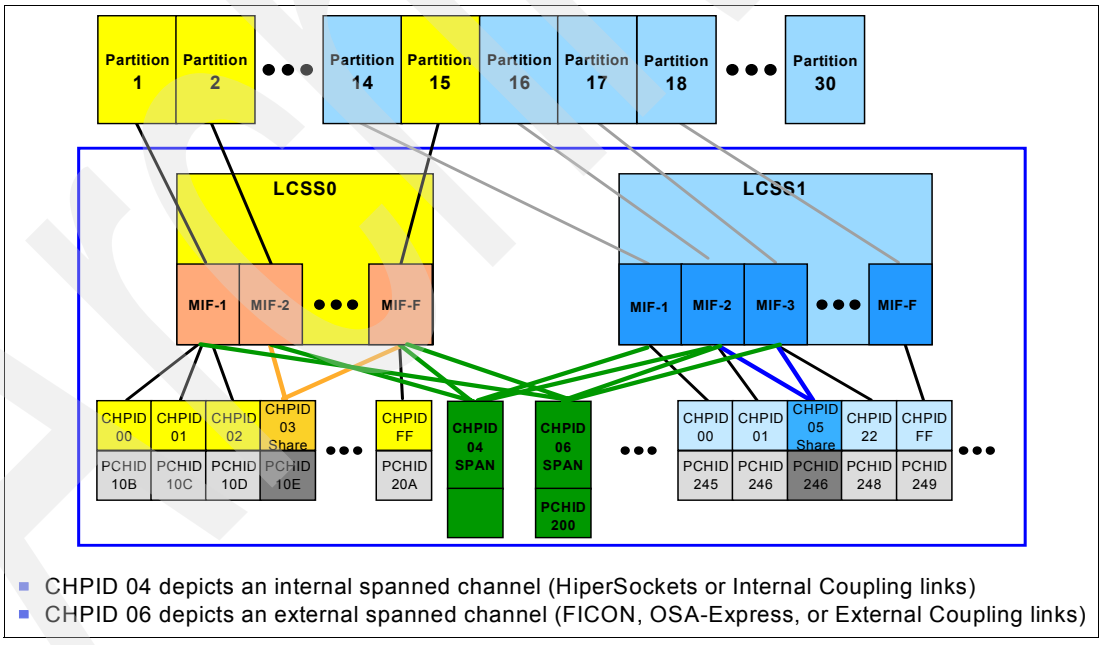

Figure 2-12 Summary of spanned, MIF’ed, shared and reconfigurable channels

2.8 Coupling links

ICs, ICBs, and ISC-3s on z890: Like its predecessor, z890 supports Internal Coupling Channels (ICs), Integrated Cluster Bus (ICB), and InterSystem Channel-3 (ISC-3) for passing

information back and forth in a Parallel Sysplex environment. These technologies are all members of the family of Coupling connectivity options available on z890.

The type of CF links you can use to connect a CF to an operating system logical partition is important because of the impact of the link performance on response times and coupling overheads. The types of links that are available to connect an operating system logical partition to a CF on a z890 are:

- *ICs*: Microcode-defined links to connect a CF to a z/OS logical partition in the same z890 processor. IC links require two CHPIDs to be defined and can only be defined in peer mode. The link bandwidth is greater than 2.0 GB/sec. ICs (CHPID type ICP) can be defined as spanned channels and can be shared among logical partitions within and across LCSSs.
- *ICB-4s*: Copper links available to connect two z890, two z990, or a combination of both; the maximum distance between the two processors is 7 meters (maximum cable length is 10 meters). The link bandwidth is 2.0 GB/sec. ICB-4 links can only be defined in Peer mode. The maximum number of ICB-4 links is 8. ICB-4s (CHPID type CBP) can be defined as spanned channels and can be shared among logical partitions within and across LCSSs.
- *ICB-3s*: Copper links available to connect a z890 to z890, z990, z900 or z800 processors; the maximum distance between the two processors is 7 meters (maximum cable length is 10 meters). The link bandwidth is 1 GB/sec. ICB-3 links can only be defined in Peer mode. Maximum number of ICB3 links is 16. ICB-3s (CHPID type CBP) can be defined as spanned channels and can be shared among logical partitions within and across LCSSs.
- *ISC-3s (Peer mode)*: Fiber links defined in peer mode available to connect z890 to z890, z990, z900 or z800 processors. The maximum distance is 10 km, 20 km with an RPQ, or 40 km with Dense Wave® Division Multiplexing (DWDM) connections. Link bandwidth is 200MB/sec (for distances up to 10 km) and 100MB/sec for greater distances (RPQ). The maximum number of ISC-3 links in peer mode is 48. If the ISC-3 link is connected to another zSeries server and defined as a peer (CFP) mode channel, the link operates in peer mode, and the link is capable of 2 Gbps. Peer mode is used between zSeries servers only. ISC-3s (CHPID type CFP - peer) can be defined as spanned channels and can be shared among logical partitions within and across LCSSs.
- *ISC-3s (Compatibility mode)*: Fiber links defined in compatibility mode are used to connect to G5/G6 systems. The maximum distance is 10 km, 20 km (RPQ), or 40 km with DWDM connections. These operate on single mode fiber at 100 MB/sec or multimode fiber at 50 MB/sec. The maximum number of ISC-3 links in compatibility mode is 32. If the link is connected to a coupling-capable server (G5/G6) which is not a zSeries and is defined as a sender/receiver (CHPID types CFS/CFR) channel, the link operates in compatibility mode, and the link is capable of 1 Gbps. Compatibility mode is used between a zSeries server and 9672 coupling-capable servers. ISC-3s (CHPID type CFS - Coupling Sender, compatibility mode) can be defined as spanned channels and can be shared among logical partitions within and across LCSSs.

Note: The number of combined ICs, ISC-3s, ICB-3s, and ICB-4s CHPIDs cannot exceed 64 per server.

There are several advantages in using peer mode links. All peer mode links operate on a higher performance than the equivalent non-peer link. A single CHPID (one side of the link) can be both sender and receiver; this means that a single CHPID can be shared between multiple OS logical partitions and one CF logical partition. The number of subchannels defined when peer mode links are used is seven (7) per logical partition per link compared to two (2) per logical partition per link in compatibility mode. This is particularly important with System-Managed CF Structure Duplexing. Peer links have 224 KB data buffer space

compared to 8 KB on non-peer links; this is especially important for long distances as it reduces the handshaking for large data transfers.

2.8.1 Coupling Link connectivity with other servers

A variety of Coupling Link connectivity is available on the z890. The z890 supports three types of coupling channels, opening the way to connect to all zSeries servers (z800, z900, z890, and z900), and Generation 5, and 6 servers. z890 supports:

- ISC-3 links
  - ISC-3 links make peer mode connections to z890, z990, z800, and z900 servers with a link bandwidth of 2 Gbps (200 MB/s).
  - ISC-3 links connect to Generation 5, and 6 servers in compatibility mode with a link bandwidth of 1 Gbps (100 MB/s).
- ICB-3 links

  ICB-3 links make 1 GBps (1000 MB/s) peer mode connections to z890, z990, z800 and z900 servers.
- ICB-4 links

  ICB-4 links make 2 GBps (2000 MB/s) peer mode connections to z890 and z990.

ICB-4 links are by far the higher bandwidth links and are preferred wherever the opportunity exist since they save an I/O slot and offer the most efficient coupling link.

Figure 2-13 shows all possibilities.

Connectivity Options	z890 ISC-3	z890 ICB-3	z890 ICB-4
G5/G6 ISC	1 Gbps Compatibility	n/a	n/a
z800/z900 ISC-3	2 Gbps Peer Mode	n/a	n/a
z890/z990 ISC-3	2 Gbps Peer Mode	n/a	n/a
G5/G6 ICB	n/a	n/a	n/a
z900 ICB-2	n/a	n/a	n/a
z990 ICB-2	n/a	n/a	n/a
z800/z900 ICB-3	n/a	1 GBps Peer Mode	n/a
z890/z990 ICB-3	n/a	1 GBps Peer Mode	n/a
z890/z990 ICB-4	n/a	n/a	2 GBps Peer Mode

Figure 2-13 z890 coupling link connectivity with other servers

2.9 OSA-Express features

The OSA-Express features available for the z890 are the same as for the z990. The OSA-Express Fast Ethernet feature has been replaced by the OSA-Express 1000BASE-T feature, although an existing Fast Ethernet feature can be used (for example, upgrading from z800 with a Fast Ethernet feature installed). The older OSA-2 features are not available on a z890. A maximum of 20 OSA-Express cards per system may be used (up to 12 OSA-Express cards on capacity setting model 110). Each card provides two ports, each port has a PCHID identifier, and each port requires a CHPID number in order to be used. This differs from the earlier OSA-2 features, where one CHPID was used by the feature and three more were blocked. No CHPIDs are blocked by OSA-Express cards.

The OSA-Express features can generally run at *line speeds*. This was not the case for earlier LAN interfaces, and is a major improvement in capability. It has been common practice to use ESCON channels to connect to external LAN routers in order to reach the nominal line speeds for LAN connections. With the OSA Express cards, this is not required, and this provides a number of advantages:

- ESCON-attached routers are expensive. Eliminating them provides a direct savings.
- External routers can be complex and often require unique personnel training (or contract services). Eliminating them can remove the need for an expensive skill, and remove a potential failure point in your system.
- External routers can require another level of IP routing. Eliminating this simplifies your logical routing structure.

Note that the 1000BASE-T, Fast Ethernet, and Token Ring ports connect to copper cables, while the Gigabit Ethernet ports connect to fiber cables.

The OSA-Express cards (and the FICON Express card) have indicator lights that may be useful during installation and for problem determination.

Note that the OSA-Express 155 ATM and the FDDI feature is not supported on the z890.

2.9.1 OSA-Express Fast Ethernet

This feature can be carried forward from a z800 system, but cannot be ordered as a new feature for the z890. Instead, new orders will use the 1000BASE-T feature.

The OSA Express Fast Ethernet card has two independent ports (channels) and uses traditional copper wiring with an RJ-45 connector. The ports provide 10 or 100 Mbps Ethernet, with auto-negotiation of the speed used. A QDIO interface[12] may be used for TCP/IP. A non-QDIO interface may be used for SNA (including APPN and HPR) and TCP/IP. By default, the feature automatically adapts to 10 or 100 Mbps operation, and to half or full duplex operation. You can set these options (to avoid the automatic selection) using Support Element panels or the OSA/SF program. Automatic operation can be a problem when used with a very lightly-loaded network with not enough activity for the feature to properly sense.

2.9.2 OSA-Express 1000BASE-T Ethernet

This feature has two independent ports (channels) and provides gigabit Ethernet (as well as 100 Mbps and 10 Mbps Ethernet) over copper cables. It replaces the OSA-Express Fast Ethernet feature, and complements the OSA-Express Gigabit Ethernet feature (which uses

[12] This is an alternative to the standard I/O interface using SSCH commands with traditional CCWs. QDIO functions have been available for some time with the z900 machines, but may not be familiar to installations with earlier systems. QDIO protocols are especially efficient for LAN interfaces. Unique operating system code is required to use QDIO, and this support exists in z/OS and other operating systems.

fiber cables). A port on the 1000BASE-T Ethernet feature can be configured as CHPID type OSC, OSD, or OSE.

The OSA Express Fast Ethernet card uses traditional copper wiring with an RJ-45 connector. The ports provide 10, 100, or 1000 Mbps Ethernet, with auto-negotiation of the speed used. A QDIO interface may be used for TCP/IP. A non-QDIO interface may be used for SNA (including APPN and HPR) and TCP/IP. By default, the feature automatically adapts to 10, 100, or 1000 Mbps operation, and to half or full duplex operation.

Integrated Console Controller function

The z890 has a new function for the OSA-Express 1000BASE-T Ethernet feature and a new Channel Path Identifier (CHPID) type, OSC. The OSA-Express Integrated Console Controller (OSA-ICC) function supports TN3270E (RFC 2355) and non-SNA DFT 3270 emulation.

Now, 3270 emulation for console session connections is integrated in the zSeries 890 via a port on the OSA-Express 1000BASE-T Ethernet feature. This can eliminate the requirement for external console controllers (2074, 3174), helping to reduce cost and complexity.

Each port can support up to 120 console session connections. OSA-ICC support on zSeries is exclusive to z890 and z990 and is available via the OSA-Express 1000BASE-T Ethernet feature, only. OSA-ICC can be configured on a port-by-port basis.

Figure 2-14 shows an example of high availability of the OSA-Integrated Console Controller in a single system configuration.

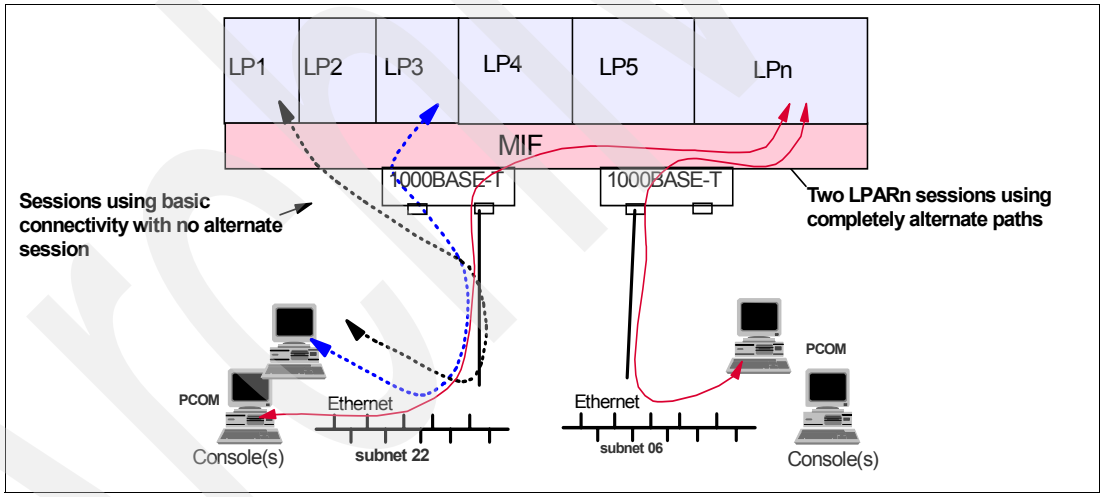

Figure 2-14 Example of high availability of OSA-ICC in a single system configuration

Figure 2-15 on page 51 compares the OSA-Integrated Console Controller function with the IBM 2074-model 003 Console Support Controller and HMC Integrated Console.

	OSA-ICC	2074-003	HMC Integrated Console Level 1.8.0 or higher
z/OS, OS/390	z/OS V1.3, z/OS.e V1.3 and higher	All, even back to MVS/ESA	No
Linux	No	No	Linux LPAR only, using ASCII
z/VM, VM/ESA	z/VM V4.4 and higher	All	z/VM 4.4 and higher LPAR only, using 3270
VSE/ESA	VSE/ESA V2.6 and higher	All	No
TPF	TPF 4.1	All	No
# sessions	120 per port. Multiple sessions to multiple LPARs on single CEC. Up to 48 OSA-Express ports per z890 and z990	48 per ESCON, 96 maximum with two ESCONs. Multiple sessions to multiple LPARs and multiple CECs	HMC can simultaneously support one session on each of multiple LPARs. HMC can switch a session to another HMC. HMC can control up to 100 Support Elements (SE). A SE can be controlled by up to 32 HMCs.
Spanned channel	Yes	No	No
User connection	LAN	LAN	LAN
Remote access	Yes, unsecured	Yes, unsecured	Yes, secured
Remote configuration	Yes, secure via HMC	No, local access only	Yes, secure via HMC
3270 device	Yes	Yes	Yes
ESCON to host	No, OSA-Express	Yes	No
Ethernet to user	10/100/1000 Mbps	10/100/1000 Mbps	10/100 Mbps
Token-Ring	No	4/16 Mbps	4/16 Mbps
S/390 G5/G6	No	Yes	Yes, driver 26. ECF99918
z800/z900	No	Yes	Yes, driver 3G. EC J11219
z890/z990	Yes, May, 2004	Yes	Yes

Figure 2-15 Comparison between OSA-ICC, 2074 and HMC Integrated Console

2.9.3 OSA-Express Gigabit Ethernet

There are two generations of this feature. The older generation may be carried forward from a z800 system. The new generation is provided for new orders. The primary difference is the connectors. The older generation uses SC Duplex fiber connectors and the newer generation uses LC Duplex fiber connectors.

These features have two independent ports (channels). Each feature is available in two versions. The LX (long wave length) version uses single mode fiber with an SC or LC Duplex connector. The SX (short wave length) version uses multimode fiber with an SC or LC Duplex connector. The LX version can use multimode fiber for shorter distances if MCP cables are added at each end of the multimode cable.

This feature uses only QDIO, and is used only for TCP/IP. That is, SNA is not supported through the Gigabit Ethernet card. However, Enterprise Extender support can be used to send SNA traffic over IP.[13] QDIO mode on Gigabit Ethernet on OS/390® or z/OS requires Release 7 or later of Communications Server. Gigabit Ethernet always operates in full duplex mode.

2.9.4 OSA-Express Token Ring

This feature has two ports (channels) and each operates independently at 4, 16, or 100 Mbps. The ports automatically adjust to the correct speed using auto-sense and auto-negotiation functions. It uses traditional copper wire connections with an RJ-45 connector. Each feature port operates in either half or full duplex, depending on the speed. This feature operates in both QDIO mode (TCP/IP traffic only) and non-QDIO mode (TCP/IP and SNA, APPN, HPR). Implementing QDIO on an OSA-Express token ring card requires Release 10 (or later) of Communications Server for OS/390. QDIO mode for Token Ring on

[13] This requires matching Enterprise Extender software at the “other end” of the connection, of course.

Linux requires Linux kernel V2.4 or later. This card replaces the OSA-2 card used in earlier systems and also provides the higher speed option.

Note: OSA migration planning information is found in “OSA-Express migration” on page 101.

2.10 ESCON channels

On the z890, ESCON channels are always delivered with the 16 channel I/O cards. The older 4-port cards cannot be used. The channels are packaged with 16 ports on a single I/O card. Up to 15 channels on each card are available for use; the last channel is reserved as a spare. In practice, any unallocated ports on the card can act as a spare.

A minimum of two ESCON channel cards are always installed if any ESCON channels are configured for the processor. ESCON channels are ordered in groups of four. For configurations greater than 28 ESCON channels, individual (not pairs) of ESCON channel cards are added as necessary. The active channels are distributed across the physical cards to provide additional redundancy.

If one of the activated ports fails, the system performs a *call home* to inform IBM. An IBM Service Representative will initiate the repair by selecting the “Perform a Repair Action” icon at the Service task panel on the SE. This will start the Repair&Verify procedure.

- If sparing can take place, then the IBM Service Representative moves the external fiber optic cable from the failing port to a spare or unconfigured port on the same card.
- If sparing cannot be performed, the card will be marked for replacement by the procedure. Upon replacement of the ESCON card, the cables that were changed are installed at the *original* port locations. Repair&Verify will recognize the unused ports on the new card as candidates for future sparing.

These ESCON cards, which are also used with z990, z900 and z800 machines, use the small MT-RJ connectors. These are different from the traditional ESCON connections that are familiar to most S/390 owners. You can use an ESCON cable with an MT-RJ connector on one end (for the channel connection) and a traditional ESCON connector on the other end (for the control unit). Or you can use conversion cables.

The conversion cables are short cables with an MT-RJ connector on one end and a duplex connector for traditional ESCON cables on the other end. This permits you to use your existing ESCON cables.

The converter cables are relatively fragile and should only be used for converting large numbers of channels if a suitable housing/cabinet is placed in the floor to anchor the connections. An optional wiring harness provides a block of quick-disconnect junctions for connecting groups of conversion cables to existing ESCON cables.

Note: ESCON conversion channel planning information is found in “Consideration for ESCON conversion channels” on page 100.

2.10.1 ESCON directors and multiple LCSSs

An ESCON channel may belong only to a single LCSS. What happens when ESCON Directors are involved, as illustrated in Figure 2-16 on page 53.

In this example we have two channel subsystems. Each LCSS might have several LPs associated with it, but no logical partition is associated with more than one LCSS. So, how does this situation appear to the ESCON director and the two disk sets? For all practical purposes, this configuration is treated as if it were two separate zSeries machines.

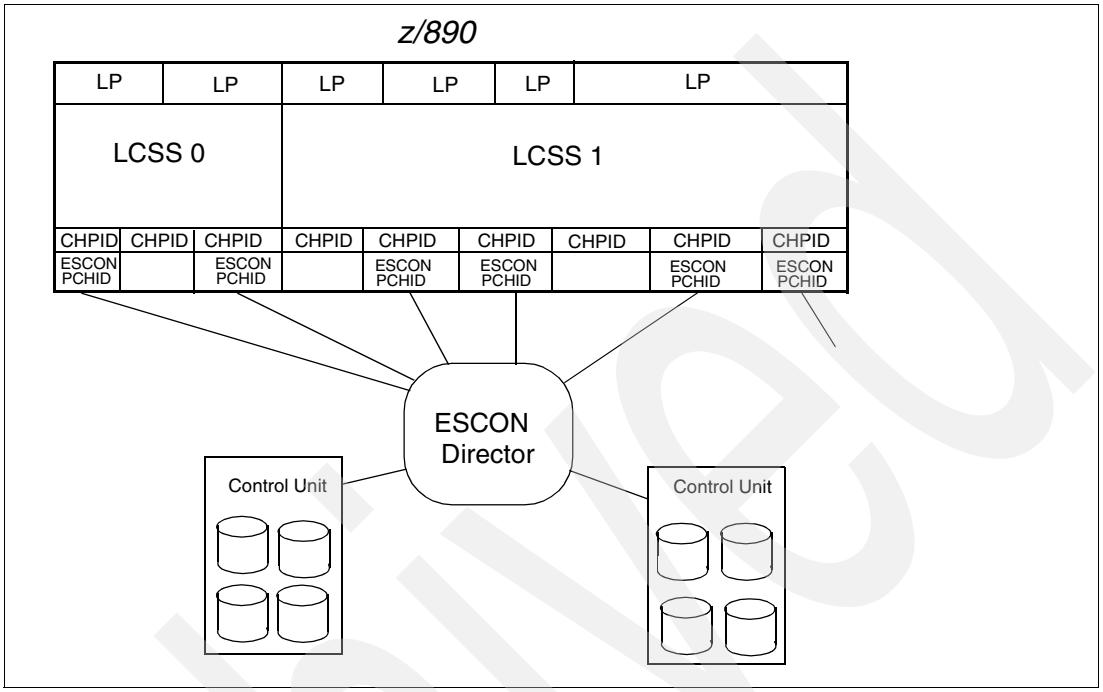

Figure 2-16 DASD connected to multiple LCSSs via ESCON Directors

2.10.2 Fiber Quick Connect (FQC) for ESCON "Quick Connect"

The Fiber Quick Connect (FQC) features are optional features for factory installation of the IBM Fiber Transport System (FTS) fiber harnesses for connection to ESCON channels in the I/O cage. Each direct-attach fiber harness connects to six ESCON channels at one end and one coupler in a Multi-Terminated Push-On Connector (MTP) coupler bracket at the opposite end. When ordered, the features support all of the installed ESCON features in all of the installed I/O cages. FQC cannot be ordered on a partial or one cage basis.

FQC supports all of the ESCON channels in the I/O cage. FQC cannot be ordered for selected channels.

2.11 FICON and FCP channels

The FICON Express features available for the z890 are the same as for the z990. In this section, we discuss the connectivity options provided with z890 for FICON and FCP environments.

2.11.1 FICON Express features

There are two types of FICON channel transceivers supported on the z890, a long wavelength (LX) laser version, and a short wavelength (SX) laser version, as explained:

- z890 FICON Express LX feature with two ports per feature, supporting LC Duplex connectors

- z890 FICON Express SX feature with two ports per feature, supporting LC Duplex connectors

Up to 40 FICON channels (20 features) can be installed in the z890. The smallest z890 sub uniprocessor (capacity setting model 110) can have up to 32 FICON channels (16 features).

2.11.2 FICON Express LX feature

The z890 FICON Express LX feature occupies one I/O slot in the z890 I/O cage. The feature has two Peripheral Component Interconnect (PCI) cards. Each PCI card has a single port supporting an LC duplex connector, with one PCHID associated to each port, and supports link bandwidths of 1 Gbps or 2 Gbps.

Each port supports attachment to the following:

- FICON LX Bridge one port feature of an IBM 9032 ESCON Director at 1Gbps *only*
- Fibre Channel Switch that supports 1Gbps/2Gbps Fibre Channel/FICON LX
- Control unit that supports 1Gbps/2Gbps Fibre Channel/FICON LX
- FICON channel in Fibre Channel Protocol (FCP) mode

Each port of the z890 FICON Express LX feature uses a 1300 nanometer (nm) fiber bandwidth transceiver. The port supports connection to a 9 micron single-mode fiber optic cable terminated with an LC Duplex connector.

Note: Mode Conditioning Patch (MCP) cables are for use with 1 Gbit/sec (100 MB/sec) links *only*.

Multimode (62.5 or 50 micron) fiber optic cable may be used with the z890 FICON Express LX feature for 1 Gbps *only*. The use of this multimode cable type requires a Mode Conditioning Patch (MCP) cable to be used at each end of the fiber optic link, or at each optical port in the link. Use of the single-mode to multimode MCP cables reduces the supported optical distance of the 1 Gbps link to an end-to-end maximum of 550 meters.

Fiber optic conversion kits and Mode Conditioning Patch (MCP) cables are not orderable as features on z890. Fiber optic cables, cable planning, labeling, and installation are all customer responsibilities for new z890 installations and upgrades.

IBM Fiber Cabling Services offer total a cable solution service to help with your cable ordering needs, and is highly recommended.

2.11.3 FICON Express SX feature

The z890 FICON Express SX feature occupies one I/O slot in the z890 I/O cage. The feature has two Peripheral Component Interconnect cards. Each PCI card has a single port supporting an LC Duplex connector, with one PCHID associated with each port, and supports link bandwidths of 1Gbps and 2 Gbps.

Each port supports attachment to the following:

- Fibre Channel Switch that supports 1Gbps/2Gbps Fibre Channel/FICON SX
- Control unit that supports 1Gbps/2Gbps Fibre Channel/FICON SX

Each port of the z890 FICON Express SX feature uses an 850 nanometer (nm) fiber bandwidth transceiver. The port supports connection to a 62.5 micron or 50 micron multimode fiber optic cable terminated with an LC Duplex connector.

2.11.4 FICON channel in Fibre Channel Protocol (FCP) mode

When configured for FCP mode, the FICON Express features can access FCP devices either:

- Via a FICON channel in FCP mode through a single Fibre Channel switch or multiple switches to an FCP device
- Via a FICON channel in FCP mode through a single Fibre Channel switch or multiple switches to a Fibre Channel-to-SCSI bridge

Note: z890 FCP channel direct attachment in point-to-point and arbitrated loop topologies is not supported as part of the zSeries FCP enablement.

z890 adapter interruptions enhancement for FCP

The z890 servers, Linux on zSeries, and z/VM V4.4 and higher work together to provide performance improvements by exploiting extensions to the Queued Direct Input/Output (QDIO) architecture. Adapter Interruptions, first added to z/Architecture with HiperSockets, provide an efficient, high-performance technique for I/O interruptions to reduce path lengths and overhead in both the host operating system and the feature FICON Express, when using the FCP CHPID type.

In extending the use of adapter interruptions to FCP channels, the programming overhead to process a traditional I/O interruption is reduced. This benefits FCP support in Linux on zSeries.

Adapter interruptions apply to a z890 FICON Express channel when in FCP mode (FCP CHPID type), which supports attachment of SCSI devices in a Linux on zSeries environment.

z890 FCP SCSI IPL feature enabler (FC 9904)

This optional z890 feature (FC 9904) allows Linux on zSeries operating system IPL from a SCSI or FCP disk. Both IPL of logical partition images and z/VM V4.4 and higher guests are supported.

Using this feature, Linux logical partitions can be started and run completely from SCSI or FCP disks. Further, a standalone dump program can be loaded from such a SCSI or FCP disk in order to dump the contents of a logical partition, and the dump data can be written to this same disk.

z890 FCP Concurrent Patch

FICON, when configured as CHPID type FCP, supports concurrent patch allowing application of new Licensed Internal Code (LIC) without requiring a configuration off/on. This is a zSeries exclusive FCP availability feature, available with FICON Express feature codes 2319 and 2320.

Cascaded FICON Directors

Some time ago, IBM made generally available the FICON Cascaded Director function. This means that a native FICON (FC) channel or a FICON CTC can connect a server to a device or other server via two (same vendor) FICON Directors.

This type of cascaded support is important for disaster recovery and business continuity solutions because it can provide high availability and extended distance connectivity, and (particularly with the implementation of 2 Gb/s Inter Switch Links) has the potential for fiber infrastructure cost savings by reducing the number of channels for interconnecting the two sites.

The following Directors are supported:

- CNT (INRANGE) FC/9000 128-port and 256-port models
- McDATA Intrepid 6064
- McDATA Intrepid 6140
- IBM TotalStorage® Director 2109-M12

FICON Cascaded Directors have the added value of ensuring high integrity connectivity. Transmission data checking, link incidence reporting, and error checking are integral to the FICON architecture, thus providing a true enterprise fabric.

For more information on Cascaded Directors, consult the I/O Connectivity Web page:

http://www.ibm.com/servers/eserver/zseries/connectivity/ficon_cascaded.html

2.12 Cryptographic support

The cryptographic functions that IBM provides is defined by the IBM Common Cryptographic Architecture. These functions, external interfaces, and a set of key management rules which pertain both to the Data Encryption Standard (DES)-based symmetric algorithms and the Public Key Algorithm (PKA) asymmetric algorithms are available through the facilities provided by the z/OS Integrated Cryptographic Service Facility/MVS™ (ICSF). ICSF uses the corresponding hardware features like the Cryptographic Coprocessor Facility (CCF), the PCI Cryptographic Coprocessor (PCICC) and the PCI Cryptographic Accelerator (PCICA). The z890 does not have CCFs and does not use PCICC features.

As mentioned, the CCF feature is not provided with the z890 processors. Some of its functions are initially replaced by the clear key DES, Triple DES (TDES), authentication code (MAC) message authentication and Security Hash Algorithm (SHA-1) Cryptographic-assist instructions available to all processors on a z890. Most of the remaining functions will be replaced by the PCIXCC card when it becomes available. The IBM CP assist for cryptographic function architecture (CPACF/CP Assist) provides a set of cryptographic functions that enhance the performance of the encrypt/decrypt functions of Secure Socket Layers (SSL), Virtual Private Network (VPN), and data storing applications not requiring Federal Information Processing Standards (FIPS) 140-2 level 4 security.

To achieve the required throughput and implement new functions while maintaining balanced usage of server resources, integrated hardware is key. zSeries introduces the Message Security Assist Architecture along with the CP Assist for Cryptographic Function (CPACF) delivering cryptographic support on every Central Processor (CP)—data encryption/decryption with support of Data Encryption Standard (DES), Triple Data Encryption Standard (TDES), and Secure Hash Algorithm-1 (SHA-1). This offers balanced use of server resources and is designed to provide unprecedented scalability—a z890 can have from one to four CPs, depending upon capacity setting, and data rates at up to 1.8X or more faster than the CMOS Cryptographic Coprocessor Facility (CCF). Since these cryptographic functions are implemented in each and every CP, the association of cryptographic functions to specific CPs, as was done with previous generations of zSeries, is eliminated.

The DES and TDES functions use clear key values. The SHA-1 function is shipped enabled. However, DES and TDES functions require enablement of the CPACF function (#3863) for export control. For IBM and customer-written programs, the CPACF for DES, TDES, and SHA-1 functions can be invoked by five new problem state instructions as defined by an

extension to the zSeries architecture. Support is available via the Integrated Cryptographic Service Facility (ICSF).

2.12.1 z990 Cryptographic processors

Three types of cryptographic hardware features are available on z990. The cryptographic features are usable only when explicitly enabled through IBM.

- CP Assist for Cryptographic Function (CPACF)

  The CP Assist for Cryptographic Function feature provides hardware acceleration for DES, TDES, MAC, and SHA-1 cryptographic services. Cryptographic keys must be protected by The application system.

- PCIX Cryptographic Coprocessor (PCIXCC)

  The PCIX Cryptographic Coprocessor provides a replacement for both the PCICC and the CMOS Cryptographic Coprocessor Facility (CCF). The PCIXCC on z890 provides equivalent PCICC functions at higher performance. It also includes functions that were implemented in the CCF. The PCIXCC supports highly secure cryptographic functions, use of secure encrypted key values and user-defined extensions.

- PCI Cryptographic Accelerator (PCICA)

  Secure Web transactions frequently employ the secure Socket Layer (SSL) protocol. The IBM e-business PCI Cryptographic Accelerator offloads your server from compute-intensive public-key cryptographic operations employed in the protocol. This cost-effective solution often enables significantly greater server throughput.

2.13 HiperSockets

HiperSockets provide high-performance “networks in a box.” A z890 has up to 16 internal HiperSocket LANs compared to the four supported on the z800 processor. Each of these uses a special shared internal CHPID that can be accessed by candidate partitions specified in the IOCDS. Full control can be exercised over which partitions can or cannot access each HiperSocket LAN. For example, you may which to designate certain HiperSocket LANs as production and restrict access to them to partitions running production application servers, database servers, firewalls, and so forth. Similarly, you could designate other HiperSocket LANs for test or development purposes only. Access to a HiperSocket LAN is via an IQD CHPID; all partitions connecting to one of these CHPIDs are in effect, sharing the same internal LAN.

The characteristics of HiperSocket LANs include:

- Excellent performance and response times because all operations are through the system memory bus.
- High availability because there are no external parts or connections involved.
- Cost savings, again because no external parts, OSA-Express features or physical network cabling is involved.
- General connectivity for z/OS, Linux on zSeries, and z/VM, with standard CPs or IFLs.
- Easy to install and implement since they operate like a traditional LAN.
- High levels of security:
  - There is no physical network that might be used for unauthorized connections.
  - Only the partitions authorized to access the IQD CHPID associated with that HiperSocket LAN may connect to it.

Because of this inherent security, it is probably not necessary to encrypt the traffic travelling over the HiperSocket LAN.

- Up to 4,096 communications queues across all 16 HiperSocket LANs (the z800 supported 4 HiperSocket LANs and 1,024 queues). Each communication queue may be shared by multiple logical partitions. Each TCP/IP stack within a logical partition must use a different communications queue, although each queue can be shared across multiple logical partitions with MIF and multiple Logical Channel Subsystems with spanning.
- No interference with normal system performance because HiperSocket data flow does not go through the processor L1 or L2 caches.
- Channel access is through QDIO programming.
- Multicasting is supported, but broadcasting is not supported.
- Not used by current TPF systems.

HiperSockets can be configured many ways. Figure 2-17 shows a simplistic arrangement (all in one LCSS, which is not shown), but illustrates key concepts. Up to 16 HiperSocket CHPIDs can be specified, and each one is, in effect, a LAN. In the illustration, two HiperSocket LANs are used. One logical partition is connected to both HiperSocket LANs and to an external LAN. This logical partition could be used as a router.

TCP/IP connections to each LAN are handled as normal TCP/IP connections. Notice that HiperSockets exist only within the system. There are no external connections, not even to another zSeries machine. Any external connection must be through other means, such as OSA-Express features, channel-to-channel connections, XCF connections, and so forth. However, a z/OS partition connected to an OSA-Express feature can act as a HiperSocket Accelerator and provide high-speed routing of traffic between an external LAN and the systems connected to a HiperSocket LAN.

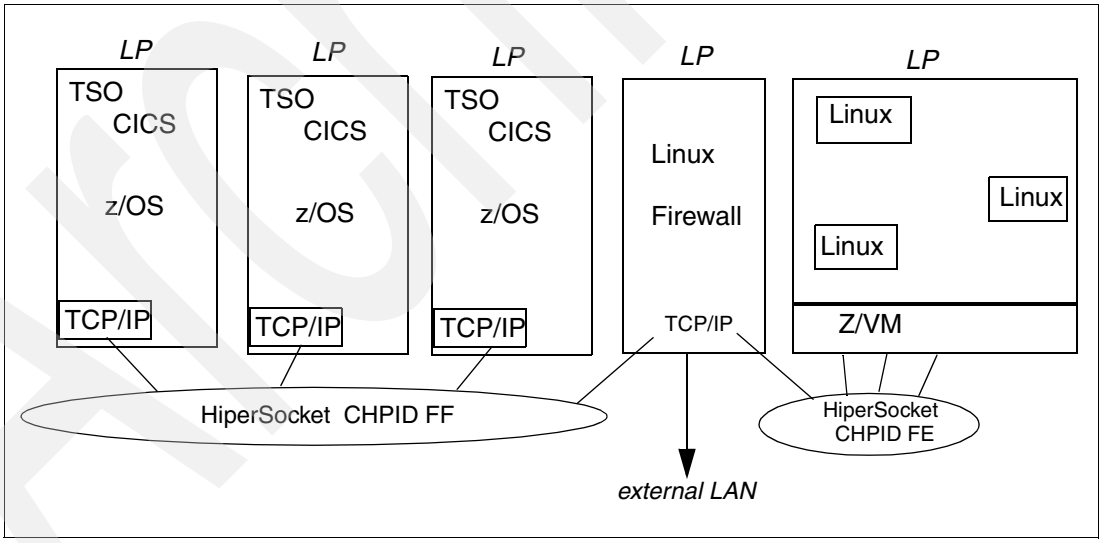

Figure 2-17 Two HiperSocket LANs in a z890 system

HiperSockets are a major addition to IBM's architectures and are expected to become a significant element in future middleware and application designs.

2.14 HMC and SE functionss

Users familiar with the existing S/390 and zSeries Hardware Management Consoles (HMCs) should have no problem in navigating the panels on the z890 HMC. In this section we highlight the most significant changes and enhancements for the z890.

2.14.1 Integrated 3270 console

A new HMC icon provides an emulated 3270. One emulated 3270 is available for each LP. The emulated 3270s do not appear as standard I/O devices. That is, there is no address (device number) associated with them and special programming is needed to address the 3270 sessions.

The Integrated 3270 Console is exploited by z/VM V4.4 and higher, removing the need for a dedicated z/VM system console and associated 2074 control unit, or Integrated Console function of the OSA-Express 1000BASE-T Ethernet feature. The Integrated 3270 Console is exploited by z/VM, the z/VM Standalone Program Loader, and the standalone DASD Dump-Restore program. It can only be used for one console session per z/VM logical partition at a time. Multiple HMCs cannot each simultaneously establish a console session with the same z/VM logical partition.

z/OS currently does not support the Integrated 3270 Console. It cannot be used for system console functions or by VTAM®.

The Integrated 3270 Console icon can be found on the “CPC Recovery” page on the right-hand side of the HMC display. It is invoked in exactly the same way as the HMC Operating System Messages display. For example, a CPC Image icon that represents the target z/VM logical partition can be dragged and dropped onto the “Integrated 3270 Console” icon.

The hardware support for this feature is also integrated into Generation 5, and 6 servers with driver 26 and z800 and z900 processors with driver 3G.

2.14.2 Integrated ASCII console

The Integrated ASCII Console can be exploited by Linux logical partitions (not Linux guests under z/VM) with a planned code drop to the Open Source community. This feature allows a Linux logical partition to establish an ASCII console session on the HMC and should simplify Linux installation and initial operation. It can only be used for one console session per Linux logical partition at a time. Multiple HMCs cannot each establish a console session with the same Linux guest.

The Integrated ASCII console icon can be found on the “CPC Recovery” page on the right-hand side of the HMC display. It is invoked in exactly the same way as for the HMC Operating System Messages display.

The hardware support for this feature is also integrated into Generation 5, and 6 servers with driver 26 and z800 and z900 processors with driver 3G.

2.14.3 Optional Strict password rules

This feature enforces strict password rules defined by the ACSADMIN userid and prompt the HMC user for a new password if the existing one has expired. Password expiration intervals can be set on the ACSADMIN “User Administration” screen.

2.14.4 Customizable HMC data mirroring

This feature allows two or more Hardware Management Consoles to be *associated* by sharing the same user data file. This eliminates the need to copy data between HMCs via diskette. The mirroring support can be customized to specify the data that is to be mirrored.

2.14.5 Extended console logging

The View Console Tasks Performed log size has been increased to show the last 500 actions. Previously this log only showed the last 100 events.

2.14.6 Operating System Messages display

The Operating System Messages display has been enhanced to display the command line on the first page. A check box has been added to indicate a Reply to Priority Message.

Note: How to configure HMCs and SEs is further explained in "SE and HMC connectivity planning" on page 94.

2.15 Comparison table - z890 versus prior servers

The Table 2-9 summarizes some of the differences between 9672 Generation 5, and 6 (G5/G6), Multiprise 3000 (MP3000), z800, z890, and z990 servers. Please note that some maximum limitations are reduced for the smallest z890 sub uniprocessor (capacity setting model 110) and these reduced limitations are not shown in the table.

Table 2-9 Summary of server differences

	MP3000	9672 G5 / G6	z800	z890	z990
Maximum PUs	3	12 / 14	5	5	48
Maximum customer PUs	2	10 / 12	4	4	32
Maximum number CPs	2	10 / 12	4	4	32
Maximum number IFLs	1	7 / 11	4	4	32
Maximum number ICFs	0	7 / 11	4	4	16
Maximum number zAAPs	N/A	N/A	N/A	2	16
Standard SAPs	1	1-2 / 2	1	1	2/book
Cycle time	N/A	2.0-2.6 ns / 1.57-1.8 ns	1.6 ns	1.0 ns	.83 ns
Memory	1-4 GB	5-32 GB	8-32 GB	8-32 GB	16-256 GB
Maximum I/O cages	N/A	3	1	1	3
Support elements	1	2	2	2	2
Maximum STIs	N/A	24@333 MB/s	6@1 GB/s	8@2 GB/s	12@2 GB/s per Book
Maximum logical partitions	15	15	15	30	30
Maximum HiperSockets	0	0	4	16	16

	MP3000	9672 G5 / G6	z800	z890	z990
Maximum FICON Express channels	0	24 / 36 FICON *non*-Express	32	40	120
Maximum OSA-Express ports	0	12	24	40	48
Maximum OSA ports	0	12	0	0	0
Maximum ESCON channels	56	256	240	420	1024
Integrated disks	Yes	No	No	No	No
Internal emulated I/O	Yes	No	No	No	No
Maximum ISC-3 Links	0	32	24	48	48
Maximum ICB-2 Links	0	18	0	0	8
Maximum ICB-3 Links	0	0	5	16	16
Maximum ICB-4 Links	0	0	0	8	16
Standard Crypto Coprocessors	1 Optional	2	2	0	0
Maximum PCICCs	0	8	16	0	0
Maximum PCICAs	0	0	12	4	12
Maximum PCIXCCs	0	0	0	4	4
L1 cache (per processor)	256K	256K	256K I/256K E	256K I/256K E	256K I/256K E
L2 cache	8 MB	8 / 16 MB	8 MB	32 MB	32 MB/book
I/O bandwidth	N/A	8 GB/sec	6 GB/sec	16 GB/sec	24-96 GB/sec

3

Software support

This chapter describes the software requirements and support considerations for the z890 (IBM zSeries 2086 model A04). It covers z/OS, z/OS.e, OS/390, z/VM, z/VSE, VSE/ESA™, TPF, and Linux on zSeries operating systems. Software migration considerations, concurrent upgrade considerations, and in addition workload license charging are discussed as well.

3.1 Operating system support

There are many significant changes in the z890 functions and features when compared to the 9672, z800, and z900 servers. For this reason, software services are available to either let the supported software level run on the z890 hardware, or exploit the new features offered by the z890. The z890 software support uses the same support structure introduced by the z990 server:

- The *Compatibility* support, which allows:
  - To define a z890 environment with HCD.
  - To run on a z890 server in a logical partition in LCSS 0, using an Logical Partition ID equal or less than 15 (x'F').
- The *Exploitation* support, which includes the compatibility support, and also allows:
  - More than 15 logical partitions (except for the z890 capacity setting model 110)
  - Two Logical Channel Subsystems

Extensive software support has been made available to supported operating levels via compatibility and exploitation support to accommodate these changes in the z/OS, z/OS.e, OS/390, z/VM, z/VSE, VSE/ESA, TPF, and Linux on zSeries operating systems. Table 3-1 summarizes supported software on the z890.

Table 3-1 z890 software support summary

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
OS/390® Version 2 Release 10	Yes	Yes	Yes	No
z/OS Version 1 Release 2	No[a]	Yes	Yes	No
z/OS and z/OS.e Version 1 Release 3	No[a]	Yes	Yes	No
z/OS and z/OS.e Version 1 Release 4	No[a]	Yes	Yes	Yes
z/OS and z/OS.e Version 1 Release 5 and 6	No	Yes	Included	Included
Linux for S/390	Yes	No	Yes	Yes
Linux® on zSeries	No	Yes	Yes	Yes
z/VM Version 3 Release 1	Yes	Yes	Yes	No
z/VM™ Version 4 Release 3	Yes	Yes	Yes	No
z/VM Version 4 Release 4	Yes	Yes	Included	Included
z/VM Version 5 Release 1	No	Yes	Included	Included
VSE/ESA™ Version 2 Release 6 and 7	Yes	No	Yes	Yes
z/VSE Version 3 Release 1[b]	Yes	No	Included	Included
TPF Version 4 Release 1 (ESA mode only)	Yes	No	Yes	No

a. 31-bit mode is only available as part of the z/OS Bimodal Migration Accommodation software program. The program is intended to provide fallback support to 31-bit mode in the event that it is required during migration to z/OS in z/Architecture mode (64-bit).

b. The z/VSE operating system can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. The z/VSE operating system is designed to exploit select features of IBM eServer zSeries hardware.

3.2 z/OS, z/OS.e, and OS/390 software support

z/OS, z/OS.e, and OS/390 software support is designed at two levels: *Compatibility* support and *Exploitation* support.

Note: Compatibility and exploitation support for OS/390, z/OS, and z/OS.e on the z890 are included in the z990 software support. So, the z990 compatibility and exploitation support *features* also apply to z890 servers.

Both the compatibility and exploitation support features are included in the z/OS and z/OS.e V1.5, and in the z/OS and z/OS.e V1.6 operating systems.

In addition, the z/OS and z/OS.e V1.6 introduce the support for the zSeries Application Assist Processor (zAAP). See “zSeries Application Assist Processor (zAAP)” on page 76 for more information.

z990 Compatibility for selected OS/390, z/OS, and z/OS.e releases

OS/390 V2.10, z/OS V1.2, and z/OS V1.3 require the Web delivered z990 Compatibility for selected Releases, to allow these releases to run on the z890 server.

This support is also required on all servers in a parallel sysplex that run OS/390 V2.10, z/OS V1.2, z/OS V1.3, and z/OS.e V1.3 when a z/OS or coupling facility image in that same parallel sysplex runs on a z890, or z990, and the LPAR ID of the operating system or coupling facility image is greater than 15 (x’F’). Note that z/OS.e V1.3 cannot run on a z990.

Attention: Compatibility and exploitation support is *not* available for z/OS 1.1.

z/OS V1.4 z990 compatibility support feature

The z/OS V1.4 z990 compatibility support feature is an unpriced optional feature required to to allow z/OS V1.4 to run on the z890 server.

This support is also required on all servers in a parallel sysplex when a z/OS or coupling facility image in that same parallel sysplex runs on a z890 or z990, and the LPAR ID of the operating system or coupling facility image is greater than 15 (x’F’). In addition, it provides support for the PCICA, and CP Assist for Cryptographic Function (CPACF) cryptographic features.

Note: This feature is no longer orderable and is replaced by the z/OS V1.4 z990 Exploitation Support feature. If the feature has been ordered earlier to support a z990 it can be used to run z/OS V1.4 on a z890.

z/OS V1.4 z990 exploitation support feature

The z/OS V1.4 z990 exploitation support feature is an unpriced optional feature that provides support to run z/OS V1.4 on the z890 and to let z/OS V1.4 make use of more than one Logical Channel Subsystem (LCSS) and up to 30 logical partitions.

This support is also required on all servers in a parallel sysplex when a z/OS or coupling facility image in that same parallel sysplex runs on a z890 or z990, and the LPAR ID of the operating system or coupling facility image is greater than 15 (x'F').

This is a mandatory feature when z/OS V1.4 is running on a z890 as of February 24, 2004.

3.2.1 z/OS.e V1.4 z990 coexistence feature

This feature provides support to allow z800 and z890 servers running z/OS.e V1.4 to coexist in a Parallel Sysplex with z890 and z990 servers when a z/OS or coupling facility image in that same parallel sysplex runs on a z890, or z990, and the LPAR ID of the operating system or coupling facility image is greater than 15 (x'F').

z/OS.e V1.4 cannot run on a z990.

Note: This feature is no longer orderable and is replaced by the z/OS V1.4 z990 exploitation support feature.

z/OS.e V1.4 z990 coexistence update feature

This feature provides support for a consistent code base for z/OS and z/OS.e, and provides exploitation support for two Logical Channel Subsystems (LCSSs) and 30 logical partitions. It is required on all servers running z/OS.e V1.4 in parallel sysplex when a z/OS or a coupling facility image in the same parallel sysplex is running on a z890 or z990 server, and the LPAR ID of the operating system or coupling facility image is greater than 15 (x'F'). It replaces the z/OS.e V1.4 z990 Coexistence feature.

This is a mandatory feature when z/OS.e V1.4 is running on a z890 as of February 24, 2004.

z/OS.e V1.4 cannot run on a z990 server.

3.3 OS/390, z/OS, and z/OS.e software considerations

z990 Compatibility for selected releases allows OS/390 V2R10, z/OS, and z/OS.e releases to run on and coexist with z890 and z990 servers. Certain functions of the z890 (particularly, a second LCSS and more than 15 logical partitions) are not supported until the exploitation support is.

Pre- z/OS V1R4 and z/OS.e V1.4 releases and OS/390 V2 R10 can run in compatibility mode.

Note: z/OS 1.1 is not supported with the z890 server.

If you run on OS/390 V2R10, z/OS V1R2, z/OS V1R3, or z/OS.e V1.3, the z990 Compatibility for selected Releases support can be downloaded from IBM Resource Link™:

`http://www.ibm.com/servers/resourcelink`

If z/OS V1R4 or z/OS.e V1.4 is being used, support for exploitation mode should be ordered through the standard software deliverable channels. Either one is required:

- z/OS V1.4 z990 Exploitation Support feature
- z/OS.e V1.4 z990 Coexistence Update feature

New orders for z/OS V1R5 and later releases will have exploitation support included.

Table 3-2 summarizes the OS/390, z/OS, and z/OS.e software support.

Table 3-2 Supported OS/390, z/OS, and z./OS.e releases on z890

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
OS/390® Version 2 Release 10	Yes	Yes	Yes	No
z/OS Version 1 Release 2	No[a]	Yes	Yes	No
z/OS and z/OS.e Version 1 Release 3	No[a]	Yes	Yes	No
z/OS and z/OS.e Version 1 Release 4	No[a]	Yes	Yes	Yes
z/OS and z/OS.e Version 1 Release 5 and 6	No	Yes	Included	Included

a. 31-bit mode is only available as part of the z/OS Bimodal Migration Accommodation software program. The program is intended to provide fallback support to 31-bit mode in the event that it is required during migration to z/OS in z/Architecture mode (64-bit).

3.3.1 Compatibility support

Compatibility support for the z890 is available for OS/390 V2R10, z/OS V1R2, z/OS V1R3 and it is part of z/OS V1R4. This is also true for all corresponding z/OS.e releases. Compatibility support allows these releases of the operating system to do the following:

- Define a z890 environment with HCD.
- Run on a logical partition in LCSS 0.
- Make dynamic I/O changes for LCSS 0 (only).
- Coexist with servers in a Sysplex.
- Coexist with servers sharing DASD (even outside of a sysplex).

These are broad statements which, in themselves, do not cover all variations of all possible environments. It may not be needed to activate all aspects of compatibility support, although it is highly desirable that they are. For example, you only need to install the HCD maintenance on the system that is defining the z890 configuration; it is not needed on the systems where HCD is not used. The precise circumstances under which compatibility support is needed are explored later in this section.

The rules for whether compatibility support is required are:

- Compatibility support is required for *all* images running on a z890.
- Compatibility support is required on *any* image that is used for defining the I/O configuration for the z890.
- Compatibility support is required on *all* images in a sysplex if a Coupling Facility logical partition has a logical partition identifier greater than X'F'. Note that this is not the same as having more than 15 logical partitions. Logical partition identifiers can be assigned arbitrarily; it is possible to have a logical partition identifier greater than X'F' even if it is the only logical partition on the server.

Compatibility support allows the operating systems to run in LCSS 0 on the z890 server. It is not possible to run with z/OS or OS/390 in LCSS 1, even with the compatibility support installed. (Running in LCSS 1 is supported with z/OS V1 R4 and later releases with exploitation support.) A Coupling Facility logical partition can reside in any LCSS.

Compatibility support is supported on any server that is already supported by that operating system. For example, you could install compatibility support on a z/OS V1 R2 system that is running on a 9672 Generation 5 server.

Figure 3-1 illustrates a situation where compatibility support is *not* required.

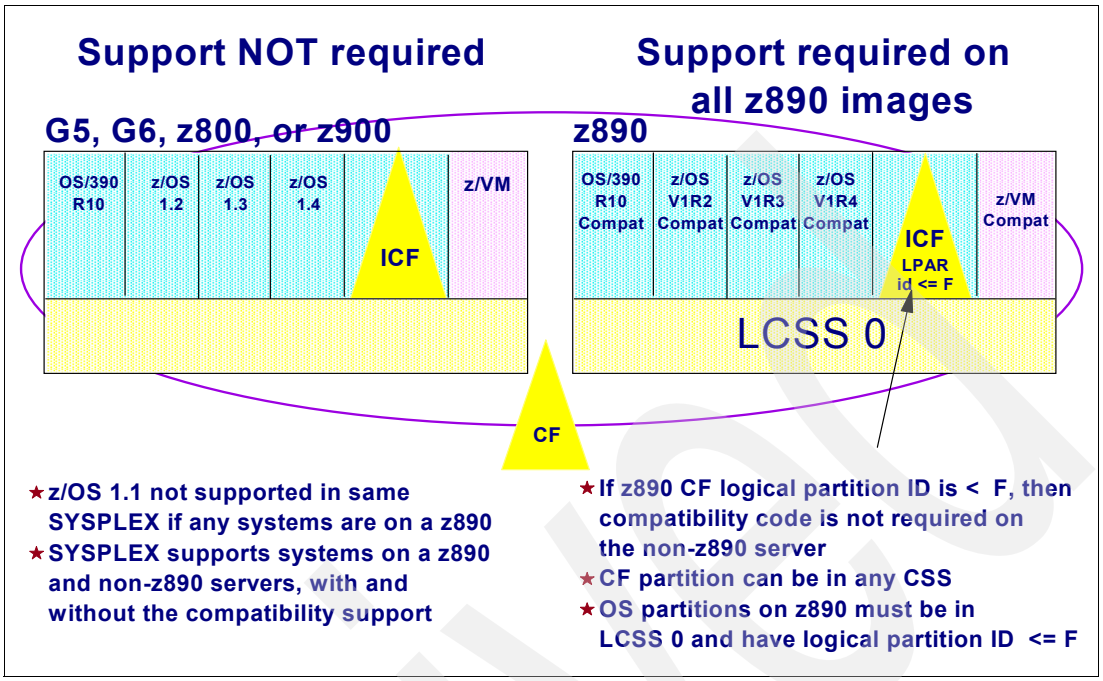

Figure 3-1 Conditions where compatibility maintenance is not required

Figure 3-2 illustrates a situation where compatibility support *is* required.

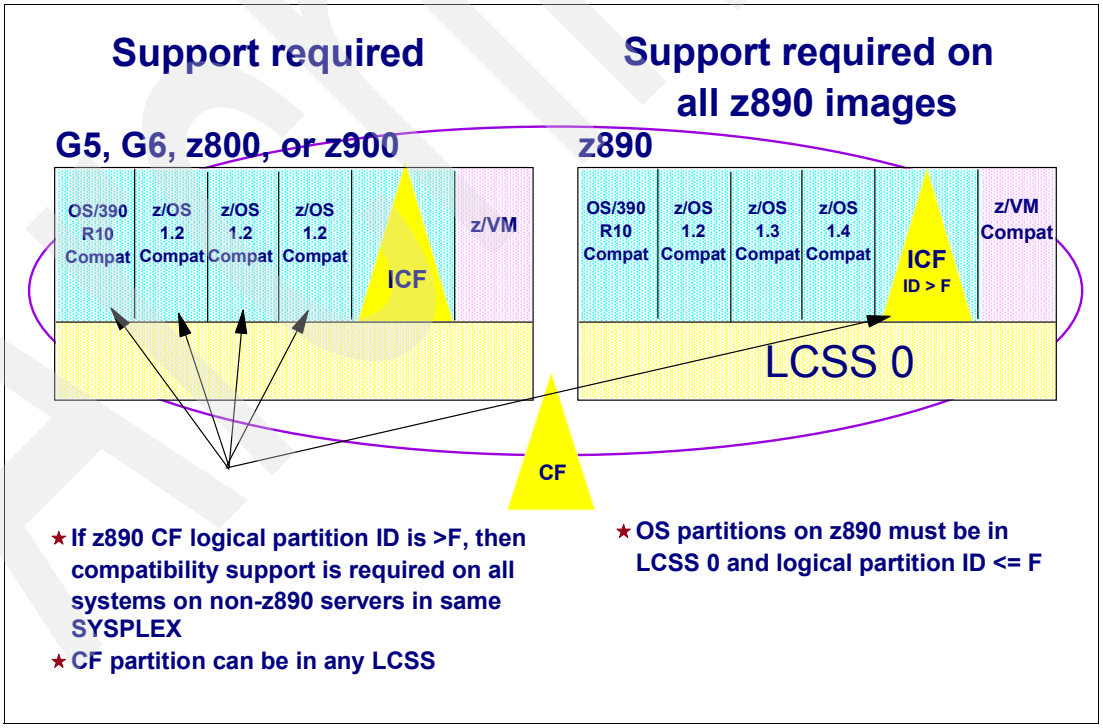

Figure 3-2 Conditions where compatibility maintenance is required

Compatibility support updates several functional areas in order to accommodate the z890 server, as a standalone system and in a multisystem environment.

HCD support

The support added to HCD[14] for both compatibility and exploitation support allows the definition of the z890 server and I/O configuration while working in another (possibly older) system. The new HCD elements apply to z/VM and z/OS.

The updated HCD function introduces a new concept: A *validated work IODF*. This is a new status for an IODF data set. It contains a complete set of validated processor, LCSS, channel, partition, control unit and I/O device definitions. A validated work IODF would not normally contain the physical channel identifiers (PCHIDs) for channels. These can be added to the validated work IODF by the output from the CHPID Mapping Tool (CMT).

Support has been added to HCD to work with the CHPID Mapping Tool for PCHID assignments. This support allows an IOCP source statement data set to be created from a validated work IODF; it also allows the data from the CHPID Mapping Tool to be merged with the validated work IODF to complete the PCHID assignment.

The compatibility support for HCD allows an installation to do the following:

- Define a z890 environment with multiple (up to two) Logical Channel Subsystems.
- Make dynamic hardware changes to LCSS 0 only. Devices cannot be added, modified or deleted if they are also defined to another LCSS. This would require a Power-on Reset. Full dynamic change support is implemented in the exploitation support for HCD.
- Software ACTIVATEs can be done regardless of how many Logical Channel Subsystems are defined.

Definition of the new server environment is discussed at length in “Hardware Configuration Dialog (HCD)” on page 109.

CFRM and CFRM policy support

The CFRM support for the z890 and z990 server now allows two-digit logical partition identifiers to be assigned in the CFRM policy. This extension is essential for allowing more than 15 logical partitions to coexist on the server.

The logical partition ID is *not* the same as the partition number used in prior servers. The logical partition identifier is assigned in the image profile on the Support Element and is unique to each logical partition. It is a two-digit value in the range x’00’ to x’3F’.

Attention: The logical partition ID for a Coupling Facility (specified in its image profile on the Support Element) must correspond with that specified on the PARTITION() keyword in the CFRM. This is different than earlier zSeries and S/390 servers. On previous systems, the value specified on the PARTITION() keyword was the partition number from the IOCDS.

This update is identical for both compatibility and exploitation support. The logical partition ID is also used in IOS and XCF messages.

Automation changes

Several commands and messages have been adapted to accommodate the new two-digit logical partition identifier.

[14] This section discusses compatibility support for HCD. Equivalent compatibility support is also needed for HCM.

SMF support

CPU and PR/SM activity data is recorded by RMF™ in the SMF type 70 subtype 1 records. Prior to the z890 and z990, these records were always shorter than 32 KB. With the increased number of logical partitions and logical processors, these records could potentially increase in size beyond the 32 KB limit. To accommodate this, each record is now broken into pieces where each piece is shorter than 32 KB. Each piece is self-contained; that is, the record can be processed without re-assembling the broken pieces. If you have any site-specific processing of this data outside of RMF, you may need to review that application to ensure that it is no longer dependent upon all this data being contained within a single record.

RMF Monitor 1 Device Activity reporting is recorded in the SMF 74 subtype 1 records. These have been updated to support an extended device data section. This section now includes the Initial Command Response time for the device. A similar change has also been made to the RMF Monitor II Device Activity recording in the SMF 79 subtype 9 records.

The SMF records for the SRM Decisions data stored in type 99 (subtypes 8 and 9) have also been extended. These records now contain the LCSS ID for the WLM LPAR management and I/O subsystem information.

RMF support

There are several changes to RMF reports to accommodate the enhanced I/O subsystem and improved collection of channel measurement data.

The I/O Activity report no longer shows the Control Unit Busy (CUB) and Director Port Busy (DPB) times. (The corresponding percentage and pending reason fields for CUB and DPB have also been removed from the Monitor III reports.) This information was already available at an LCU level and is much more useful than figures broken out for individual devices. Furthermore, the Director Port Busy field would only show a non-zero value for the events when *all* director ports were busy. If an individual director port was found to be busy but a connection was established through an alternate path, then this figure was not updated. With FICON connections, a Director Port was never reported busy since this type of channel allows multiple data transfers to occur simultaneously.

A better measure of fabric contention has now been provided with the Initial Command Response Time (CMR). This is a measure of the time taken from sending a command to a device, to a response that it has accepted the command. This new metric (AVG CMR DLY) has now replaced the older AVG CUB and AVG DPB columns on the Monitor I Device Activity Report. Corresponding DELAY CMR% and PENDING REASON CMR displays have been introduced into the Monitor III reports. Monitor III exception reporting has also been updated to replace the previous CUBDL and DPBDL conditions with the new condition, CMRDL.

ICKDSF requirements

ICKDSF Release 17 is now needed on all systems that share DASD with a z890 server. This applies to z/VM and z/OS systems. The need for this ICKDSF release applies even to systems that are not part of the same sysplex, or that are running a non-MVS-based operating system, such as z/VM.

Compatibility support restrictions

Compatibility support provided does not allow you to make full use of all the capabilities of the z890. Specific restrictions include:

- z/OS must be IPLed in a partition defined in LCSS 0. If it is IPLed in a partition in LCSS 1, it will terminate with a 07C-01 wait state.

- z/OS must be IPLed in a partition that has a logical partition identifier in the range 0-F. The logical partition identifier is specified in the Image profile on the HMC. If the logical partition identifier is outside of this range, then z/OS will terminate with a 07C-02 wait state.
- Dynamic activates for hardware changes can only be done for LCSS 0. A Power-on Reset is required for changes to other Logical Channel Subsystems. Dynamic activates for hardware changes within LCSS 0 cannot be done if the resource is also defined in any other LCSS. For example, if a DASD control unit has connections to LCSS 1, then additional connections cannot be added to LCSS 0 dynamically while in compatibility mode.

3.3.2 Exploitation support

z/OS exploitation support for the z890 is available for z/OS and z/OS.e V1R4 and later releases. All references to z/OS in the following topics refer to z/OS and z/OS.e V1R4 and later releases.

Exploitation support allows:

- z/OS to run in a partition defined to any Logical Channel Subsystem
- z/OS to run in a partition with a logical partition identifier greater than X'F'
- z/OS dynamic activation for hardware changes to any LCSS

The following areas need to be considered when running z/OS with exploitation mode.

SMF

The SMF type 89 record used for recording Product Usage data has been extended for exploitation mode support. A new field, SMF89LP3, allows an 8-bit logical partition ID to be stored. This field is marked valid by the new flag bit SMF89LPM. When a logical partition ID is less than or equal to X'F', the logical partition ID is stored in both the new field and the old 4-bit SMF89LP2 field to maintain compatibility.

Standalone dump

The z/OS systems must generate a new version of the standalone dump program. This standalone dump program cannot be used for dumping systems at earlier releases of z/OS.

Automation

The output from the D M=CPU command has been enhanced to show the two-digit logical partition ID (set in the Image profile), the LCSS associated with the logical PUs associated with the logical partition in that LCSS, and the MIF Image ID (see Example 3-1). The logical PU address no longer appears in the first digit of the serial number as a result of the change to the STIDP instruction.

Example 3-1 Changes to the D M=CPU command output

```
D M=CPU
IEE174I 12.26.42 DISPLAY M 002
PROCESSOR STATUS
ID  CPU                  SERIAL
0    +                   036A3A2086
1    +                   036A3A2086

CPC ND = 002086.220.IBM.02.00000023FADE
CPC SI = 2086.220.IBM.02.000000000023FADE
CPC ID = 00
```

```
CPC NAME = A03
LP NAME = A03        LP ID =  3
CSS ID  = 0
MIF ID  = 3

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE

CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
CSS ID   CHANNEL SUBSYSTEM IDENTIFIER
MIF ID   MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER
```

The output from the D IOS,CONFIG(HSA) or D IOS.CONFIG(ALL) commands has been changed to remove all references to SHARED and UNSHARED control units. Previously, this command would have been used to determine the HSA space available for dynamically adding control units and I/O devices. On the z890 server, the number of additional devices that can be added dynamically is determined by the MAXDEV value associated with each LCSS. This parameter is specified via HCD and is set in the IOCDS. The new output from the D IOS command is shown in Example 3-2.

Example 3-2 Output from the D IOS command on the z890

```
D IOS,CONFIG(HSA)
IOS506I 13.37.12 I/O CONFIG DATA 032
HARDWARE SYSTEM AREA AVAILABLE FOR CONFIGURATION CHANGES
PHYSICAL CONTROL UNITS                 7710
CSS  0 - LOGICAL CONTROL UNITS       3892
             SUBCHANNELS                          55804
CSS  1 - LOGICAL CONTROL UNITS       3981
             SUBCHANNELS                          56262
```

The IEE174I, IOS050I and IOS051I messages have also been changed to display not just the CHPID number, but also the associated PCHID. This addition assists diagnosis of hardware problems. The new output is shown in Example 3-3.

Example 3-3 PCHID support for channel messages

```
D M=CHP(3F)
IEE174I 13.51.30 DISPLAY M 063
CHPID 3F:  TYPE=05, DESC=ESCON SWITCHED POINT TO POINT, ONLINE
DEVICE STATUS FOR CHANNEL PATH 3F
    0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
510 +  +  +  +  .  .  .  .  .  .  .  .  .  .  .  .
SWITCH DEVICE NUMBER = 001C
PHYSICAL CHANNEL ID = 01E0
********************** SYMBOL EXPLANATIONS ***********************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL

IOS050I CHANNEL DETECTED ERROR ON dddd,yy,op,stat,PCHID=pppp

IOS051I INTERFACE TIMEOUT DETECTED ON dddd,yy,op,stat,PCHID=pppp
```

EREP

The PCHID value associated with a particular CHPID is now displayed in the EREP Subchannel Logout record, as shown in Example 3-4.

Example 3-4 EREP support of PCHIDs

```
DEVICE NUMBER:    0A02      REPORT:  SLH EDIT                    DAY YEAR  JOB
                            SCP:     VS 2 REL.  3          DATE: 323  03
 DEVICE TYPE:      CACA
                            CPU MODEL:    2084                  HH MM SS.TH
 CHANNEL PATH ID:  3E        LOGICAL  CPU ID: 232920        TIME: 12 50 29.84
 PHYSICAL CHAN ID: XXXX      PHYSICAL CPU ID: 612920        PHYSICAL CPU ADDRESS: 00

             CC   CA   FL  CT
 FAILING CCW     02 01DA2500 24 5000           VOLUME SERIAL      N/A
                                         SUBCHANNEL ID NUMBER  000108B2
              K  FLAGS   CA    US SS  CT     ERROR TYPE         OTHER
 SCSW         64 C24017 01DA2318 00 02 5000

 ---UNIT STATUS----  SUB-CHANNEL STATUS  ----------------SCSW FLAGS-----------------
                                                FLAG 0     FLAG 1     FLAG 2
```

Extended Channel Measurement Block (ECMBs)

The z890 server supports Extended Channel Measurement Blocks for the new I/O architecture. It also supports the original XA I/O architecture CMBs for compatibility reasons.

Under z/OS with the exploitation support, the ECMBs will now be placed in a system common area dataspace. Because of this, you might want to review your setting of the MAXCAD value in IEASYSxx. The CMB parameter in IEASYSxx is now redundant and will be ignored.

Dynamic activates for hardware changes

If a z890 server is running multiple z/OS systems with a mixture of exploitation and compatibility support and a new hardware configuration is being activated, the activate must be done using one of the z/OS systems with exploitation support installed if:

- There is more than one Logical Channel Subsystem defined.
- A non-zero LCSS is being changed, or resources affected by the change are defined to non-zero LCSSs.

Dynamic CHPID management

On a z890, systems that are part of the same LPAR cluster may be in different LCSSs. Figure 3-3 on page 74 helps to illustrate that.

Dynamic CHPID management is supported even if the LPAR cluster spans multiple Logical Channel Subsystems. Movement of CHPIDs within the LPAR cluster is confined to movement within that LCSS.

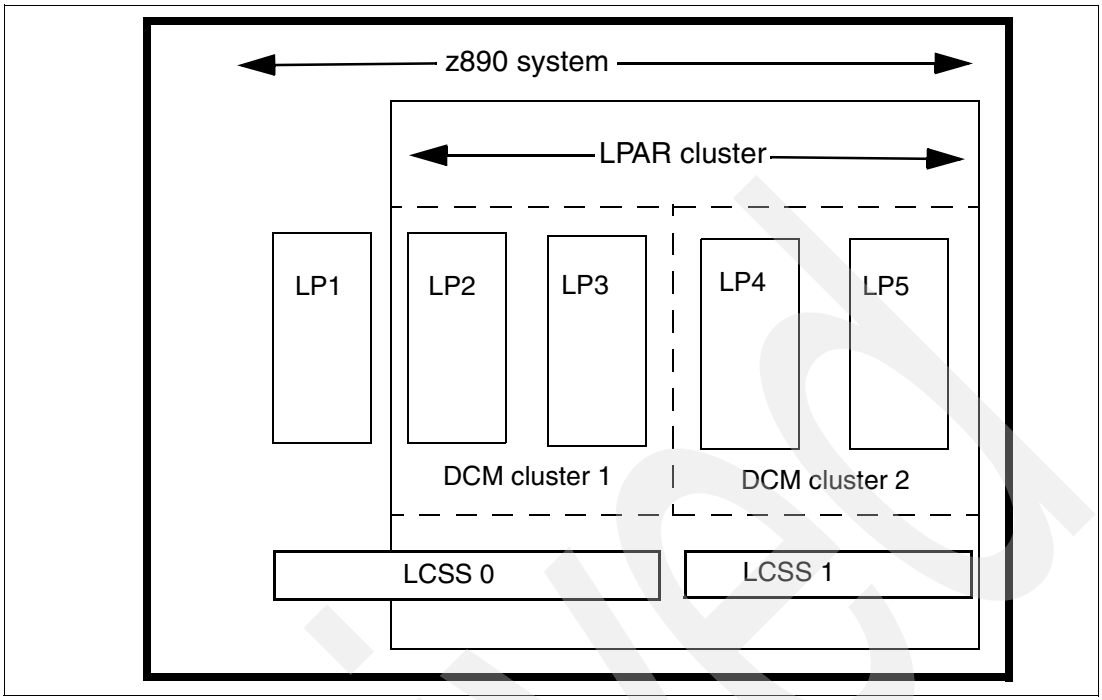

Figure 3-3 DCM clusters

If an LPAR cluster consists of multiple LCSSs, then the DCM command VARY SWITCH must be issued from one of the z/OS systems with exploitation support installed. The DCM command SETIOS DCM=ON/OFF can be issued from any system, whether in compatibility mode or exploitation mode.

Greater than 15 logical partitions

You must define a second LCSS if you plan to use more than 15 logical partitions. An individual Logical Channel Subsystem can only support up to 15 logical partitions. z/OS and OS/390 V2R10 logical partitions running with compatibility support can only reside in LCSS 0. Coupling Facilities, z/VM V4R4, and Linux on zSeries can reside in any Logical Channel Subsystem.

ISV software

You should check with your ISV about any required maintenance for ISV products running on the z890. This is particularly important for any software that uses the results of the STIDP instruction to control where that software can run.

3.4 z/VM software considerations

z/VM support on the z890 hardware is provided by z/VM V3R1 and z/VM V4R3 in compatibility mode, and z/VM V4R4, z/VM V5R1 and later releases, in exploitation mode.

Table 3-3 Supported z/VM releases on z890

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
z/VM Version 3 Release 1	Yes	Yes	Yes	No
z/VM™ Version 4 Release 3	Yes	Yes	Yes	No

z/VM Version 4 Release 4	Yes	Yes	Included	Included
z/VM Version 5 Release 1	No	Yes	Included	Included

Almost all new functions of the z890 server can be addressed with z/VM in compatibility support. But the use of the complete set of facilities are exclusive to z/VM with exploitation support applied. One example is Dynamic I/O configuration. Dynamic I/O is the ability to add, change or delete channel paths, control units and devices in multiple LCSSs without the need of a POR (power on/restart) or IPL. This is restricted to z/VM with exploitation support only. However, if only one LCSS is being used, any z/VM supported on this machine can do a Dynamic I/O configuration. See also “HCD support” on page 69.

z/VM exploitation support that is available in z/VM V4R4, z/VM R5V1 and later releases will allow operation in all Logical Channel Subsystems. It also covers use of extended I/O measurements, Linux for zSeries (guests) directed SCSI IPL, adapter interruptions, and new performance assist functions. z/VM V5R1 includes support for installation and IPL of z/VM on a FCP attached disk, support for PCIXCC and PCICA cryptographic features for z/OS and Linux guests, and the ability to install z/VM from HMC using the DVD or CD-ROM console driver.

If you are using z/VM 3.1 or z/VM 4.3, it is highly recommended that you check the last preventive service planning (PSP) prior to install a z890 server. You should also reference the IBM announcement letter for the z890, and subsequent announcements, for more detailed information concerning z/VM support.

3.5 Linux on zSeries software considerations

Linux on zSeries support is delivered through the June 2003 stream for kernel 2.4. The open source code is an effort of the Open Source Community and can be downloaded from the following URL:

`http://www10.software.ibm.com/developerworks/opensource/linux390/index.shtml`

Linux on zSeries support provides both compatibility and exploitation support, and will support superscalar performance, two LCSSs, and enhanced network performance and functionality.

Table 3-4 Supported Linux versions on z890

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
Linux for S/390	Yes	No	Yes	Yes
Linux® on zSeries	No	Yes	Yes	Yes

The following distributions already include the June 2003 stream: Red Hat AS 3.0, SuSE SLES 8, Turbolinux TLES 8 and Conectiva CLEE; the last three are Linux Enterprise Edition powered by UnitedLinux for zSeries. Also there are two different flavors for each distribution: 31-bit or 64-bit mode, and both modes are supported on a z890 server.

3.6 z/VSE and VSE/ESA software considerations

Compatibility support for z/VSE and VSE/ESA on the z890 is provided for VSE/ESA V2R6, VSE/ESA V2R7, z/VSE V3R1 and later releases. z/VSE and VSE/ESA support two LCSSs

and up to 30 logical partitions. z/VSE V3R1 is planned to have support for the Fiber Channel Protocol (FCP attached to a SCSI disk). It is recommended that you check the VSE subset of the 2086DEVICE Preventive Service Planning (PSP) bucket prior installing a z890 server.

Table 3-5 Supported VSE/ESA and z/VSE releases on z890

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
VSE/ESA™ Version 2 Release 6 and 7	Yes	No	Yes	Yes
z/VSE Version 3 Release 1	Yes	No	Included	Included

Important: The z/VSE operating system can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. The z/VSE operating system is designed to exploit select features of IBM eServer zSeries hardware.

3.7 TPF software considerations

The TPF operating system is supported on the z890 server. It requires TPF V4R1 to run. Note that there is no exploitation mode for TPF and it must run in 31-bit mode.

Table 3-6 Supported TPF releases on z890

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
TPF Version 4 Release 1 (ESA mode only)	Yes	No	Yes	No

3.8 Hardware feature support

This section describes the operating system levels needed to exploit the various functionalities on the z890 server. These include:

- z/Series Application Assist Processor (zAAP)
- Cryptographic features
- FICON Express features
- OSA-Express features
- Internal and External HiperSockets support
- Spanned channels

3.8.1 zSeries Application Assist Processor (zAAP)

Support for the zSeries Application Assist Processor is introduced in z/OS V1.6 and z/OS.e V1.6. This z/OS release is planned to be available in September 2004.

A zAAP reduces the standard processor (CP) capacity requirements for Java applications, freeing up capacity for other workload requirements. zAAPs do not increase the MSU value of the processor and therefore do not affect the software license fee.

zAAPs only run Java code. The IBM SDK for z/OS, Java 2 Technology Edition (the Java Virtual Machine) in cooperation with z/OS, and PR/SM directs JVM processing from CPs to zAAPs. Apart from the cost savings this may realize, the integration of Java-based applications with their associated database systems such as DB2, IMS, or CICS may simplify the infrastructure, for example, reducing the number of TCP/IP programming stacks and server interconnect links. Furthermore, processing latencies that would occur if Java application servers and their database servers were deployed on separate server platforms are prevented.

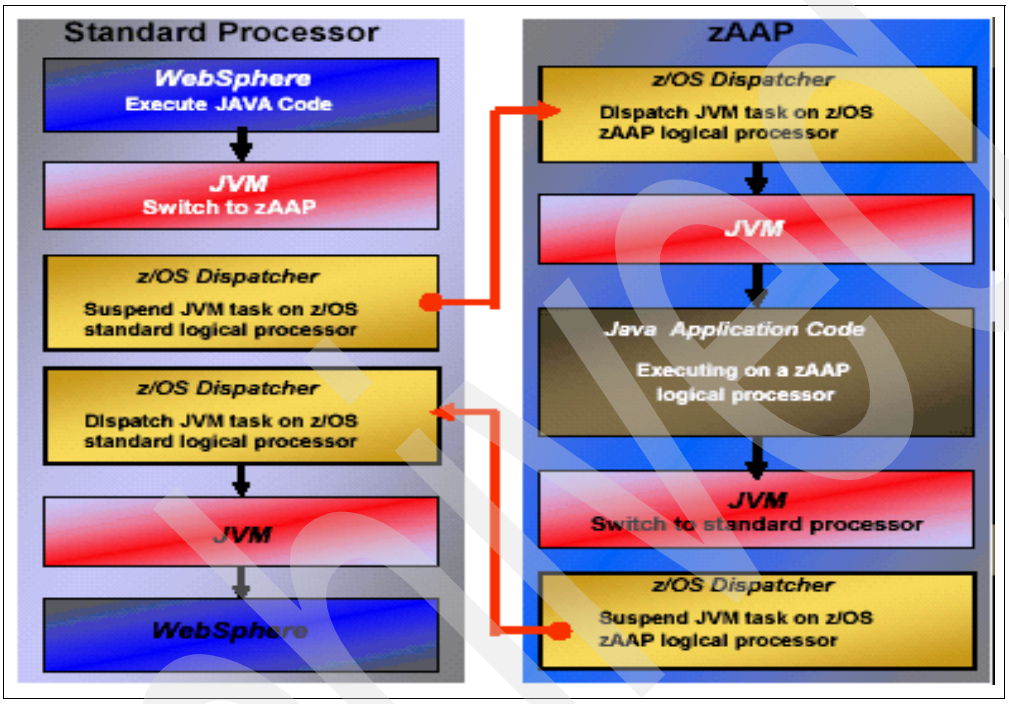

Figure 3-4 Logical flow of Java code execution on a zAAP

Figure 3-4 shows the logical flow of Java code running on a z890 server that has zAAPs available. The Java Virtual Machine (JVM), when it starts execution of a Java program, passes control to the z/OS dispatcher that will verify the availability of a zAAP:

- If a zAAP is available (not busy), the dispatcher will suspend the JVM task on the CP, and assign the Java task to the zAAP. When the task returns control to the JVM, it passes control back to the dispatcher that will reassign the JVM code execution to a CP.
- If there is no zAAP available at that time, the z/OS dispatcher may allow a Java task to run on a standard CP (depending on the option used in the OPT statement in the IEAOPTxx member of SYS1.PARMLIB).

zAAPs do not affect the overall MSU or capacity rating of a system or logical partition. That is, adding zAAPs to a system or defining to a logical partition does not affect the software license charges. There is no additional z/OS charge for zAAPs.

Subsystems that exploit zAAPs are:

- WebSphere Application Server (WAS) V5.1
- CICS/TS V2.3
- DB2 V8
- IMS V8

- WebSphere WBI for z/OS

The functioning of a zAAP is transparent to all IBM and ISV Java programming on JVM 1.4.1 and later.

Three execution options for Java code execution are available in z/OS. These options are user specified in IEAOPTxx and can be dynamically altered by the SET OPT command.

Option 1: Java dispatching by priority (honor_priority=yes)

Option 1 is the default option and specifies that standard CPs execute both Java and non-Java work in priority order when zAAPs are not configured. When zAAPs are configured they only execute Java work in priority order, while the CPs execute normal tasks and JVM tasks in priority order. This option is oriented towards servicing work with the highest priority first, regardless of the type of work.

Option 2: Java discretionary crossover (crossover=yes)

Standard CPs execute Java work in priority order only if no non-Java (standard) work is available to be dispatched. This way Java work may execute on a CP as if it has a lower priority than the non-Java work. This option is oriented towards environments where not enough zAAP capacity is available and the Java work has no need for priority over non-Java work. When executed on a zAAP, Java work is handled in priority order.

Option 3: No Java crossover (crossover=no)

This option is set to prevent Java work to be executed on a CP. If, for example, Sub Capacity Workload License Charging is applicable, Java work that is executed on a CP will increase CP utilization, and consequently may increase the software charges for non-Java work. Also, if ample zAAP capacity is available, this option assures that all Java work is done on a zAAP. Only if the last available zAAP would fail, crossover to a standard CP is enforced.

3.8.2 Cryptographic feature

Three cryptographic features are available on the z890 server:

- CP Assist for Cryptographic Function (CPACF) - feature #3863
- PCI Cryptographic Accelerator (PCICA) - optional feature #0862
- PCI X Cryptographic Coprocessor (PCIXCC) - optional feature #0868

Cryptographic functions that IBM provides for the z890 are defined by the IBM Common Cryptographic Architecture. Note that the Cryptographic Coprocessor Facility (CCF) and PCICC cryptographic features used on the z900 and z800 are not supported and are not available on z890 and z990 servers.

The CPACF feature has built-in circuitry in each physical processor on the z890 that partially replaces earlier CCF functions. In addition, additional *crypto-assist instructions* are supported. CPACF enabling is ordered separately with feature code 3863. The CPACF feature performs clear key DES, Triple DES (TDES), message authentication code (MAC) message authentication, and Security Hash Algorithm (SHA-1).

The IBM CP Assist for Cryptographic Function (CPACF) architecture provides a set of cryptographic functions that enhance the performance of encryption and decrypt functions for Secure Socket Layers (SSL), Virtual Private Network (VPN), and data storing applications not requiring Federal Information Processing Standards (FIPS) 140-2 level 4 security.

The PCICA feature is designed specifically for Secure Sockets Layer (SSL) processing, handling any key size of up to 2048 bits. It does not support symmetric cryptography, cannot

manage keys, and is not tamper proof. Each PCICA processor (there are two per feature) can support up to 2100 SSL handshakes per second.

The PCIXCC feature provides improved scalability, performance, and reliability over the cryptographic implementations of previous servers, due to:

- Use of a more powerful processor
- Extensive hardware error checking
- Faster RSA/SHA/DES engines
- More memory and flash memory
- A hardware-assisted communications protocol

The PCIXCC feature is the only cryptographic hardware for the z890 that supports Secure Keys, that is, the master key used is stored inside the hardware. Most of the functions that were available on CCF are replaced by the PCIXCC feature. Without the PCIXCC feature, cryptographic functions are supported by the CP Assist functions in each PU, and PCICA features with Clear Key support only.

Via the dynamic add/delete of a logical partition name, a logical partition can be renamed. Its name can be changed from 'NAME1' to ' * ' and then changed again from ' * ' to 'NAME2'. In this case, the logical partition number and MIFID are retained across the logical partition name change. However, the master keys in PCIXCC that were associated with the old logical partition 'NAME1' are retained. There is no explicit action taken against a cryptographic component for this.

Attention: Cryptographic cards are not tied to partition numbers or MIFIDs. They are set up with AP numbers and domain indices. These are assigned to a partition profile of a given name. The customer assigns these "lanes" to the partitions now and continues to have responsibility to clear them out if he changes who is using them.

z/OS, z/OS.e, and OS/390 support

Cryptographic functions, external interfaces, and a set of key management rules that pertain to the Data Encryption Standard (DES)-based symmetric algorithms and the Public Key Algorithm (PKA) asymmetric algorithms are available through the facilities provided by z/OS Integrated Cryptographic Service Facility/MVS (ICSF). ICSF uses the corresponding hardware features like the CPACF, PCIXCC and PCICA.

The support that ICSF for z/OS provides for z890 comprises:

- Clear key DES, TDES, and SHA instructions on all processor units (PUs) on a z890 processor. In other words, the z890 PUs are a very fast DES, TDES, or SHA engine.
  - Crypto-assisted instructions can be used by programs running in problem state; these instructions are described in the *Principle of Operations*, SA22-7832.
  - These DES/TDES functions use clear keys only; they do not use enciphered keys under a master key.
  - The ICSF Common Cryptographic Architecture (CCA) services CSNBSYE and CSNBSYD provide access to these instructions; the path length is minimal and there is no SAF or EXIT support. An assembly language programmer can code the instructions directly, but we suggest using ICSF interfaces for wider compatibility.
  - Key management support is limited to Rivest-Shamir-Adleman algorithm (RSA) Public-Key Cryptography Standards (PKCS) 1.2 key distribution via CSNDPKE, and CSNDPKD services; there is no support for the Cryptographic Key Data Set (CKDS).
  - There is no CPU affinity issue with these instructions.
- SSL and IP Security (IPSEC) acceleration is supported:
  - The RSA operations are driven by ICSF to the PCICA cards.

  - The DES operations are driven by ICSF to the z890 CP engines and use the crypto-assisted instructions.
- With PCIXCC support installed, it will be possible to use:
  - Secure key cryptographic functions
  - Secure encrypted key values
  - User Defined Extensions (UDX)

The following releases of z/OS and z/OS.e have Cryptographic Support available today: V1R2, V1R3, V1R4 and later releases. OS/390 V2R10 also has support. The Web deliverable z890 and z990 Enhancements to Cryptographic Support will be available on May 28, 2004, and can be downloaded from:

http://www.ibm.com/eserver/zseries/zos/download

Web deliverables are unpriced and the end of service coincides with the end of service for the release on which it runs.

z/VM support

z/VM's cryptographic support allows guests operating systems such as Linux on zSeries and z/OS to exploit the z890 cryptographic features.

Support for CPACF is available in z/VM 3.1, z/VM 4.3 and later releases.

Support for PCICA is available in z/VM 4.3 and later (for Linux guests) and z/VM 5.1 (for Linux and z/OS guests).

Support for PCIXCC is planned to become available in z/VM 5.1 (z/OS and Linux guests).

Linux support

The support for Linux on zSeries and Linux for S/390 is delivered at the following URLs:

http://www-124.ibm.com/developerworks/projects/libica
http://www-124.ibm.com/developerworks/projects/openCryptoki

z/VSE and VSE/ESA support

VSE support is available in VSE/ESA V2R7, z/VSE V3R1 and later releases for PCICA and can be used by IBM TCP/IP for VSE.

3.8.3 FICON Express features

FICON Express channels provide increased channel bandwidth and connectivity. It increases the data rate from 20 MBps for ESCON channels to 200 MBps.

FICON Express channels can be defined with either of three CHPID types:

- FICON in native mode (FC). Native FICON devices can be attached to servers via FICON channels with CHPID type FC, either directly or through a Fibre Channel Director. FC also supports Channel-to-Channel (CTC) communication.
- FICON bridge (FCV). ESCON devices can be attached only through a 9032 model 5 via the FICON bridge card (CHPID type FCV).
- Fibre Channel Protocol (FCP). FCP is for support of SCSI devices in Linux and z/VM environments.

z/OS, z/OS.e, and OS/390 support

For this support:

- FICON in native mode (CHPID type FC): Requires at minimum OS/390 V2R10, z/OS V1R2 (or later), or z/OS.e V1R3 (or later). In the case of cascaded FICON Directors, including CTCs with cascading, a minimum of z/OS V1R3 or V1R4 with PTFs, or z/OS.e V1R3 (or later) is required. For OS/390 V2R10, it is necessary that z/OS V1R3 (or later) be in another logical partition to dynamically define cascaded directors for dynamic I/O changes and to use enhanced display functions. More detailed information can be obtained in the 9032/9042 PSP buckets.
- FICON Bridge (CHPID type FCV): Attachment to ESCON Director Model 9032 with the FICON Bridge feature requires z/OS V1R2 or later, or z/OS.e V1R3 or later.
- Fibre Channel Protocol (CHPID type FCP): z/OS does not exploit FCP.

z/VM support

For this support:

- FICON in native mode (CHPID type FC): Requires z/VM V3R1 and V4R3 or later releases. FICON cascaded Directors, including CTCs with cascading, requires z/VM V4R4 or later.
- FICON bridge (CHPID type FCV): Requires z/VM V3R1 and z/VM V4R3 or later releases.
- Fibre Channel Protocol (CHPID type FCP): Requires z/VM V4R3 or later for Linux or zSeries or Linux for 390 running as a guest under z/VM. The Performance Assist for Adapter Interruptions and Performance Assist for V=V Guests require z/VM V4R4 or later. Native support of SCSI disks for installation, IPL, and operations of z/VM operating system require z/VM V5R1 or later.

Linux support

Fibre Channel Protocol (CHPID type FCP) support is included on the following releases of Linux for 390 and Linux on zSeries: SuSE SLES 8, Turbolinux TLES 8, and Conectiva CLEE, all of them powered by United Linux V1 and RedHat AS V3. SCSI IPL for FCP also requires z/VM V4R4 or later. Also available is FCP SAN management delivered in January, 2004, as an Open Source contribution. All required drivers can be downloaded from the following URL:

`http://www10.software.ibm.com/developerworks/opensource/linux390`

z/VSE and VSE/ESA support

FICON in native mode (CHPID type FC) and FICON Bridge (CHPID type FCV) are supported on VSE/ESA V2R6, VSE/ESA V2R7, and z/VSE V3R1 and later releases. Fibre Channel Protocol (CHPID type FCP) is planned to be supported by z/VSE V3R1 and later releases.

TPF support

FICON in native mode (CHPID type FC) and FICON Bridge (CHPID type FCV) are supported by TFP V4R1 at PUT level 16. The support of cascaded FICON Directors is also included.

3.8.4 OSA-Express features

The following OSA-Express features are orderable with the z890 server:

- OSA-Express Gigabit Ethernet (GbE)
- OSA-Express 1000BASE-T Ethernet
- OSA-Express 1000BASE-T Ethernet used for the implementation of the OSA-Express Integrated Console Controller (OSA-ICC)

- OSA-Express Token Ring

The current OSA-Express Fast Ethernet feature as used on the z800 is also supported on z890.

z/OS, z/OS.e, and OS/390 support

z/OS, z/OS.e, and OS/390 support for OSA-Express features on z890:

- *OSA-Express Gigabit Ethernet* (CHPID type OSD): Requires z/OS V1R2 and later, z/OS.e V1R3 and later, OS/390 V2R10 with Communication Server. OSA-Express Gigabit Ethernet has QDIO support only.
- *OSA-Express 1000BASE-T Ethernet* (CHPID type OSD in QDIO mode, CHPID type OSE for non-QDIO mode): Requires z/OS V1R2 and later for both QDIO and non-QDIO mode. If running non-QDIO mode, the appropriate release of the Communications Server (ACF/VTAM) must be used. For OS/390 V2R10 and Communication Server with PTFs for QDIO and non-QDIO mode, z/OS.e V1R3 and later.
- *OSA-Express Integrated Console Controller (OSA-ICC)* (CHPID type OSC): This feature is exclusive for z890 and z990 servers. It requires z/OS V1R3 and later, with PTF for APAR OA05738. HCD also requires PTF for APAR OA03689. z/OS.e V1R3 and later.
- *OSA-Express Token Ring* (CHPID type OSD in QDIO mode, CHPID type OSE for non-QDIO mode): Requires z/OS V1R2 and later releases, for both QDIO and non-QDIO mode. For non-QDIO mode, Communication Server and the appropriate release of ACF/VTAM. OS/390 V2R10 for QDIO mode and non-QDIO mode require PTFs. The appropriate release of Communication Server and ACF/VTAM is also required if running in non-QDIO mode.

OSA-Express functionality for z/OS and OS/390 includes:

- *Checksum offload for IPv4 Packets:* Applies to OSA-Express Gigabit Ethernet (Feature Codes 1364 and 1365) and 1000BASE-T Ethernet (FC 1366) on z890 and z990 when in QDIO mode, and requires z/OS V1R5 or z/OS.e V1R5 and later.
- *z/OS full VLAN (802.1q) support*: Applies to OSA-Express 1000BASE-T Ethernet (FC 1366), Fast Ethernet (FC 2366) and Gigabit Ethernet (FCs 2364, 2365, 1364, 1365) on z800, z900, z890 and z990 when in QDIO mode, and requires z/OS V1R5 or z/OS.e V1R5 Communications Server.
- *Intrusion Detection Services*: Applies to OSA-Express features supported on z890 and z990 servers when in QDIO mode, and requires z/OS V1R5 or z/OS.e V1R5 Communications Server.
- *160 TCP/IP stacks per OSA-Express port:* Applies to all OSA-Express features supported on z890 and z990 when in QDIO mode and is transparent to the operating systems.
- *OSA/SF Java GUI:* Is a complete replacement of OSA/SF V2R1 (5655-B57). It is integrated in the operating systems and shipped as a PTF where applicable. This support applies to all of the OSA-Express features and to all servers that support them. The new version of OSA/SF is supported by z/OS V1R2 and later, z/OS.e V1R3 and later, and OS/390 V2.R10. Refer to the appropriate PSP bucket for more information.
- *OSA-Express direct SNMP subagent support*:
  - *Traps and Set*: Applies to all OSA-Express features supported on z890 and z990 servers when configured in QDIO mode (CHPID type OSD), and requires z/OS V1R5 or z/OS.e V1R5 and later releases.
  - *Direct SNMP for LCS:* Applies to all OSA-Express features supported on z890 and z990 servers when configured in non-QDIO mode (CHPID type OSE), using TCP/IP passthru and LAN Channel Station (LCS). Direct SNMP for LCS supports the same

SNMP commands and alerts currently offered in QDIO mode, that is, Get, GetNext, Trap, Set. Support is planned to be available for z/OS V1R6 and z/OS.e V1R6 and later releases.

- *Performance data*: Applies to all OSA-Express features supported on z800, z900, z890 and z990 when configured in QDIO mode (CHPID type OSD). It requires z/OS V1R4, or z/OS.e V1R4 and later releases.
- *Get and GetNext:* Applies to all OSA-Express features supported on z800, z900, z890 and z990 when configured in QDIO mode (CHPID type OSD). It requires z/OS V1R4 with co-requisite (refer to PSP bucket), or z/OS.e V1R4 with co-requisite, and later releases.
- *Ethernet data for dot3StatsTable:* Applies to all OSA-Express features supported on z800, z900, z890 and z990 servers when configured in QDIO mode. It requires z/OS V1R4 and z/OS.e V1R4, both with co-requisite (refer to PSP), and later releases.

z/VM support

For this support:

- *OSA-Express 1000BASE-T Ethernet* (CHPID type OSD in QDIO mode, CHPID type OSE for non-QDIO mode): Supported by z/VM V3R1, z/VM V4R3 and later. If running in QDIO mode check for the appropriate level of TCP/IP. If running in non-QDIO mode, you must check also the appropriate release of ACF/VTAM.
- *OSA-Express Integrated Console Controller (OSA-ICC)* (CHPID type OSC): Supported by z/VM V4R4 with PTF for APAR VM63405.
- *OSA-Express Gigabit Ethernet* (CHPID type OSD): z/VM V3R1, z/VM V4R3 and later, with the appropriate level of TCP/IP.
- *OSA-Express Token Ring* (CHPID type OSD in QDIO mode, CHPID type OSE in non-QDIO mode): Supported by z/VM V3R1, z/VM V4R3 and later releases with the appropriate level of TCP/IP if using it in QDIO mode. Requires also the appropriate ACF/VTAM release to support it in non-QDIO mode.

OSA-Express functionality for z/VM includes:

- *z/VM VLAN (802.1q) support:* Applies to OSA-Express 1000BASE-T Ethernet (FC 1366), Fast Ethernet (FC 2366) and Gigabit Ethernet (FCs 2364, 2365, 1364, 1365) on z890 and z990 servers when in QDIO mode. Support is provided by z/VM V4R4 and later for one global VLAN ID for IPV4. For one global VLAN ID for IPV6, support is provided by z/VM V5R1 and later releases.
- *160 TCP/IP stacks per OSA-Express port:* Applies to all of the OSA-Express features supported on z890 and z990 when in QDIO mode and is transparent to the operating systems.
- *OSA/SF Java GUI:* Supported is integrated within z/VM V4R4, and later releases.
- *OSA port name relief:* Applies to all OSA-Express features supported on z800, z900, z890 and z990 when in QDIO mode. Support is provided by z/VM V4R3 with co-requisites (please refer to z/VM subset of the 2086DEVICE PSP bucket). On z/VM V4R4 support is integrated.

Linux support

For this support:

- *OSA-Express 1000BASE-T Ethernet:* Requires Linux kernel 2.2.16 and later. Note that only QDIO mode is supported.
- *OSA-Express Gigabit Ethernet* (CHPID type OSD): Linux with kernel 2.2.16 and later.

- *OSA-Express Token Ring:* Linux with kernel 2.4 and later.

OSA-Express functionality for includes:

- *Checksum offload for IPv4 Packets:* Applies to OSA-Express Gigabit Ethernet (Feature Codes 1364 and 1365) and 1000BASE-T Ethernet (FC 1366) on z890 and z990 when in QDIO mode. Linux on zSeries support was included in the qeth driver delivered in June, 2003.
- *OSA port name relief:* Linux on zSeries support was delivered in June, 2003.
- *OSA-Express Direct SNMP subagent support:*
  - *Performance data:* Support for Linux on zSeries was delivered in June, 2003.
  - *Get and GetNext:* Linux on zSeries support is available for SUSE SLES8, Turbolinux TLES 8 and Conectiva CLEE distributions.
  - *Ethernet data for dot3StatsTable:* Support for Linux on zSeries was delivered in June 2003.

  Note that for updated information about OSA-Express drivers for Linux on zSeries, refer to:

  `http://www10.software.ibm.com/developerworks/opensource/linux390`

z/VSE and VSE/ESA support

For this support:

- *OSA-Express 1000BASE-T Ethernet* (CHPID type OSD in QDIO mode, CHPID type OSE for non-QDIO mode): Support is delivered by VSE/ESA V2.6, VSE/ESA V2.7, and z/VSE V3.1 and later releases, with the appropriate release of TCP/IP. Support for non-QDIO mode also requires the appropriate release of ACF/VTAM.
- *OSA-Integrated Console Controller (OSA-ICC)* (CHPID type OSC): Supported by VSE/ESA V2.6, VSE/ESA V2.7, and z/VSE V3.1 and later releases.
- *OSA-Express Gigabit Ethernet* (CHPID type OSD): VSE/ESA V2.6, VSE/ESA V2.7, and z/VSE V3.1 and later releases with the appropriate release of TCP/IP.
- *OSA-Express Token Ring* (CHPID type OSD in QDIO mode, CHPID type OSE for non-QDIO mode): Support is delivered by VSE/ESA V2.6, VSE/ESA V2.7, and z/VSE V3.1 and later releases with appropriate release of TCP/IP. Support for non-QDIO mode also requires the appropriate release of ACF/VTAM.

OSA-Express functionality for z/VSE and VSE/ESA includes *OSA/SF Java GUI;* support is provided for VSE/ESA V2.6, VSE/ESA V2.7, and z/VSE V3.1 and later releases.

TPF support

For this support:

- *OSA-Express Gigabit Ethernet* (CHPID type OSD): TPF V4R1 at PUT level 13 with PTF for APAR JP2733.
- *OSA-Integrated Console Controller (OSA-ICC)* (CHPID type OSC): Supported by TPF V4R1.

3.8.5 HiperSockets

HiperSockets (also known as internal Queued Direct Input/Output (iQDIO)) provides network capability within the boundaries of the server. This functionality allows high speed any-to-any connectivity among logical partitions, that is, operating system images, within the z890 server without requiring any physical cabling. This “network within the box” concept minimizes

network latency and maximizes bandwidth capabilities between z/VM, Linux on zSeries, z/VSE, VSE/ESA, and z/OS images, or combinations of these. There is also the possibility to have HiperSocket under z/VM which permits internal network between guests operating system, like many Linux servers, for example.

Up to 16 separate internal LANs with up to 4096 TPC/IP stacks can be configured within a server, and the support is transparent to the operating systems. If you want the internal LANs shared between partitions in different LCSSs then the channel must be spanned. For more information on spanned channels, refer to “Spanned channels” on page 45.

z/OS support

All z/OS versions have support for HiperSockets. However, *Broadcast for IPv4 packets*, a new function, requires z/OS V1R5, or z/OS.e V1R5, and later releases.

z/VM support

Support for HiperSockets is provided by z/VM V4R3 and later releases.

HiperSocket functionality include:

- *Broadcast for IPv4 Packets:* Applies to HiperSockets on z890 and z990. The support is provided by z/VM V4R4 and later releases.
- *HiperSockets Network Concentrator:* Applies to HiperSockets on z890 and z990. Requires z/VM V4R3 with PTF for APAR VM63397, or z/VM V4R4 and later releases. This facility is exploited by Linux on zSeries guests.

Linux support

Support for Linux on zSeries is delivered through the *qeth* device driver in all Linux 2.4 kernels, that include August, 2001, May, 2002, and June 2003 “streams” that were made available by a contribution to the Open Source community. For up-to-date information, refer to the following Web site:

`http://www10.software.ibm.com/developerworks/opensource/linux390`

HiperSocket functionality include:

- *VLAN (802.1q):* Applies to HiperSockets on z890 and z990. Linux on zSeries support is included in the qeth device driver delivered in June, 2003.
- *Broadcast for IPv4 packets:* Applies to HiperSockets on z890 and z990. Linux on zSeries support is included in the qeth device driver in all Linux 2.4 kernels.
- *HiperSockets Network Concentrator:* Applies to HiperSockets on z890 and z990. Linux on zSeries support is included in the qeth driver dated 2003-10-31 or later. Also s390-tools-1.2.3, dated 2003-11-28 or later is needed. Please refer to the instructions accompanying the October 31, 2003, updates on DeveloperWorks about how to activate HiperSockets Network Concentrator unicast, broadcast, and multicast support.

z/VSE and VSE/ESA support

HiperSockets support is available in z/VSE and VSE/ESA and requires VSE/ESA V2R7 or z/VSE V3R1 and later releases.

3.8.6 Spanned channels

Spanning channels is the ability to configure a channel to more than one Logical Channel Subsystem (LCSSs). The channel is transparently shared by any or all of the configured

logical partitions regardless of the Logical Channel Subsystem to which the partition is configured.

Spanned channel support is an extension of Multiple Image Facility (MIF). MIF allows sharing of channel resources across logical partitions. MIF Spanned channel support allows sharing of channel resources across LCSSs. There are two types of channels that can be MIF spanned across LCSSs:

- Internal spanned channels, that is, internal coupling channels (ICs) and HiperSockets (CHPID type IQD).
- External spanned channels, that is, FICON Express channels (CHPID types FC, and FCP), ICB-3 links, ICB-4 links, ISC-3 links (CHPID type CFS, and CFP), and OSA-Express ports (CHPID types OSC, OSD, and OSE).

Note: ESCON channels cannot be spanned.

z/OS support

For this support:

- *Internal spanned channels:* Require z/OS V1R2, z/OS V1R3, or z/OS V1R4 with PTFs, z/OS.e V1R3 with PTFs, and later releases
- *External spanned channels:* Require z/OS V1R4, or z/OS.e V1R4, and later releases

z/VM support

For this support:

- *Internal spanned channels:* Require z/VM V3R1 or z/VM V4R3 with compatibility support, or z/VM V4R4 and later releases
- *External spanned channels:* Require z/VM V3R1 or z/VM V4R3 with compatibility support, or z/VM V4R4 and later releases

z/VSE and VSE/ESA support

For this support:

- *Internal spanned channels:* Require VSE/ESA V2R7 or z/VSE V3R1 for HiperSockets (CHPID type IQD) only

For more details on these and other zSeries connectivity features, refer to *zSeries Connectivity Handbook,* SC24-5444.

3.8.7 Summary of z/OS, z/OS.e and OS/390 software requirements

The table below is a summary of z/OS, z/OS.e and OS/390 software requirements.

Table 3-7 Minimum z/OS, z/OS.e and OS/390 software requirements

Software requirements / Function	z/OS and z/OS.e V1.6[a]	z/OS and z/OS.e V1.5	z/OS V1.4 Ex-ploit	z/OS.e V1.4 Co-exist Update	z/OS and z/OS.e V1.3	z/OS V1.2	OS/390 V2.10
16 to 30 logical partitions	X	X	X	X			
Two Logical Channel Subsystems (LCSSs)	X	X	X	X			
Dynamic I/O support for multiple LCSSs	X	X	X	X			

Software requirements Function	z/OS and z/OS.e V1.6[a]	z/OS and z/OS.e V1.5	z/OS V1.4 Ex-ploit	z/OS.e V1.4 Co-exist Update	z/OS and z/OS.e V1.3	z/OS V1.2	OS/ 390 V2.10
Dynamic Add/Delete Logical Partition Name	X						
zSeries Application Assist Processor (zAAP)	X[b]						
Internal spanned channels	X	X	X	X	X		
External spanned channels	X	X	X	X			
HiperSockets	X	X	X	X	X	X	
Broadcast for IPv4 packets	X	X					
16-port ESCON feature	X	X	X	X	X	X	X
FICON Express (type FCV)	X	X	X	X	X	X	X
FICON Express (type FC)	X	X	X	X	X	X	X
Cascaded FICON Directors (CHPID types FC, and FCP)	X	X	X	X	X		X
OSA-Express GbE (CHPID type OSD)[c]	X	X	X	X	X	X	X
OSA-Express 1000BASE-T Ethernet	X	X	X	X	X	X	X
OSA-Express Integrated Console Controller (OSA-ICC)	X	X	X	X	X		
OSA-Express Token Ring	X	X	X	X	X	X	X
Checksum offload for IPv4 packets[d]	X	X					
z/OS Full VLAN (IEEE 802.1q) support	X	X					
Intrusion Detection Services (CHPID type OSD)	X	X					
OSA/SF Java GUI[e]	X	X	X	X	X	X	X
OSA-Express Direct SNMP subagent support[f]	X	X	X				
CP Assist for Cryptographic Function (CPACF)[g]	X	X	X	X	X	X	X
PCI Cryptographic Accelerator (PCICA)[g]	X	X	X	X	X	X	X
PCIX Cryptographic Coprocessor (PCIXCC)[g]	X	X	X	X	X	X	X
PCIXCC User Defined Extensions (UDX)[g]	X	X	X	X	X	X	X

a. z/OS and z/OS.e V1.6 is planned to be available September 2004.
b. z/OS and z/OS.e V1.6 planned to be available September 2004, plus IBM SDK for z/OS, Java 2 Technology Edition, V1.4.
c. CHPID type OSE (non-QDIO) supports TCP/IP and SNA.
d. For OSA-E GbE and 1000BASE-T EN with CHPID type OSD.
e. With z990 Compatibility for Selected Releases, a Web deliverable.
f. z/OS and z/OS.e V1.5 for 'Traps and Set', z/OS and z/OS.e V1.6 for 'Direct SNMP for LCS' support all other SNMP subagent support z/OS and z/OS.e V1.4 and later.
g. With z990 Cryptographic Support, or z990 and z890 Enhancements to Cryptographic Support.

3.8.8 Summary of z/VM, z/VSE, VSE/ESA, TPF, and Linux software requirements

The following table is a summary of z/VM, z/VSE, VSE/ESA, TPF, and Linux software requirements.

Table 3-8 Minimum z/VM, z/VSE, VSE/ESA, TPF and Linux on zSeries requirements

Software requirements / Function	z/VM V5.1	z/VM V4.4	z/VM V3.1 & V4.3	VSE/ESA V2.7 and z/VSE V3.1	VSE/ESA V2.6	TPF V4.1	Linuxon zSeries [a]
16 to 30 logical partitions	X	X		X	X	X	X
Two Logical Channel Subsystems (LCSSs)	X	X	X[b]	X	X		X
Dynamic I/O support for multiple LCSSs	X	X	X[b]				X
Dynamic Add/Delete LP Name							X
Internal spanned channels	X	X	X	X[c]			X
External spanned channels	X	X	X				X
Adapter Interruption (CHPID types FCP and OSD)	X	X		X[d]			X
V=V support for CHPID types FCP, OSD, and IQD	X	X					
HiperSockets	X	X	X[e]				X
VLAN (IEEE 802.1q)							X
Broadcast for IPv4 packets	X	X					X
HiperSockets Network Concentrator	X	X	X[f]				X
16-port ESCON feature	X	X	X	X	X	X	X
FICON Express (CHPID type FCV)	X	X	X	X	X	X	
FICON Express (CHPID type FC)	X	X	X	X	X	X	X
FICON Express (CHPID type FCP)	X[g]	X	X[h]	X[i]			X
FCP SAN Management							X[j]
SCSI IPL for FCP	X[k]	X[l]					X
Cascaded FICON Directors (CHPID types FC and FCP) including CTC	X	X		X	X	X	X
OSA-Express GbE (CHPID type OSD)	X	X	X	X	X	X[m]	X
OSA-Express 1000BASE-T Ethernet	X	X	X	X	X		X
OSA-Express Integrated Console Controller (OSA-ICC)	X	X[n]		X	X	X	
OSA-Express Token Ring	X	X	X	X	X		X
Checksum offload for IPv4 packets							X
z/VM VLAN (IEEE 802.1q) support	X[o]	X[p]					

Software requirements Function	z/VM V5.1	z/VM V4.4	z/VM V3.1 & V4.3	VSE/ ESA V2.7 and z/VSE V3.1	VSE/ ESA V2.6	TPF V4.1	Linuxon zSeries [a]
Linux on zSeries VLAN (IEEE 802.1q) support							X[q]
Intrusion Detection Services							X[r]
OSA/SF Java GUI	X	X	X	X	X		
OSA port name relief	X	X	X[s]				
OSA-Express Direct SNMP subagent support							X[t]
CP Assist for Cryptographic Function (CPACF)	X	X	X				X
PCI Cryptographic Accelerator (PCICA)	X	X	X[e]	X			X
PCIX Cryptographic Coprocessor (PCIXCC)	X						X

a. Current distributions are SUSE SLES7, and SLES8, Red Hat RHEL 3.0, Turbolinux TLES 8, and Conectiva CLEE.
b. Dynamic I/O configuration for LCSS 0 only.
c. HiperSockets only (CHPID type IQD).
d. CHPID type OSD only.
e. z/VM V4.3 only.
f. z/VMV4.3 only with PTF for APAR VM63397.
g. For z/VM install, IPL, and operation from SCSI disks.
h. z/VM V4.3 only. For Linux as a guest.
i. Planned for z/VSE V3.1 only (for FCP attached SCSI disks on the IBM ESS).
j. See //http:www10.software.inm.com/developerworks/opensource/linux390.
k. z/VM IPL from SCSI disks.
l. For Linux as a guest.
m. PUT 13 with PTF for APAR PJ2733.
n. With PTF for APAR VM63405.
o. For one global VLAN ID for IPv6 (applies to OSA-Express 1000BASE-T Ethernet, Fast Ethernet, and GbE).
p. For one global VLAN ID for IPv4 (applies to OSA-Express 1000BASE-T Ethernet, Fast Ethernet, and GbE).
q. Applies to OSA-Express 1000BASE-T Ethernet, Fast Ethernet, and GbE with CHPID type OSD.
r. Applies to all OSA-Express features in QDIO mode (CHPID type OSD).
s. z/VM V4.3 only, applies to all OSA-Express features in QDIO mode (CHPID type OSD).
t. Applies to performance data, Get and GetNext, and Ethernet data for dot3StatsTable in QDIO mode, CHPID Type OSD.

4

Planning and migration considerations

This chapter is aimed to help you with some challenges that may arise when migrating to a z890 server. It is impossible to cover all that you may come across, but you may be able to find some references that can be useful to you, especially if you are running on earlier s/390 servers such as Multiprise 2000, Multiprise 3000, or 9672 CMOS servers of all generations.

This chapter also offers planning information when appropriate, and migration information, where migration issues may arise when moving from where you are today to the implementation of a z890. It is not a replacement for other specific material on this subject such as the z890 Physical Planning manual IMPP or z890 SAPR Guide. It is meant to be complementary. Also some thought is given to the implications of moving from one operating system level (OS) to another, often necessitated by the level of OS support required on the z890.

4.1 Power and cooling

A z890 systems uses 50/60 Hz three-phase power or (optionally) single-phase power. Power requirements vary from 2.5 KW for a z390 used as dedicated Coupling Facility with only ICB-4 (cluster bus) links for communication (that is, no I/O feature cards) to 4.2 KW for a system with a full load of I/O features installed.

Power connection requirements are one of the following:

```
200 - 240 volts (three phase or single phase)      24 amps
380 - 415 volts (three phase or single phase)      12 amps
480 volts (three phase)                            10 amps
```

Dual power cords are used and these are intended to be connected to independent power sources.

The z890 is an air cooled server. There is no internal refrigeration unit as in the z990. Chilled air must be provided from under a raised floor. The z890 contains two fans. It is important to use the internal air flow controls, as designed, to route sufficient cooling to all areas of the frame.

The optional battery feature (IBF) provides approximately 10 minutes of operation for a fully loaded system and up to 24 minutes for a z890 dedicated to coupling functions only. (These numbers assume the batteries are less than three years old and have experienced no more than two complete discharge cycles per year.)

The table below provides a summary of physical and environmental information for the z890. However, the reader should refer to the more authoritative manual, *zSeries 890 Installation Manual - Physical Planning*, GC28-6828, for planning and installation of the z890.

	G5 / G6 Minimum 1 Frame System	G5 / G6 Maximum 2 Frame System	Multiprise 3000 1 Frame System Maximum	z800 Maximum	z890 Minimum	z890 Maximum
Power 50/60 Hz, kVA	0.6 / 1.0	5.5 / 5.5	1.32	2.95KW	1.5	4.7
Heat Output KBTU/hr	2.0 / 2.5	18.8 / 18.8	4.5	10.0	5.12	16.05
Air Flow CFM Air Flow m*3/min	290 / 290 7.1 / 7.1	1400 / 1400 38.6 / 38.6		400 11.1	640 17.64	640 17.64
Floor Space - Sq. meters - Sq. feet	1.0 / 1.6 10.4 / 16.4	1.8 / 1.8 19.7 / 19.7		0.83 8.9	1.24 13.33	1.24 13.33
Including service clearance - Sq. meters - Sq. feet	2.5 / 2.5 27.4 / 27.4	4.8 / 4.8 51.9 / 51.9		6.0 64.5	3.03 32.61	3.03 32.61
Approximate weight - kg - lbs	612 / 612 1346 / 1346	938 / 938 2057 / 2057	236 520	545 1201	674 1482	785 1730
Approximate height - cm - inches	199.8 78.7	199.8 78.7	80 31.5	181.1 71.3	194.1 76.4	194.1 76.4

Figure 4-1 Power and cooling requirements and other environmental properties of the z890

4.2 Hardware Management Console (HMC)

The z890 server requires a HMC to be operated from. Some older HMC features can be used with the z890, see Figure 4-2. There are some enhancements for HMC that can be useful and may avoid acquisition of other additional hardware. An example is the Integrated 3270 Console. If you are planning to run z/VM on z890 for example, there is no need to acquire a IBM 2074 Console Controller or an IBM 3174 control unit for console attachment. Another feature is the Integrated ASCII console that can be used with Linux running in under logical partition (not under z/VM). For a description of the console functions see “Integrated 3270 console” on page 59 and “Integrated ASCII console” on page 59.

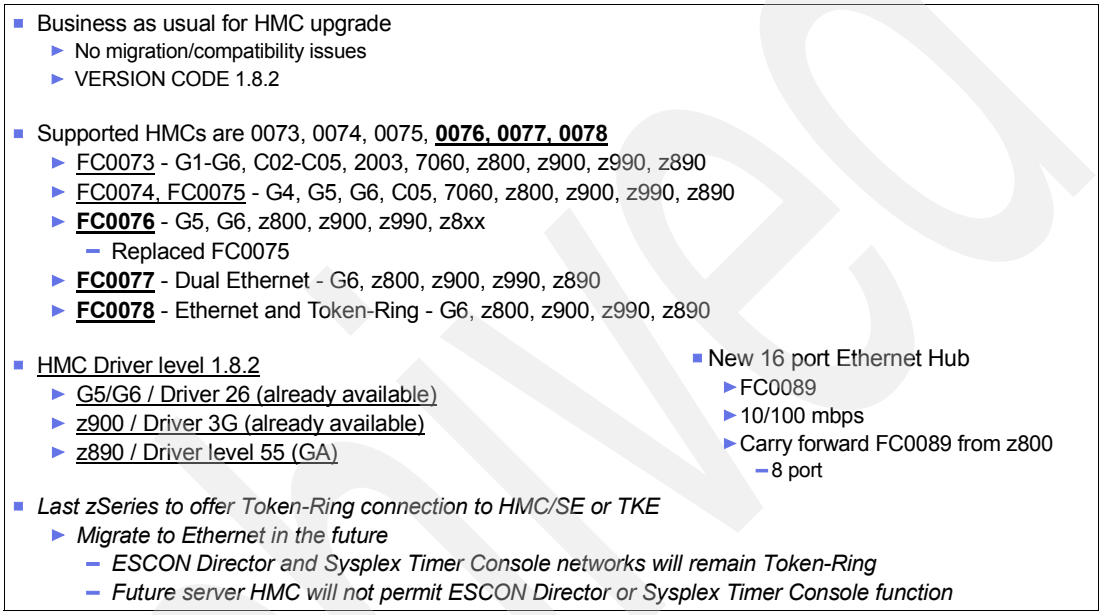

Figure 4-2 HMC feature codes

4.2.1 Token ring planning

The z890 system accepts OSA-Express Token Ring features and also supports token ring attachment of HMCs and SEs. Concurrently with the z890 announcement, IBM released planning information for future systems (*after* the z890). This planning information includes the following:

- Future systems will not offer Token Ring features for Hardware Management Consoles (HMCs), Support Elements (SEs), or Trusted Key Entry (TKE) workstations.
- While IBM intends to offer the OSA-Express Token Ring features for nearer term future systems, the long-term direction is clearly away from token ring use.
- The ESCON Director Model 5 and the Sysplex Timer Model 2 continue to support *only* Token Ring connectivity to their respective components. IBM suggests that customers procure appropriate independent workstations (with token ring features) to continue operation of these components. Future operation of these components must not assume the availability of HMCs as consoles for ESCON Directors or Sysplex Timers.
- Future HMCs will be closed platforms that run *only* the HMC application. It will not be possible to install additional applications on these HMCs.
- When available, the next generation HMC may communicate only with G5 servers and later (that is, Multiprise 3000, Generation 5, and 6, z800, z900, z890, and z990).
- TCP/IP is intended to be the only communications protocol supported.

These comments are for future systems, not the z890. This planning information is subject to change without notice. It is recommended that for future installations token ring environments be migrated to Ethernet.

4.2.2 SE and HMC connectivity planning

Configuring and ordering SEs and HMCs requires an understanding of the LAN interfaces involved. Both units normally have two LAN interfaces. The general rules are:

- An SE is connected to a given HMC by only one LAN.
- In simple situations, the second LAN interfaces included with SEs and HMCs are not used.
- The two LAN interfaces on an SE may be used to connect to different sets of HMCs, using two independent LANs.
- The two SE LAN interfaces may be both Ethernet, or one token ring and one Ethernet. An option of two token ring interfaces is not available.[15] Both SEs in a system will have the same LAN configuration.
- A default HMC configuration has one Ethernet and one token ring interface.

SEs and HMCs may be connected through public LANs, but this is typically not done for several reasons:

- It is an obvious security exposure.
- It could result in connection losses, which would probably impact customer operation.
- MCL distribution from HMC to SE cannot tolerate connection losses.
- An SE upgrade or restore is quite sensitive, and could be impacted if there is much traffic on the LAN.

A very common arrangement is to use at least two HMCs, one near the system and the other near the general operations control area.

The following discussion assumes that you purchase a new HMC with your z890. You can purchase any reasonable number of HMCs to use with the system.[16] You have the following choices when you order your system:

- SEs with a Token Ring and an Ethernet feature are supplied with two 50 feet Ethernet cables.
- SEs with two Ethernet features are supplied with four 50 feet Ethernet cables.
- HMC with one Token Ring and one Ethernet feature. The HMC comes with a 75 feet Token Ring cable, a 7 feet Token Ring cable, and a 50 feet Ethernet cable.
- A display for the HMC.
- An MAU for Token Ring is included in the z890 frame if the Support Elements include Token Ring ports. You do not need to order the MAU in this case.
- An Ethernet switch. This is ordered automatically if your order does not include Token Ring interfaces.

The SE and HMC Ethernet ports can run at 10 Mbps or 100 Mbps and autosense the LAN speed.

[15] There is no obscure theoretical reason for the mixture of LAN interface types offered. It is based on the practical observation that an Ethernet port is now included on the planar board of most PCs, whether you use it not.
[16] Up to 32 HMCs can be used to control your zSeries systems.

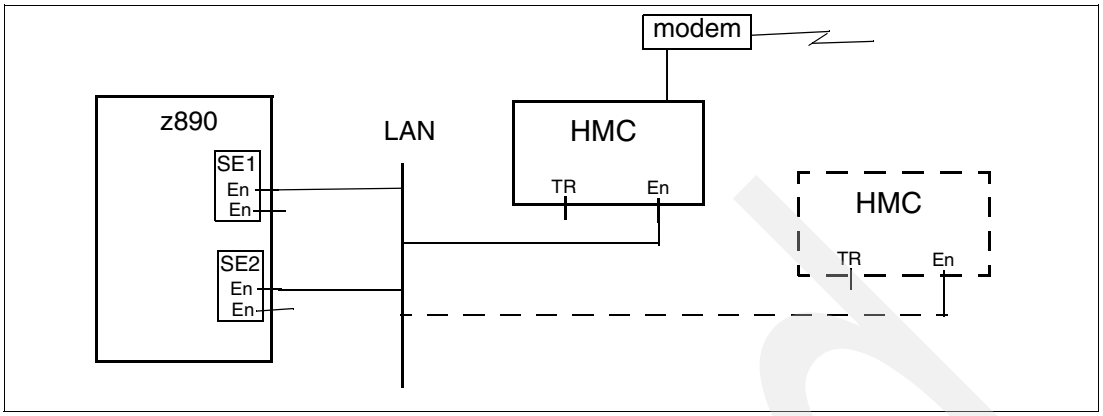

Figure 4-3 Basic SE - HMC connections

Figure 4-3 illustrates a basic SE/HMC configuration. This example uses Ethernet. The mixed token ring/Ethernet SEs would work just as well. A second HMC might be connected to permit operator actions from a different location.

What is *not* shown in this illustration is important—the second Ethernet interfaces are not connected to the LAN. OS/2 (the operating system for SEs and HMCs), in the implementation used for these functions, will not automatically use a second interface to the same LAN as an alternate path if the first interface fails.

The modem shown in the sketch is required if a LAN connection from the HMC to the IBM Network is not available.[17] It is used for Remote Support Facility (RSF) connections to transfer microcode updates (MCLs) and system status information via automated scheduled transmissions. Also, if a system fault occurs, the HMC places a call to the IBM Support System, alerting IBM and transferring error information and logs to reduce the impact of unplanned outages.

Figure 4-4 illustrates use of both LAN features in the SEs. This illustration uses Ethernet, but the principles are the same if token ring is used for one of the LANs.

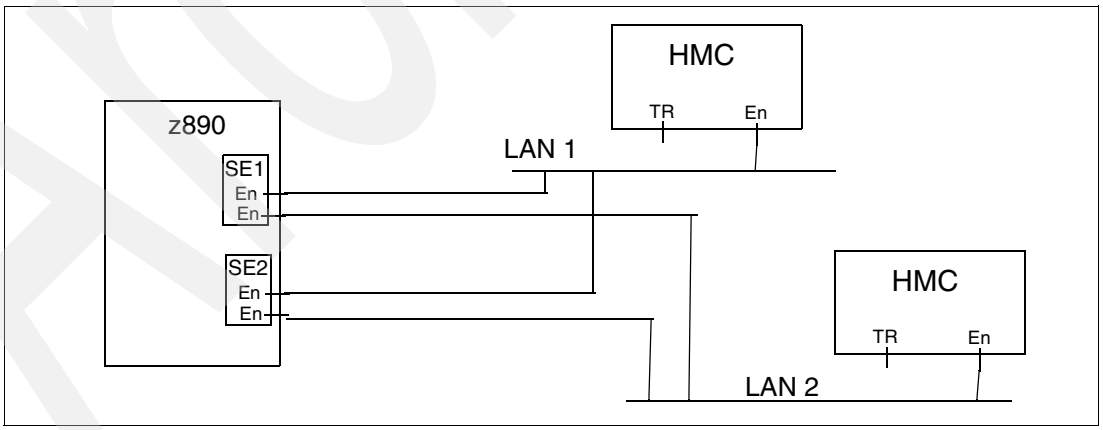

Figure 4-4 Alternate LAN use

In the examples in both figures, other zSeries systems might be connected to the same LANs to share the HMCs.

[17] This option is not available in *secure accounts*, where a connection to the IBM Support System via modem is not allowed. A LAN connection to a public Internet Service Provider is sufficient for connection to IBM facilities.

4.3 Cabling services and cabling migration

Fiber optic cables, cable planning, labeling, and installation are customer responsibilities for new installations and upgrades. Fiber optic conversion kits and Mode Conditioning Patch (MCP) cables are not orderable as features on z890.

To better serve the cabling needs of zSeries customers, IBM Networking Services has enhanced their fiber optic cabling services to better match your requirements, and has a set of services available. These services take into consideration the requirements for all of the protocols/media types supported on zSeries (for example, ESCON, FICON, Coupling Links, and OSA), whether the focus is the data center, the storage area network (SAN), local area network (LAN), or the end-to-end enterprise.

There are three options to provide individual fiber optic cables (jumper cables, conversion kits, MCP cables) for connecting to z800, z900, z990, or z890.

- Option 1 - Fiber optic jumper cabling package

  IBM does the detailed planning. This option includes planning, new cables, installation, and documentation. An analysis of the zSeries channel configuration, I/O devices, and any existing fiber optic cabling is required to determine the appropriate fiber optic cables.

- Option 2 - Fiber optic jumper cable migration and reuse for a zSeries upgrade

  This option includes planning, reuse of existing cables, and documentation. IBM organizes the existing fiber optic cables based upon the new z890 connection details. Relabeling, rerouting, and reconnection to the appropriate z890 channels is performed. New cables are not offered as a part of this option.

- Option 3 - Fiber optic jumper cables and installation

  The customer tells IBM what they need, but the customer does the detailed planning. The service includes new cables, installation, and documentation. Planning and providing the list of required cables are customer responsibilities.

Options 1 and 2 can be combined within one contract to provide complete upgrade coverage.

Under the Enterprise Fiber Cabling Services umbrella there are two options to provide IBM Fiber Transport System (FTS) trunking commodities (fiber optic trunk cables, fiber harnesses, panel-mount boxes) for connecting to the z800, z900, z890, and z990:

- Option 1 - zSeries fiber optic trunk cabling package

  IBM reduces the cable clutter under the floor. An analysis of the zSeries (z800, z900, z890, and z990) channel configuration and any existing fiber optic cabling is performed to determine the required FTS fiber optic trunking commodities (trunk cables, harnesses, panel-mount boxes). This option includes zSeries planning, FTS fiber optic trunking commodities, installation, and documentation.

- Option 2 - Enterprise fiber cabling services

  IBM organizes the entire enterprise. This option includes enterprise planning, new cables, fiber optic trunking commodities, installation, and documentation. This is the most comprehensive set of services.

Under the zSeries Fiber Cabling Services umbrella there are two options to provide individual fiber optic cables (jumper cables, conversion kits, MCP cables) for connecting to z800, z900, z890, and z990:

- Option 1 - Fiber optic jumper cabling package

  IBM does the detailed planning. This option includes planning, new cables, installation, and documentation. An analysis of the zSeries channel configuration, I/O devices, and any existing fiber optic cabling is required to determine the appropriate fiber optic cables.

- Option 2 - Fiber optic jumper cable migration and reuse for a zSeries upgrade

  This option includes planning, reuse of existing cables, and documentation. IBM organizes the existing fiber optic cables based upon the new z890 connection details. Relabeling, rerouting, and reconnection to the appropriate z890 channels is performed. New cables are not offered as a part of this option.

Tip: Additional information is documented in the IBM Redbook: *IBM eServer zSeries Connectivity Handbook*, SG24-5444.

Cabling migration

The z890 does not automatically come with cables. Cabling is a customer responsibility except for ICB-3 and ICB-4 connections.

On z890 server parallel channels are not supported. If you have a control unit that has only parallel channel features and you are planning to use it, you must use an ESCON converter.

Fiber cable connectors used on older servers (Multiprise 2000, Multiprise 3000, and IBM 9672), but also zSeries servers, have been changed on z890. Figure 4-5 shows a subset of the connectors that have been modified. Reuse of existing cables on z890 server demands conversion kits. Figure 4-6 on page 98 presents a sample of conversion kits that are available through the IBM fiber cabling services offering.

Feature	z890 Connector	Previous Connector
FICON Express LX Single mode (SM) fiber	 LC Duplex SM	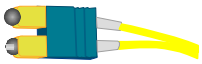 SC Duplex SM
FICON Express SX Multimode (MM) fiber	 LC Duplex MM	 SC Duplex MM
ISC-3 Single mode (SM) fiber	 LC Duplex SM	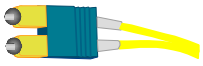 SC Duplex SM
ESCON (16 ports) Multimode (MM) fiber	 MTRJ MM	 ESCON Duplex MM

Figure 4-5 z890 connectors compared to previously used connectors

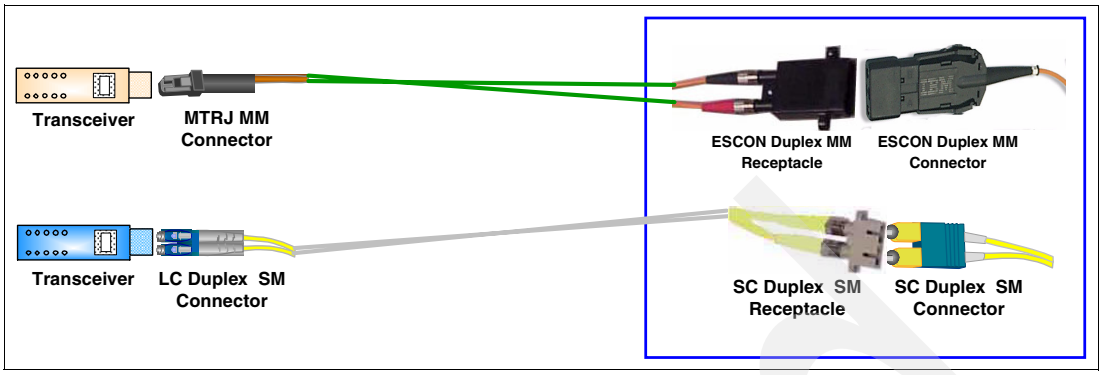

Figure 4-6 Under floor connector conversion kits

Reuse of a multimode trunk cable is also something to consider. If you have a multimode fiber infrastructure, the use of Mode Conditioning Patch (MCP) cables allow reuse of the existing multimode fiber infrastructure with LX features that normally attach to single mode fiber connectors. Two MCP cables are required per link, that is, one at each end. There are some restrictions when using MCP:

- Restricted to 1 Gigabit (100 Mbyte/sec) links only
- Not supported with 2 Gigabit ISC-3 or FICON Express features
- Feature-to-feature distance limited to 550 meters (1804 feet)

As a recommendation, you should consider migrating from a multimode infrastructure to a single mode infrastructure.

Customers with the resources and personnel to plan and implement their own connectivity, or those with less complex system configurations, may consult the following manuals to help them determine and order the required fiber optic cabling:

- *IBM zSeries 890 Installation Manual for Physical Planning*, GC28-6828
- *IBM zSeries 990 Installation Manual for Physical Planning*, GC28-6824
- *IBM zSeries 900 Installation Manual for Physical Planning*, 2064-IMPP
- *IBM zSeries 800 Installation Manual for Physical Planning*, 2066-IMPP
- *IBM zSeries G5/G6 Installation Manual for Physical Planning*, GC22-7106
- *Planning for Fiber Optic Links*, GA23-0367

These manuals are available on the IBM Resource Link Web site:

`http://www.ibm.com/servers/resourcelink`

4.4 I/O and network connectivity

This section provides information related to I/O and network connectivity when moving to a z890.

4.4.1 Parallel channel migration

As stated, z890 systems do not have parallel channels.[18] If you have parallel channel devices, you need to do some planning. Plans typically involve the following choices:

- Retire the devices.[19] This is not as trivial as it may sound. Parallel channel tape and DASD devices are probably quite old and ready for retirement. Maintenance costs may exceed

the cost of replacement devices using ESCON or FICON channels, and newer DASD devices have far greater capacity and performance.

- Purchase converter boxes to connect parallel devices to ESCON channels. IBM made such converters some time ago (IBM 9034), but these are no longer manufactured. IBM recommends the use of the Optica ESCON® Converter. The IBM 9034 units may also be used, but these cannot be ordered through normal IBM ordering channels.
- Keep only selected parallel devices and consolidate them onto as few channels as possible.

The most typical parallel devices that installations might want to keep include:

- IBM 34xx tape drives ("round tapes"), kept for archival functions.
- IBM 3174 control units, used for 3270 operating system consoles and TSO/CICS 3270 terminals. However, the new OSA-Express ICC function may eliminate the need for these control units.
- Various line printers.

There are often valid reasons for keeping these devices. However, we suggest that they can be consolidated onto a reduced number of channels. For example, a single channel could probably handle a mixture of all of the devices mentioned here. You might not normally mix tape drives with other devices on the same channel, but this depends on how often the tape drives are used. Typically, the older drives are so rarely used that their channel sharing characteristics can be ignored.

4.4.2 Byte Multiplexor channel migration

A few parallel devices may require byte multiplexor channels.[20] Typically, there are no modern direct replacements for these devices. Few customers still use these devices, and each might be considered as a special case. Possibly EP and PEP lines are the most common of these.[21] The ESCON-to-parallel converter units may support such devices; you should check with the manufacturer for the latest information. The older IBM 9034 units do support byte multiplexor modes, but may not have been tested with your particular device types. Again, we suggest you consider these as special cases and discuss them with your marketing representatives.

A 9034 unit used with byte multiplexor control units must have serial number 41-53345 or higher and must contain a logic card with part number 42F8047. If use of older 9034 units is required, RPQ 8P1767 should be investigated.

4.4.3 ESCON channel planning

On the z890, ESCON channels are always delivered with the 16 channel I/O feature cards. The older 4-port cards cannot be used. The channels are packaged with 16 ports on a single I/O card. Up to 15 channels on each card are available for use; the last channel is reserved as a spare. In practice, any unallocated ports on the card can act as a spare.

A minimum of two ESCON channel feature cards are always installed if any ESCON channels are configured for the processor. ESCON channels are ordered in groups of four. For configurations greater than 28 ESCON channels, individual (not pairs) ESCON channel

[18] Parallel channels are often known as *bus and tag* or OEMI channels, named for the functions of the two cables used for these channels.
[19] We use the term *devices* here to include the associated control units. It is the control units, of course, that actually use the channels.
[20] Other devices were traditionally used with byte multiplexor channels, but can also be used with block multiplexor channels.
[21] If you do not recognize these names, you need not worry about them.

cards are added as necessary. The active channels are distributed across the physical cards to provide additional redundancy.

If one of the activated ports fails, the system performs a *call home* to inform IBM. An IBM Service Representative will initiate the repair by selecting the “Perform a Repair Action” icon at the Service task panel on the SE. This will start the Repair&Verify procedure.

- If sparing can take place, then the IBM Service Representative moves the external fiber optic cable from the failing port to a spare or unconsidered port on the same card.
- If sparing cannot be performed, the card will be marked for replacement by the procedure. Upon replacement of the ESCON card, the cables that were changed are installed at the *original* port locations. Repair&Verify will recognize the unused ports on the new card as candidates for future sparing.

These ESCON cards, which are also used with z990, z900, and z800 servers, use the MT-RJ connectors. These are different from the traditional ESCON connections that are familiar to most S/390 owners. You can use an ESCON cable with an MT-RJ connector on one end (for the channel connection) and a traditional ESCON connector on the other end (for the control unit). Or you can use conversion cables. See “Cabling services and cabling migration” on page 96 for more details.

Consideration for ESCON conversion channels

There are considerations for ESCON conversion channels, each of which should be connected to an ESCON convertor. When one of the conversion channel types (CHPID type CVC, or CBY) is defined, the channel hardware expects that an ESCON convertor is connected to the channel. If the convertor is not connected, a permanent hardware error may be reported at POR. We recommend that you do not define a conversion channel type until the convertor is actually connected.

4.4.4 OSA considerations

You should be aware of the current OSA features supported on z890 server. Earliest OSA features such as OSA-2, OSA-2 FDDI and OSA ATM are no longer supported on z890, and z990 machines. The recommendations are:

- OSA-2: Can be replace by the equivalent OSA-Express feature. Eventually a new connector may be required. All LAN Channel Stations (LCS) functions are still available, that is, you can use the OSA-Express feature in non-QDIO mode and have TCP/IP passthru and SNA on the same port.
- OSA-2 FDDI: Can be replaced by the equivalent OSA-Express feature. If you are willing to keep the FDDI infrastructure, you should use a multi protocol switch or router with the appropriate network interface. Please refer to the following Web site for more details.

  `http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100340`
- OSA ATM: Can be replaced by OSA-Express 1000BASE-T Ethernet or OSA Express Gigabit Ethernet and a multi protocol switch or router with the appropriate network interface.

The new function available on OSA-Express 1000BASE-T Ethernet, that is, the OSA-Express Integrated Console Controller (OSA-ICC) should also be considered once it can replace or avoid the acquisition of 2074 or 3174 non-SNA console control units. The OSA-ICC is a standard option within the OSA-Express 1000BASE-T Ethernet feature and is made available through the CHPID=OSC coding on HCD. You should consider having two OSA-ICCs on different OSA-Express features for availability reason.

OSA-Express migration

Users migrating from earlier S/390 machines often find OSA-Express concepts to be quite different from previous LAN interface hardware. For example, OSA-Express features combine the traditional channels, control units, and device concepts into a single element—the OSA-Express feature. The OSA-Express feature has two channels (one for each port).

A full description of OSA-Express features is beyond the scope of this publication. We strongly recommend that you consult the IBM Redbook *OSA-Express Implementation Guide,* SG24-5948-01 or later.

A configuration program, known as OSA/SF, is *sometimes* needed to customize an OSA-Express feature. Earlier versions of this program (used with earlier versions of OSA features) required some effort to understand and use. The current OSA-Express features, when used for TCP/IP traffic, have greatly reduced the need to use OSA/SF. Also, newer versions of the OSA/SF program can be used at a workstation and provide easier-to-use GUI interfaces. In addition to feature configuration, OSA/SF can display very useful statistics about LAN usage.

Any 155 ATM or FDDI features must be replaced with different modes of connections, such as an Ethernet or token ring link to a switch or router that provides the ATM or FDDI interfaces.

Table 4-1 LAN interface summary

Feature	CHPID type	SNA/APPN/HPR	TCP/IP	OSA/SF needed
Gigabit Ethernet	OSD (QDIO)	No[a]	Yes	No
Fast Ethernet	OSD (QDIO) OSE (non-QDIO)	No[a] Yes	Yes Yes	No Yes
1000Base-T EN	OSD (QDIO) OSE (non-QDIO)	No[a] Yes	Yes Yes	No Yes
Token Ring	OSD (QDIO) OSE (non-QDIO)	No[a] Yes	Yes Yes	No Yes[b]

a. However, this support can be provided by the Enterprise Extender function.
b. OSA-Express Token Ring and Ethernet requires OSA/SF for non-QDIO except for when it uses the default OAT without port sharing.

FICON Express planning

FICON channels have many benefits over the traditional ESCON and Parallel Channels. To mention some:

- Increased Data Transfer Rate: Up to 200 MB/sec for FICON Express.
- Improved Performance: Native FICON attachment can offer much higher data transfer rates than ESCON attachments. This is important for data-intensive applications that can get significant improvements in overall elapsed time.
- Reduced Backup Windows®: Working together with the appropriated DASD subsystem and Tape subsystem can slash elapsed times for backup operations by up to half.
- Increased Distance: Up to 100 Km, with repeaters, or DWDMs, without data rate droop.
- Channel Aggregation: A single FICON channel can replace multiple ESCON Channels.

For more details about FICON and FICON Express performance refer to *FICON and FICON Express Channel Performance Version 2.0* white paper.

FICON channels can also work as fiber channel (Fiber Channel Protocol) and access SCSI devices, allowing this way the use of Storage Area Network devices for example. This feature can be helpful specially for Linux environments as well for storage consolidation aspects.

4.4.5 I/O feature configuration rules

The following general rules apply for z890 systems:

- A single I/O cage is used, with 28 feature slots. It is not possible to add additional I/O cages. Each domain (4 slots) with an I/O feature card requires an STI connection. The capacity setting model 110 has 16 slots.
- OSA-Express feature cards include:
  - Token Ring.
  - Gigabit Ethernet (GbE).
  - Fast Ethernet (FE) for features carried forward from a z800 system.
  - 1000BASE-T Ethernet, which replaces the OSA Express Fast Ethernet feature. One or both ports on the feature can be used for 3270 function.
- OSA-Express features are ordered in increments of two ports (equal to one feature).
- For FICON Express, OSA-Express, PCICA, and PCIXCC cards, taken together, the maximum is 20 of these feature cards per I/O cage (maximum of 16 features on the capacity setting model 110).
- There is a maximum of 20 FICON Express cards per system, less any restrictions imposed by the number of OSA-Express, PCICA, and PCIXCC cards installed. The capacity setting model 110 can have up to 16 FICON Express cards.
- There is a maximum of 2 PCICA cards.
- There is a maximum of 4 PCIXCC cards.
- ESCON features:
  - Only the 16-port ESCON features may be used.
  - The first group of 4 ESCON channel(s) require two features.
  - After the first two features, single features are added.
  - ESCON channels are ordered in increments of 4.
  - IBM will determine the number of features needed, based on the number of ESCON channels ordered. All the potential channels on a feature might not be enabled, depending on the total number of channels ordered.
  - A maximum of 15 of the 16 ports on a feature will be enabled. The 16th port is a spare.
- ICB links:

  ICB-4 links are for direct z890 or z990 connections and operate at 2.0 GB/second. These links use direct connections to a book and have no connections to I/O cages. There maximum number of links is 8 per system, but may be further limited by the number of available STI connections on the book. A special cable is needed for ICB-4 connections.

  ICB-3 links are for connections to zSeries machines and operate at 1 GB/second. An STI-3 feature card is used in an I/O cage, and this provides two ICB-3 ports. There is a maximum of 8 STI-3 feature cards (providing 16 ICB-3 links) per system.

  - ICB (also known as ICB-2) links are not available for the z890.
- ISC-3 links:

  ISC-3 links are for connections to z890. z990, z900, z800, and 9672 processors (generation 5 and 6 only), and run at 100 or 200 MB/second. A maximum of 48 ISC-3 links in peer mode may be used (maximum of 32 ISC-3 links in compatibility mode). These are packaged on I/O feature cards, and each feature can have two daughter cards. Each daughter feature has two ISC links.

- ICB and ISC links:

  The maximum number of external coupling links (ICB-3, ICB-4, active ISC-3s) and internal links (ICs) cannot exceed 64 per server.

The complete I/O configuration rules are more complex than outlined here. Consult your IBM representative for more specific and timely information.

You can address 256 CHPIDs (that is, channels) within a Logical Channel Subsystem. Two Logical Channel Subsystems may be used with the z890, allowing up to 512 CHPIDs (channels) to be addressed.[22] (However, the maximum number of channels that can be installed is 421, based on 28 ESCON features with 15 channel ports per feature and one ICB-4; the capacity setting model 110 can have up to 240 channels). The physical channel (which, in practical terms, is a connector jack on an I/O feature card) is specified by its PCHID number. You can have more installed PCHIDs (physical channels) than CHPIDs. Your IOCDS definitions connect each CHPID to a PCHID. You must have all your defined CHPIDs connected to PCHIDs to have a valid IOCDS. However, you need not define all your installed PCHIDs (physical channels) in your IOCDS.

Logical Channel Subsystems are, in a sense, entities provided by the system firmware. They are discussed in more detail in *IBM eServer zSeries 990 Technical Guide*, SG24-6947.

Other relevant information about I/O and networking connectivity can be obtained from the *IBM eServer zSeries Connectivity Handbook*, SG24-5444-03.

4.5 Parallel Sysplex considerations

This section provides information applicable to the Parallel Sysplex environment when moving to a z890.

4.5.1 Coupling Facility planning

The z890 and z990 servers do not provide a special model for a CF-only processor (like the model 100 for z900 processors). You can, however, have a z890 or z990 processor with all enabled PUs defined as ICFs. Some precautions should be taken related to:

- The differences between z890 and z990 coupling link connectors and previous zSeries models.
- The ICB link distances and their incompatibility with non-zSeries servers.

The Coupling Facility Level 13 (CFLEVEL 13) introduces the application of patches to Coupling Facility Control Code (CFCC) concurrently, or better, in “rolling” mode. As shown in Figure 4-7 on page 104, on a single CEC you can first apply the patch on a test internal coupling logical partition, perform all necessary tests, and later on, roll the patch to the production logical partition. To use the updated CFCC code to a CF logical partition, simply deactivate and activated the partition. When the CF comes up, it displays its version on the OPRMSG panel for that partition.

With this facility there is no need to stop the entire machine for CFCC patches; however, CFLevel upgrades will still be disruptive to the entire CEC.

[22] There are additional limitations if you use spanned channels.

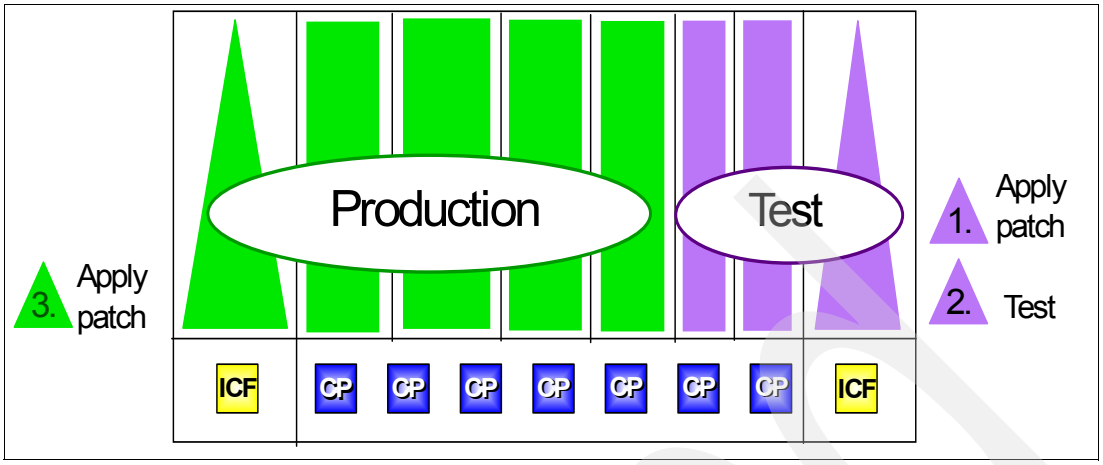

Figure 4-7 Concurrent CFCC Patch

4.5.2 Coupling Links connectivity

Coupling link connectivity should be planned carefully, especially if you have Generation 4 or earlier machines participating on a Parallel Sysplex environment. For this there is no coupling link connectivity to the z890 server. Figure 2-13 on page 48 shows the available connectivity options between z890, and earlier servers.

The connection between z890 and G5/G6 machines can be done with ICS-3 links. If it is the case, an STI-3 feature card with two ICS-3 ports should be ordered. The STI-3 card requires an I/O slot on z890.

4.5.3 Sysplex Timer ETR Network ID

As part of the installation of a Sysplex Timer in either Basic or Expanded Availability configuration, each IBM 9037 Sysplex Timer Unit is assigned a ETR Network ID (0 to 31decimal) and ETR Unit ID (0 to 31decimal). The ETR network ID and ETR Unit ID values are arbitrary (within the valid range) and can be chosen by the customer to uniquely identify an ETR network and a unique ETR unit (Sysplex Timer) within the ETR Network.

For example, in a Sysplex Timer Expanded Availability configuration there are two IBM 9037 Sysplex Timer Units. The first IBM 9037 may have a ETR Network ID value of 0, and a ETR Unit ID value of 0. The second IBM 9037 will have the same ETR Network ID value of 0, but a different ETR Unit ID value of 1. The two Sysplex Timer Units in an Expanded Availability configuration must have the same ETR Network ID value defined. However, their ETR Unit ID values must be unique within the Sysplex Timer ETR Network.

A Sysplex Timer Unit's ETR Network ID, ETR Unit ID, and Port number are transmitted along with timing signals to an attached server's ETR port. This information is available to sysplex systems running on the server and can be displayed using the D ETR command.

A function, introduced with the z990 server, and now also on the z890 server, is implemented in the server's Support Element code, which now requires the ETR Network ID of the attached Sysplex Timer Network to be manually set in the Support Element at installation time. In addition, on the same panel, the ETR ports have to be enabled for stepping the TOD. This function checks that the ETR Network ID being received in the timing signals via each of the server's two ETR ports matches the ETR Network ID manually set in the server's Support Element (SE).

This function provides greater checking, helping eliminate cabling errors where a server's ETR port may be incorrectly connected to a Sysplex Timer Unit in an incorrect Sysplex Timer ETR Network, and allows verification of cabling connectivity from the Sysplex Timer to the z890 server prior to IPL of z/OS or OS/390.

If the ETR Network ID received on an ETR port does not match the value set in the server's Support Element, that ETR port state is made semi-operational by the server. In the semi-operational state, timing signals are still received by the ETR port, but are not used for stepping the server TOD clock. This has some important operational considerations at system IPL time as well as for running sysplex systems:

- If one server ETR port receives a ETR Network ID which does not match the server defined value, that ETR port state is made semi-operational by the server:
  - Systems running with ETRMODE=YES on the server will continue to operate normally using the server TOD clock, which is stepping to the timing signals being received on the second server ETR port. A server hardware error message is generated.
  - At each IPL of a system with ETRMODE=YES on the server, the operating system will alert the server of the state of the ETR port as semi-operational, generating a server hardware error message.
- If both server ETR ports receive a ETR Network ID which does not match the server defined value, both ETR port states are made semi-operational by the server:
  - Systems with ETRMODE=YES already running on the server will immediately enter a non-restartable disabled wait state. A server hardware error message is generated.
  - At each IPL of a system with ETRMODE=YES on the server, the operating system will not be able to enter the sysplex configuration. The customer may choose to continue the IPL in local mode using the server's local TOD clock. The operating system will alert the server of the state of the ETR ports as semi-operational, generating a server hardware error message.

Important: Changing the ETR Network ID to an incorrect value on a z890 server's SE will cause all operating systems with ETRMODE=YES on the server to immediately enter a non-restartable disabled wait state.

Ordering ETR cables

Two ESCON MT-RJ cables for the ETR feature on the zSeries z890 system will automatically be added to any zSeries Fiber Cabling Service contract offered by IBM Global Services. Contact your local IBM Installation Planning Representative, IBM zSeries z890 Product Specialist, or IBM Connectivity Services Specialist for details. If you choose not to use this Service, you may purchase the two ESCON MT-RJ cables separately, or provide them yourself from another source. An ESCON MTRJ-to-ESCON Duplex conversion kit will also have to be provided to connect the MT-RJ cables to the Duplex connectors on the 9037.

Sysplex Timer attachment planning

ETR ports on z890 are optional features and must be ordered separately. However, if you order any CF links it will automatically generate an order for ETR ports.

4.6 Capacity upgrades planning

Capacity upgrades on z890 servers can be either:

- Permanent, via:
  - A regular order (MES), or
  - The Customer Initiated Upgrade (CIU)
- Temporary, via:
  - The On/Off Capacity on Demand (On/Off CoD), or
  - The Capacity Backup (CBU)

Capacity Upgrade on Demand (CUoD) functions are available to provide concurrent capacity upgrades.

While capacity upgrades to the processor itself are concurrent, your software may not be able to take advantage of the increased capacity without performing an Initial Programming Load (IPL), for example, when:

- There are no reserved logical processors defined to the logical partition.
- The operating system does not support dynamic changes of processors' capacity (engine size). Example: z/OS does not support a concurrent upgrade from the capacity model 210 to the model 220, as the capacity of CPs is increased.

Attention: With proper planning and the operating system support, horizontal capacity upgrades can be non-disruptive; as such, upgrades only add more CPs of the same capacity. However, vertical or diagonal capacity upgrades change the CP capacity, and are disruptive for some operating systems, as z/OS.

Note: Only CPs can have sub-capacity; all other z890 processors (zAAPs, IFLs and ICFs) always run at full capacity.

4.6.1 Capacity Upgrade on Demand (CUoD)

Capacity Upgrade on Demand (CUoD) allows for the non-disruptive addition of additional Central Processors (CPs) capacity, Internal Coupling Facilities (ICFs), Integrated Facilities for Linux (IFLs), zSeries Application Assist Processors (zAAPs). CUoD can quickly add processor capacity up to the maximum number of available inactive engines. This provides you with the capacity for much needed dynamic growth in an unpredictable e-business world. The CUoD function, combined with Parallel Sysplex technology, enables virtually unlimited capacity upgrade capability.

The CUoD functions are:

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

Plan Ahead and Concurrent Conditioning

Concurrent Conditioning configures a system for hot plugging of I/O based on a future specified target configuration. Concurrent Conditioning of the z/Series I/O is minimized by the fact that all I/O cards plugging into the zSeries I/O cage are hot pluggable. This means that the only I/O to be conditioned is the I/O cage itself. The question of whether to concurrently condition a cage is an important consideration, especially with the rapid change in the IT environment (e-business) as well as the technology. Migration to FICON Express or

additional OSA-Express networking is exceptionally easy and non-disruptive with the appropriate LIC and if the cage space is available.

The z890 supports concurrent memory upgrade. This capability allows a processor's memory to be increased without disrupting the processor operation. To take advantage of this capability, a customer should not plan processor storage on the 8 or 16 GB increments. If you have a Model A04 with 24 GB of storage you will be able to concurrently upgrade to 32 GB, as there is spare memory capacity on the card. However, if you have 8 or 16 GB of memory installed, to get to the next increment would be disruptive:

- 8 GB to 16 GB (disruptive card change)
- 16 GB to 24 GB (disruptive card change)
- 24 GB to 32 GB (concurrent LIC CC upgrade)

The Plan Ahead process can easily identify your 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, non-disruptive way.

4.6.2 Customer Initiated Upgrade (CIU)

Customer Initiated Upgrade (CIU) is designed to allow you to respond to sudden increased capacity requirements by downloading and applying a Processor Unit (PU) and/or memory upgrade via the Web, using IBM Resource Link and the Remote Support Facility. IBM now has a faster process for upgrading your server. With the Express option on CIU, an upgrade may be installed as fast as within a few hours after order submission.

Orders (MESs) of Processor Units (PUs) and memory can be delivered by Licensed Internal Code; Control Code (LIC CC) is eligible for CIU delivery. This includes the upgrade of PUs and limited memory upgrades (from 24 GB to 32 GB only) for the z890, up to the maximum available on the installed server.

4.6.3 On/Off Capacity on Demand (On/Off CoD)

On/Off Capacity on Demand (On/Off CoD) is used when you need short term additional capacity. On/Off CoD is designed to temporarily turn on previously uncharacterized Processor Units (PUs), or any unassigned Integrated Facilities for Linux (IFLs) that are available within the current machine, as CPs, Integrated Coupling Facilities (ICFs), IFLs, and zSeries Application Assist Processors (zAAPs).

On/Off CoD capacity setting is restricted to the base capacity. For example, capacity setting 110 is eligible for On/Off CoD upgrades to only capacity settings 120 and 210. All other target capacity settings would result in more than two times the base capacity.

The maximum number of On/Off CoD zAAPs available for z890 cannot exceed the number of zAAPs, with the additional restriction that the sum of zAAPs and On/Off CoD zAAPs cannot exceed the number of CPs.

Activation of this capability is mutually exclusive with Capacity Backup Upgrade (CBU) activation. Both On/Off CoD and CBU can reside on the server, but only one can be activated at a time. On/Off CoD is delivered through the function of Customer Initiated Upgrade (CIU). To participate in this offering, you must have installed CIU Enablement (FC 9898) and On/Off CoD Enablement (FC 9896). Subsequently, you may concurrently install temporary capacity by ordering On/Off CoD Use Days (one of FC 6121 through FC 6471), On/Off CoD Active IFLs (FC 9888), On/Off CoD Active ICFs (FC 9889), or On/Off CoD Active zAAP (FC 9893) up to the current capacity or the number of IFLs, ICFs, and zAAPs, respectively, and use the additional capacity for an indeterminate time.

You will be billed for each On/Off CoD capacity turned on in any given 24-hour period continuing until such On/Off CoD capacity is turned off. Each month your bill will be calculated for the sum of all orders installed within the prior month. Monitoring will occur through the server call home facility and a bill will be generated if the capacity has been enabled for any portion of a calendar month. You will continue to be billed for use of temporary capacity until you return the server to the original state. After concurrently returning to the original state, you may choose to activate a new On/Off CoD upgrade which can be different from the previous upgrade. When you dispose of the server, or decide that you want to disable future temporary upgrades, you are required to remove the enablement feature, On/Off CoD Enablement (FC 9896).

4.6.4 Capacity Backup (CBU)

The z890 Capacity Backup Upgrade (CBU) capability is the non-disruptive temporary addition of Central Processors (CPs) on servers with the CBU feature installed. It enables enterprises to provide flexible, cost-effective Disaster Recovery on the z890. The CBU feature provides the ability to concurrently increment the capacity of your processor, using LIC CC, in the event of an unforeseen loss of substantial zSeries computing capability at one or more of your eligible sites.

A Special Bid Contract for CBU must be approved before the CBU features can be ordered. The feature identifies how many CPs are in the Capacity Backup to IBM's vital product data base for the customer system. This contract is required and the model must be configured for the system when additional capacity is invoked. I/O cage channel requirements must be planned in advance and should already be in the configuration prior to the test or CBU being invoked.

CBU is done concurrently with system operations and performs the following:

- Activating CBU via Password Panel
- CBU Automatic Activation using the IBM Service Support System
- Viewing CBU Feature Information Panel
- Automatic Enablement of CBU for the GDPS®
- Concurrent Undo CBU

CBU upgrades apply to whole CP engines additions and only to the largest capacity configuration (full engine). This means that a z890 capacity setting model 120 (FC 6120, one sub-capacity CP) with 2 CBU features (FC 6800) installed will result on a capacity model 370 (FC 6370, three full capacity CPs) after the CBU activation. So, any CBU activation will result on an upgraded capacity setting model 270, 370, or 470, as shown in Table 4-2.

Table 4-2 Capacity setting models after CBU activation

Capacity Setting Model Range	Plus 1 CBU (one FC 6800)	Plus 2 CBUs (two FC 6800)	Plus 3 CBUs (three FC 6800)
110 to 170	270	370	470
210 to 270	370	470	N/A
310 to 370	470	N/A	N/A
410 to 470	N/A	N/A	N/A

CBU activations of any capacity setting model having sub-capacity CPs result on processor capacity (engine size) increases. Some operating systems (like z/OS) do not support dynamic changes of CPs capacity, and will require an Initial Programming Load (IPL) after the hardware upgrade.

4.7 I/O configuration definition and management

The architectural enhancements of the z890 enforce a new approach to configuration management. Every CHPID is mapped to a PCHID; it is mandatory that every CHPID has a PCHID associated with it. For internal channels, such as IC links, and HiperSockets, CHPIDs are not assigned a PCHID.

The z890 does *not* have default CHPIDs assigned to channel ports as part of the initial configuration process. CHPIDs are assigned to physical channel path identifiers (PCHIDs) in the IOCP input file. It is the customer's responsibility to perform these assignments by using the HCD/IOCP definitions or by importing the output of the CHPID Mapping Tool.

4.7.1 I/O configuration definition

Tools are provided to maintain and optimize the I/O configuration of a z890.

Hardware Configuration Dialog (HCD)

HCD supplies an interactive dialog to generate your I/O definition file (IODF) and subsequently your Input/Output Configuration Data Set (IOCDS). It is strongly recommended that HCD or HCM be used to generate your I/O configuration, as opposed to writing your own IOCP. The validation checking that HCD performs as you enter data helps eliminate errors before you implement your I/O configuration.

CHPID Mapping Tool (CMT)

The CHPID Mapping Tool provides a mechanism to map CHPIDs onto PCHIDs as required on a z890. Additional enhancements have been built into the CMT to cater for the new requirements of the zSeries; it provides the best availability recommendations for the installed z890 features and defined configuration.

When building an IOCP, even under the stand alone version of IOCP, you must match the physical channels IDs (PCHIDs) with the channel IDs (CHPID). Recall that it is possible to have up two sets of 240 channels on a z890 server, that is, in two logical channel subsystems. Each set of channels could have, for example, two CHPIDs X'10': one for LCSS 0 and other for LCSS1. The manner to identify each CHPID exclusively is by associating a different PCHID to each CHPID. The CHPID mapping tool available from IBM Resource Link, can help you with this task.

For further details on the CMT, refer to "I/O configuration management" for a brief introduction, and to *IBM eServer zSeries Connectivity Handbook,* SG24-5444, for a comprehensive explanation.

Stand Alone IOCP

Stand alone IOCP programs now run in a logical partition. On earlier machines, IOCP stand alone requires basic mode to run. If you are moving from an older machine to z890 and you are running an earlier operating system version not supported on z890 server, maybe you will have to create a stand alone IOCP to start up the new environment. You can obtain more information about stand alone IOCP in the *z890 IOCP Users Guide.*

4.7.2 I/O configuration management

It is recommended that the CHIPD Mapping Tool (CMT) be used for all new build z890 configurations, or when upgrading from a z800. It can also be used as part of standard hardware changes to your installed z890.

The CHIPD Mapping Tool takes input from two sources:

1. The Configuration Report file (CFreport) produced by the IBM order tool (e-Config) can be obtained from your IBM representative, or the Hardware Configuration File produced by IBM manufacturing can be obtained from IBM Resource Link.
2. An IOCP statement file.

The following output is produced by the CMT:

- Tailored reports. All reports should be saved for reference. The Port Report sorted by CHPID number and location should be supplied to your IBM hardware service representative for the z890 installation.
- An IOCP input file with PCHIDs mapped to CHPIDs. This IOCP input file can then be migrated back into HCD from which a production IODF can be built.

Important: When an IOCP statement file is exported from a Validated Work IODF using HCD, it must be imported back to HCD for the process to be valid. The IOCP file cannot be used directly by the IOCP program.

The configuration management process is reflected in Figure 4-8.

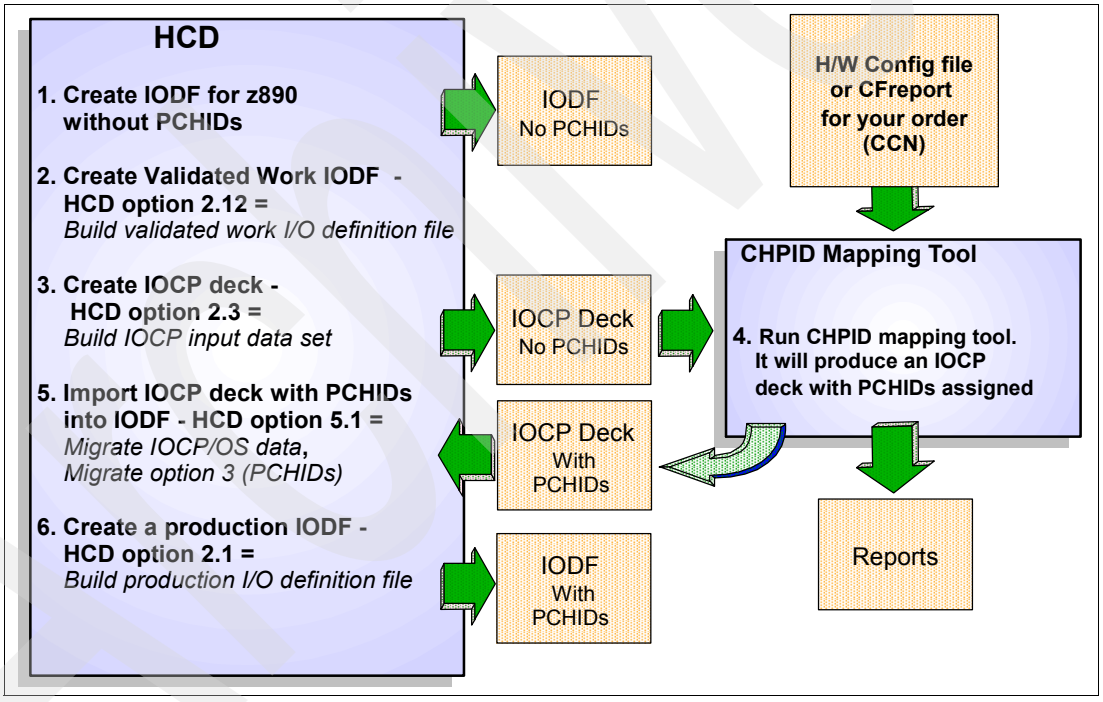

Figure 4-8 z890 I/O configuration definition flow

The z890 elements, LCSS, CHPID, and PCHID must be included in an IOCDS. An IOCDS is created through the HCD utility program or by use of a standalone IOCP file. In practical use, HCD is normally used. For illustration purposes we will examine an IOCP file and will assume the reader is generally familiar with such files.

4.7.3 IOCP example

Figure 4-9 on page 111 illustrates the relationship between logical partitions, LCSSs, CHPIDs, and PCHIDs. This illustration also shows spanned FICON channels.

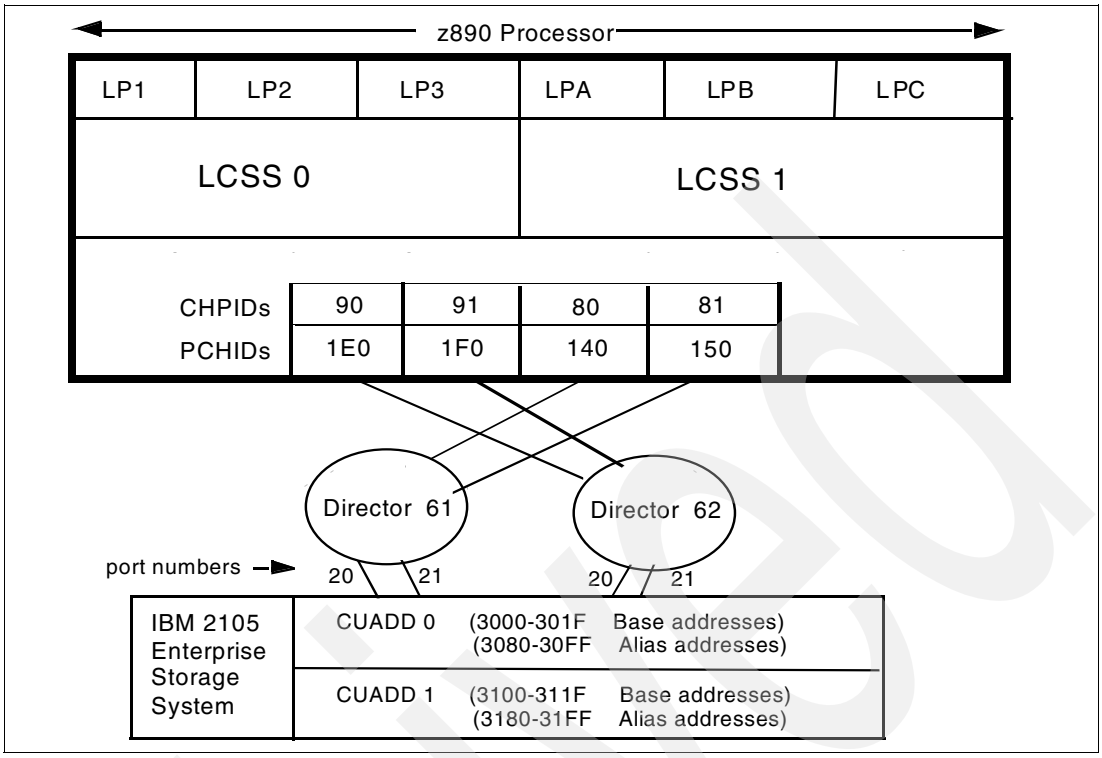

Figure 4-9 Overview of LPARs, LCSSs, CHPIDs, and PCHIDs

Depicted in the figure are a number of important points, such as:

- Two LCSSs are shown, which is the maximum supported with the z890.
- A logical partition (LP) is associated with a specific LCSS. (Also note that LPs have unique names across the complete system. LP naming and numbering has become a little more complex and this is discussed in “LPAR mode updates” on page 116.)
- Multiple LPs (up to 15) may be associated with an LCSS.
- A CHPID is associated to a PCHID, and PCHID numbers are unique across the server and are based on the physical location of the I/O port.
- A spanned channel is connected to more than one LCSS and occupies the same CHPID and PCHID in all LCSSs in which it is used. For example, all our FICON channels (CHPIDs and PCHIDs) defined to CSS 0 are also defined to CSS 1.

The IOCP definitions shown in Example 4-3 describe the system shown in Figure 4-9.[23] This is not intended to represent a practical system—it has no consoles, for example—but it illustrates the z890 definition elements.

Table 4-3 IOCP definitions for spanned channels

```
ID MSG1='BILLIOCP',MSG2='z890 2 LCSS',SYSTEM=(2086,1)
RESOURCE PARTITION=((CSS(0),(LP1,1),(LP2,2),LP3,3)),                    X
  (CSS(1),(LPA,1),(LPB,2),LPC,3))),                                      X
  MAXDEV=((CSS(0),64512),(CSS(1),64512))
CHPID PATH=(CSS(0,1),80),SHARED,PARTITION=((LP1,LP2,LP3),(LPA,LPB,LPC)), X
  SWITCH=61,TYPE=FC,PCHID=140
CHPID PATH=(CSS(0,1),81),SHARED,PARTITION=((LP1,LP2,LP3),(LPA,LPB,LPC)), X
```

[23] We cheated a little in this listing. Some lines are too long for the required IOCP format and should be broken into continuation lines. Also, we may have gone to a continuation line at an inappropriate point. We did this in the interest of readability. The “X” continuation indicators should be in column 72 of a proper IOCP file. Also, the PCHID number associated with a CHPID is placed on a separate line by the CHPID Mapping Tool and is typically listed that way; however, it can be entered as shown here.

```
  SWITCH=61,TYPE=FC,PCHID=150
CHPID PATH=(CSS(0,1),90),SHARED,PARTITION=((LP1,LP2,LP3),(LPA,LPB,LPC)),        X
  SWITCH=62,TYPE=FC,PCHID=1E0
CHPID PATH=(CSS(0,1),91),SHARED,PARTITION=((LP1,LP2,LP3),(LPA,LPB,LPC)),        X
  SWITCH=62,TYPE=FC,PCHID=1F0
CNTLUNIT CUNUMBR=3000,                                                          X
  PATH=((CSS(0),80,81,90,91),(CSS(1),80,81,90,91)),                             X
  UNITADD=((00,256)),CUADD=0,UNIT=2105,                                         X
  LINK=((CSS(0),20,21,20,21),(CSS(1),20,21,20,21))
CNTLUNIT CUNUMBR=3100,                                                          X
  PATH=((CSS(0),80,81,90,91),(CSS(1),80,81,90,91)),                             X
  UNITADD=((00,256)),CUADD=1,UNIT=2105,                                         X
  LINK=((CSS(0),20,21,20,21),(CSS(1),20,21,20,21))
IODEVICE ADDRESS=(3000,032),CUNUMBR=(3000),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(3080,128),CUNUMBR=(3000),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(3100,032),CUNUMBR=(3100),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(3180,128),CUNUMBR=(3100),STADET=Y,UNIT=3390A
END
```

The z890 parameters elements in this IOCP listing include:

- *CSS definitions.* The LCSSs used in the IOCDS created from this IOCP are stated in the RESOURCE statement. In the example, we define two LCSSs and indicate which LPs are associated with each one. This *defines* the LCSSs.
- *Subchannel maximums.* The MAXDEV parameter (in the RESOURCE statement) specifies the maximum number of subchannels (device numbers) that can be defined in an LCSS. (The MAXDEV value includes the devices you have defined plus whatever expansion number you need.) This affects the HSA storage used.
- *CSS number for CHPID.* Each CHPID statement specifies its LCSS number as part of the PATH parameter.
- *PCHID number.* Each CHPID statement *must* include a PCHID parameter to associate the CHPID with a physical channel. (The PCHID parameters can be added to the IOCP definitions by the CHPID Mapping Tool, through HCD definitions, or defined in a standalone IOCP file, but the PCHID parameters must be present in all CHPID statements in order to create an IOCDS.)
- *CSS number for PATH and LINK parameters.* PATH and LINK parameters in CNTLUNIT statements must indicate the LCSS number associated each path and link.
- *A CHPID is designated as being spanned*, if it has multiple CSS values assigned on the CHPID IOCP statement and also has the SHARED keyword specified. Within HCD on z/OS, a CHPID is designated as spanned if the Operation Mode is specified as SPAN when it is defined.

As can be seen from this example, the basic concepts and definitions for multiple logical channels subsystems and spanned channels are straightforward and mesh with existing IOCP parameters. Equivalent fields have been added to HCD panels for entering LCSS and PCHID numbers.

Channel definitions in the IOCP statement

The following channel types (as defined in an IOCDS) are used with the z890:

- FICON channel types
  - FC - Native FICON channel (both for native FICON devices and FICON CTCs)
  - FCV - FICON bridge channel
  - FCP - Fibre Channel Protocol (SCSI)

- ESCON channel types
  - CNC - Native ESCON channel
  - CTC - ESCON CTC channel
  - CVC - ES conversion channel, which connects to a converter in block multiplexer mode
  - CBY - ES conversion channel, which connects to a converter in byte multiplexer mode

  Note that ESCON channels cannot be spanned.
- CF link channel types
  - CBP - Integrated Coupling Bus (ICB-3) channel, for both OS and CF partitions, to connect to z900 or z800 servers
  - CBP - An ICB-4 connection to another z890 or z990
  - CFP - InterSystem Coupling (ISC-3) peer mode channel, for both OS partitions and a CF partition
  - CFS - ISC-3 compatibility mode sender channel, for OS partitions
  - CFR - ISC-3 compatibility mode receiver channel, each must be used for only 1 CF partition
  - ICP - Peer mode Internal Coupling (IC-3) channel, for both OS and CF partitions, to connect among LPARs within a z890 system internally
- OSA-Express channel types
  - OSC - OSA-ICC (Integrated Console Controller)
  - OSD - OSA-Express (QDIO)
  - OSE - OSA-Express (LCS)
- HiperSockets channel type

  IQD - HiperSockets channel, QDIO mode only

Each of these channel types requires that a CHPID be defined, even if it is an internal channel and no physical hardware (channel card) exists. Each channel, whether a “physical” channel device or a virtual device (such as a HiperSocket) must be assigned a unique CHPID within the LCSS. There are no default CHPID numbers for the z890 and you can arbitrarily assign whatever number you like (within the X’00’ to X’FF’ range, of course).

Most of these channel types can be shared and used concurrently among multiple LPs within the same LCSS. This capability is known as the Multiple Image Facility (MIF). Exceptions are for ESCON conversion channels (CVC and CBY) and CF receiver channels (CFR). These channel types cannot be shared concurrently, but can be defined as reconfigurable channels by specifying the REC parameter on the channel definition. The channel can be reassigned to another LP after the former owning LP configures the channel offline.

Parallel channels, ICB-2 compatibility link to G5/G6 systems, and OSA-2 features are not supported on the z890. Therefore, CHPID types BY, BL, CBS, CBR and OSA are not allowed.

For more information on IOCP definitions, refer to *Input/Output Configuration Program User’s Guide,* SB10-7037.

4.8 LPAR planning

The first noticeable difference between the z890 and its predecessors is that it requires that all operating systems run in a logical partition. Meaning that the z890 (as the z990) does not support basic mode. The standalone IOCP program was modified to only support LPAR mode for the z890 and z990 servers.

4.8.1 LPAR concepts

Since the concept of logical partitions may be new for many potential users of the z890, a short description plus some usability information about logical partitioning follows.

Processor Resource/System Manager (PR/SM)

Processor Resource/System Manager (PR/SM)[24] is a standard hardware facility (microcode) that permits partitioning of a physical server into several logical servers. This is known as logical partitioning or simply LPAR mode (see Figure 4-10). When running in LPAR mode, as a z890 always does, you can isolate workloads in different logical images, each using their own operating system. For example, production workloads can run separated from test workloads.

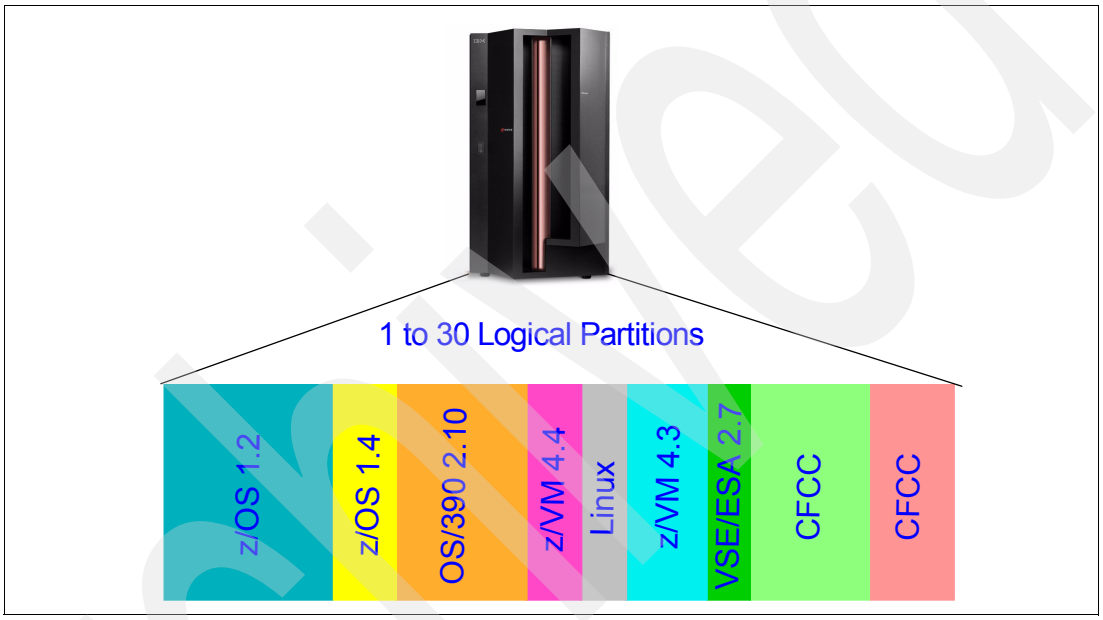

Figure 4-10 z890 in LPAR mode

LPAR mode has the following properties:

- Each logical partition is defined as a set of (logical) resources that are mapped onto a set of physical resources (processor units, memory, and channels). In each logical partition an operating system can be activated. Current operating systems supported to run in a logical partition are z/OS, OS/390, z/VM, Linux, VSE/ESA, and the Coupling Facility Control Code (CFCC).
- Up to 30 logical partitions can be defined on a z890 (except for capacity setting model 110, which is limited to 15 logical partitions).
- A logical partition is always related to a Logical Channel Subsystem (LCSS), of which two can be specified on a z890. Up to 15 logical partitions can be defined per LCSS. Each logical partition is defined through IOCP/HCD. For example, the IOCP statement RESOURCE PARTITION=((CSS(0),(PROD,1),(TEST,2)) statement defines two logical partitions (PROD and TEST). A Power-on-Reset (POR) is required to add or remove logical partitions from the system. However, a dynamic add/delete of a logical partition *name* can be used if reserved logical partition slots have been previously defined (see “Dynamic Add/Delete of a logical partition name” on page 117).

[24] Most people use the terms LPAR and PR/SM interchangeably although it is not strictly accurate. Similarly, many people use the term LPAR when referring to an individual logical partition.

- There are many options that can be assigned to a logical partition, for example:
  - The number of logical processors, and the amount of memory and channels.
  - The logical partition weight, that is, the relevance of a logical partition to the others in the same server when scarce resources must be shared.
  - LPAR capping for a logical partition set a cap on processor resource usage by the logical partition.
  - Security, and other characteristics.

  These parameters are defined and stored in the logical partition Image Profile on the HMC.
- Individual physical processors can be shared between multiple logical partitions, or they can be dedicated for use by a single logical partition.
- Channels can be dedicated, reconfigurable (dedicated to one logical partition, but able to be switched manually between logical partitions), or shared (except for ESCON channel that are connected to a “parallel channel converter”).
- Memory defined to a logical partition is a dedicated resource, but can be reconfigured from one logical partition to another.

Physical processors can be dedicated or shared (see Figure 4-11).

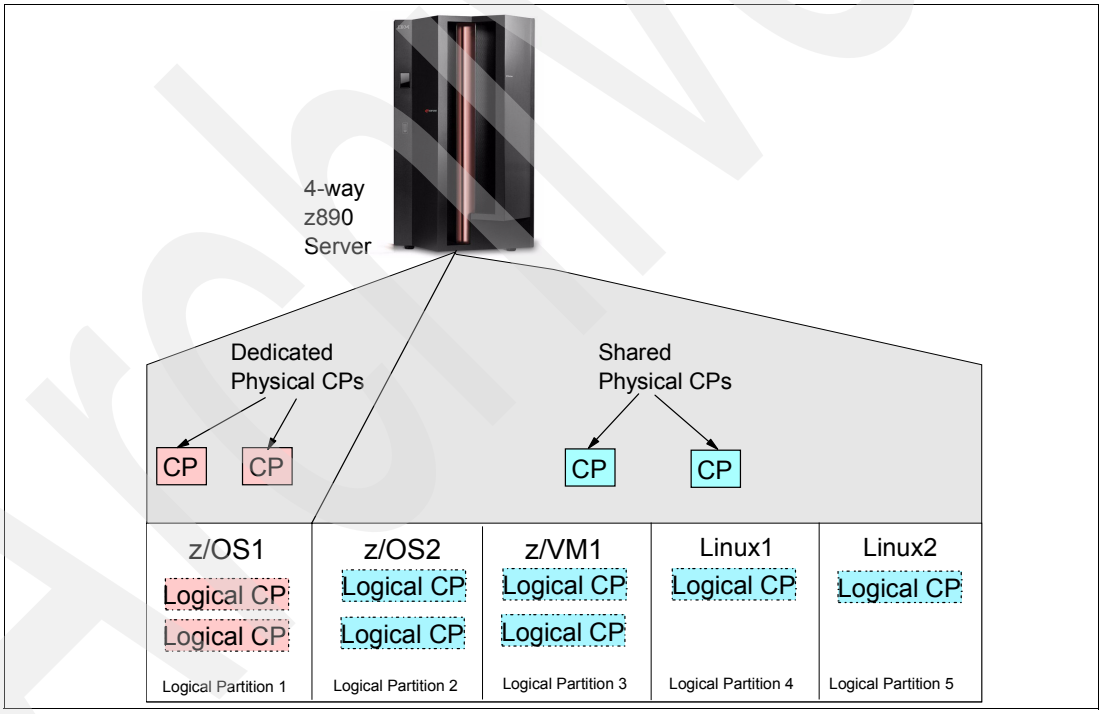

Figure 4-11 Shared and dedicated processors

When dedicated, the logical processors of a partition are permanently assigned to a physical processor. When shared, a logical processor is temporarily associated with a physical processor during an interval of time when the logical partition is being dispatched by PR/SM.

The maximum number of logical processors for any logical partition cannot exceed the number of physical processors on the server. If dedicated logical processors are defined to a partition, the maximum number for logical processors for any other partition cannot exceed the number of the server’s physical processors *minus* the number of dedicated logical processors.

For simplicity, we show processor units (PUs) that are characterized as CPs. In an environment where only z/VM logical partitions with Linux guests are running, IFLs may be used. Refer to “Processor unit functions” on page 22 for a description of the different characterizations of PUs.

Detailed information about PR/SM can be obtained from *zSeries PR/SM Planning Guide,* SB10-7036.

4.8.2 LPAR mode updates

Except for the model with the smallest capacity setting (model 110), a z890 can have up to 30 logical partitions. The key thing here is that the previous limit of 15 logical partitions has been extended and that identification of more than 15 logical partitions no longer fits in a 4-bit field. Therefore, the logical partition ID now is registered in a 1-byte field. The logical partition ID is specified by the user as a part of the LPAR activation profile on the HMC. There are a few additional things that are important to remember, such as:

- A z890 only runs in LPAR mode. Basic mode is not supported.
- A maximum of 15 logical partitions may be associated with a single Logical Channel Subsystem.
- Operating systems with compatibility support can only be activated in a logical partition related to LCSS 0 with a logical partition ID equal or smaller than 15 (x’F’).
- Partition identifiers, names and numbers associated with a logical partition:
  - *Logical partition identifier*. This is a identifier in the range 0 - X’3F’ and is assigned by the user when defining logical partition Image profiles through the SE or HMC. It is unique across the z890 server. This identifier is returned by an STIDP or STSI instruction. It is also known as the user logical partition ID (UPID).
  - *MIF Image identifier*.[25] This identifier is defined through HCD or IOCP and is the partition number defined in the RESOURCE PARTITION statement in the IOCP source. It is in the range x’1’ to xX’F’ and is unique *within* an LCSS. It does not need to be unique within a z890 system. The MIF Image identifier is also known as the IID.
  - *Logical partition name*. This name is defined through HCD or an IOCP and is the partition name in the RESOURCE statement in the IOCP source. A logical partition name must be unique across the z890 server.
  - *Logical partition number*. This number is not specifiable by the user for the z890. It is assigned during Power-On Reset POR), by PR/SM.

Figure 4-12 on page 117 shows where the LCSS, logical partition, and MIF definitions are defined.

[25] MIF is Multiple Image Facility. It was formerly known as EMIF for ESCON Multiple Image Facility when introduced for ESCON channels. This function now relates to all I/O, so the “E” has been dropped.

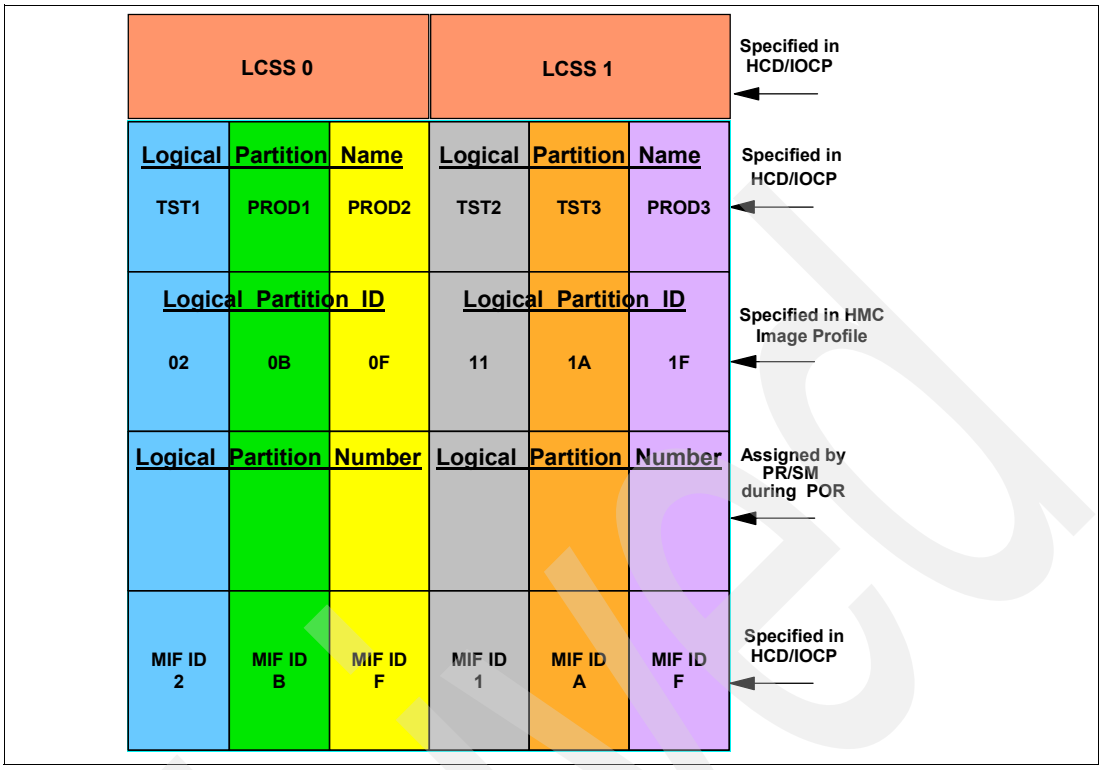

Figure 4-12 Where to define LCSS, logical partition, and MIF definitions

The logical partition ID is specified when defining activation profiles. The HCD field for the MIF Image ID has the label “Partition Number” and this is equivalent to the MIF Image ID used in IOCP RESOURCE statement. We suggest you establish a naming convention for the logical partition identifiers. For example, you could use the LCSS number concatenated to the MIF Image ID, which means logical partition ID 1A is in LCSS 1 with MIF ID A. This fits within the allowed range of Logical partition IDs and conveys useful information to the user.

Dynamic Add/Delete of a logical partition name

The ability to add meaningful logical partition names to the configuration without a Power-On Reset is being introduced. Prior to this support, extra logical partitions were defined by adding reserved names in the Input/Output Configuration Data Set (IOCDS), but one may not have been able to predict what might be meaningful names in advance.

Dynamic add/delete of a logical partition name allows reserved logical partition 'slots' to be created in an IOCDS in the form of extra logical channel subsystem (CSS), Multiple Image Facility (MIF) image ID pairs. A reserved partition is defined with the partition name placeholder ‘ * ’, and cannot be assigned to access or candidate list of channel paths or devices. These extra logical channel subsystem MIF image ID pairs (CSSID/MIFID) can be later assigned an logical partition name for use (or later removed) via dynamic I/O commands using the Hardware Configuration Definition (HCD). The IOCDS still must have the extra I/O slots defined in advance since many structures are built based upon these major I/O control blocks in the Hardware System Area (HSA). This support is exclusive to the z890 and z990 and is applicable to z/OS V1.6, which is planned to be available in September 2004.

When a logical partition is renamed, its name can be changed from ’NAME1’ to ‘ * ’ and then changed again from ‘ * ’ to ‘NAME2’. The logical partition number and MIFID are retained across the logical partition name change. However, the master keys in PCIXCC that were

associated with the old logical partition ‘NAME1’ are retained. There is no explicit action taken against a cryptographic component for this.

Attention: Cryptographic cards are not tied to partition numbers or MIFIDs. They are set up with AP numbers and domain indices. These are assigned to a partition profile of a given name. The customer assigns these "lanes" to the partitions now and continues to have responsibility to clear them out if he changes who is using them.

4.9 Operating Systems migration

Operating systems support for z890 is summarized in “Operating system support” on page 64. The summary is repeated here for convenience reasons.

Table 4-4 z890 software support summary

Operating system	ESA/390 (31-bit)	z/Arch. (64-bit)	Compatibility	Exploitation
OS/390® Version 2 Release 10	Yes	Yes	Yes	No
z/OS Version 1 Release 2	No[a]	Yes	Yes	No
z/OS and z/OS.e Version 1 Release 3	No[a]	Yes	Yes	No
z/OS and z/OS.e Version 1 Release 4	No[a]	Yes	Yes	Yes
z/OS and z/OS.e Version 1 Release 5 and 6	No	Yes	Included	Included
Linux for S/390	Yes	No	Yes	Yes
Linux® on zSeries	No	Yes	Yes	Yes
z/VM Version 3 Release 1	Yes	Yes	Yes	No
z/VM™ Version 4 Release 3	Yes	Yes	Yes	No
z/VM Version 4 Release 4	Yes	Yes	Included	Included
z/VM Version 5 Release 1	No	Yes	Included	Included
VSE/ESA™ Version 2 Release 6 and 7	Yes	No	Yes	Yes
z/VSE Version 3 Release 1[b]	Yes	No	Included	Included
TPF Version 4 Release 1 (ESA mode only)	Yes	No	Yes	No

a. 31-bit mode is only available as part of the z/OS Bimodal Migration Accommodation software program. The program is intended to provide fallback support to 31-bit mode in the event that it is required during migration to z/OS in z/Architecture mode (64-bit).
b. The z/VSE operating system can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. The z/VSE operating system is designed to exploit select features of IBM eServer zSeries hardware.

If you are running an operating systems other than the ones listed above, that is, an unsupported operating system version, you must plan the software migration steps to z890 and the corresponding supported OS carefully. As a general rule, it is not recommended that you migrate software and hardware at same time. However, there are circumstances where this is unavoidable. Careful planning is required.

4.9.1 z/OS and z/OS.e migration

z/OS or z/OS.e migration requires some attention especially if you are running earlier versions of the operating system not supported on z890 server. Presented here are some general cases and ideas to migrate to z/OS, that may be helpful. However, as a recommendation, if you are really back-level in terms of the operating system, ask your IBM representative help to create an accurate and responsible migration plan. Keep in mind that there may be other issues, that may not be addressed here. Also, IBM Global Services has services offerings that can help you with the migration efforts.

Migrating from earlier versions of OS/390 usually is done in two steps:

1. Move to OS/390 V2R10 and after that to z/OS. This is especially interesting if you have a Generation 4 or earlier S/390 CMOS server that *can not* run z/OS, as shown in Figure 4-13. If this is the case, you should migrate first to OS/390 V2R10.
2. Move to the z890 server and migrate from OS/390 V2R10 to z/OS V1R4 or z/OS.e V1R4 or later. This scenario also addresses coexistence, that is, you can keep your S/390 server with OS/390 V2R10 and the z890 server with z/OS for a certain period of time until a stable environment is established.

If you run on a S/390 server capable to run the z/OS operating systems and you are on OS/390 V2R9 or earlier, you should consider moving to z/OS prior to moving to z890. For example, you could first install z/OS V1R4 on your current S/390 server and later move to the z890. In this case, coexistence of both environments is also addressed.

Figure 4-13 summarizes z/OS and z/OS.e support for z890 and other servers.

		G3-G4	G5/G6 Multiprise 3000	z800	z890	z900	z990	End of Service	Coexists with z/OS...	Planned Ship Date
OS/390	2.8	x	x	x		x		9/02	1.2	
	2.9	x	x	x		x		9/03	1.3	
	2.10	x	x	x	x_c	x	x_c	9/04	1.4	
z/OS	1.1		x	x		x		3/04	1.4	
	1.2		x	x	x_c	x	x_c	10/04	1.5	
	1.3*		x	x	x_c	x	x_c	3/05	1.6	
	1.4*		x	x	x	x	x	3/07	1.7	9/02
	1.5*		x	x	x	x	x	3/07**	1.8	3/04
	1.6*			x	x	x	x	9/07**	TBD	9/04

x_c - Compatibility support - does not exploit new z890/z990 features: 30 Logical Partitions and multiple Logical Channel SubSystems
IBM Bimodal Accommodation Offering is available for z/OS 1.2, 1.3, and 1.4 (not z/OS.e). It will not be provided for z/OS 1.5
z/OS 1.4 will remain orderable until Sept 9, 2004
z/OS 1.5 will be orderable between March13 and Sept. 9, 2004
*z/OS.e - z800 and z890 only
**Planned

Figure 4-13 z/OS and zOS.e support summary

If you are already running z/OS, the migration is easier except for those that are running z/OS V1R1. Remember that z/OS V1R1 is not supported on a z890, so you should first install a later version of z/OS or z/OS.e prior to migrating to the z890 server. If you have a z/OS or z/OS.e version that is supported, you have to install the compatibility support if this is not yet included in the current version of z/OS or z/OS.e. Refer to the 2086DEVICE preventive service planning (PSP) bucket for more detailed information. New orders of z/OS or z/OS.e operating systems already include compatibility support z/OS V1R5 and later.

If you want to exploit two logical channel subsystems and some other features of the z890, you should install exploitation support on z/OS or z/OS.e V1R4 or install V1R5 or later that has exploitation support included.

It is always important to recall that z/OS, and z/OS.e V1R5 and later will run only in 64 bits mode and will require a z/Series server. These releases do not support the bimodal Accommodation Offering (see caption in Table 4-4 on page 118).

4.9.2 z/VM migration

z/VM support on the z890 starts with z/VM V3R1 and later releases in compatibility mode, from z/VM V4R4 and later, exploitation mode is supported.

If you are running any other version of VM, like VM/ESA®, for example, you must install z/VM prior to moving to the z890 server. If you have a Generation 4 S/390 CMOS server or earlier, the only alternative you have is to make an intermediate step with z/VM V3R1, since z/VM V4R3 and later releases require at least a Generation 5 S/390 CMOS server to run on. Another alternative would be to install z/VM V4R4 on the z/890, and have the old VM operating system running as a guest. Performance may become an issue, but you can migrate smoothly to z/VM V4R4 this way.

It is important to note that z/VM V5R1 will run only on z/Series architecture, that is, there will be no support for 31 bits. You can, however, have 31-bit operating systems run as guests under z/VM V5R1 without problem. Figure 4-14 on page 121 summarizes z/VM support on z890 and other servers.

4.9.3 VSE/ESA migration

In Figure 4-14 on page 121 is a summary of VSE operating system support on z890 server. VSE/ESA V2R6 and later releases support two Logical Channel Subsystems (LCSSs), up to 30 logical partitions, and many other features available on z890.

If you are running earlier releases of VSE/ESA you should migrate at least to VSE/ESA V2R6 prior to moving to z890.

The planned future release of VSE, z/VSE V3R1, will execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. z/VSE is designed to exploit select features of IBM zSeries hardware only. Figure 4-14 on page 121 summarizes VSE/ESA and z/VSE support on z890 and other servers.

		G3-G4	G5/G6 MP3000	z800	z890	z900	z990	End of Mkt	End of Service	Ship Date
VSE/ESA	2.5	x	x	x		x	x_c	12/01	12/03	9/00
	2.6	x	x	x	x_c	x	x_c	3/03		12/01
	2.7*		x	x	x	x	x			3/03
z/VSE	3.1*		x	x	x	x	x			1Q05**
z/VM	3.1*	x	x	x	x_c	x	x_c	8/04	tbd	02/01
	4.1		x	x	x_c	x	x_c	10/01	6/03	7/01
	4.2		x	x	x_c	x	x_c	5/02	12/03	10/01
	4.3		x	x	x_c	x	x_c	8/03	5/05	5/02
	4.4*		x	x	x	x	x	tbd	9/06	8/03
	5.1*			x	x	x	x	tbd	xx/07**	3Q04**

x_c - Compatibility support
*Releases currently orderable
**Planned

Figure 4-14 z/VM, z/VSE and VSE/ESA support summary

4.9.4 Linux migration

Linux on z/Series provides both compatibility and exploitation mode support. The RedHat AS 3.0, SUSE Linux Enterprise Server 8.0, Turbolinux Enterprise Server 8.0 and Conectiva Linux Enterprise Edition (powered by UnitedLinux for z/Series) distributions provide the necessary support for z890 servers.

If you run on earlier versions of Linux for 390 or Linux on z/Series, be sure to apply the latest patches available on the DeveloperWorks Web site or, at least, have the June 2003 stream for kernel 2.4 installed.

4.10 Migration to 64-bit hardware

The most significant architectural change on zSeries hardware relative to S/390 hardware is the ability to run in 64-bit addressing mode.

When moving to a z890 server, you should have no major concerns about 64-bit feature once:

- There are no incompatibility API changes.
- The differences in the architectural changes is absorbed by Operating System.
- Even low level authorized services remain compatible.
- Three different addressing modes can coexist on a single image: 24-, 31- and 64-bit applications.

The z/OS Bimodal Accommodation is an offering, available until September 2004, where you can have z/OS, up to V1R4, run in 31-bit mode until you feel comfortable to go to 64-bit mode. Once installed, this offering option has a validity period of six months.

There are minor changes to the operating system required to enable 64-bit mode:

- All processor memory for the logical partition image must be configured as Central Storage. Expanded Storage is not supported in 64-bit mode.

- The LOADxx parameters must be reviewed.
- An image must be re-IPL'ed.

As a general rule there is no need for application changes. However, you should pay attention to ISV software and check if there are no special requirements related to real storage addressing when going to 64-bit.

4.11 Cryptographic migration considerations

The following considerations should be noted by users of cryptographic functions to recognize the difference between their current cryptographic environment and that on the z890.

4.11.1 Unsupported cryptographic functions

The following functions are no longer be supported by ICSF when the PCIX Cryptographic Coprocessor is installed on the z890:

- Digital Signature Algorithm (DSA) signature and key generation.
- American National Standard Institute (ANSI) x9.17 services (offset and notarization), and associated key types.
- Ciphertext_translate (CSNBCTT).
- German bank Pool-Pin offset.
- CSFUDK: This support is replaced by CSNBDKG.
- Commercial Data Masking Facility algorithm (CDMF), commonly known as 40-bit encryption.

These functions are thought to be 'not used' and should not impact the use of cryptography in most installations.

4.11.2 Functions changed and coexistence considerations

The support added to ICSF for double-length Message Authentication Code (MAC) keys has the following implications:

- Current DATAM and DATAMV keys have identical Crypto-Values (CVs) on left and right keys for internal token and CCA-compliant CVs (different) for external tokens. Such internal tokens will continue to work on PCIXCC features so that existing CKDS entries can be shared with z890 processors.
- Generation, import, or export of DATAM and DATAMV on a z890 creates a key with true CCA CVs; that is, different CVs for left and right keys in internal and external tokens. These tokens will work, with restrictions, on PCICC features with OS/390 V2R10 if APAR OW46382 is applied.
- z890 handling of keys with the Prohibit Export CV bit set *off* has the following migration implications: An external token having a CV with the *prohibit export* bit set *off* will be imported to an internal token with identical CV; the prohibit export CV bit will not be set.
- If the CKDS is to be shared, then all CKDS management should be performed from a non-z890 processor to avoid possible sharing problems.

Almost all ICSF application programming interface (API) calls that would execute on CCF, PCICC, or PCICA will be moved over, without change, to PCIXCC, PCICA, or the Cryptographic assist instructions available on the z890. There will normally be no need for application reworking, including recompile or re-linked, to move it to a z890.

4.12 Migration from Multiprise 3000 to z890

It is expected that many Multiprise 3000 users will migrate to z890 servers. The wide range of capacity settings for the z890 fits very well with the three Multiprise 3000 models. There are a number of considerations for this migration:

- The z890 does not use emulated I/O. This has several side effects:
  - There is no DEVMAP function on the z890. Definitions that may have involved a DEVMAP (such as LAN interfaces) are completely within the IOCDS on the z890.
  - There is no equivalent for emulated DASD and no way to load z/OS software from CD-ROM to disks on the z890. However, Linux systems can be installed from the DVD drive in the z890 HMC. Also z/VM V5R1 can be installed from the DVD drive.
  - The Support Element, 3270 console sessions that are possible on the Multiprise 3000 have no direct equivalent on the z890. The 3270 sessions (one per LPAR) available on the z890 HMC are not equivalent to the Multiprise 3000 emulated 3270 sessions on the Support Element.
  - There is no 4mm tape drive on the z890 and no way to use 4-mm tapes. Likewise there is no way to move external SCSI-attached tape drives to a z890 for z/OS, z/VM, or VSE/ESA use.
- The emulated I/O LAN interfaces on the Multiprise 3000 are replaced with OSA-Express features on the z890. The OSA-Express features have a *much* greater effective bandwidth than the emulated I/O LAN interfaces of the Multiprise 3000. The differences are so great that you might want to rethink your network implementation. Some Multiprise 3000 users have ESCON-connected routers to overcome the limited throughput of the emulated I/O LAN features. This approach is no longer needed with OSA-Express implementations.

  The Multiprise 3000 is limited to four LAN feature ports and these could not be shared among LPARs. The z890 can have many more LAN feature ports and each can be shared among multiple LPARs for typical TCP/IP usage.
- The z890 does not have internal disks. This can affect your physical installation plans once there will be more floor space needed to accommodate a DASD and/or Tape control unit.
- The I/O configuration of a practical z890 system is likely sufficiently different from that of a typical Multiprise 3000 system. This implies that the HCD (or IOCP) definitions will need manual reconfiguration. Also, z890 systems can have multiple LCSSs and this requires changes to HCD and IOCP definitions.
- The Multiprise 3000 used “normal wall outlets” for power. This is not the case for a z890.
- A raised floor environment was not needed for most Multiprise 3000 installations. For practical reasons, it is required for a z890.
- While an Multiprise 3000 could use an HMC, most systems did not have one. An HMC is required for a z890.
- The z890 is a z/Architecture machine, that is, 64-bits capable, while the Multiprise 3000 was not. This has implications for older versions of operating systems and software subsystems. Existing applications should run without changes provided they work with current releases of the operating systems and subsystems. Refer to “Operating Systems migration” on page 118 for more details.
- The cryptographic coprocessor of the Multiprise 3000 is not available on the z890. The z890 has much more extensive cryptographic capabilities, but the implementation is different.

4.13 Migration from S/390 CMOS servers to z890

There are many considerations regarding IBM S/390 servers migrations. To mention some of them:

- On average, power and cooling requirements are bigger on z890, unless you have a two frames S/390 CMOS server installed.
- There is no parallel channel support. Use ESCON converters if parallel channels are still required.
- The current ESCON fiber cable infrastructure will require connector conversion kits to be reused.
- Evaluate the use of FICON to “consolidate” ESCON channels.
- There is no OSA-2 ATM or FDDI on z890. Migrate to OSA-Express Ethernet and multi-purpose switches or routers if needed.
- There is no ICB-2 connection on z890. For Generation 5 and 6 servers only the use of ISC-3 links in compatibility mode (Feature Code 0008) is available to connect to z890. Earlier than Generation 5 CMOS servers cannot be part of parallel sysplex with z890.
- Check if your current HMC can be used; otherwise order a new one.
- Check for unsupported operating systems. Consider migration to a supported one before going to z890.
- Check for coexistence of current and z890 servers when running with different operating system releases. Use compatibility support if necessary.
- The impact of fewer, faster CPs versus more slower CPs must be verified to see whether this phenomenon will have an impact on the behavior of the workload.

5

Frequently asked questions

This chapter lists some of the more commonly asked questions concerning the new functions, features, and technology improvements offered by the IBM eServer zSeries 890. Questions regarding migrating from previous smaller servers or the more recent zSeries 800 server are also covered.

The questions and answers touch on topics, such as:

- General hardware features
- Connectivity enhancements
- zSeries Application Assist Processor
- Cryptographic support
- Parallel Sysplex support
- On/Off Capacity on Demand (CoD)
- Miscellaneous

5.1 z890 general hardware feature FAQs

This section lists questions that are more general in nature concerning the announcement of the new IBM eServer z890. Questions such as what was announced, general hardware functions and some comparisons to the previous severs are addressed here. More detailed FAQs regarding specific areas like networking and I/O, zSeries Application Assist Processor (zAAP) and cryptographic functions are found in separate sections.

Question
What enhancements does the IBM eServer zSeries 890 offer over the previous IBM eSeries zSeries 800?
Answer
Persistent innovation and improved value have been important elements of IBM's zSeries mainframes and both of these elements are a major part of IBM's eServer zSeries 890 server announcement. These include the following as compared to the z800:

- zSeries Application Assist Processor (zAAP)
- OSA - Integrated Console Controller (OSA-ICC)
- Parallel Sysplex Coupling Facility Control Code (CFCC) level 13 and enhanced Patch Apply
- Cryptographic enhancements
- Two Logical Channel SubSystems (LCSSs) supporting 256 CHPIDs each
- Internal and external channel spanning
- Parallel Sysplex enhancements
- Four times the number of HiperSockets available compared to the z800
- On/Off Capacity on Demand (On/Off CoD) enhancements
- Twenty eight capacity settings, offering a lower entry point than the smallest z800 with scalability of up to 123 percent of the largest z800

Question
What are the 28 capacity settings for the z890?
Answer
With the z890, you can start at the capacity setting you need to meet your IT infrastructure requirements, then easily add additional capacity to meet your changing needs as you grow. Each of the four processors on the z890 can be divided into seven sub-units. This creates a 7 by 4 matrix of settings. As long as upgrades are positive MIPS growth, you can move around anywhere within the matrix when adding capacity.

Question
How can I determine the z890 capacity setting from the feature code?
Answer
There is only one model of the z890, the IBM 2086 model A04. Each of the 28 capacity settings will have a 4 digit associated feature code. The first digit is always a '6' and can be ignored. The second digit relates to the number of active Central Processor (CP) engines defined to your machine. The third number describes the number of the capacity setting for the engines. The last digit is an engine identifier called 'use day', used only when On/Off Capacity on Demand is executed. Using the above information, the following table shows the capacity settings for the z890 Model A04.

Table 5-1 Capacity settings for z890 Model A04

1-way	2-way	3-way	4-way
110	210	310	410
120	220	320	420

1-way	2-way	3-way	4-way
130	230	330	430
140	240	340	440
150	250	350	450
160	260	360	460
170 Full 1-way	270 Full 2-way	370 Full 3-way	470 Full 4-way

Question
What is different about Capacity Setting 110?
Answer
The entry model z890 is Capacity Setting 110. The 110 has most of the same features as the other 27 capacity settings; however, it has smaller I/O configuration/expansion capabilities. The 110 will only be allowed to run 15 LPs, it can have a maximum of 240 ESCON channels, 32 FICON Express, and 24 OSA-Express. The 110 is also different from the other models because it offers zELC pricing.

Question
How is the z890 different from previous generations of mainframes?
Answer
For starters, the z890 is very different from past generations of mainframes because of this radical new look at granularity and the 28 capacity setting matrix. You will be able to install the size machine you want, and as long as they have positive capacity upgrades, they can move anywhere on the matrix when they upgrade. This will offer new options for running traditional work, if desired, on more, smaller engines, and additional engines can be acquired for Linux (IFL) or the new zSeries Application Assist Processor.

Question
How does the z890 report capacity setting, CPU ID information?
Answer
Capacity setting information is returned via the STSI instruction.

Question
Are there upgrade paths from z800 to z890?
Answer
There are upgrade paths from z800 models 002 and 004 to any z890 Capacity Settings.

Question
Are there upgrade paths from z890 to z990?
Answer
There are upgrade paths from z890 to any z990 Model A08 (0-8 Central Processors) from the following z890 Capacity Settings: 170, 250 - 270, 330 - 370, and 430 - 470.

Question
How many spare processing units are on the z890?
Answer
There are four Processing Units on the z890. They can be assigned as central processor (CP), Integrated Facility for Linux (IFL), zSeries Application Assist Processor (zAAP) Internal Coupling Facility (ICF) engines. If all four characterizable Processor Units are utilized there are no spares.

Question
How many System Assist Processors (SAPs) are on the z890?
Answer
There is one standard SAP on the z890. No optional SAPs are offered. A spare engine may be used as a SAP if there were to be a SAP failure.

Question
On the z890, if I add an IFL to a sub-uniprocessor z890, will I receive a sub-uniprocessor IFL?
Answer
No, when adding an IFL, zAAP, or ICF to a subcapacity machine, the machine is able to run with mixed capacities. The IFL, zAAP or ICF are full capacity engines. Only Central Processors are tuned down (for performance and software granularity).

5.2 z890 connectivity enhancements FAQs

This section lists FAQs directed toward networking and I/O enhancements, providing more detailed responses. Questions addressed here deal with channel spanning, OSA Express, and the new Integrated Console Controller.

Question
What does channel spanning mean to me?
Answer
Channel spanning is when channels can be shared between LPs in different LCSSs. Prior to this they could only be shared within a single LCSS.

Question
Can channel spanning help me to simplify and reduce my zSeries infrastructure?
Answer
It can, if the load placed on the shared channel can meet your performance requirements. Channel spanning brings a higher level sharing to z890, allowing multiple LPs associated with different LCSS to attach to and share a common port. This can reduce the number of additional channels that might be required to attach multiple LPs without the spanning capability.

Question
Can you give me an example of reducing my zSeries infrastructure with channel spanning?
Answer
If you configure your z890/z990 for multiple LCSSs, and you have a requirement to share resources (I/O devices, networks, CFs) between LPs in those LCSSs, you can now share the channels between those LCSSs. This can reduce the total number of needed physical channels, switch ports, and so on.

Question
What channels can not be spanned?
Answer
ESCON, FICON Bridge (CHPID type FCV), and ISC-3 receiver (CHPID type CFR).

Question
Can you give me an example of how spanning OSA-Express adapters can help my environment?
Answer
Channel spanning brings a higher level of OSA-Express MIF sharing to z890, allowing multiple LPs associated with different LCSS to attach to and share a common OSA-Express port. Each port can be separately configured to provide either QDIO, non-QDIO or new OSC (OSA-ICC) support. This can reduce the number of separate or additional OSA-Express channels that

might be required for connecting to multiple LPs without the spanning capability.

Question
Can you give me an example of how OSA-Express Integrated Console Controller (OSA-ICC) can help my current environment?
Answer
Each port of the OSA-Express 1000BASE-T Ethernet feature configured for OSA-ICC is designed to support up to one hundred and twenty (120) IPL and operating systems console sessions in z/OS, z/OS.e, z/VM, VSE/ESA and TPF environments. This console session attachment capacity is a significant increase over prior IBM 2074 (model 3 attachment maximum of 96) and 3174 (one console per LP per controller) capabilities and is designed to provide the capacity spanning both medium and large enterprise requirements.

Question
Can the OSA-ICC connect to multiple separate z890 servers?
Answer
No. Since the OSA-ICC is a function operating on a port of the OSA-Express 1000BASE-T Ethernet feature which is installed into the frame of a single z890, a single OSA-ICC port can only support the mainframe to which it is attached.

Question
What desktop or laptop 3270 emulator does the OSA-ICC support?
Answer
IBM has tested the eNetwork Personal Communications V5.6 emulator running under Microsoft® Windows XP + for compatibility, however there is no reason to believe that other compatible emulators will not work. Please contact the emulator product vendor directly for terms, conditions, prices and other product details related to such vendor's emulator.

Question
What level of security does the OSA-ICC provide for session connections?
Answer
There are different types of remote connections supported on the OSA-ICC. The configuration and diagnostic support for OSA-ICC is provided via the traditional zSeries HMC or SE which do provide the ability for a secure remote connection capability. The user client console connections are LAN attached using workstations or laptops supporting 3270 emulation. Given the importance of operational consoles, clients typically attach workstations and laptops to dedicated, local LAN(s) to minimize exposure to unsecure access. Userid and IP address access checking is provided in this case.

5.3 z890 zSeries Application Assist Processor (zAAP) FAQs

The zSeries Application Assist Processor (zAAP) is a completely new function that can only be enabled on the z890 and z990 servers. zAAPs offer improved efficiencies for certain software applications when exploited properly on the zSeries platform.

This section lists FAQs directed toward the new zSeries Application Assist Processor.

Question
What is a zSeries Application Assist Processor (zAAP)?
Answer
zSeries Application Assist Processors (zAAPs) are specialized assist processors that are designed to execute Java programming when jointly configured with central processors within logical partitions (LPs). zAAPs operate asynchronously from the central processors and reduce processor consumption for Java-based applications. The central processors use the SDK 1.4.1 for zSeries to direct the Java Virtual Machine processing cycles to the configured

zAAPs.

Question
What is the value of zAAPs?
Answer
zAAPs enable you to run e-business Java Web applications next to mission critical data for high performance, reliability, availability and security at a competitive cost.
zAAPs carry no traditional zSeries software charges. In addition, zAAPs reduce central processor consumption for Java-based applications, thereby increasing the capacity of the central processors.

Question
How do I order the zSeries Application Assist Processors (zAAPs)?
Answer
The zSeries Application Assist Processor (zAAP) (FC 0520) can be ordered with a new-build IBM eServer zSeries 890 (z890) or 990 (z990) machine, or with any MES, either directly from IBM or through an authorized IBM Business Partner.

Question
How many zSeries Application Assist Processors (zAAPs) can a machine have?
Answer
A machine may have zAAPs up to the number of active CPs on a given machine. For z890 the maximum number of zAAPs would be two with two CPs.

Question
What are the requirements for ordering and running the zSeries Application Assist Processors (zAAPs) on z890?
Answer
To order zAAPs, you must be WLC subcapacity qualified and utilize the SCRT tool and provide Call Home data. In order to exploit zAAPs on z990 or z890, the z/OS operating system must be migrated to z/OS 1.6 and the WebSphere must be Version 5.1 or later.

5.4 z890 cryptographic FAQs

There have been many improvements in the hardware cryptographic capabilities with the first zSeries 900 and 800 as compared to the previous 9672 servers. Now with new functions and capabilities introduced first with the z990 and continuing with the z890 even more functions and capabilities are available.

This section lists FAQs directed toward the cryptographic functions provided by the z890 server.

Question
What is the new CP Assist for Cryptographic Function (CPACF) and is it standard on every z890?
Answer
Like the z990, the z890 comes with the new Message Security Assist Architecture along with the new CPACF, delivering cryptographic support on every Processor Unit (PU), along with DES and TDES data encryption/decryption and SHA-1 hashing. By having the capability already on the PU, this means that association of cryptographic functions to a specific PU (such as the CCF) is eliminated. The SHA-1 function is shipped enabled, but the DES and TDES functions require enablement of the CPACF function (feature code 3863) for export control.

Question
Will the PCI Cryptographic Coprocessor (PCICC), available on the z800, be available on the z890?
Answer
No. The new PCIX Cryptographic Coprocessor (PCIXCC) feature code 0868 is a replacement for the PCICC and the CMOS Cryptographic Coprocessor Facility (CCF) that were offered on the z800. The PCIXCC supports all of the equivalent PCICC functions. In addition, the functions of the CMOS CCF used by known applications have also been implemented by the PCIXCC feature. It supports secure cryptographic functions, use of secure encrypted key values, and user-defined extensions.

Question
Will the PCI Cryptographic Accelerator (PCICA) feature, supported on the z800, be available to the z890?
Answer
Yes. The PCICA can be carried over from a z800. This hardware-based cryptographic solution continues to address the high Secure Sockets Layer (SSL) performance needs of an on demand e-business.

Question
Can you describe the PKE Service Enhancements for z890 servers?
Answer
The Public Key Encrypt (PKE) service has been enhanced to support the Mod_Raised_to_Power (MRP) function. MRP has the same functionality as Zero-Pad with the addition of supporting cryptographic keys with both even and odd exponents. The MRP function can be used to off load the compute-intensive portion of the Diffie-Hellman protocol to the PCIXCC and PCICA features for improved performance and more efficient use of CP resources.

Question
What are the benefits of Double Length Derived Unique Key Per Transaction (DUKPT)?
Answer
The PCIXCC feature supports Derived Unique Key Per Transaction (DUKPT) for double length keys. The previous ANSI standard supported single length keys for transactions. DUKPT allows you to write applications that implement the Double Length Derived Unique Key per Transaction (DUKPT) algorithm as defined by the 2002 version of ANSI X9.24 standard and implemented by many financial institutions.

Question
Will zServers support the Europay Mastercard and Visa (EMV) 2000 standard?
Answer
Yes, the PCIXCC feature supports Europay Mastercard and Visa (EMV) 2000 standard. EMV, Europay Mastercard and Visa, is a standard that is developed and implemented for financial transactions between heterogeneous hardware and software. This support applies only to the PCIXCC feature on z890 and z990 servers.

Question
What services will be available in support of the Europay Mastercard and Visa (EMV) 2000 standard?
Answer
The Pin Change and Unblock service are included in EMV 2000 support.
The service allows you to write applications to change or unblock personal identification numbers (PIN) on smart cards, as required by the Visa specifications. This support applies only to the PCIXCC feature on z890 and z990 servers.

Question
Are there enhancements to the Public Key Decrypt (PKD) service in May 2004?
Answer
Yes. The PCIXCC and PCICA features support Public Key Decrypt (PKD) service enhancement. The Public Key Decrypt (PKD service has been upgraded to support a Zero-Pad option for Clear RSA private keys on both the PCIXCC feature and the PCICA feature.

Question
Can you describe the benefits of TKE 4.1 workstation operational key entry?
Answer
Trusted Key Entry (TKE) 4.1 operational key entry allows you to securely and remotely load operational keys into the Integrated Cryptographic Service Facility (ICSF) Cryptographic Key Data Set (CKDS) via Cryptographic Coprocessor hardware features on the host.

Question
What are the recent changes incorporated in TKE code 4.1 level on the z890 server in May 2004?
Answer
This level of code includes support for Operational Key Support. Functionally TKE with 4.1 code is capable of controlling 9672/G5 and G6, 7060, z800, z900, z890 and z990 processors.

Question
Are new TKE hardware features available to support the z890 server in May 2004?
Answer
Yes. There are two new TKE workstations: Feature code 0896 is available for token ring support and Feature code 0899 for ethernet support.

Question
Why is z990 and z890 enhancements to Cryptographic support, an ICSF Web deliverable, required?
Answer
Cryptographic support for z890 and z990 for OS/390 V2.10, z/OS V1.2 and later, z/OS.eV1.3 and later, an ICSF Web deliverable, FMID: HCR 770B, is needed to enable and manage access to new and/or enhanced cryptographic functions on the PCIXCC feature. Specifically, this Web deliverable supports the following new Cryptographic functions:

- PKE/PKD Service Enhancements
- TKE V4.1 Operational Key Entry
- Double Length Unique Key Per Transaction (DUKPT)
- EMV 2000 Standard

5.5 z890 Parallel Sysplex FAQs

The IBM zSeries 890 server can participate in a multisystem Sysplex with other zSeries servers or even with previous servers. The z890 offers the latest link technology with speeds at the same level as the z990 servers. New Coupling Facility Control Code (CFCC) available with the z890 and also the z990 offers additional benefits.

This section lists FAQs directed toward the z890 and Parallel Sysplex operation.

Question
Is there a new CFCC level associated with the z890?
Answer
Yes, CFCC level 13 is part of the z890 Licensed Internal Code (LIC).

Question
Is there software support associated with the new CFCC level?
Answer
Yes, z/OS APAR OA01517 enables CF exploiters, such as DB2, to request placement of their structures in a CFCC level 13 CF, and provides the required z/OS software support for the use of CFCC level 13.

Question
Are there structure sizing changes associated with the new CFCC level?
Answer
Yes, there are. Additional “facility” storage is consumed by CFCC level 13, leaving less storage available for use by structures. Additional “structure” storage is also consumed in individual structures.

Question
Are there any implications for System-Managed CF Structure Duplexing for the new CFCC level?
Answer
Yes, there are. Several of the CFCC enhancements are expected to benefit duplexed structures particularly, but their benefit will be obscured unless BOTH of the duplexed structure instances are located in CFCC level 13 CF images.

Question
Now I can install up to 48 ISC-3 links. Am I still free to define them all in compatibility mode?
Answer
No. You are still restricted to a maximum of 32 ISC-3 links in compatibility mode (CHPID types CFS/CFR). If you install more than 32 ISC-3 links, those greater than 32 must be defined for peer mode (CHPID type CFP).

Question
Can you provide an example of how spanning links can reduce the infrastructure in my Parallel Sysplex environment?
Answer
If you have z/OS images in one LCSS and a CF image in another LCSS, you can span a single peer or sender coupling link to an external CF. This will allow that link to be used for normal message traffic and for CF-to-CF duplexing. Likewise, if you have z/OS images in different LCSS's, a single spanned coupling link can be shared by them all.

Question
Can I attach the z890 into a 9672 Parallel Sysplex®?
Answer
Yes, using the InterSystem Coupling-3 (ISC-3) link operating in compatibility mode (100 MB/sec) you can attach to a 9672, however it is recommended to use all zSeries processors in a Parallel Sysplex for best performance results.

5.6 z890 On/Off Capacity on Demand (CoD) FAQs

On/Off Capacity on Demand was first introduced with the z990 server. It has been extended to the z890 server and new functionality has been since its original introduction. The z890 offers many capacity settings, several options exist to permanently or temporally add capacity, with these options capacity can be added dynamically and concurrently.

This section lists FAQs directed toward the cryptographic functions provided by the z890 server.

Question
What is On/Off Capacity on Demand (On/Off CoD)?
Answer
On/Off CoD is the latest addition to the IBM zSeries existing Capacity on Demand offerings such as Customer Initiated Upgrade (CIU) and Capacity Backup Upgrade (CBU). On/Off CoD gives z990 and z890 customers the ability to rent hardware capacity by the day. With On/Off CoD you can turn on processors from your original configuration when your business needs them and then turn them off when the need subsides, and only pay for the days the processors are turned on.

Question
How do I order the zSeries On/Off Capacity on Demand Offering?
Answer
On/Off CoD Enablement (FC 9896) can be ordered with a new-built IBM eServer zSeries 990 (z990) machine, with a new-built IBM eServer zSeries 890 (z890) machine or with any MES, either directly from IBM or through an authorized IBM Business Partner. This first step includes the signing of contracts, establishment of upgrade pricing and creation of Resource Link™ profiles. Once this set-up is completed, you may activate and deactivate temporary CP, IFL, ICF, and beginning 4Q04, zAAP capacity via the CIU application on ResourceLink at ibm.com/servers/resourcelink.

Question
What are the requirements for ordering this On/Off CoD Offering?
Answer
On/Off CoD requires a) Customer Initiated Upgrade (CIU) Enablement on the same machine (FC 9898), b) executed On/Off CoD / CIU end user agreements, c) an active RSF connection between the z990 machine and IBM.
For customers without an RSF connection between the z990 or z890 machine and IBM, the On/Off CoD Offering may be ordered with an RPQ request.

Question
What kinds of workloads is On/Off CoD intended for?
Answer
On/Off CoD is intended for any workloads that run on the permanent z990 or z890 CPs, IFLs, or zAAPs.

Question
How much temporary capacity may be ordered on a z890 with On/Off CoD?
Answer
You may order CP capacity up to the amount of permanently purchased CP capacity on a given machine. Temporary capacity on a z890 can not exceed permanently purchased capacity. The CP capacity ordered with On/Off CoD on the z890 may not exceed the limit of available engines in the machine model.

Question
How many temporary IFLs or ICFs can you have with On/Off CoD?
Answer
You can have IFLs up to the number of permanently purchased IFLs on a given machine or ICFs up to the number of permanently purchased ICFs on a given machine. The number of IFLs or ICFs with On/Off CoD may not exceed the limit of available engines in the machine model.

Question
How many temporary zAAPs can you have with On/Off CoD?
Answer
Beginning October 1, 2004, when On/Off CoD begins support of zAAPs, you may have

zAAP(s) up to the number of permanently purchased zAAPs on a given machine. The number of zAAPs ordered with On/Off CoD may not exceed the limit of available engines in the machine model.

Question
What is the relationship between CBU and On/Off CoD?
Answer
A machine can have CBU (FC6800) and On/Off CoD Enablement (FC 9896) at the same time. You can now install/activate a temporary capacity upgrade without physically removing the CBU record from the machine. Note that only one type of temporary capacity (CBU or On/Off CoD) can be activated at a time.

5.7 z890 miscellaneous FAQs

The questions listed in this section are wide ranging, addressing various areas such as migrating from older technology to architectural related items. They did not quite fit in the other sections, therefore they listed here.

Question
Do I need to make program changes for superscalar operation on the z890?
Answer
No, but…. existing programs work correctly (with the usual disclaimers). However, some programs may run *faster* when written for optimum superscalar usage. This is normally a function of a compiler, which should produce instructions in an optimum order. Programs written in assembly language often do not use optimum instruction sequences or memory reference patterns. There is typically no compelling need to rework such programs, but assembly programmers might want to be aware of these issues when writing new programs.

Question
Does the z890 use the normal OS/2-based HMCs?
Answer
Yes, the z890 uses the normal (current) OS/2-based HMCs.

Question
Can I order a Linux-only model?
Answer
There is no special model for this configuration, but you can order a z890 with only IFL PUs.

Question
Is ICMF operation supported?
Answer
No, this is not supported in the z800, z890, or the z990.

Question
Can I order cryptographic coprocessors compatible with older S/390 machines?
Answer
No, these are not available on the z890 or the z990.

Question A minimal z/VM system can be run in an IFL. Can a minimal z/OS or z/OS.e be run in an zAAP?
Answer
No.

Question
Can I connect my internal Multiprise 3000 disks to the z890?
Answer
No.

Question
Can I move any of the adapter cards from my MP 3000 to the z890?
Answer
No.

Question
Can I move ESCON-attached devices from ESCON channels on my MP 3000 to the z890? Can I connect both my MP 3000 and a z890 to the same ESCON device?
Answer
Yes, you can move the ESCON-attached devices to the z890. To connect both systems you need an ESCON Director.

Question
My MP 3000 is a 31-bit system, but the z890 is a 64-bit system. Do I need to recompile all my programs? Do I need new compilers?
Answer
There is no need to recompile anything. All your 31-bit programs continue to run correctly (with the usual disclaimers for really unusual programs).

Question
Can I reuse the internal MP 3000 disk drives in any way?
Answer
No, not with the z890.

Question
Do 64-bit programs run approximately as fast as 31-bit programs? I understand this is not the case for emulated S/390 solutions.
Answer
To a first approximation, yes. The z890 (like all zSeries machines) has native 64-bit processors. The emulated solutions you may have looked at probably emulate 64-bit zSeries operation on a 32-bit processor. Such emulation works correctly, but will obviously be slower than emulation of 31/32-bit operations.

A

Instruction set

The z890 and z990 includes many new and changed instructions. The details are the same for the z890 and the z990. The changes can be placed into three categories:

- Long-displacement instructions
- Cryptographic instructions
- Other instructions

Long-displacement instructions

A long displacement is a signed 20-bit displacement value, instead of the unsigned 12-bit displacement used with most existing instructions. There are two groups of long-displacement instructions:

- New instructions. There are 44 of these.
- Existing instructions that are changed to include long displacements. There are 69 of these.

The new instructions are available only when operating in zArchitecture mode (and when the STFL instruction results indicate that the long displacement facility is installed). The modified instructions operate with a long displacement only in zArchitecture mode. In ESA/390 mode, they operate as they did previously.

All of the new (and modified) instructions are 6 bytes long. The general formats are shown in Figure A-1. Existing instructions in RXE and RSE format have an 8-bit field that is unused. These instructions are now considered to be in RXY and RSY formats, and the unused field (now the DH field) is concatenated to the left of the 12-bit D field (now the DL field), producing a 20-bit displacement.[1] This is considered to be a signed value and will be added or subtracted from the value in the base register (plus index register, if used). A number of new instructions in the RXY and RSY format have also been added.

While these instructions are available to any assembly language programmer, their design was based primarily on the needs of the C and PL/X compilers, DB2, and Java functions.

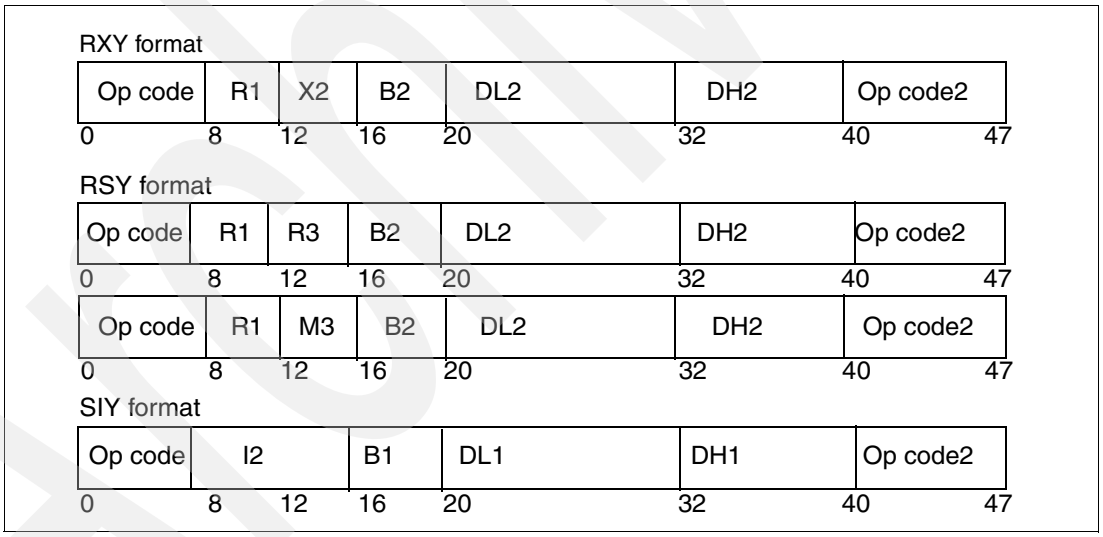

Figure A-1 Long displacement formats

An attempt to execute any of these instructions on a system that does not have the long-displacement facility installed will produce an operation exception (interruption).

Do not confuse long displacement instructions with relative addressing, as used in the *relative and immediate* instructions added to S/390 architecture a few years ago. An instruction using relative addressing for an operand requires no base register for that operand, while a long displacement instruction needs a base register. However, programs using long displacement instructions may need fewer base registers. Compilers can use these instructions to better optimize the use of registers (and assembly language programmers can do the same thing, of course).

[1] The naming of these fields is generally DH (for Displacement High order) and DL (for Displacement Low order).

The following list contains the new and changed instructions. An asterisk (*) in the first column denotes a new instruction. The letters PRIV indicates a privileged instruction. These are all six-byte instructions with part of the operation code in the first byte and part in the last byte.

	Format	Op codes	Mnemonic	Name....	Notes.....
*	RXY	E3 5A	AY	Add	32 bits
	RXY	E3 08	AG	Add	64 bits
	RXY	E3 18	AGF	Add	64<--32
*	RXY	E3 7A	AHY	Add halfword	16 bits
	RXY	E3 0A	ALG	Add logical	64 bits
	RXY	E3 1A	ALGF	Add logical	64<--32
*	RXY	E3 5E	ALY	Add logical	32 bits
	RXY	E3 98	ALC	Add logical with carry	32 bits
	RXY	E3 88	ALCG	Add logical with carry	64 bits
*	SIY	EB 54	NIY	AND	8 bits
	RXY	E3 80	NG	AND	64 bits
	RXY	E3 46	BCTG	Branch on count	64 bits
	RSY	EB 45	BXLEG	Branch on index low or equal	64 bits
	RSY	EB 44	BXHG	Branch on index high	64 bits
*	RRE	B9 2E	KM	Cipher message	
*	RRE	B9 2F	KMC	Cipher message with chaining	
*	RXY	E3 59	CY	Compare	32 bits
	RXY	E3 20	CG	Compare	64 bits
	RXY	E3 30	CGF	Compare	64:32
*	RSY	EB 14	CSY	Compare and swap	32 bits
	RSY	EB 30	CSG	Compare and swap	64 bits
*	RSY	EB 31	CDSY	Compare double and swap	32 bits
	RSY	EB 3E	CDSG	Compare double and swap	64 bits
*	RXY	E3 79	CHY	Compare halfword	16 bits
*	RXY	E3 55	CLY	Compare logical	32 bits
	RXY	E3 31	CLGF	Compare logical	64:32
	RXY	E3 21	CLG	Compare logical	64 bits
*	SIY	EB 55	CLIY	Compare Logical (immediate)	8 bits
*	RSY	EB 21	CLMY	Compare logical characters under mask (low)	
	RSY	EB 20	CLMH	Compare logical characters under mask (high)	
	RSY	EB 8F	CLCLU	Compare logical long unicode	
*	RRE	B9 3E	KIMD	Compute intermediate message digest	
*	RRE	B9 3F	KLMD	Compute last message digest	
*	RRE	B9 1E	KMAC	Compute message authentication code	
*	RXY	E3 06	CVBY	Convert to binary	32 bits
	RXY	E3 0E	CVBG	Convert to binary	64 bits
*	RXY	E3 26	CVDY	Convert to decimal	32 bits
	RXY	E3 2E	CVDG	Convert to decimal	64 bits
	RXY	E3 0D	DSG	Divide single	64 bits
	RXY	E3 1D	DSGF	Divide single	
	RXY	E3 97	DL	Divide logical	32<--64
	RXY	E3 87	DLG	Divide logical	64<--128
*	SIY	EB 57	XIY	Exclusive OR (immedicate)	8 bits
	RXY	E3 82	XG	Exclusive OR	64 bits
*	RXY	E3 57	XY	Exclusive OR	32 bits
*	RXY	E3 73	ICY	Insert Character	8 bits
*	RSY	EB 81	ICMY	Insert characters under mask (low)	
	RSY	EB 80	ICMH	Insert characters under mask (high)	
	RXY	E3 04	LG	Load	64 bits
	RXY	E3 14	LGF	Load	64<--32
*	RXY	E3 58	LY	Load	32 bits
*	RXY	ED 65	LDY	Load	floating long
*	RXY	ED 64	LEY	Load	floating short
*	RSY	EB 9A	LAMY	Load access multiple	
*	RXY	E3 71	LAY	Load address	32 or 64 bits
*	RXY	E3 76	LB	Load byte	signed byte to 32 bits

*	RXY	E3 77	LGB	Load byte	signed byte to 64 bits
	RSY	EB 2F	LCTLG	Load control	64 bits PRIV
	RXY	E3 15	LGH	Load halfword	64<--16
*	RXY	E3 78	LHY	Load halfword	16 to 32 bits
	RXY	E3 16	LLGF	Load logical	64<--32
	RXY	E3 90	LLGC	Load logical character	64<--8
	RXY	E3 91	LLGH	Load logical halfword	64<--16
	RXY	E3 17	LLGT	Load logical 31 bits	
	RSY	EB 04	LMG	Load multiple	64 bits
*	RSY	EB 98	LMY	Load multiple	32 bit words
	RSY	EB 96	LMH	Load multiple high	
	RXY	E3 8F	LPQ	Load pair from quadword	
*	RXY	E3 13	LRAY	Load real address	32 bits PRIV
	RXY	E3 03	LRAG	Load real address	64 bits PRIV
	RXY	E3 1E	LRV	Load reversed	32 bits
	RXY	E3 0F	LRVG	Load reversed	64 bits
	RXY	E3 1F	LRVH	Load reversed	16 bits
*	SIY	EB 52	MVIY	Move (immediate)	8 bits
	RSY	EB 8E	MVCLU	Move long unicode	
	RXY	E3 86	MLG	Multiply logical	128<--64
	RXY	E3 96	ML	Multiply logical	64<--32
*	RXY	E3 51	MSY	Multiply single	32<--32x32
	RXY	E3 1C	MSGF	Multiply single	
	RXY	E3 0C	MSG	Multiply single	64 bits
*	SIY	EB 56	OIY	OR (immediate)	8 bits
	RXY	E3 81	OG	OR	64 bits
*	RXY	E3 56	OY	OR	32 bits
	RSY	EB 1C	RLLG	Rotate left single logical	64 bits
	RSY	EB 1D	RLL	Rotate left single logical	32 bits
	RSY	EB 0A	SRAG	Shift right single	64 bits
	RSY	EB 0B	SLAG	Shift left single	64 bits
	RSY	EB 0C	SRLG	Shift right single logical	64 bits
	RSY	EB 0D	SLLG	Shift left single logical	64 bits
*	RXY	ED 67	STDY	Store	floating long
*	RXY	ED 66	STEY	Store	floating short
*	RXY	E3 50	STY	Store	32 bits
	RXY	E3 24	STG	Store	64 bits
*	RSY	EB 9B	STAMY	Store access multiple	32 bit words
*	RXY	E3 72	STCY	Store character	
*	RSY	EB 2D	STCMY	Store characters under mask (low)	
	RSY	EB 2C	STCMH	Store characters under mask (high)	
	RSY	EB 25	STCTG	Store control	64 bits PRIV
*	RXY	E3 70	STHY	Store halfword	16 bits
	RSY	EB 24	STMG	Store multiple	64 bits
	RSY	EB 26	STMH	Store multiple high	
*	RSY	EB 90	STMY	Store multiple	
	RXY	E3 8E	STPQ	Store pair to quadword	
	RXY	E3 2F	STRVG	Store reversed	64 bits
	RXY	E3 3E	STRV	Store reversed	32 bits
	RXY	E3 3F	STRVH	Store reversed	16 bits
*	RXY	E3 5B	SY	Subtract	32 bit words
	RXY	E3 09	SG	Subtract	64 bits
	RXY	E3 19	SGF	Subtract	64<--32
*	RXY	E3 7B	SHY	Subtract halfword	16 bit integers
	RXY	E3 0B	SLG	Subtract logical	64 bits
*	RXY	E3 5F	SLY	Subtract logical	32 bit words
	RXY	E3 1B	SLGF	Subtract logical	64<--32
	RXY	E3 89	SLBG	Subtract logical with borrow	64 bits
	RXY	E3 99	SLB	Subtract logical with borrow	32 bits
*	SIY	EB 51	TMY	Test under mask	8 bits

```
RSY     EB 0F      TRACG   Trace                          64 bits PRIV
```

Not included in this list are:

- The STFL instruction, which has its output extended to indicate the presence of new facilities and assist functions
- The STIPD, described in the text
- The STCPS instruction, which functions as a NO OP on this system

The exact format and function of the instructions are explained in the Principles of Operation manual associated with the z990.

The long displacement instructions have been made available on z800 and z900 systems.

Cryptographic assist instructions

Five new instructions are included with the cryptographic assist function that is standard on every z990 PU. As a group, these instructions are known as the *Message Security Assist* (MSA). Briefly, the instructions are:

```
Cipher message
Cipher message with chaining
Compute intermediate message digest
Compute last message digest
Compute message authentication code
```

These are all problem-mode instructions and are all in RRE format.

The Store Facility List (STFL) instruction results have been expanded to include:

```
Bit 17: The Message Security Assist is installed.
Bit 18: The long-displacement facility is installed (in zArch mode)
```

The STFL instruction is a privileged instruction, normally used only by the operating system.

STIDP instruction changes

The results of the STIDP (Store CPU ID) instruction is different on z890 and z990 systems than on earlier systems. The logical CPU address is absent from the results, and the LPAR identifier is now a full byte. This is illustrated in Figure A-2 on page 142.

In the new format, the combination of the LPAR number (8 bits) and the extracted portion of the system serial number (16 bits) provide a 24-bit value that will be unique across all z890 machines and their LPARs.

A program requiring the logical CPU address can use an STSI instruction (instead of the STIDP instruction) to obtain it, although the STAP instruction is the recommended method.

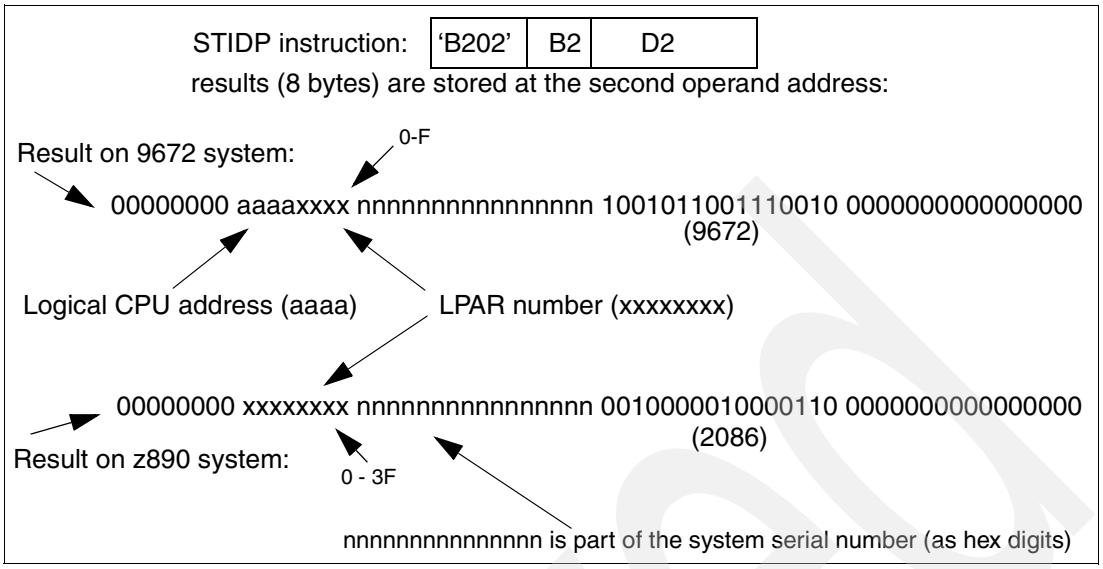

Figure A-2 Changes to STIDP instruction results

Extended Translation Facility

The Extended Translation Facility adds 10 new instructions to the zSeries instruction set. These new instructions may enhance performance for data conversion operations done supporting data encoded in Unicode, improving the ability to efficiently support applications enabled for Unicode and/or Globalization. The instructions add hardware support for conversions to and from the various Universal Character Set (UCS) Transformation Formats (UTF) encodings used to support Unicode. These formats for encoding data are used in a number of important technologies emerging in the Web Services, Grid, and on demand environments, such as XML and SOAP, as well as being supported in DB2 for data storage. The initial programming support for the new instructions are intended to be provided in High Level Assembler.

B

Geographically Dispersed Parallel Sysplex (GPDS)

IBM Installation Services for GDPS is a total end-to-end solution that manages availability within a site and across multiple sites. It provides the automation to manage not only unplanned exception conditions, but also the many planned exception conditions that are faced as a part of normal everyday processing in any IT environment. The GDPS solution can be tailored to specific Business Continuance requirements, and is based on either the synchronous Peer to Peer Remote Copy (PPRC) or the asynchronous Extended Remote Copy (XRC).

GDPS also supports the Peer-to-Peer Virtual Tape Server (PtP VTS) form of remote copying tape data. By extending GDPS support to data resident on tape, the GDPS solution is designed to provide continuous availability and near transparent business continuity benefits for both disk- and tape-resident data. Enterprises should no longer be forced to develop and utilize processes that create duplex tapes and maintain the tape copies in alternate sites.

GDPS is application independent and is enabled by means of key IBM technologies and architectures:

- Parallel Sysplex
- Tivoli® Netview for z/OS or OS/390
- System Automation for z/OS or OS/390
- Enterprise Storage Server™ (ESS)
- Peer-to-Peer Virtual Tape Server (PtP VTS)
- Optical Dense or Coarse Wavelength Division Multiplexer
- PPRC (Peer-to-Peer Remote Copy) architecture
- XRC (Extended Remote Copy) architecture
- Virtual Tape Server Remote Copy architecture

All GDPS images are running GDPS automation based upon Tivoli Netview for z/OS or OS/390 and System Automation for z/OS or OS/390. Each image will monitor the base or Parallel Sysplex cluster, Coupling Facilities, and storage subsystems; and maintain GDPS status. GDPS automation can coexist with an enterprise existing automation product.

For more detailed information on GDPS, see the white paper *GDPS: The e-business Availability Solution,* GF22-5114, at:

http://www.ibm.com/servers/eserver/zseries/library/whitepapers/gf225114.html

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 146. Note that some of the documents referenced here may be available in softcopy only.

- *IBM @server zSeries Connectivity Handbook,* SG24-5444
- *IBM @server zSeries 990 Technical Guide,* SG24-6947

Other publications

These publications are also relevant as further information sources:

- *Hardware Management Console Operations Guide*, SC28-6819
- *Support Element Operations Guide*, GC38-0608
- *PR/SM Planning Guide*, SB10-7036
- *zSeries 890 Installation Manual for Physical Planning*, GC28-6828
- *zSeries 890 and 990 CHPID Mapping Tool User’s Guide*, GC28-6825
- *Input/Output Configuration Program User’s Guide*, SB10-7037
- *zSeries 890 System Overview*, SA22-6832
- *Stand-alone IOCP User’s Guide*, SB10-7040
- *Hardware Configuration Definition: User’s Guide*, SC33-7988

Online resources

These Web sites and URLs are also relevant as further information sources:

- IBM zSeries server product line

  http://www.ibm.com/servers/eserver/zseries/
- IBM zSeries connectivity options

  http://www.ibm.com/servers/eserver/zseries/connectivity
- IBM Large Systems Performance Reference for zSeries

  http://www.ibm.com/servers/eserver/zseries/lspr
- IBM Resource Link

  http://www.ibm.com/servers/resourcelink

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

N

O

P

Q

R

S

T

U

V

X

Z

IBM eserver zSeries 890 Technical Introduction

Redbooks

Hardware description and software support

Planning and migration considerations

Frequently asked questions

This IBM Redbook introduces the IBM eServer zSeries 890, which represents the continuation of the scalable servers featured with the IBM eServer zSeries 990. The z890 is based on z/Architecture, the zSeries building blocks of the z990, and the virtualization technology of passed sever families. It is designed to be resilient in the unpredictable on demand world.

The z890 is a single model server with a wide range of capacity settings, delivering significantly improved granularity and enriched functions over its predecessor. At the same time, the z890 is also introducing the new eServer zSeries Application Assist Processor (zAAP), which provides a Java execution environment.

This publication provides information on the hardware and software features available with the z890. It also includes planning and migration considerations.

This technical introduction is intended for hardware planners, system engineers, and consultants who need to understand the capabilities of the z890.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6310-00 ISBN 0738497649