IBM System z9 Enterprise Class (EC) Reference Guide

April 2007

Table of Contents

System z9 EC Overview

Data alone is really just the bits and bytes that get loaded every day into databases. By being able to integrate and merge data from different databases, businesses are able to provide timely, current, correct and secure information to their clients. This type of delivery can allow a business to be responsive and flexible to its clients' requests. Taking this a step further, a business may gain a competitive edge by being able to analyze the information for insight, research or simulations. This may help set corporate strategy and allow a business to have a competitive edge.

But managing data can be a big project. It needs to be available globally, it must be kept up in real time, and accessible 24x7. At the same time, it needs to be protected from unauthorized access and in compliance with new regulations. The IBM System z™ mainframe family has a strong heritage in data serving, and its architecture is designed for massive data access, whether across the Internet, to storage devices or across to remote backup sites. Scalability, virtualization, availability, security – these are all core competencies of the mainframe.

The IBM System z9™, having classic strengths on which its reputation was formed, continues to be a leader in areas such as data management, availability, security and resiliency, virtualization, and integration. Its ability to support open and industry standards, with support for Service-Oriented Architecture makes it an ideal platform for deploying new workloads, or for interoperating with new workloads on alternative technologies.

The IBM System z9 Enterprise Class (z9 EC) provides a collaboration with IBM storage and software to help achieve advanced I/O function and performance. And while the z9 EC can scale to offer a large capacity mainframe in a single footprint, some customers have requested the z9 EC innovative features with more flexibility in customizing and sizing the capacity of the general purposes processors (PUs) that reside in the server. These customers want more options as they weigh business demands and costs in an effort to achieve the right balance. The new subcapacity models of the z9 EC help satisfy these requests. These enhanced capabilities illustrate that the z9 EC is a classic that doubles as an advanced solution designed to support the IT infrastructure for today's on demand business.

System z and IBM DB2 for z/OS – synergy and strength for Data Serving

The foundation for data serving on the mainframe is provided by System z and DB2® for z/OS®. DB2 for z/OS is written to exploit the System z platform and as a result can offer advanced features and function. IBM DB2 for z/OS delivers rich function for highly scalable, industry-leading high availability IT infrastructure for your enterprise data and on demand business applications. The combined power and capacity of the z9 EC with the high performance and availability of the z/OS operating system, and the strength of the DB2 for z/OS data server can help expand and extend your IT infrastructure and the business value of your data. The combination of DB2 and System z can provide an advantage for on demand environments by enabling a flexible, cost effective and optimized foundation for Information on Demand. This foundation can allow you to manage risk, support your efforts to demonstrate compliance with policies and standards, and help to simplify management of your information infrastructure. These capabilities are important to enable customers to use their core business data to drive insight and help gain competitive advantage.

In today's on demand world, providing the support for internal and external clients requires an increasing dependence on IT, with more users and more applications seeking access to a common set of data in real time that is accessible anytime, from anywhere across the globe. So how do you help meet these needs? The answer will, of course, be dependent on your individual situation and requirements; there is no 'one size fits all' answer for data serving. However, if you feel that your individual situation can include consolidating to fewer copies of data, running on highly scalable servers, while also providing a high level of security and availability, then the IBM System z9, which builds upon the inherent strengths of the IBM mainframe to deliver industry-leading data and transaction serving capabilities, may well be the right choice for your data serving needs.

Enhancing flexibility with z9 EC and SOA

There is a growing recognition in the IT industry of the potential benefits of Service-Oriented Architecture (SOA) for building new applications. Over the years, many IBM customers have employed and developed business applications running on z/OS, using a combination of CICS®, IMS™, and DB2 for z/OS. Consequently, the inherent strengths and capabilities of a z/OS environment running on a z9 EC makes it an ideal platform from which to develop, deploy and manage applications.

The use of the IBM SOA products, such as the new IBM WebSphere® Developer for zSeries® V6.0, may help assist in the generation of Web and user interfaces for core business. IBM WebSphere Process Server for z/OS V6.0 is designed to help enable the integration of diverse "services" such as multiple core applications, new applications or packaged applications within the same workspace. The IBM WebSphere, Rational® and Tivoli® products feature technology in middleware and management tools designed to help reduce operational overhead.

Choosing to deploy SOA on the z9 EC may help enhance application re-use, and may help reduce the cost and risk of new development projects and bring flexibility and responsiveness to the way customers are able to tackle business challenges or opportunities.

Many enterprises are realizing that the mainframe, which is at the core of their infrastructure today, is a critical element of their on demand operating environment. Its core strengths of scalability, security and resiliency, availability, and its data serving capabilities can work together to enhance the role of the mainframe as the data hub of the enterprise.

Yet strength is not measured by the power and hardware features alone. Strength is also derived from the ability of the z9 EC to utilize open computing standards to allow integration of existing resources and to build and deploy effective applications. It's also derived from the teaming of I/O capabilities with IBM's storage products to create an environment that is optimized to work together.

Helping Secure Your Enterprise – IBM Mainframe Encryption

Protecting sensitive data is a growing concern for companies around the globe. The importance of securing critical business data and customer information reaches to the corporate boardroom, because failure to protect these assets may result in high out-of-pocket costs and, more importantly, may also result in lost customer and investor confidence. Data protection may also be required by stringent government regulations and contractual obligations with business partners. Whether the data moves across the network or across town on a tape in a truck, the object is to make it usable to those who are authorized and inaccessible to those who are not.

With IBM Encryption Facility for z/OS software and Integrated Cryptographic Service Facility (ICSF), IBM offers a solution for encrypting data at rest that exploits the existing strengths of the mainframe. The Encryption Facility for z/OS software allows you to exchange encrypted tapes across the enterprise and with partners even if the recipient does not have access to IBM software.

The encryption capabilities provided in IBM mainframe servers, the z/OS operating system, and the Encryption Facility for z/OS are designed to provide a comprehensive approach for data encryption to tape and disk.

z/Architecture

The System z9 platform is based on the z/Architecture®, designed to reduce bottlenecks associated with the lack of addressable memory and built to automatically direct resources to priority work through Intelligent Resource Director. The z/Architecture is a 64-bit superset of ESA/390.

z/Architecture is implemented on the System z9 platform to allow 64-bit real and virtual storage support. A maximum 512 GB of real storage is available on z9 EC. z9 EC can define any logical partition as having 31-bit or 64-bit addressability.

z/Architecture has:

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

- *IEEE Floating Point architecture adds twelve new instructions for 64-bit integer conversion*
- *The 64-bit SIE architecture allows a z/Architecture server to support both ESA/390 (31-bit) and z/Architecture (64-bit) guests and Zone Relocation is expanded to 64-bit for LPAR and z/VM®*
- *64-bit operands and general registers are used for all Cryptographic instructions*
- *The implementation of 64-bit z/Architecture can help reduce problems associated with lack of addressable memory by making the addressing capability virtually unlimited (16 Exabytes)*

z/Architecture operating system support

The z/Architecture is a tri-modal architecture capable of executing in 24-bit, 31-bit, or 64-bit addressing modes. Operating systems and middleware products have been modified to exploit the new capabilities of the z/Architecture. Immediate benefit may be realized by the elimination of the overhead of Central Storage to Expanded Storage page movement and the relief provided for those constrained by the 2 GB real storage limit of ESA/390. Application programs can run unmodified on the System z9 platform.

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

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

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

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

Operating System	ESA/390 31-bit mode	z/Architecture 64-bit mode
z/OS.e[1] and z/OS V1R6, 7, 8	No	Yes
z/OS V1R9 (Planned*)	No	Yes
Linux® on System z, 64-bit distribution	No	Yes
Linux on System z, 31-bit distribution	Yes	No
z/VM V5R1[2], 2, 3[3]	No	Yes
z/VSE™ V4R1[5]	No	Yes
z/VSE V3R1[4]	Yes	No
z/TPF V1R1	No	Yes
TPF V4R1 (ESA mode only)	Yes	No

1. z/OS.e - z800, z890 and z9 BC only. Release 1.8 will be the last release of z/OS.e.
2. Support for z/VM 5.1 will end September 30, 2007.
3. z/VM 5.3 is planned to GA in June 2007.
5. z/VSE V3 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 System z9 and zSeries hardware.
7. z/VSE V4 is designed to exploit 64-bit real memory addressing, but will not support 64-bit virtual addressing.

Note: Please refer to the latest PSP bucket for latest PTFs for new functions/features.

* All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

IBM System z9 EC

IBM System z9 Enterprise Class (z9 EC) brings the power of IBM to deliver intelligent and adaptive virtualization capabilities, leveraged to both protect the business and designed to dynamically orchestrate and integrate end-to-end computing power so that it can be optimized to the priorities of the business. z9 EC offers high levels of security and resiliency, giving you the peace of mind and confidence in a foundation designed to avoid costly downtime, create new business opportunities, and maximize user productivity. System z9 virtual system solutions are designed to be dynamically optimized, providing the capability to orchestrate resources according to business priorities and provide the business flexibility that enables peak levels of efficiency and effectiveness. By applying its intelligence and adaptability to advanced technologies, z9 EC creates a resilient resource pool that combines a solid foundation with flexibility that helps enable the integration of business and information management.

z9 EC builds upon the structure introduced on the IBM eServer™ zSeries 990 (z990) – scalability and z/Architecture. z9 EC expands upon a key attribute of the platform – availability – to help ensure you have a resilient infrastructure designed to satisfy the demands of your business. With the potential for increased performance and capacity, you have an opportunity to continue to consolidate diverse applications on a single platform. With a modular book design, the z9 EC is designed to provide up to 95% more total system capacity than the z990 Model D32, and has up to double the available memory. The maximum number of Processor Units (PUs) has grown from 32 to 54, and memory capacity has doubled, up to 128 GB per book and up to 512 GB per system. To improve the system virtualization capacity, the number of logical partitions has increased from 30 to up to 60 partitions, and the new N_Port ID Virtualization (NPIV) for FCP channel virtualization has been implemented. The I/O cage connectivity has been increased by up to 40%, now allowing you to populate an I/O cage with up to 28 features in any combination. You can now install up to 336 FICON® Express4 or FICON Express2 channels per server, which represents up to 40% more FICON channels than z990.

z9 EC provides increased addressable storage, with up to 768 additional I/O devices to the system, and also implements support for multiple subchannel sets (MSS), doubling the I/O device addressing capability and benefiting Parallel Access Volumes (PAV). The exclusive Modified Indirect Data Address Word (MIDAW) facility for FICON and ESCON® channels is also introduced to help reduce latency and system overhead for I/O requests using extended format datasets. Other System z9 exclusive functions are Program Directed re-IPL and native FICON link incident reporting.

PUs defined as Internal Coupling Facilities (ICFs), Integrated Facility for Linux (IFLs), System z Application Assist Processor (zAAPs) and System z9 Integrated Information Processor (zIIPs) are no longer grouped together in one pool as on the z990, but are grouped together in their own pool, where they can be managed separately. This simplifies capacity planning and management for LPAR significantly. The separation also can have an effect on weight management since CP weights and zAAP and zIIP weights can now be managed separately. Capacity BackUp (CBU) features are available for IFLs, ICFs, zAAPs and zIIPs.

For LAN connectivity, z9 EC now provides the OSA-Express2 1000BASE-T Ethernet feature, and supports IP version 6 (IPv6) on HiperSockets™. OSA-Express2 OSN (OSA for NCP) is also available on z9 EC to support the Channel Data Link Control (CDLC) protocol, providing direct access from the host operating system images to the Communication Controller for Linux on System z9 (CCL) using OSA-Express2 to help eliminate the requirement for external hardware for communications.

z9 EC now offers a configurable Crypto Express2 feature, with PCI-X adapters that can be individually configured as a secure coprocessor or an accelerator for SSL, a new TKE 5.0 workstation with optional Smart Card Reader, and provides the following enhancements to the CP Assist for Cryptographic Function (CPACF):

- *Advanced Encryption Standard (AES)*
- *Secure Hash Algorithm – 256 (SHA-256)*
- *Pseudo Random Number Generation (PRNG)*

To help reduce planned and unplanned server outages, the z9 EC provides the following functions and features:

- *Enhanced Book Availability*
- *Redundant I/O Interconnect*
- *Flexible Memory Option*
- *Enhanced Driver Maintenance*
- *Concurrent MBA fanout card replacement*
- *Dynamic Oscillator switchover*

Enterprises with IBM System z9 Business Class (z9 BC) Model S07, z990s and IBM eServer zSeries 900 (z900) (except z900 Model 100) may upgrade to z9 EC to help ensure a resilient infrastructure. When you demand the highest level of availability, you will want z9 EC for its nondisruptive memory, book, and I/O repair and upgrade capabilities, achievable with proper planning, particularly for planned outages. If you desire a consolidation platform for your mainframe and Linux capable applications, you can add capacity and even expand your current application workloads in a cost-effective manner. If your traditional applications and new applications are growing, you will find z9 EC a good fit with its base qualities of service and its specialty processors designed for assisting with new workloads. If you have G5 and G6 servers (9672s) you should consider the z9 EC to migrate your workloads to z/Architecture in order to take advantage of on demand functions like On/Off Capacity on Demand and Intelligent Resource Director (IRD). Innovation on z9 EC can provide increases in capacity, bandwidth, and logical partitions in a single server. Value for all is leveraged with improved hardware price/performance and System z9 software pricing strategies.

With a superscalar microprocessor based on CMOS 10S-SOI technology, the z9 EC is designed to further extend and integrate key platform characteristics such as dynamic flexible partitioning and resource management in mixed and unpredictable workload environments, providing scalability, high availability and Quality of Service to emerging applications such as WebSphere, Java™ and Linux.

With the logical partition (LPAR) group capacity limit on z9 EC and z9 BC, you can now specify LPAR group capacity limits allowing you to define each LPAR with its own capacity and one or more groups of LPARs on a server. This is designed to allow z/OS to manage the groups in such a way that the sum of the LPARs' CPU utilization within a group will not exceed the group's defined capacity. Each LPAR in a group can still optionally continue to define an individual LPAR capacity limit.

LPAR group capacity limit requires that all LPARs managed in the group are running at z/OS or z/OS.e 1.8 or later. LPAR group capacity limits may help provision a portion of a System z9 server to a group of LPARs allowing the CPU resources to float more readily between those LPARs, resulting in more productive use of "white space" and higher server utilization.

The z9 EC has five models with a total of 78 capacity settings available as new build systems and as upgrades from the z990 and z900 (except for the Model 100).

The five z9 EC models are designed with a multi-book system structure that provides up to 54 Processor Units (PUs) that can be characterized as either Central Processors (CPs), IFLs, ICFs, zAAPs or zIIPs.

The System z Application Assist Processor (zAAP), available also on the z9 BC, z990 and IBM eServer zSeries 890 (z890) servers, is an attractively-priced specialized processor that provides a strategic z/OS application Java execution environment for those who desire the powerful integration advantages and traditional Qualities of Service of the platform. Java is the first application exploiter of the zAAP.

When configured with general purpose Central Processors (CPs) within logical partitions running z/OS, zAAPs can help you to extend the value of your existing investments and strategically integrate and run Java workloads on the same server as your database. This may help to simplify and reduce the infrastructure required for Web applications while helping to lower your overall total cost of ownership.

zAAPs are designed to operate asynchronously with general purpose CPs and can execute Java programming under control of the IBM Java Virtual Machine (JVM). This can help reduce the demands and capacity requirements on general purpose CPs which may then be available for reallocation to other System z9 workloads. The amount of general purpose CP savings may vary based on the amount of Java application code executed by zAAPs. Best of all, IBM JVM processing cycles can be executed on the configured zAAPs with no anticipated modifications to the Java applications. Execution of the JVM processing cycles on a zAAP is a function of the IBM Software Developer's Kit (SDK) for z/OS Java 2 Technology Edition, z/OS 1.6 and later and the innovative Processor Resource/Systems Manager™ (PR/SM™).

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

zAAPs allow you to purchase additional processing power exclusively for z/OS Java application execution without affecting the total MSU rating or server model designation. Conceptually, zAAPs are very similar to a System Assist Processor (SAP); they cannot execute an Initial Program Load and only assist the general purpose CPs for the execution of Java programming. IBM does not impose software charges on zAAP capacity. Additional IBM software charges will apply when additional general purpose CP capacity is used.

The System z9 Integrated Information Processor (zIIP) is the latest customer-inspired specialty engine designed to help improve resource optimization and lower the cost of portions of eligible workloads. The zIIP can help to strengthen the System z9 mainframe as the data serving hub, helping customers to more fully leverage their valuable assets.

The zIIP's execution environment will accept eligible work from z/OS 1.6 and above, which will manage and direct the work between the general purpose processor and the zIIP. The zIIP is designed so that a program can work with z/OS (z/OS.e) to have all or a portion of its Service Request Block (SRB) dispatched work directed to the zIIP. The zIIP is available on the System z9 mainframe, and its introduction can help increase the value that customers may derive from the System z9 mainframe over previous generations of the IBM mainframe.

DB2 for z/OS V8 (and DB2 9 for z/OS) will exploit the zIIP capability for portions of eligible workload. Types of eligible work that access DB2, such as Customer Relationship Management (CRM), Enterprise Resource Planning (ERP), Business Intelligence (BI), and data warehousing applications,can have portions of their work be directed to the zIIP.

Other types of eligible DB2 for z/OS V8 (and DB2 9 for z/OS) workloads executing in SRB mode, all or a portion of that can be sent to the zIIP include:

1. *For Network Connected Applications – An application (running on UNIX®, Linux, Intel®, Linux on System z or z/OS) may access a DB2 for z/OS database that is hosted on a z9 EC. Eligible work that can be directed to the zIIP are portions of those requests made from the application server, to the host, via SQL calls over a DRDA® over TCP/IP connection. DB2 for z/OS gives z/OS the necessary information to direct portions of the eligible work to the zIIP. Examples of workloads that may be running on the server connected via DRDA over TCP/IP to the z9 EC may include BI, ERP, or CRM application serving. IBM's DB2 Connect™ is an example of IBM software that uses DRDA over TCP/IP.*
2. *For Data Warehousing Applications – Applications may execute queries to a DB2 for z/OS database that is hosted on a z9 EC. Eligible work that can be directed to the zIIP are portions of requests that utilize star schema parallel queries. DB2 for z/OS gives z/OS the necessary information to direct portions of these queries to the zIIP. Examples of these applications may include BI applications.*
3. *There are some DB2 for z/OS utility functions (LOAD, REORG and REBUILD INDEX), written in SRB mode, that are used to maintain index maintenance structures. Eligible work that can be directed to the zIIP are portions of those utility functions that execute in SRB mode. DB2 for z/OS gives z/OS the necessary information to direct a portion of these functions to the zIIP.*

If a star schema parallel query comes in remotely via DRDA over TCP/IP, a portion of the workload coming across the DRDA over TCP/IP connection can be redirected to the IBM zIIP, along with the portion of the star schema parallel query processing that is redirected.

In addition to DB2 exploitation of the zIIP, the z/OS Communications Server will be enhanced to allow IPSec processing to take advantage of zIIPs. The new zIIP Assisted IPSec function is designed to move most of the IPSec processing from general purpose processors to the zIIPs. In addition to performing eligible encryption processing, the zIIP will also handle cryptographic validation of message integrity, and IPSec header processing. This is designed to allow you to take advantage of the cost saving benefits of the zIIP when you implement IPSec to secure your valuable business transactions and bulk data movement and to protect your host. Specifically, the z/OS Communications Server is designed to interact with z/OS Workload Manager to have all of its enclave Service Request Block (SRB) work made eligible to run on the zIIP. This capability is planned to be available in August 2007 with z/OS 1.8 and PTFs, and native in z/OS 1.9, when available.

The zIIP is intended to help customers integrate and secure their database workloads, better leverage the data that they have on the mainframe and free up capacity on the general purpose processor which may make it available for use by other workloads running on the server.

The use of zAAPs and zIIPs by a single transaction flow is not mutually exclusive. The two specialty engines are designed to run two different types of work. The zAAP processor provides a specialized engine for running new application technologies (such as Java) and the zIIP is designed to redirect workload that is typically more closely associated with operating system processing.

Customers may order zIIPs up to the number of permanently purchased general purpose processors (CPs) on a given z9 EC model. This requirement is at a server level, so a customer could have an LPAR with more zIIPs than general purpose processors (CPs), as long as there are enough general purpose processors (CPs) in the entire server to meet the one for one requirement. For example, if a customer currently has a general purpose processor (CP) and one zAAP, they can order one zIIP without needing to order another general purpose processor (CP) to meet the ordering restriction.

IBM does not impose software charges on zIIP capacity. Additional IBM software charges will apply when general purpose processor (CP) capacity is used. The amount of general purpose processor savings will vary based on the amount of workload executed by the zIIP, among other factors. Combining the qualities of service provided by the z9 EC, z/OS and DB2 for z/OS V8, with the cost effectiveness of data access via zIIP may help reduce the need for many local copies of data and additional IT complexity that scenario brings.

z/VM 5.3 is designed to provide new guest support for zAAPs and zIIPs and includes:

- *Simulation support — z/VM simulates specialty processors for guest virtual machines by dispatching the virtual specialty processors on real CPs. Simulating specialty processors provides a test platform for z/VM guests to exploit mixed-processor configurations. This allows users to assess the operational and CPU utilization implications of configuring a z/OS system with zIIP or zAAP processors without requiring the real specialty processor hardware. This simulation also supports z/VM's continuing role as a disaster-recovery platform, since a virtual configuration can be defined to match the real hardware configuration even when real zIIP or zAAP processors are not available on the recovery system. z/VM simulates specialty processors using real CPs if the underlying hardware is capable of supporting the real specialty processor. zIIPs can be simulated only on System z9 (z9 EC and z9 BC) servers. zAAPs can be simulated only on z9 EC, z9 BC, z990, and z890 servers.*
- *Virtualization support — z/VM can create virtual specialty processors for virtual machines by dispatching the virtual processors on corresponding specialty processors of the same type in the real configuration. Guest support for zAAPs and zIIPs may help improve your total cost of ownership by allowing available zAAP and zIIP capacity not being used by z/OS LPARs to be allocated to a z/VM LPAR hosting z/OS guests running Java and DB2 workloads. zAAPs and zIIPs cost less than standard CPs, so this support might enable you to avoid purchasing additional CPs, thereby helping to reduce your costs both for additional hardware and for software licensing fees.*

You are encouraged to contact your specific ISVs and USVs directly to determine if or how your charges will be affected.

Some of the significant enhancements in the z9 EC that help bring improved performance, availability and function to the platform have been identified. The following sections highlight the functions and features of the z9 EC.

z9 EC Design and Technology

The z9 EC is designed to provide balanced system performance. From processor storage to the system's I/O and network channels, end-to-end bandwidth is provided and designed to deliver data where and when it is needed.

The z9 EC provides a significant increase in system scalability and opportunity for server consolidation by providing five models, from one to four MultiChip Modules (MCMs), delivering up to a maximum 54-way configuration. The MCMs are configured in a book package, with each book comprised of a MultiChip Module (MCM), memory cards and Self-Timed Interconnects. The MCMs, which measure approximately 95 x 95 millimeters, contain the Processor Unit (PU) chips, the cache structure chips and the processor storage controller chips. The MCMs contain 102 glass ceramic layers to provide interconnection between the chips and the off-module environment. A 12-PU MCM is used on z9 EC models from one to four books providing up to 38 PUs, and a 16-PU MCM is used on a four-book model providing up to 54 PUs on a single z9 EC. Each MCM provides support for 40 MB level 2 cache. Each PU measures 15.78 mm x 11.84 mm and has level 1 cache sizes of 256 KB for instructions and 256 KB for data. The design of the MCM technology on the z9 EC provides the flexibility to configure the PUs for different uses; two of the PUs on each MCM are reserved for use as System Assist Processors (SAPs), and two PUs of the first book's MCM are reserved as spares. The remaining inactive PUs on each installed MCM are available to be characterized as either CPs, ICF processors for Coupling Facility applications, or IFLs for Linux applications and z/VM hosting Linux as a guest, System z Application Assist Processors (zAAPs), System z9 Integrated Information Processors (zIIPs) or as optional SAPs provide you with tremendous flexibility in establishing the best system for running applications. Each model of the z9 EC must always be ordered with at least one CP, IFL or ICF.

The PU, which uses the latest chip technology from IBM semiconductor laboratories, is built on CMOS 10S-SOI with copper interconnections and has a cycle time of 0.58 nanoseconds. Implemented on this chip is the z/Architecture with its 64-bit capabilities including instructions, 64-bit General Purpose Registers and translation facilities.

Each book can support up to 128 GB of memory, delivered on four or eight memory cards, and up to 16 Self-Timed Interconnects (STIs) per book, delivering up to 512 GB of memory and up to 64 STIs on the S38 model. The memory is delivered on 4 GB, 8 GB or 16 GB memory cards which can be purchased in 16 GB increments. The minimum memory is 16 GB. Each book has up to 8 Memory Bus Adapter (MBA) fanout cards and each MBA fanout card supports 2 STIs. The bandwidth of each STI is up to 2.7 GigaBytes per second (GB/sec) for I/O and 2 GB/sec for ICB-4s.

All books are interconnected with a super-fast bi-directional redundant ring structure which allows the system to be operated and controlled by PR/SM operating in LPAR mode as a symmetrical, memory coherent, multiprocessor. PR/SM provides the ability to configure and operate as many as 60 Logical Partitions which may be assigned processors, memory and I/O resources from any of the available books. The z9 EC supports LPAR mode only (i.e. basic mode is no longer supported).

The MultiChip Module (MCM) is the technology cornerstone for flexible PU deployment in the z9 EC models. For most models, the ability of the MCM to have inactive PUs allows such features as Capacity Upgrade on Demand (CUoD), Customer Initiated Upgrade (CIU), and the ability to add CPs, ICFs, IFLs, zAAPs and zIIPs dynamically, providing nondisruptive upgrade of processing capability. Also, the ability to add CPs, ICFs, IFLs, zAAPs and zIIPs lets a z9 EC with spare PU capacity become a backup for other systems in the enterprise, expanding the z9 EC to meet emergency outage situation. This is called Capacity

BackUp (CBU). The greater capacity of the z9 EC offers you even more flexibility for using this feature to back up critical systems in their enterprise.

The z9 EC has been designed to offer high performance and efficient I/O structure. All z9 EC models ship with two frames: an A-Frame and a Z-Frame, which together support the installation of up to three I/O cages. Each I/O cage has the capability of plugging up to 28 I/O features. When used in conjunction with the software that supports Logical Channel SubSystems, it is possible to have up to 360 ESCON channels in a single I/O cage and a maximum of 1024 channels across 3 I/O cages. Alternatively, three I/O cages will support up to 336 FICON Express4 and FICON Express2 channels. Each book will support up to 16 STIs for I/O connectivity. Eight STIs are required to support the 28 channel slots in each I/O cage, so in order to support a fully configured three I/O cage system, 24 STIs are required. To achieve this maximum I/O connectivity requires at least an S18 model, which provides 32 STIs.

To increase the I/O device addressing capability, the I/O subsystem has been enhanced by introducing support for multiple subchannels sets (MSS), which are designed to allow improved device connectivity for Parallel Access Volumes (PAVs). To support the highly scalable multi-book system design, the z9 EC I/O subsystem uses the Logical Channel SubSystem (LCSS) which provides the capability to install up to 1024 CHPIDs across three I/O cages (256 per operating system image). The Parallel Sysplex® Coupling Link architecture and technology continues to support high speed links providing efficient transmission between the Coupling Facility and z/OS systems. HiperSockets provides high-speed capability to communicate among virtual servers and logical partitions. HiperSockets is now improved with the IP version 6 (IPv6) support; this is based on high-speed TCP/IP memory speed transfers and provides value in allowing applications running in one partition to communicate with applications running in another without dependency on an external network. Industry standard and openness are design objectives for I/O in z9 EC.

The z9 EC has five models offering between 1 to 54 processor units (PUs), which can be configured to provide a highly scalable solution designed to meet the needs of both high transaction processing applications and On Demand Business. Four models (S08, S18, S28 and S38) have 12 PUs per book and the high capacity model, S54, offers 16 PUs in each of its four books. The PUs can be characterized as either CPs, IFLs, ICFs, zAAPs or zIIPs. An easy-to-enable ability to “turn off” CPs or IFLs is available on z9 EC, allowing you to purchase capacity for future use with minimal or no impact on software billing. An MES feature will enable the “turned off” CPs or IFLs for use where you require the increased capacity. There are a wide range of upgrade options available which are indicated in the z9 EC models chart.

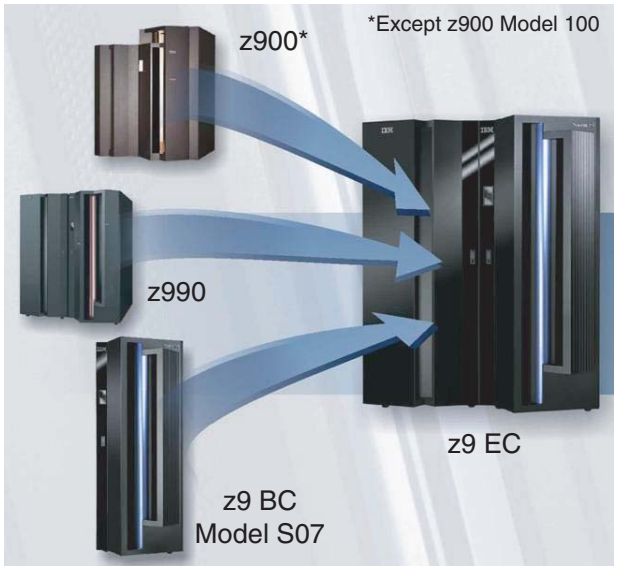

The z9 EC hardware model numbers (S08, S18, S28, S38, S54) on their own do not indicate the number of PUs which are being used as CPs. For software billing purposes only, there will be a Capacity Indicator associated with the number of PUs that are characterized as CPs. This number will be reported by the Store System Information (STSI) instruction for software billing purposes only. There is no affinity between the hardware model and the number

of CPs. For example, it is possible to have a Model S18 which has 5 PUs characterized as CPs, so for software billing purposes, the STSI instruction would report 705.

z9 EC model upgrades

There are full upgrades within the z9 EC models and upgrades from z9 BC Model S07 (to S08 only), z990 or z900 (except Model 100) to z9 EC. There are any-to-any upgrades from any of the z990 models and the z900 general purpose models. A z900 Coupling Facility Model 100 must first be upgraded to a z900 general purpose model before upgrading to a z9 EC. There are no direct upgrades from IBM eServer zSeries 890 (z890), or IBM eServer zSeries 800 (z800), or 9672 G5 and G6.

Announced in April 2006, IBM is introducing subcapacity engines on the z9 EC. A total of 24 new settings, each with less capacity than the full capacity 8-way (z9 EC 708) will offer outstanding granularity for customers' workload and cost management.

For the z9 EC 701 to 708, there are four capacity settings per engine for central processors (CPs) providing a new lower point of entry. The smallest capacity model is approximately 66% smaller than the z9 EC 701. Subcapacity processors will have availability of z9 EC features/functions and any to any upgradeability is available within the new subcapacity matrix, as well as to current z9 EC capacity settings. All CPs must be the same capacity setting size within one z9 EC. Specialty engines will continue to run as full size engines.

z9 EC software models:

- *700, 401 to 408, 501 to 508, 601 to 608 and 701 to 754*
- *Capacity setting 700 does not have any CP engines*
- *Nxx, where n = the capacity setting of the engine, and xx = the number of PU characterized as CPs in the CEC*
- *Once xx exceeds 08, then all CP engines are full capacity*

Upgrades from Models S08, S18, S28, or S38 to Model S54 are disruptive, requiring the replacement of the 12-PU MCM book(s) by the 16-PU MCM books.

z9 EC Granular Capacity for up to 8 CPs

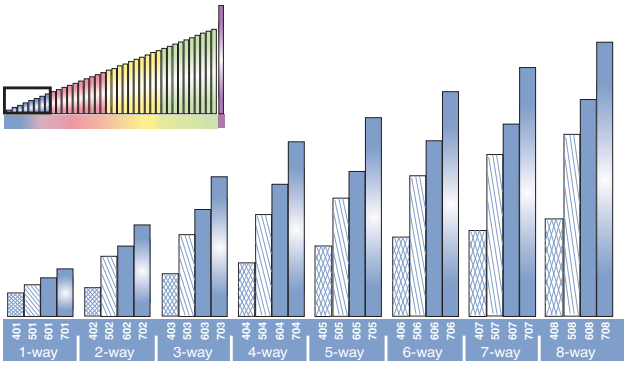

- *The z9 EC now has 24 additional capacity settings within the first eight CPs*
- *Entry point is approximately 66% smaller than the 701*
- *All CPs must be the same capacity within one z9 EC*
- *Combined zAAPs and/or zIIPs can not be more than 2x the number of CPs*
- *Only 8 CPs can have granular capacity, other PUs must be characterized as specialty engines*

z9 EC Performance

The performance design of the z/Architecture can enable the server to support a new standard of performance for applications through expanding upon a balanced system approach. As CMOS technology has been enhanced to support not only additional processing power, but also more PUs, the entire server is modified to support the increase in processing power. The I/O subsystem supports a greater amount of bandwidth than previous generations through internal changes, providing for larger and faster volume of 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 operating system and hardware support of 64-bit addressing. The combined balanced system design allows for increases in performance across a broad spectrum of work. However, due to the increased flexibility in the z9 EC model structure and resource management in the system, it is expected that there will be greater performance variability than has been previously seen by our traditional customer set. This variability may be observed in several ways.

Large System Performance Reference

The Large System Performance Reference (LSPR) should be referenced when considering performance on the z9 EC. The range of performance ratings across the individual LSPR workloads is likely to have a large spread. 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 and additional PUs available on the z9 EC. The impact of this increased variability is expected to 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 z9 EC has been optimized to run many workloads at high utilization rates.

With a modular book design, the z9 EC is designed to provide up to 95% more total system capacity than the z990 Model D32, and has up to double the available memory. The performance of the z9 EC (2094) 701 is expected to be 1.35 times the z990 (2084) 301 (LSPR mixed workload).

It is also important to notice that the LSPR workloads have been updated to reflect more closely your current and growth workloads. The traditional Commercial Batch Short Job Steps (CB-S) workload (formerly CB84) is dropped and a new Java batch (CB-J) workload is added. The remainder of the LSPR workloads are the same as ones as were used for the z990 LSPR. The new LSPR provides two tables: the single image z/OS from 1-way to 32-way, and the typical LPAR configuration from 1-way to 54-way, which is based on customer profiles. The typical LPAR configuration table is used to establish single-number-metrics such as MIPS and MSUs. The z9 EC LSPR will rate all z/Architecture processors running in LPAR mode and 64-bit mode. The existing zSeries servers have all been re-measured using the new workloads running in LPAR mode and 64-bit mode.

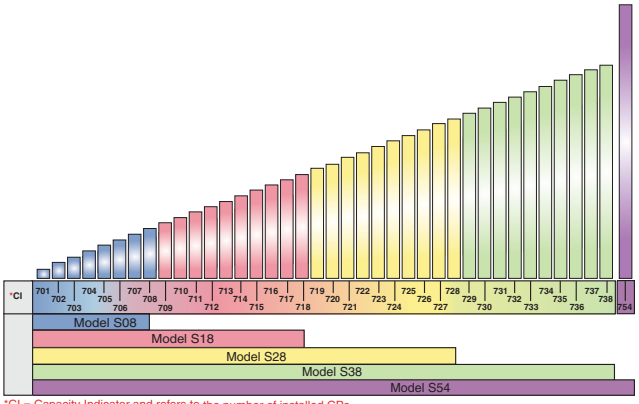

*CI = Capacity Indicator and refers to the number of installed CPs.
Reported by STSI instruction. Model 700 does not have any CPs

For more information about performance comparisons, please consult the LSPR at **ibm.com**/servers/eserver/zseries/lspr/, MSU Values; www-1.**ibm.com**/servers/eserver/zseries/library/swpriceinfo.

z9 EC Modified Indirect Data Address Word (MIDAW) Facility

The System z9 I/O architecture supports a new facility for indirect addressing, Modified Indirect Data Address Word (MIDAW) facility, for both ESCON and FICON channels. This is a new system architecture and software exploitation designed to improve the performance of many applications by reducing FICON channel, director, and control unit overhead.

The MIDAW facility is exclusive to System z9 and is supported by ESCON using CHPID type CNC and by FICON using CHPID types FCV and FC. The MIDAW facility is exploited by z/OS and z/VM 5.3 for z/OS guest exploitation.

FCV is available for configuring on System z servers when a FICON Express LX feature #2319 exists. Feature #2319 is only available as carry forward on z9 EC and z9 BC. It cannot be ordered.

The z9 EC contains an I/O subsystem infrastructure which uses an I/O cage that provides 28 I/O slots and the ability to have one to three I/O cages delivering a total of 84 I/O slots. ESCON, FICON Express4, FICON Express2, FICON Express, OSA-Express2, OSA-Express and Crypto Express2 features plug into the z9 EC I/O cage along with any ISC-3s and STI-3 distribution cards. All I/O features and their support cards can be hot-plugged in the I/O cage. Installation of an I/O cage remains a disruptive MES, so the Plan Ahead feature remains an important consideration when ordering a z9 EC system. Each model ships with one I/O cage as standard in the A-Frame (the A-Frame also contains the Central Electronic Complex [CEC] cage where the books reside) and any additional I/O cages are installed in the Z-Frame. Each STI has a bandwidth up to 2.7 GigaBytes per second (GB/sec) for I/O domains and 2.0 GB/sec for ICBs-4.

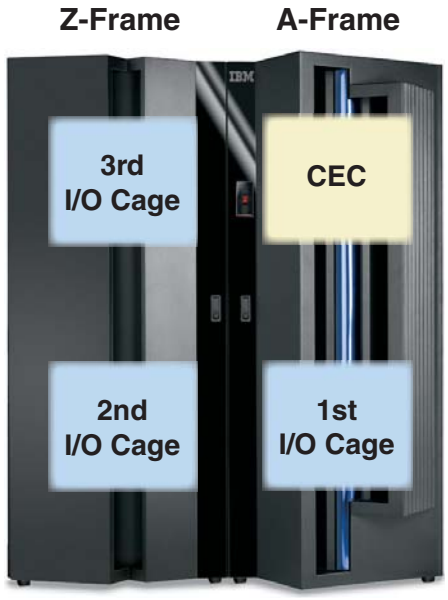

z9 EC Logical Channel SubSystems (LCSSs) and support for up to 60 logical partitions

In order to provide the channel connectivity required to support the scalability of the z9 EC, the z9 EC channel I/O subsystem provides up to 4 Logical Channel Subsystems (LCSSs) per CEC, each of which can support up to 256 Channel Path Identifiers (CHPIDs). This implementation is provided in such a way that it intended to be transparent to the programs operating in the logical partition. Each LCSS may have from 1 to 256 CHPIDs and may in turn be configured with 1 to 15 logical partitions. Up to 60 logical partitions can be defined, with up to 15 running under a single LCSS. As with previous mainframes, Multiple Image Facility (MIF) channel sharing as well as all other channel subsystem features are available to each Logical Partition configured to each LCSS.

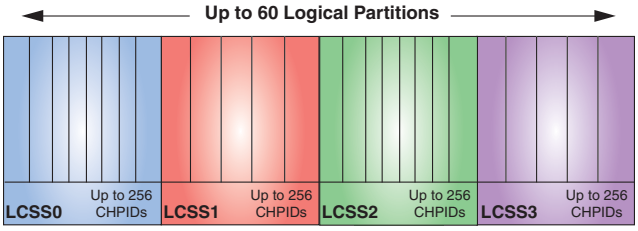

z9 EC Increased number of Subchannels (63.75K)

A subchannel represents an I/O device to the hardware and is used by operating systems to pass an I/O request to the channel subsystem. System z9 technology uses a maximum of 65,536 (64K) subchannels with 1,024 (1K) of these previously reserved for system use. IBM has now made available 768 of these 1K reserved subchannels for customer use, which can result in 65,280 (63.75K) subchannels for I/O devices addressability. The potential increased addressable storage this represents can be significant: using 3390-3 volume sizes and 768 volumes of 54 GB/volume, this represents 41 TeraBytes (TB) of increased storage addressability (54 GB/volume * 768 volumes = 41 TB). This is exclusive to System z9.

z9 EC Multiple Subchannel Sets (MSS)

Multiple Subchannel Sets (MSS) are designed to provide more subchannels on System z9 servers. Two subchannel sets are now available per LCSS and are designed to enable a total of 63.75K subchannels in set-0 and adding 64K-1 subchannels in set-1. z/OS will allow Parallel Access Volume Alias (PAV-alias) devices in the subchannel set 1. Only one set of 63K subchannels are available with z990, z890, z900, and z800 servers. This capability can provide greater I/O device configuration capability of large enterprises. This function is exclusive to System z9. Refer to the System z9 Operating Systems Support section for further information.

z9 EC Redundant I/O Interconnect

Redundant I/O Interconnect is designed to allow you to replace a book or respond to a book failure and retain connectivity to resources. In the event of a failure or customer initiated action such as the replacement of a book or an MBA fanout card, the z9 EC is designed to provide access to your I/O devices in the affected I/O domain through another STI. This function is exclusive to System z9.

Physical Channel IDs (PCHIDs) SubSystem

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

System-initiated CHPID Reconfiguration

System-initiated CHPID reconfiguration function is designed to reduce the duration of a repair action when an ESCON or FICON channel, an OSA port, or an ISC-3 link is shared across logical partitions. This is designed to minimize operator interaction to configure channels offline and online.

Logical Channel SubSystem (LCSS) Spanning

The concept of spanning channels provides the ability for a channel to be configured to multiple LCSSs and therefore be transparently shared by any or all of the logical partitions in those LCSSs. Normal Multiple Image Facility (MIF) sharing of a channel is confined to a single LCSS. The z9 EC supports the spanning of these channel types: IC, HiperSockets, FICON Express4, FICON Express2, FICON Express, OSA-Express2, OSA-Express, ISC-3, ICB-3, ICB-4. Note: ESCON architecture does not support spanning of ESCON channels.

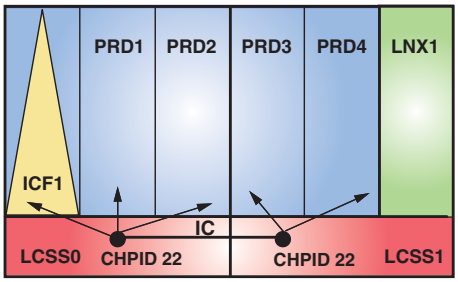

A z9 EC with all three I/O cages installed has a total of 84 I/O slots. These slots can be plugged with a mixture of features providing the I/O connectivity, networking connectivity, coupling and cryptographic capability of the server.

Multipath Initial Program Load (IPL) is designed to help eliminate manual problem determination when executing an IPL. If an error occurs, an alternate path is selected. Multipath IPL is applicable to ESCON channels (CHPID type CNC) and FICON channels (CHPID type FC). z/OS and z/OS.e provide support for this enhancement.

ESCON Channels

The System z9 supports up to 1,024 ESCON channels. The high density ESCON feature has 16 ports, 15 of which can be activated for customer use. One port is always reserved as a spare which is activated in the event of a failure of one of the other ports. For high availability the initial order of ESCON features will deliver two 16-port ESCON features and the active ports will be distributed across those features. After the initial install the ESCON features are installed in increments of one. ESCON channels are available in four-port increments and are activated using IBM Licensed Internal Code, Configuration Control (LIC CC).

Fibre Channel Connectivity

The on demand operating environment requires fast data access, continuous data availability, and improved flexibility, all with a lower cost of ownership. The four port FICON Express4 and FICON Express2 features available on the z9 EC helps distinguish the z9 EC, further setting it apart as enterprise class in terms of the number of simultaneous I/O connections available.

FICON Express4 Channels

The z9 EC supports up to 336 FICON Express4 channels, each one operating at 1, 2 or 4 Gb/sec auto-negotiated. The FICON Express4 features are available in long wavelength (LX) and short wavelength (SX). For customers exploiting LX, there are two options available for unrepeated distances of up to 4 kilometers (2.5 miles) or up to 10 kilometers (6.2 miles). Both LX features use 9 micron single mode fiber optic cables. The SX feature uses 50 or 62.5 micron multimode fiber. Each FICON Express4 feature has 4 independent channels (ports) and can be configured to carry native FICON traffic or Fibre Channel (SCSI) traffic. LX and SX cannot be intermixed on a single feature. The receiving devices must correspond to the appropriate LX or SX feature. Customers will see an approximate 35% improvement in FICON channel throughput. The maximum number of FICON Express4 features is 84 using three I/O cages.

FICON Express2 Channels

The z9 EC supports up to 336 FICON Express2 channels, each one operating at 1 or 2 Gb/sec auto-negotiated. The FICON Express2 features are available in long wavelength (LX) using 9 micron single mode fiber optic cables and short wavelength (SX) using 50 and 62.5 micron multimode fiber optic cables. Each FICON Express2 feature has four independent channels (ports) and each can be configured to carry native FICON traffic or Fibre Channel (SCSI) traffic. LX and SX cannot be intermixed on a single feature. The maximum number of FICON Express2 features is 84, using three I/O cages.

FICON Express Channels

The z9 EC also supports carrying forward to z9 EC FICON Express LX and SX channels from z990 (up to 120 channels) and z900 (up to 96 channels), each channel operating at 1 or 2 Gb/sec auto-negotiated. Each FICON Express feature has two independent channels (ports).

The Model S08 is limited to 64 features – any combination of FICON Express4, FICON Express2 and FICON Express LX and SX features.

The FICON Express4, FICON Express2 and FICON Express feature conforms to the Fibre Connection (FICON) architecture and the Fibre Channel (FC) architecture, providing connectivity between any combination of servers, directors, switches, and devices in a Storage Area Network (SAN). Each of the four independent channels is capable of 1 gigabit per second (Gb/sec), 2 Gb/sec, or 4 Gb/sec depending upon the capability of the attached switch or device. The link speed is auto-negotiated, point-to-point, and is transparent to users and applications. Not all switches and devices support 2 or 4 Gb/sec link data rates.

FICON Express4 and FICON Express2 Performance

Your enterprise may benefit from FICON Express4 and FICON Express2 with:

- *Increased data transfer rates (bandwidth)*
- *Improved performance*
- *Increased number of start I/Os*
- *Reduced backup windows*
- *Channel aggregation to help reduce infrastructure costs*

For more information about FICON, visit the IBM Redbook™ Web site at: http://www.redbooks.**ibm.com**/ search for SG24-5444. There are also various FICON I/O Connectivity at: www-03.**ibm.com**/systems/z/connectivity/.

Concurrent Update

The FICON Express4 and FICON Express2 SX and LX features may be added to an existing z9 EC concurrently. This concurrent update capability allows you to continue to run workloads through other channels while the new FICON Express4 and FICON Express2 features are being added. This applies to CHPID types FC and FCP.

Continued Support of Spanned Channels and Logical Partitions

The FICON Express4 and FICON Express2, FICON and FCP (CHPID types FC and FCP) channel types, can be defined as a spanned channel and can be shared among logical partitions within and across LCSSs.

Modes of Operation

There are two modes of operation supported by FICON Express4 and FICON Express2 SX and LX. These modes are configured on a channel-by-channel basis – each of the four channels can be configured in either of two supported modes.

- *Fibre Channel (CHPID type FC), which is native FICON or FICON Channel-To-Channel (server-to-server)*
- *Fibre Channel Protocol (CHPID type FCP), which supports attachment to SCSI devices via Fibre Channel switches or directors in z/VM, z/VSE, and Linux on System z9 environments*

Native FICON Channels

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

The FICON Express4, FICON Express2 and FICON Express channels support native FICON and FICON Channel-to-Channel (CTC) traffic for attachment to servers, disks, tapes, and printers that comply with the FICON architecture. Native FICON is supported by all of the z9 EC operating systems. Native FICON and FICON CTC are defined as CHPID type FC.

IBM has a full range of disk, SAN, tape, software, and services for the System z9. The IBM System z9 and IBM Storage 4 Gb FICON/FCP connectivity may help to support faster link speeds and shorter backup windows, enable channel and link consolidation to help simplify management and reduce the cost of the storage infrastructure.

IBM has a full range of Disk, SAN, Tape, Software, & Services for System z9

Disk	DS8000 – 4 Gbps FICON/FCP Planned 2Q06 DS6000 – 2 Gbps FICON/FCP
SAN	IBM SAN256B and SAN32B-2 FICON/FCP IBM SAN256M (Planned 2006) and SAN32M, and SAN140M 4 Gbps FICON/FCP Cisco MDS 9500 and 9216 4 Gbps FICON/FCP Planned 2006
Virtualization	IBM SVC 4 Gb FCP for Linux on System z Planned 2Q06 VTS 2 Gbps FICON/FCP TS7510 Virtualization Engine – 2 Gbps FCP for Linux on System z Planned 2Q06
Tape	IBM TS1120 4 Gbps FCP Tape Drive IBM TS1120 Tape Controller 4 Gbps FICON Planned 2Q06 IBM LTO Gen 3 – 4 Gbps FCP for Linux on System z Planned 2006 IBM 3494 and 3584 Tape Libraries IBM TS3310 Tape Library – 4 Gbps FCP for Linux on System z Planned 2Q06

FICON CTC function

Native FICON channels support channel-to-channel (CTC) on the z9 EC, z9 BC, z990, z890, z900 and z800. G5 and G6 servers can connect to a FICON CTC as well. This FICON CTC connectivity will increase bandwidth between z9 EC, z9 BC, z990, z890, z900, z800 G5, and G6 systems.

Because the FICON CTC function is included as part of the native FICON (FC) mode of operation, FICON CTC is not limited to intersystem connectivity (as is the case with ESCON), but will also support multiple device definitions.

For example, ESCON channels that are dedicated as CTC cannot communicate with any other device, whereas native FICON channels are not dedicated to CTC only. Native can support both device and CTC mode definition concurrently, allowing for greater connectivity flexibility.

FICON Support for Cascaded Directors

Native FICON (FC) channels support cascaded directors. This support is for a single hop configuration only. Two-director cascading requires a single vendor high integrity fabric. Directors must be from the same vendor since cascaded architecture implementations can be unique. This type of cascaded support is important for disaster recovery and business continuity solutions because it can help provide high availability, extended distance connectivity, and (particularly with the implementation of 2 Gb/sec Inter Switch Links), has the potential for fiber infrastructure cost savings by reducing the number of channels for interconnecting the 2 sites.

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

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

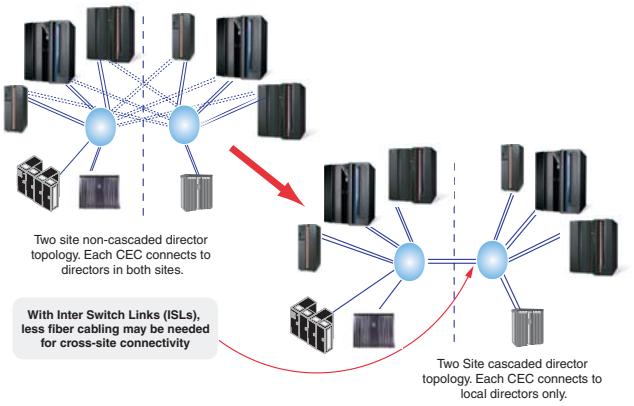

FCP Channels

System z9 supports FCP channels, switches and FCP/SCSI disks with full fabric connectivity under Linux on System z9 and z/VM 4.4 for Linux as a guest under z/VM, under z/VM 5.1, and under z/VSE 3.1 for system usage including install and IPL. Support for FCP devices means that z9 EC servers are capable of attaching to select FCP-attached SCSI devices and may access these devices from Linux on System z9 and z/VSE. This expanded attachability means that enterprises have more choices for new storage solutions, or may have the ability to use existing storage devices, thus leveraging existing investments and lowering total cost of ownership for their Linux implementations.

The same FICON features used for native FICON channels can be defined to be used for Fibre Channel Protocol (FCP) channels. FCP channels are defined as CHPID type FCP.

The 4 Gb/sec capability on the FICON Express4 channel means that 4 Gb/sec link data rates are available for FCP channels as well.

FCP performance*

An enhancement to the Queued Direct Input/Output (QDIO) architecture is designed to help increase performance for a FICON channel configured as CHPID type FCP (supporting communication with SCSI devices). Laboratory measurements for half-duplex data transfers yielded 31,500 starts per second for reading or writing 4k block sizes (up to twice the number of start I/Os per second previously measured). Measurements yielded 30,500 starts per second for reading or writing 8k block sizes.

Note: Performance projections are based on a controlled environment using standard IBM benchmarks; individual user performance may vary.

These results were achieved in a laboratory environment with an I/O driver program using one channel configured as CHPID type FCP with no other workload on the System z9 and do not represent actual field measurements.

The FCP performance enhancement applies to the FICON Express4 features (CHPID type FCP) on z9 EC and z9 BC and is transparent to operating systems.

FCP throughput improvements for half-duplex and full-duplex data transfers

The enhancement to the QDIO architecture is also designed to improve throughput. An FCP channel is now capable of achieving full data rate, 400 MegaBytes per second (MBps), for all reads and all writes (half-duplex data transfers). An FCP channel is capable of achieving 550 MBps for a mix of reads and writes (full-duplex data transfers). These measurements apply to large sequential data transfers.

The FCP throughput improvements were achieved in a laboratory environment with an I/O driver program using one channel configured as CHPID type FCP with no other workload on the System z9 and do not represent actual field measurements.

The FCP throughput improvements applies to the FICON Express4 features (CHPID type FCP) on z9 EC and z9 BC and is transparent to operating systems.

FCP performance metrics

For an FCP channel, I/O information is being made available using Linux on System z. This data, relating to FCP performance (latencies) and FCP channel usage, may help with the analysis of FCP channels.

Linux on System z can extract the hardware statistics (time spent in the fabric and time spent in the channel) to assist with the preparation of graphics to help analyze the performance and usage of FCP channels.

FCP performance metrics applies to the FICON Express4 and FICON Express2 features (CHPID type FCP) on z9 EC and z9 BC in the Linux on System z9 environment.

FCP Full fabric connectivity

FCP full fabric support means that any number of (single vendor) FCP directors/ switches can be placed between the server and FCP/ SCSI device, thereby allowing many "hops" through a Storage Area Network (SAN) for I/O connectivity. This support, along with 2 Gb/sec link data rates, is available together with IBM switch vendors IBM, McDATA, and Cisco. FCP full fabric connectivity enables multiple FCP switches/ directors on a fabric to share links and therefore provides improved utilization of inter-site connected resources and infrastructure. Further savings may be realized in the reduction of the number of fiber optic cabling and director ports.

When configured as CHPID type FCP, the z9 EC FICON Express4, FICON Express2, and FICON Express features support the industry standard interface for SAN management tools.

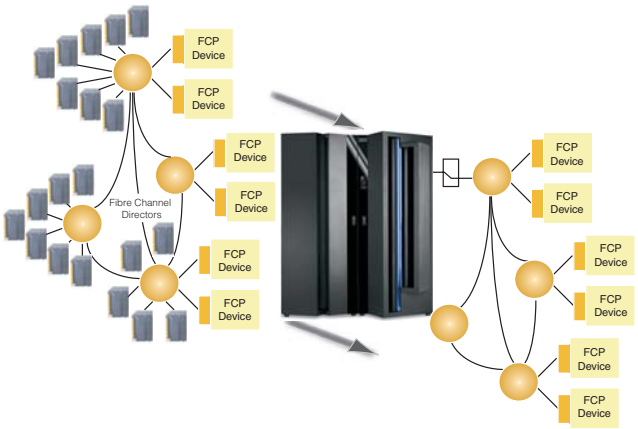

FICON Express enhancements for Storage Area Networks

N_Port ID Virtualization

N_Port ID Virtualization is designed to allow for sharing of a single physical FCP channel among multiple operating system images. Virtualization function is currently available for ESCON and FICON channels, and is now available for FCP channels. This new function offers improved FCP channel utilization due to fewer hardware requirements, and can reduce the complexity of physical FCP I/O connectivity.

Program Directed re-IPL

Program Directed re-IPL is designed to enable an operating system to determine how and from where it had been loaded. Further, Program Directed re-IPL may then request that it be reloaded again from the same load device using the same load parameters. In this way, Program Directed re-IPL allows a program running natively in a partition to trigger a re-IPL. This re-IPL is supported for both SCSI and ECKD™ devices. z/VM 5.3 provides support for guest exploitation.

FICON Link Incident Reporting

FICON Link Incident Reporting is designed to allow an operating system image (without operating intervention) to register for link incident reports, which can improve the ability to capture data for link error analysis. The information can be displayed and is saved in the system log.

Serviceability Enhancements

Requests Node Identification Data (RNID) is designed to facilitate the resolution of fiber optic cabling problems. You can now request RNID data for a device attached to a native FICON channel.

Connectivity for LANs – Open Systems Adapter-Express2

Open Systems Adapter-Express2 (OSA-Express2), for connectivity to Local Area Networks (LANs), supports 1000BASE-T Ethernet, Gigabit Ethernet (GbE) LX and SX, and 10 GbE LR. The OSA-Express2 features are hot-pluggable, support the Multiple Image Facility (MIF) sharing of channels across logical partitions, and can be defined as a spanned channel to be shared among logical partitions within and across LCSSs. The maximum combined number of OSA-Express2 and OSA-Express features supported per server is 24 on the z9 EC (up to 48 ports). OSA-Express2 and OSA-Express features can be carried forward on an upgrade from a z990 or z900 server. The OSA-Express Token-Ring feature is not supported on z9 EC servers.

The new OSA-Express2 1000BASE-T Ethernet feature and the OSA-Express2 Gigabit Ethernet (GbE) feature support the new IBM Communication Controller for Linux (CCL) on the System z9 platform and introduces the OSA-Express2 OSN (OSA for NCP) to support the Channel Data Link Control (CDLC) protocol, which provides direct access from the host operating system (such as z/OS and TPF) to the CCL.

With the large volume and complexity of today's network traffic, the System z9 offers systems programmers and network administrators the ability to more easily solve network problems. With the introduction of the OSA-Express Network Traffic Analyzer and QDIO Diagnostic Synchronization on the System z9, customers will have the ability to capture trace/trap data and forward it to z/OS 1.8 tools for easier problem determination and resolution.

An enhancement to the QDIO architecture, OSA-Express Network Traffic Analyzer is designed to allow trace records to be sent to the host operating system to improve the capability to capture data for both the system programmer and the network administrator.

This function is designed to allow the operating system to control the sniffer trace for the LAN and capture the records into host memory and storage (file systems), using existing host operating system tools to format, edit, and process the sniffer records.

OSA-Express Network Traffic Analyzer is exclusive to z9 EC and z9 BC, is applicable to the OSA-Express2 features when configured as CHPID type OSD (QDIO), and is supported by z/OS and z/OS.e.

Dynamic LAN idle for z/OS and z/OS.e

Dynamic LAN idle is designed to reduce latency and improve network performance by dynamically adjusting the inbound blocking algorithm. When enabled, the z/OS TCP/IP Stack is designed to adjust the inbound blocking algorithm to best match the application requirements.

For latency sensitive applications, the blocking algorithm is modified to be "latency sensitive." For streaming (throughput sensitive) applications, the blocking algorithm is adjusted to maximize throughput. The z/OS TCP/IP stack can dynamically detect the application requirements, making the necessary adjustments to the blocking algorithm. The monitoring of the application and the blocking algorithm adjustments are made in real-time, dynamically adjusting the application's LAN performance.

System administrators can authorize the z/OS TCP/IP stack to enable a dynamic setting, which was previously a static setting. The z/OS TCP/IP stack is able to help determine the best setting for the current running application, based on system configuration, inbound workload volume, CPU utilization, and traffic patterns.

Dynamic LAN idle is exclusive to z9 EC and z9 BC, is applicable to the OSA-Express2 features (CHPID type OSD), and is supported by z/OS and z/OS.e.

Dynamic LAN idle (performance improvement) for the OSA-Express2 features (CHPID type OSD) on z9 EC and z9 BC requires at a minimum:

- *z/OS or z/OS.e 1.8 with PTFs planned to be available third quarter 2007*
- *z/VM 5.1 for guest exploitation*

Link aggregation for z/VM in Layer 2 mode

z/VM Virtual Switch-controlled (VSWITCH-controlled) link aggregation (IEEE 802.3ad) allows you to dedicate an OSA-Express2 port to the z/VM operating system when the port is participating in an aggregated group when configured in Layer 2 mode. Link aggregation (trunking) is designed to allow you to combine multiple physical OSA-Express2 ports into a single logical link for increased throughput and for nondisruptive failover in the event that a port becomes unavailable.

- *Aggregated link viewed as one logical trunk and containing all of the Virtual LANs (VLANs) required by the LAN segment*
- *Load balance communications across several links in a trunk to prevent a single link from being overrun*
- *Link aggregation between a VSWITCH and the physical network switch*
- *Point-to-point connections*
- *Up to eight OSA-Express2 ports in one aggregated link*
- *Ability to dynamically add/remove OSA ports for "on demand" bandwidth*
- *Full-duplex mode (send and receive)*
- *Target links for aggregation must be of the same type (for example, Gigabit Ethernet to Gigabit Ethernet)*

The Open Systems Adapter/Support Facility (OSA/SF) will provide status information on an OSA port – its "shared" or "exclusive use" state. OSA/SF is an integrated component of z/VM.

Link aggregation is exclusive to z9 EC and z9 BC, is applicable to the OSA-Express2 features in Layer 2 mode when configured as CHPID type OSD (QDIO), and is supported by z/VM.

OSA Layer 3 Virtual MAC for z/OS and z/OS.e environments

To simplify the infrastructure and to facilitate load balancing when an LPAR is sharing the same OSA Media Access Control (MAC) address with another LPAR, each operating system instance can now have its own unique "logical" or "virtual" MAC (VMAC) address. All IP addresses associated with a TCP/IP stack are accessible using their own VMAC address, instead of sharing the MAC address of an OSA port. This applies to Layer 3 mode and to an OSA port shared among Logical Channel Subsystems. This support is designed to:

- *Improve IP workload balancing*
- *Dedicate a Layer 3 VMAC to a single TCP/IP stack*
- *Remove the dependency on Generic Routing Encapsulation (GRE) tunnels*
- *Improve outbound routing*
- *Simplify configuration setup*
- *Allow WebSphere Application Server content-based routing to work with z/OS in an IPv6 network*
- *Allow z/OS to use a "standard" interface ID for IPv6 addresses*
- *Remove the need for PRIROUTER/SECROUTER function in z/OS*

VMACs are currently available for Layer 2 mode in the z/VM and Linux on System z9 environments. OSA Layer 3 VMAC is exclusive to z9 EC and z9 BC, is applicable to the OSA-Express2 features when configured as CHPID type OSD (QDIO), and is supported by z/OS and z/OS.e.

OSA-Express2 Ethernet features on z9 EC

The OSA-Express2 features provide you with the function and scalability required to help satisfy the demands of your global businesses. With data rates of 10 or 100 Megabits per second (Mb/sec), 1 Gigabit per second (Gb/sec), and 10 Gb/sec, you can select the features that best suit your current and your future application requirements.

- *OSA-Express2 Gigabit Ethernet LX*
- *OSA-Express2 Gigabit Ethernet SX*
- *OSA-Express2 1000BASE-T Ethernet*
- *OSA-Express2 10 Gigabit Ethernet LR*

The OSA-Express2 Ethernet features support the following CHPID types:

CHPID Type	OSA-Express2 Features	Purpose / Traffic
OSC	1000BASE-T	TN3270E, non-SNA DFT, IPL CEC's and logical partitions Operating system console operations
OSD	1000BASE-T GbE 10 GbE	QDIO, TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
OSE	1000BASE-T	Non-QDIO, SNA/APPN®/HPR and/or TCP/IP
OSN	1000BASE-T GbE	OSA for NCP providing support for IBM Communication Controller for Linux (CCL)

The OSA-Express2 1000BASE-T Ethernet

IBM is expanding the family of OSA-Express2 features to include 1000BASE-T Ethernet, supporting a link data rate of 10, 100, or 1000 Mb/sec over a copper infrastructure. The OSA-Express2 1000BASE-T Ethernet feature continues to provide support for:

- *OSA-Integrated Console Controller (OSA-ICC) – TN3270E and non-SNA DFT 3270 emulation*
- *Queued Direct Input/Output (QDIO), CHPID type OSD, for TCP/IP traffic when using Layer 3, and protocol-independent packet forwarding when using Layer 2 (z/VM and Linux on System z9)*
- *Non-QDIO, CHPID type OSE, for SNA/APPN/HPR and/or TCP/IP traffic*
- *Checksum Offload (exclusive to QDIO mode, CHPID type OSD)*
- *Spanned channels and sharing among logical partitions*
- *Jumbo frames in QDIO mode (when operating at 1 Gb/sec)*
- *Auto-negotiation (the target device must also be set to auto-negotiate)*
- *Category 5 Unshielded Twisted Pair (UTP) cabling*

The OSA-Express2 1000BASE-T Ethernet feature supports the following modes of operation:

- *OSA-ICC (CHPID type OSC), for 3270 data streams*
- *QDIO (CHPID type OSD), for TCP/IP traffic when Layer 3, and for protocol-independent when Layer 2*
- *Non-QDIO (CHPID type OSE), for TCP/IP and/or SNA/APPN/HPR traffic*
- *OSA for NCP (CHPID type OSN), to provide channel connectivity between operating systems and CCL*

The OSA-Express2 1000BASE-T Ethernet feature is a dual-port feature occupying a single I/O slot and utilizes one CHPID per port; two CHPIDs per feature. Each port can be independently configured as CHPID type OSC, OSD, OSE, or OSN. The OSA-Express2 1000BASE-T Ethernet feature is offered on new builds while the OSA-Express 1000BASE-T Ethernet feature can be carried forward on an upgrade from a z990 or z900 server. Refer to the System z9 Operating Systems Support section for further information.

OSA-Express2 Gigabit Ethernet

The third generation of Gigabit Ethernet features is designed to support line speed – 1 Gb/sec in each direction or 2 Gb/sec full duplex and support the following functions:

- *QDIO architecture*
- *Layer 2*
- *Spanned channels*
- *SNMP*
- *IPv4 and IPv6*
- *640 TCP/IP stacks per CHPID*
- *Jumbo frames (8992 byte frame size)*
- *Large send, for TCP/IP traffic and CPU efficiency, offloading the TCP segmentation processing from the host TCP/IP stack*
- *Concurrent LIC update*
- *OSA-Express2 OSN (OSA for NCP)*

OSA-Express2 Gigabit Ethernet (GbE) operates in QDIO mode only and supports full duplex operation, and jumbo frames (8992 byte frame size).

The OSA-Express2 GbE features continue to be dual-port features occupying a single I/O slot and utilize one CHPID per port; two CHPIDs per feature. Each port can be independently configured as CHPID type OSD or OSN. The OSA-Express2 Gigabit Ethernet SX and LX features are offered on new builds while the OSA-Express Gigabit Ethernet features can be carried forward on an upgrade from a z990 or z900 server.

The OSA-Express2 GbE features are exclusive to the System z9, z990 and z890.

OSA-Express2 10 Gigabit Ethernet LR

The OSA-Express2 10 Gigabit Ethernet Long Reach (LR) can be use in an enterprise backbone, between campuses, to consolidate file servers and to connect server farms with System z9, z990, and z890 servers. The OSA-Express2 10 GbE LR supports:

- *Queued Direct Input/Output (QDIO)*
- *One port per feature*
- *A link data rate of 10 Gb/sec*
- *Full duplex mode*
- *Spanned channels*
- *SNMP*
- *IPv4 and IPv6*
- *Jumbo frames (8992 bytes frame size)*
- *Checksum Offload for IPv4 packets*
- *Layer 2 support*
- *Large send*
- *640 TCP/IP stacks*
- *Concurrent LIC update*
- *SC Duplex connector*
- *Single mode fiber (9 micron)*
- *An unrepeated distance of 10 km (6.2 miles)*

The 10 Gigabit Ethernet (10 GbE) feature does not support auto-negotiation to any other speed. The 10 GbE feature supports 64B/66B coding, whereas the GbE supports 8B/10B coding.

The OSA-Express2 10 Gigabits per second (Gb/sec) link data rate does not represent the actual throughput of the OSA-Express2 10 GbE feature. Actual throughput is dependent upon many factors, including traffic direction, the pattern of acknowledgment traffic, packet size, the application, TCP/IP, the network, disk subsystem, and the number of clients being served.

The OSA-Express2 10 GbE feature is exclusive to the System z9, z990 and z890.

IBM Communication Controller for Linux (CCL) on the System z9 platform

CCL is designed to help eliminate hardware dependencies, such as 3745/3746 Communication Controllers, ESCON channels, and Token-Ring LANs, by providing a software solution that allows the Network Control Program (NCP) to be run in Linux on System z9. CCL helps preserve mission critical SNA functions, such as SNI, and z/OS applications workloads which depend upon these functions, allowing you to collapse SNA inside a z9 EC while exploiting and leveraging IP.

The OSA-Express2 GbE and 1000BASE-T Ethernet features provide support for CCL with OSA-Express2 OSN (Open Systems Adapter for NCP). This support is designed to require no changes to operating systems (does require a PTF to support CHPID type OSN) and also allows TPF to exploit CCL.

If you continue to need SNA solutions that require NCP functions, you can now consider CCL as a migration strategy to replace your IBM Communication Controllers (374x).

OSA-Express2 OSN (OSA for NCP)

The OSA-Express2 OSN (OSA for NCP) can help to eliminate the requirement to have any form of external medium, and all related hardware, for communications between the host operating system and the CCL image. Traffic between the two images (operating system and CCL) is no longer required to flow on an external Local Area Network (LAN) or ESCON channel.

CHPID type OSN supports both SNA PU Type 5 and PU Type 2.1 channel connectivity.

Utilizing existing SNA support (multiple transmission groups), OSA-Express2 OSN support permits multiple connections between the same CCL image and the same host operating system image. It also allows multiple CCL images to communicate with multiple operating system images, supporting up to 180 connections (3745/3746 unit addresses) per CHPID type OSN. CHPID type OSN can also span LCSSs. The CCL image connects to the OSA-Express2 feature using QDIO architecture and uses the Linux QDIO (qeth) support updated to support OSN device types.

OSA-Express2 OSN (OSA for NCP) support is exclusive to System z9, to the OSA-Express2 GbE SX, GbE LX, and 1000BASE-T Ethernet features, and requires the port to be configured as CHPID type OSN, which can be configured on a port-by-port basis.

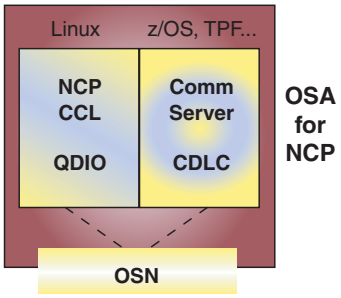

OSA-Express2 concurrent LIC update – an availability enhancement

The OSA-Express2 features have increased memory in comparison to the OSA-Express features and are designed to be able to facilitate concurrent application of Licensed Internal Code (LIC) updates, allowing the application of LIC updates without requiring a configuration off/on of the features. This can help minimize the disruption to network traffic during the update.

OSA-Express2 concurrent LIC update applies to CHPID type OSD and is exclusive to the System z9, z990 and z890.

OSA Integrated Console Controller

The Open Systems Adapter Integrated Console Controller function (OSA-ICC), which is exclusive to the System z9, z990 and z890 servers since it is based on the OSA-Express2 and OSA-Express 1000BASE-T Ethernet features, supports the attachment of non-SNA 3270 terminals for operator console applications. Now, 3270 emulation for console session connections (TN3270E [RFC 2355] or non-SNA DFT 3270 emulation) is integrated in the System z platforms which can help eliminate the requirement for external console controllers (2074, 3174), helping to reduce cost and complexity.

The OSA-ICC can be individually configured on a port-by-port basis. The OSA-ICC is enabled using CHPID type OSC.

The OSA-ICC supports up to 120 client console sessions per port either locally or remotely.

Support for this function is provided with z/OS, z/VM, z/VSE, and TPF.

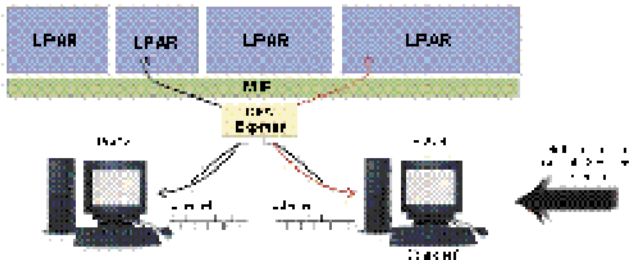

OSA Enhancements

Remove L2/L3 LPAR-to-LPAR Restriction

OSA port sharing between virtual switches can now communicate whether the transport mode is the same (Layer 2 to Layer 2) or different (Layer 2 to Layer 3). This enhancement is designed to allow seamless mixing of Layer 2 and Layer 3 traffic, helping to reduce the total cost of networking. Previously, Layer 2 and Layer 3 TCP/IP connections through the same OSA port (CHPID) were unable to communicate with each other LPAR-to-LPAR using the Multiple Image Facility (MIF).

This enhancement is designed to facilitate a migration from Layer 3 to Layer 2 and to continue to allow LAN administrators to configure and manage their mainframe network topology using the same techniques as their non-mainframe topology.

OSA/SF Virtual MAC and VLAN id Display Capability

The Open Systems Adapter/Support Facility (OSA/SF) now has the capability to support virtual Medium Access Control (MAC) and Virtual Local Area Network (VLAN) identifications (IDs) associated with OSA-Express2 and OSA-Express features configured as a Layer 2 interface. This information will now be displayed as a part of an OAT entry. This information is independent of IPv4 and IPv6 formats. There can be multiple Layer 2 VLAN Ids associated to a single unit address. One group MAC can be associated to multiple unit addresses.

For additional information, view IBM Redbook *System z9 and zSeries Connectivity Handbook* (SG24-5444) at: www.redbooks.**ibm.com**/.

HiperSockets

The HiperSockets function, also known as internal Queued Direct Input/Output (iDQIO) or internal QDIO, is an integrated function of the z9 EC server that provides users with attachments to up to sixteen high-speed "virtual" Local Area Networks (LANs) with minimal system and network overhead.

HiperSockets eliminates the need to utilize I/O subsystem operations and the need to traverse an external network connection to communicate between logical partitions in the same z9 EC server. HiperSockets offers significant value in server consolidation connecting many virtual servers, and can be used instead of certain XCF link configurations in a Parallel Sysplex cluster.

HiperSockets can be customized to accommodate varying traffic sizes. Since HiperSockets does not use an external network, it can free up system and network resources, eliminating attachment costs while improving availability, performance and security.

For additional information, consult IBM Redbook, *zSeries HiperSockets* (SG24-6816) at: www.redbooks.**ibm.com**/.

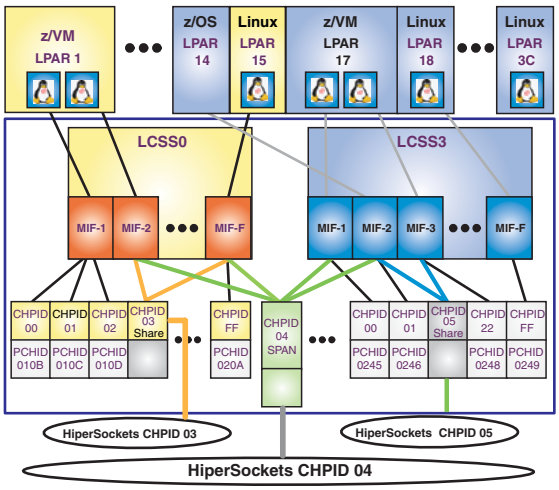

Cryptography

In the on demand era, security is a requirement. System z9 technology continues to address security, enhancing the feature and providing additional functions.

The main focus in cryptography continues to be very high, scalable performance for SSL algorithms, and second, to provide security-rich, symmetric performance for financial and banking applications using PIN/POS type encryption. As in the past, the System z9 platform is designed to deliver seamless integration of Integrated Cryptographic Service Facility (ICSF). Use of ICSF is designed to enable applications to work without change, regardless of how and where the cryptographic functions are implemented, and is also designed to enable the cryptography elements to be load balanced across the hardware resources.

CP Assist for Cryptographic Function (CPACF)

CP Assist for Cryptographic Function (CPACF) supporting clear key encryption, is activated using a no charge enablement feature and offers the following on every Processor Unit (PU) identified as a Central Processor (CP) or Integrated Facility for Linux (IFL).

With enhanced scalability and performance, the z9 EC is designed to provide a set of symmetric cryptographic functions, synchronously executed, which may enhance the performance of the encrypt/decrypt function of SSL, VPN and data storing applications which do not require FIPS 140-2 Level 4 security. The on-processor crypto functions are designed to run at z9 EC processor speed, an order of magnitude faster than the CMOS Crypto Coprocessor in the zSeries 900. As these crypto functions are implemented in each and every CP and IFL, the affinity problem of pre-z990 systems (which had only two CMOS Crypto Coprocessors) is expected to be virtually eliminated. To conform with US Export and Import Regulations of other countries, an SE panel is provided for proper enable/disable of 'strong' cryptographic functions.

- *Advanced Encryption Standard (AES) for 128-bit keys (new)*
- *Data Encryption Standard (DES)*
- *Triple Data Encryption Standard (TDES)*
- *Pseudo Random Number Generation (PRNG) (new)*
- *SHA-1*
- *SHA-256 (new)*

Performance is designed to scale with PU performance improvements. SHA-1 and SHA-256 are shipped enabled on all servers and do not require the enablement feature. Support for CPACF is also available via the Integrated Cryptographic Service Facility (ICSF). ICSF is a component of z/OS, and is designed to transparently use the available cryptographic functions, whether CPACF or Crypto Express2 to balance the workload and satisfy the bandwidth requirements of the applications.

The enhancements to CPACF are exclusive to System z9 and are supported by z/OS, z/VM, z/VSE, and Linux on System z9. Refer to the System z9 Operating Systems Support section for further information.

A third generation Cryptographic feature – Crypto Express2

The Crypto Express2 feature is designed for On Demand Business in a security-rich environment. Crypto Express2 provides the functions of PCICA and PCIXCC in a single feature that is expected to provide improved secure key and system throughput. Like its predecessors, the Crypto Express2 feature has been designed to satisfy the security requirements of an enterprise server.

The Crypto Express2 feature, with two PCI-X adapters, is configurable and can be defined for secure key encrypted transactions (Coprocessor – the default) or SSL acceleration (Accelerator). The PCIXCC, PCICC, and PCICA features are not supported on z9 EC.

The Integrated Cryptographic Service Facility (ICSF), a component of z/OS and z/OS.e, is designed to transparently use the available cryptographic functions, the CP Assist for Cryptographic Function (CPACF) as well as the Crypto Express2 features to balance the workload and satisfy the requirements of the applications.

The Crypto Express2 feature is designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 Certification. A performance benefit is expected with multitasking applications. A performance benefit may not be realized with single-threaded applications, which can utilize only one of the two coprocessors.

The Crypto Express2 feature supports the following:

- *Consolidation and simplification via a single crypto coprocessor feature on System z9, z990, and z890*
- *Compute-intensive public key cryptographic functions designed to help reduce CP utilization and increase system throughput*
- *Card Validation Value (CVV) generation and verification services for 19-digit Personal Account Number (PANs)*
- *Enabling use of less than 512-bit keys for clear key RSA operations*
- *2048-bit key RSA management capability*
- *Functions previously supported by the PCICA and PCIXCC features offered on z990 and z890 including:*
  - *Compute-intensive public key cryptographic functions to help reduce CP usage and increase system throughput*
  - *Hardware acceleration for Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols to support secure On Demand Business applications and transactions*
  - *SSL performance equivalent to the PCICA feature*
  - *The functional enhancements announced in April 2004, namely: PKE MRP support, PKD zero pad support, TDES DUKPT, and EMV2000*
  - *User Defined Extension (UDX) Service Offering – programmable to deploy standard functions and algorithms*
- *Up to eight features per server*
  - *With Crypto Express2, the z9 EC, z9 BC, z990, and z890 can have up to sixteen secure key coprocessors in comparison to the four coprocessors with the PCIXCC features. This is expected to translate into increased secure key and system throughput.*
  - *With Crypto Express2, the z9 EC, z9 BC, z990, and z890 servers can utilize up to sixteen cryptographic coprocessors for clear key SSL acceleration in comparison to twelve accelerators with the PCICA features.*
  - *A mixture of both secure and clear key applications can run on the same Crypto Express2 feature*
  - *Based on the increased throughput, the ability to consolidate both secure key and clear key crypto workloads and I/O slots on the same feature*

All logical partitions in all Logical Channel SubSystems (LCSSs) have access to the Crypto Express2 feature, up to 60 LPARs per feature. The Crypto Express2 feature occupies a card slot but does not use CHPIDs.

The Crypto Express2 feature is exclusive to System z9, z990, and z890, requires the October 2004 level of Licensed Internal Code, and is supported by z/OS, z/OS.e, z/VM, z/VSE, and Linux on System z9. z/VSE and Linux on System z9 offer support for clear key SSL transactions only.

Configurable Crypto Express2 feature

The Crypto Express2 feature has two PCI-X adapters. Each of the PCI-X adapters can be defined as either a Coprocessor or an Accelerator.

- *Crypto Express2 Coprocessor – for secure key encrypted transactions (default)*
  - *Designed to support security-rich cryptographic functions, use of secure encrypted key values, and User Defined Extensions (UDX)*

  - *Designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification*
- *Crypto Express2 Accelerator – for Secure Sockets Layer (SSL) acceleration*
  - *Designed to support clear key RSA operations*
  - *Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol*

When a System z9 platform with four CPs and both PCI-X adapters on a Crypto Express2 feature are configured as accelerators, the Crypto Express2 feature is designed to perform up to 6,000 SSL handshakes per second. This represents, approximately, a 3X performance improvement compared to z990 when using either a PCI Cryptographic Accelerator (PCICA) feature, with two PCI accelerators per feature, or the current Crypto Express2 feature, with two PCI-X adapters per feature. The SSL performance was achieved using z/OS 1.7 with Cryptographic Support for z/OS 1.6/1.7 Web deliverable, and ICSF FMID HCR7730.

Since the performance enhancements are implemented in Licensed Internal Code, current Crypto Express2 features carried forward from the z990 to the System z9 platform may take advantage of increased SSL performance and the new configuration capability.

These measurements are examples of the maximum handshakes per second achieved in a laboratory environment with no other processing occurring and do not represent actual field measurements. Details are available upon request.

The configurable Crypto Express2 feature is exclusive to the System z9 and is supported by z/OS, z/VM, z/VSE, and Linux on System z9. z/VSE and Linux on System z9 offer support for clear key SSL transactions only. z/VM 5.1 and later supports clear and secure key operations.

TKE 5.1 workstation for security and convenience

The Trusted Key Entry (TKE) capability is an optional feature of System z9 technology that provides a basic security key management system. The key management system provides authorized persons a method of security key identification, exchange, separation, update, and management.

The Trusted Key Entry (TKE) workstation feature is a combination of hardware and software, network-connected to the server, and designed to provide a security-rich, flexible method for master and operational key entry as well as local and remote management of the cryptographic coprocessor features. This optional feature provides key identification, exchange, separation, update, backup, as well as security administration. The TKE workstation is supplied with one Ethernet port capable of operating at 10, 100, or 1000 Mb/sec. Included is one mouse, one keyboard, a selectable display (small or large flat panel), and a DVD RAM drive to install Licensed Internal Code (LIC). It has a serial port for attaching a Smart Card Reader.

The Trusted Key Entry (TKE) workstation with Ethernet and the 5.1 level of Licensed Internal Code are optional features on z9 EC. Up to three TKE workstations can be used per z9 EC server. The TKE 5.1 LIC is loaded on the TKE workstation prior to shipment. For other than new shipments TKE 5.1 LIC is shipped on CD-ROM. The TKE workstation offers secure local and remote key management providing authorized persons a method of operational and master key entry, identification, exchange, separation, and update. The TKE workstation ships with a 4764 Model 001 PCI-X Cryptographic Coprocessor (4764-001).

The optional TKE features are:

- *TKE 5.1 LIC and TKE workstation with Ethernet*
- *TKE Smart Card Reader*
- *TKE additional smart cards*

TKE 3.x workstations can be used to control z900, z800, and prior servers. TKE 4.x workstations can be used to control z990, z890, and prior servers. TKE 5.x workstations must be used to control z9 EC. TKE 5.x workstations can also be used to control z9 BC, z990, z890, z900, and z800 servers.

The benefits of TKE5.1 LIC include service mode support and usability enhancements including a service user to improve access to operations and a new task layout display to maintain TKE console consistency. TKE 5.1 is a no-charge enablement feature which is loaded prior to shipment when aTKE workstation is ordered.

TKE Smart Card Reader

Support for an optional Smart Card Reader attached to the TKE 5.1 workstation allows for the use of smart cards that contain an embedded microprocessor and associated memory for data storage. Access to and the use of confidential data on the smart cards is protected by a user-defined Personal Identification Number (PIN).

The TKE 5.1 support does not remove any of the mechanisms available in the current TKE LIC with the Smart Card Reader support. It is still possible to store key parts on diskettes or paper, or optionally on smart cards, or to use a TKE authority key stored on a diskette, or optionally on a smart card, and to log on to the 4764-001 using a passphrase, or optionally a logon key pair. One feature includes two Smart Card Readers, and two cables to connect to the TKE 5.1 workstation, and 20 smart cards.

The Smart Card Reader, which can be attached to a TKE workstation with the 5.1 level of LIC, is available on z9 EC, z9 BC, z990, z890, z900, and z800.

TKE additional Smart Cards

The TKE additional smart cards are Java technology-based smart cards which provide a highly efficient cryptographic and data management application built in to read-only memory for secure storage of keys, certificates, passwords, applications, and data. The TKE blank smart cards are compliant with FIPS140-2 Level 2.

Cryptographic support for 19-digit PANs

Crypto Express2 feature offers Card Validation Value (CVV) generation and verification services for 19-digit PANs. Industry practices for use of CVV are moving to base CVV computations on a 19-digit PAN instead of the 13-digit and 16-digit PANs currently in use and supported by ICSF. ICSF and Crypto Express2 support use of the 19-digit PAN in the CVV generation and verification services (CSNBCSG and CSNBCSV, respectively).

Support of CVV generation and verification services for 19-digit PANs, an anti-fraud security feature, is supported by the Crypto Express2 feature on the z9 EC, z9 BC, z990, and z890 servers and by z/OS and z/OS.e.

Enabling use of less than 512-bit keys for clear key RSA operations

The Crypto Express2 feature supports applications that require clear key RSA operations using keys less than 512-bits, including ICSF Callable services and their corresponding verbs: Digital Signature Verify (CSNDDSV), Public Key Encrypt (CSNDPKE), and Public Key Decrypt (CSNDPKD). All other ICSF Callable services that require a Crypto Express2 feature continue to require keys of more than 511-bits.

Enabling the lower limit for clear key RSA operations may allow the migration of some additional cryptographic applications to z9 EC, z9 BC, z990, and z890 servers without requiring the applications to be rewritten.

Support of applications that require clear key RSA operations using keys less than 512 bits applies to the Crypto Express2 feature, is exclusive to z9 EC, z9 BC, z990, and z890, and is supported by z/OS, z/OS.e, and z/VM. Refer to the System z9 Operating Systems Support section for further information.

Further Cryptographic Enhancements

Remote Loading of Initial ATM Keys

Typically, a new ATM has none of the financial institutions keys installed. Remote Key Loading refers to the process of loading Data Encryption Standard (DES) keys to Automated Teller Machines (ATMs) from a central administrative site without the need for personnel to visit each machine to manually load DES keys. This has been done by manually loading each of the two clear text key parts individually and separately into ATMs. Manual entry of keys is one of the most error-prone and labor-intensive activities that occur during an installation, making it expensive for the banks and financial institutions.

Remote Key Loading Benefits

- *Provides a mechanism to load initial ATM keys without the need to send technical staff to ATMs.*
- *Reduces downtime due to key entry errors.*
- *Reduces service call and key management costs.*
- *Improves the ability to manage ATM conversions and upgrades.*

Integrated Cryptographic Service Facility (ICSF), together with Crypto Express2, support the basic mechanisms in Remote Key Loading. The implementation offers a secure bridge between the highly secure Common Cryptographic Architecture (CCA) environment and the various formats and encryption schemes offered by the ATM vendors. The following are new ICSF services are offered for Remote Key loading:

- *Trusted Block Create (CSNDTBC)*
  This callable service is used to create a trusted block containing a public key and some processing rules. The rules define the ways and formats in which keys are generated and exported.
- *Remote Key Export (CSNDRKX)*
  This callable service uses the trusted block to generate or export DES keys for local use and for distribution to an ATM or other remote device.

Refer to *Application Programmers Guide*, SA22-7522, for additional details.

Improved Key Exchange With Non-CCA Cryptographic Systems

IBM Common Cryptographic Architecture (CCA) employs Control Vectors to control usage of cryptographic keys. Non-CCA systems use other mechanisms, or may use keys that have no associated control information. This enhancement provides the ability to exchange keys between CCA systems, and systems that do not use Control Vectors. Additionally, it allows the CCA system owner to define permitted types of key import and export which can help to prevent uncontrolled key exchange that can open the system to an increased threat of attack.

These enhancements are exclusive to System z9 and supported by z/OS operating system.

ISO 16609 CBC Mode T-DES Enhancement

ISO 16609 CBC Mode T-DES MAC supports the requirements for Message Authentication, using symmetric techniques. The Integrated Cryptographic Service Facility

(ICSF) will use the following callable services to access the ISO 16609 CBC Mode T-DES MAC enhancement in the Cryptographic coprocessor:

- *MAC Generate (CSNBMGN)*
- *MAC Verify (CSNVMVR)*
- *Digital Signature Verify (CSNDDSV)*

ISO 16609 CBC mode T-DES MAC is accessible through ICSF function calls made in the Cryptographic Adapter Segment 3 Common Cryptographic Architecture (CCA) code. This enhancement is exclusive to System z9 and supported by z/OS 1.6 or higher.

IBM Encryption Facility for z/OS

The IBM Encryption Facility for z/OS 1.1 can apply mainframe encryption services, which have been helping to protect ATMs for close to 15 years, to the creation of encrypted tapes. Customers can use z/OS centralized key management to provide a highly secure exchange of encryption keys, so that if important content intended for a trusted business partner were to fall into the wrong hands, it is designed so that the tape can only be decrypted with the use of the private encryption key of the partner.

With the Encryption Facility you can encrypt data to tape on your z/OS system for transport to your partners or customers, even if they don't have a z/OS system. Partners with z/OS can use this facility on their system to decrypt the tape. For partners without z/OS, a Java technology-based decrypting and encrypting program is also be available. This separately licensed program can be used to decrypt the data encrypted by the Encryption Services feature running on z/OS, and to encrypt data for transfer back to your z/OS system, which can then be decrypted using the Encryption Services feature running on z/OS.

Encryption Facility for z/OS consists of two priced optional features:

- *The Encryption Services feature supports encrypting and decrypting certain file formats on z/OS. This can allow you to transfer them to remote sites within your enterprise, transfer them to partners and vendors, and archive them. This feature supports hardware-accelerated compression before encryption.*
- *The DFSMSdss™ Encryption feature enables the encryption of DFSMSdss dump data sets. This feature supports hardware accelerated compression before encryption.*
- *Both features can use the state-of-the-art encryption and centralized key management capabilities provided by functions of z/OS and features of System z servers to help secure data stored to tape and other removable media.*

Encryption Services feature

The Encryption Services feature can allow you to encrypt data written to tape and other removable media. This can help you share sensitive information across platforms with partners, vendors, and customers. You can also use the Encryption Services feature to encrypt certain files for archival. This feature can use the z/OS key management and access authentication capabilities provided within the Integrated Cryptographic Services Facility (ICSF) and the hardware compression and the hardware cryptographic capabilities of System z9 and zSeries servers.

The Encryption Services feature supports data encryption using TDES triple-length keys or 128-bit AES keys. RSA public/private keys can be specified to wrap and unwrap the AES and TDES data keys used to encrypt the file. The wrapped keys will be stored in a file header. With this technique, many files can be generated using different encryption keys, and each is expected to be able to be read even after years of archived storage. The Encryption Services feature also supports using a password key derivation scheme. The Encryption Services feature supports

inputs from physical sequential input files, from members of partitioned data sets (PDS) and partitioned data set extended (PDSE) data sets, and from files stored in z/OS UNIX System Services file systems. It can optionally compress input files before encrypting them and writing the output files. Also, it can use the large block interface for output files written to tape, to help optimize performance and media space.

DFSMSdss Encryption feature

The DFSMSdss Encryption feature can allow you to encrypt DFSMSdss dump data sets written to tape and disk. This feature is designed to use the z/OS key management and access authentication capabilities and the hardware cryptographic and compression capabilities of System z servers.

DFSMSdss Encryption supports encryption of data using TDES triple length keys or 128-bit AES keys. Like the Encryption Services feature, this feature supports the use of RSA public/private keys to wrap and unwrap the AES and TDES data keys used to encrypt files as well as AES and TDES key generation using a specified password. You can also specify that DFSMSdss is to compress data before encrypting it. The DFSMSdss Encryption feature includes two functions, one to encrypt data while processing DUMP commands, and the other to decrypt it while processing RESTORE commands.

Encryption Facility for z/OS Client

The Encryption Facility for z/OS Client, a separately licensed program (which is offered "as is," with no warranty), is written in Java and can be used on multiple platforms. It is designed to enable the exchange of encrypted data between z/OS systems that have the Encryption Facility installed and systems running on other platforms that provide the needed supported functions.

The Encryption Facility for z/OS Client is designed to:

- *Decrypt data that was created on a z/OS system using the Encryption Facility*
- *Encrypt data to be sent to a z/OS system, where the file will be decrypted using the Encryption Facility*

Note: Data that is to be processed using the Encryption Facility Client cannot be created using compression.

The value of mainframe encryption services

IBM mainframe encryption services are based on hardware and software integration—encryption and compression technologies found in the mainframe servers, and the centralized key management capabilities found in the z/OS operating system.

Mainframe encryption hardware provides two key functions — provides for the acceleration of encryption compared to software-based encryption, and with the appropriate features, it also provides Secure Key services. High performance encryption acceleration is provided in the CP Assist for Cryptographic Function which is built into the central processors of the System z9, z990, and z890 servers. On the System z9 EC server, new enhancements include support for the SHA-256 hashing algorithm and the 128-bit Advanced Encryption Standard (AES-128) which is rapidly becoming the de facto encryption standard.

For additional information on the Encryption Facility for z/OS view: www-03.**ibm.com**/servers/eserver/zseries/zos/encryption_facility/.

For additional information on security, visit: www-03.**ibm.com**/systems/z/security/.

On Demand Capabilities

Capacity on Demand

It may sound revolutionary, but it's really quite simple. In the highly unpredictable world of On Demand Business, you should have access to the resources you need, when you need them. This is the basic principle underlying Capacity on Demand for System z9.

The System z9 server provides for concurrent upgrade capability of Central Processors (CPs), Internal Coupling Facilities (ICFs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), and the IBM System z9 Integrated Information Processors (zIIPs), through Customer Initiated Upgrade with the following offerings:

- *Customer Initiated Upgrade (CIU) – The CIU feature enables you to order permanent capacity upgrades rapidly and download them without disrupting applications already running on the server. When extra processing power becomes necessary, an administrator simply navigates to Resource Link to order the upgrade and uses the Remote Service Facility on the Hardware Management Console to download and activate pre installed inactive processors or memory for an additional charge.*
- *On/Off Capacity on Demand (On/Off CoD) – Available through CIU, use On/Off CoD for temporary increases in processor capacity. With temporary processor capacity, you can help your business manage both predictable and unpredictable surges in capacity demands. You can activate and deactivate quickly and efficiently as the demands on your organization dictate – obtain additional capacity that you need, when you need it, and the server will keep track of your usage. On/Off CoD provides a cost-effective strategy for handling seasonal or period-end fluctuations in activity and may enable you to deploy pilot applications without investing in new hardware. The z9 provides the ability to store up to 100 On/Off CoD LIC records on the Support Element, at any given time, giving greater flexibility to quickly enable needed temporary capacity. Although there is a fee for On/Off CoD, a one time no-charge 24-hour test is available.*
- *Capacity Backup (CBU) – CBU is used to add temporary processing capacity to a backup server in the event of an unforeseen loss of server capability because of an emergency. With CBU, you can divert entire workloads to backup servers for up to 90 days. Although free tests are available for CBU, there is an up front fee and an emergency-use fee, which includes testing; software fees are not affected.*
- *Automatic CBU Enablement for Geographically Dispersed Parallel Sysplex® (GDPS®) – The intent of GDPS CBU is to enable automatic management of the reserved PUs provided by the CBU feature in the event of a server or a site failure. When a site failure or planned disaster test is detected, GDPS concurrently adds CPs to the servers in the take-over site to restore processing power for mission-critical production workloads. GDPS automation is intended to:*
  - *Perform the analysis required to determine the scope of the failure to minimize operator intervention and the potential for errors.*
  - *Automate authentication and activation of the reserved CPs.*
  - *Automatically restart the critical applications after reserved CP activation.*
  - *Reduce the outage time to restart the critical workloads from several hours to minutes.*

Extended Staging for CIU Express and On/Off CoD

All CIU Express and On/Off CoD orders may be staged for greater than 30 days. In fact, the orders may be staged for an extended period of time, unless one of the following conditions occurs:

- *Order is canceled by customer*
- *Machine is no longer under warranty or IBM Maintenance Service Agreement*
- *Permanent PU and/or memory configurations are changed outside of CIU process*

For additional information, contact your System z Sales Representative. For education and additional information, there is a System z Capacity on Demand User's Guide. Go to www.**ibm.com**/servers/resourcelink and select Customer Initiated Upgrade on the navigation bar.

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 System z9 server I/O is reduced by the fact that all I/O cards plugging into the System z9 I/O cage are hot-pluggable. But I/O cages cannot be installed concurrently to a z9 EC server. This means that the only I/O to be conditioned is the I/O cage itself. The question of whether or not to concurrently condition a cage is a very important consideration, especially with the rapid change in the IT environment as well as the technology.

The Plan Ahead process can easily identify the customer configuration that is required to meet future needs. The result of concurrent conditioning is the capability to enable a flexible IT infrastructure that can accommodate unpredictable growth in a low risk, nondisruptive way. Depending on the required Concurrent Conditioning, there should be minimal cost associated with dormant z9 EC capacity. This creates an attractive option for businesses to quickly respond to changing environments, bringing new applications online or growing existing applications without disrupting users.

Enhanced Book Availability

With proper planning, z9 EC is designed to allow a single book, in a multi-book server, to be nondisruptively removed from the server and re-installed during an upgrade or repair action. To minimize the effect on current workloads and applications, you should ensure that you have sufficient inactive physical resources on the remaining books to complete a book removal.

Customers configuring for maximum availability are recommended to purchase models with one additional book and configure a maximum of 8 processor units (PU) on a Model S18, a maximum of 18 PUs on a Model S28, or a maximum of 29 PUs on a Model S38. To ensure you have the appropriate level of memory, you may want to consider the selection of the Flexible Memory Option features to provide additional resources when completing Enhanced Book Availability action or when considering plan ahead options for the future. Enhanced Book Availability may also provide benefits, should you choose not to configure for maximum availability. In these cases, you should have sufficient inactive resources on the remaining books to contain critical workloads while completing a book replacement. Contact your IBM representative to help you determine and plan the proper configuration to support your workloads when using nondisruptive book maintenance.

Enhanced Book Availability is an extension of the support for Concurrent Book Add (CBA) delivered on z990. CBA makes it possible to concurrently upgrade a server by integrating a second, third, or fourth book into the server without necessarily affecting application processing.

The following scenarios prior to the availability of EBA would require a disruptive customer outage. With EBA these upgrade and repair procedures can be performed concurrently without interfering with customer operations.

Concurrent Physical Memory Upgrade

Allows one or more physical memory cards on a single book to be added, or an existing card to be upgraded increasing the amount of physical memory in the system.

Concurrent Physical Memory Replacement

Allows one or more defective memory cards on a single book to be replaced concurrent with the operation of the system.

Concurrent Defective Book Replacement

Allows the concurrent repair of a defective book when that book is operating degraded due to errors such as multiple defective processors.

Enhanced Book Availability is exclusive to z9 EC.

Flexible Memory Option

To plan for Enhanced Book Availability on z9 EC, Flexible Memory Option features are offered in 16 GB increments from 32 GB to 384 GB. Flexible Memory Option features provide additional physical memory to support activation of the actual purchased memory increment in the event of a single book failure or to be available during an Enhanced Book Availability action. Flexible Memory Option features are exclusive to z9 EC, and available only on S18, S28, S38 and S54 models.

Redundant I/O Interconnect

z9 EC with Redundant I/O Interconnect is designed to allow you to replace a book or respond to a book failure and retain connectivity to resources. In the event of a failure or customer initiated action such as the replacement of an MBA fanout card or book, the z9 EC is designed to provide access to your I/O devices through another STI to the affected I/O domains. This is exclusive to System z9.

Enhanced Driver Maintenance

One of the greatest contributors to downtime during planned outages is Licensed Internal Code (LIC) updates. When properly configured, z9 EC is designed to permit select planned LIC updates to the z9 EC at specified Driver Sync Points. Contact your IBM representative to help you determine the proper configuration and planning to utilize this function. Enhanced Driver Maintenance is exclusive to System z9.

Dynamic Oscillator Switchover

z9 EC has two oscillator cards, a primary and a backup. For most cases of failures, should a failure occur on the primary oscillator card, the backup can detect it, switch over, and provide the clock signal to the system transparently, with no system outage. Previously, in the event of a failure of the active oscillator, a system outage would occur, the subsequent system Power On Reset (POR) would select the backup, and the system would resume operation. Dynamic Oscillator Switchover is exclusive to System z9.

Transparent Sparing

z9 EC offers 2 PUs reserved as spares per server. In the case of processor failure, these spares are used for transparent sparing.

Enhanced Dynamic Memory Sparing

The z9 EC provides a robust recovery design with 16 times more chips available for sparing. This may eliminate the need to replace a memory card due to a DRAM failure.

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

Concurrent Maintenance

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

Upgrade for Coupling Links: z9 EC has concurrent maintenance for the ISC-3 daughter card. Also, Coupling Links can be added concurrently. This eliminates a need for scheduled downtime in the demanding sysplex environment.

Cryptographic feature: The Crypto Express2 feature plugs in the I/O cage and can be added or replaced concurrently with system operation.

Redundant Cage Controllers: The Power and Service Control Network features redundant Cage Controllers for Logic and Power control. This design enables nondisruptive service to the controllers and virtually eliminates customer scheduled outage.

Auto-Switchover for Support Element (SE): The z9 EC has two Support Elements. In the event of failure on the Primary SE, the switchover to the backup is handled automatically. There is no need for any intervention by the customer or Service Representative.

Concurrent Capacity Backup Downgrade (CBU Undo)

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

Concurrent Memory Upgrade

This function allows adding memory concurrently, up to the maximum amount physically installed. In addition, the new Enhanced Book Availability function also enables a memory upgrade to an installed z9 EC book in a multi-book server.

The System z9 offers new tools for power planning and monitoring.

Power monitoring

System z9 now provides a capability designed to monitor power consumption and temperature of the system. The System Activity Display on the Hardware Management Console will display the current total power consumption in watts and Btu/hour and will also display the input temperature.

zPower estimation tool

System z9 now provides a tool on IBM Resource Link which provides the user an estimate as to the anticipated power consumption of a particular machine model and its associated configuration. A user will input the machine model, memory, and I/O configuration and the tool will output an estimate of the power requirements needed for this system.

Parallel Sysplex Cluster Technology

Parallel Sysplex clustering is designed to bring the power of parallel processing to business-critical System z9 applications. A Parallel Sysplex cluster consists of up to 32 z/OS images coupled to one or more Coupling Facilities (CFs or ICFs) using high-speed specialized links for communication. The Coupling Facilities, at the heart of the Parallel Sysplex cluster, enable high speed, read/write data sharing and resource sharing among all the z/OS images in a cluster. All images are also connected to a Sysplex Timer® so that all events can be properly sequenced in time.

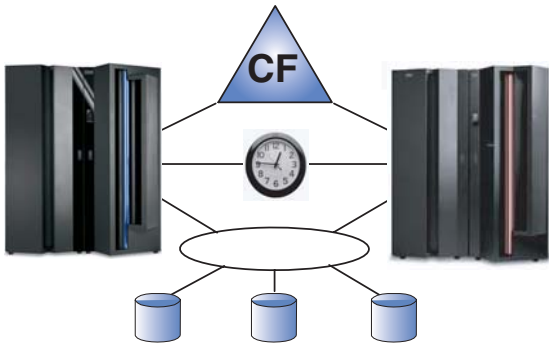

Parallel Sysplex Resource Sharing enables multiple system resources to be managed as a single logical resource shared among all of the images. Some examples of resource sharing include JES2 Checkpoint, GRS "star," and Enhanced Catalog Sharing; all of which provide simplified systems management, increased performance and/or scalability. For more detail, please see the Parallel Sysplex home page at www-03.**ibm.com**/systems/z/pso.

Although there is a significant value in a single footprint and multi-footprint environment with resource sharing, those customers looking for high availability must move on to a database data sharing configuration. With the Parallel Sysplex environment, combined with the Workload Manager and CICS TS or IMS, incoming work can be dynamically routed to the z/OS image most capable of handling the work. This dynamic workload balancing, along with the capability to have read/write access data from anywhere in the Parallel Sysplex cluster, provides scalability and availability. When configured properly, a Parallel Sysplex cluster is designed with no single point of failure and can provide customers with near continuous application availability over planned and unplanned outages. For detailed information on IBM's Parallel Sysplex technology, visit our Parallel Sysplex home page at www-03.**ibm.com**/systems/z/pso.

Parallel Sysplex Enhancements

Coupling Facility Control Code (CFCC) Level 15 is being made available on System z9 EC and BC.

Enhancement includes:

- *Increasing the allowable tasks in the CF from 48 to 112.*

Note: When migrating CF levels, lock, list and cache structure sizes may need to be increased to support the new function. This adjustment can impact the system when it allocates structures or copies structures from one coupling facility to another at different CF levels. The coupling facility structure sizer tool is designed to size structures for you, and takes into account the amount of space needed for the current CFCC levels.

Coupling Facility Configuration Alternatives

IBM offers different options for configuring a functioning Coupling Facility:

- *Standalone Coupling Facility: z900 Model 100 provides a physically isolated, totally independent CF environment. There is no unique standalone coupling facility model offered with the z9 EC, z9 BC, z990, and z890. Customers can achieve the same physically isolated environment as on prior mainframe families by ordering a z9 EC, z9 BC, z990, and z890 with PUs characterized as ICFs*

and general purpose PUs. There are no software charges associated with such configuration. An ICF or CF partition sharing a server with any operating system images not in the sysplex acts like a logical standalone CF.

- *Internal Coupling Facility (ICF): Customers considering clustering technology can get started with Parallel Sysplex technology at a lower cost by using an ICF instead of purchasing a standalone Coupling Facility. An ICF feature is a processor that can only run Coupling Facility Control Code (CFCC) in a partition. Since CF LPARs on ICFs are restricted to running only CFCC, there are no IBM software charges associated with ICFs. ICFs are ideal for Intelligent Resource Director and resource sharing environments as well as for data sharing environments where System-Managed CF Structure Duplexing is exploited.*
- *Coupling Facility partition on a System z9 platform using standard LPAR: A CF can be configured to run in either a dedicated or shared CP partition. IBM software charges apply. This may be a good alternative for test configurations that require very little CF processing resource or for providing hot-standby CF backup using the Dynamic Coupling Facility Dispatching function.*

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

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

System-Managed CF Structure Duplexing

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

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

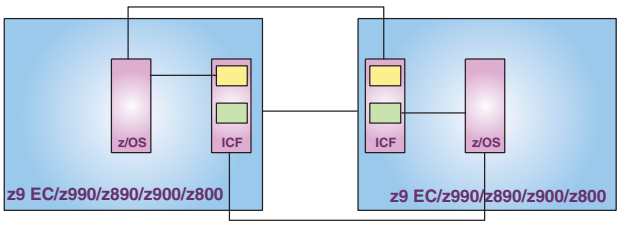

A robust failure recovery capability

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

Customers who are interested in testing and/or deploying System-Managed CF Structure Duplexing in their sysplex, please review ZSW01975, System-Managed CF Structure Duplexing at www-03.**ibm.com**/systems/z/pso or ftp://ftp.software.**ibm.com**/common/ssi/rep_wh/n/ZSW01975USEN/ZSW01975USEN.PDF to understand the performance and other considerations of using this feature.

Parallel Sysplex Coupling Connectivity

The Coupling Facilities communicate with z/OS images in the Parallel Sysplex environment over specialized high-speed links. For availability purposes, it is recommended that there be at least two links connecting each z/OS image to each CF in a Parallel Sysplex cluster. As processor performance increases, it is important to also use faster links so that link performance does not become constrained. The performance, availability and distance

requirements of a Parallel Sysplex environment are the key factors that will identify the appropriate connectivity option for a given configuration.

When connecting between System z9 and zSeries servers the links can be configured to operate in Peer Mode. This allows for higher data transfer rates to and from the Coupling Facilities. In Peer Mode, the 2 Gb/sec ISC-3 link carries traffic over single mode fiber optic cables up to an unrepeated distance of 10 km (6.2 miles). Greater distances are supported with an RPQ. ICB-3 at 1 GB/sec and ICB-4 at 2 GB/sec carry traffic over 7 meter (23 feet) copper cables. ICs are for memory-to-memory data transfers. Additional Peer Mode benefits are obtained by enabling the link to be MIFed between z/OS and CF LPARs. The peer link acts simultaneously as both a CF Sender and CF Receiver link, reducing the number of links required. Larger and more data buffers and improved protocols may also improve long distance performance.

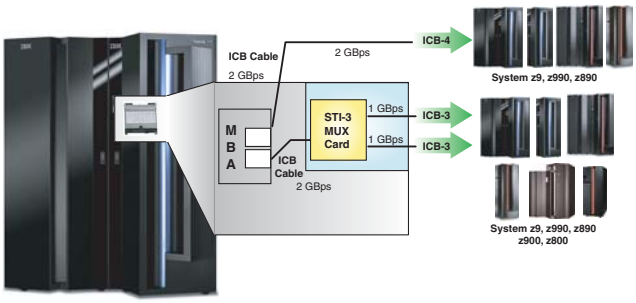

The z9 EC servers do not support ISC-3 links in Compatibility Mode. ISC-3 links are supported only in Peer Mode on z9 EC servers and so cannot be connected to 9672 G5/G6 servers. ICB-2 links are also not supported on z9 EC.

System z9 CF Link Connectivity

Connectivity Options	System z9 ISC-3	System z9 ICB-3	System z9 ICB-4
z900/z800 ISC-3	2 Gb/sec Peer Mode	N/A	N/A
System z9/z990/ z890 ISC-3	2 Gb/sec Peer Mode	N/A	N/A
z900/z800 ICB-3	N/A	1 GB/sec Peer Mode	N/A
System z9/z990/ z890 ICB-3	N/A	1 GB/sec Peer Mode Recommend use ICB-4	N/A
System z9/z990/ z890 ICB-4	N/A	N/A	2 GB/sec Peer Mode

- *ISC-3. InterSystem Channel-3 provides the connectivity required for resource or data sharing between the Coupling Facility and the systems directly attached to it.*
- *ISC-3s are point-to-point connections that require a unique channel definition at each end of the channel. ISC-3 channels operating in Peer Mode provide connection between z9 EC, z9 BC, z990, z890, z900, and z800 general purpose models and Coupling Facilities. ISC-3 channels operating in Compatibility Mode are not supported on System z9 servers, and cannot be used to provide connection between z9 EC and HiPerLink (ISC-2) channels on 9672 G5 and G6 and the 9674 R06 Models. A four port ISC-3 feature structure is provided on the z9 EC. It consists of a mother card with two daughter cards which have 2 ports each. Each port on the daughter card supports a link data rate of 2 Gb/sec in Peer Mode up to a distance of 10 km. From 10 to 20 km, an RPQ card which comes in 2 port increments is available which runs at 1 Gb/sec in Peer Mode. ISC-3 ports are purchased in one port increments.*
- *ICB-3. The Integrated Cluster Bus-3 is used to provide high-speed coupling communication between a System z9 server or CF and a z900/z800 server or CF or between two z900/z800s over short distances. (~7 meters, the ICB cable is 10 meters [33 feet] in length, but must accommodate connectivity inside the processor in addition to the distance between the machines). For longer distances, ISC-3 links must be used. When*

using ICB-3 on the z9 EC, an STI-3 distribution card is required. It resides in an I/O cage and provides 2 ICB-3 ports each capable of up to 1 GB/sec. The ports are activated in one port increments. Up to 8 STI-3 cards, 16 ICB-3 links are available on the z9. ICB-3 links operate in "Peer Mode."

- *ICB-4. The Integrated Coupling Bus-4 is a "native" coupling connection available for connecting a z9 EC server or CF to another z9 EC/z9 BC/z990/z890 server or CF over short distances (~7 meters, the ICB cable is 10 meters [33 feet] in length, but must accommodate connectivity inside the processor in addition to the distance between the machines). Capable of up to 2.0 GB/sec, the ICB-4 is the fastest external coupling connection available for the z9 EC. The ICB-4 connection consists of one link that directly attaches to an STI port on the system and does not require connectivity to a card in the I/O cage. One feature is required for each end of the link. Up to 16 ICB-4 features can be configured on a z9 EC.*
- *IC. The Internal Coupling channel emulates the Coupling Links providing connectivity between images within a single server. No hardware is required, however a minimum of 2 CHPID numbers must be defined in the IOCDS. IC links provide the fastest Parallel Sysplex connectivity.*

Continuous Availability Recommended Configuration for Parallel Sysplex Cluster

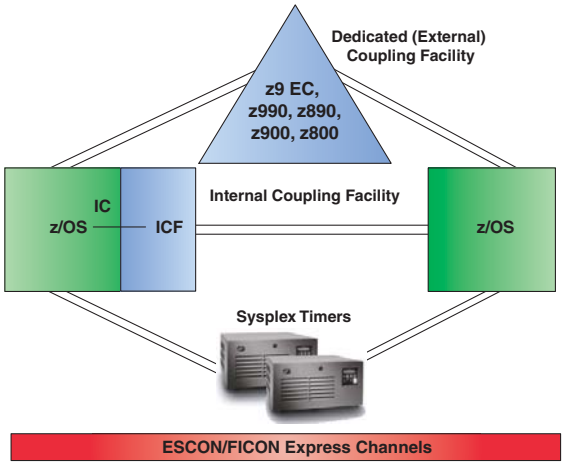

Components and assumptions

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

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

Key attributes can include

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

Message Time Ordering (Sysplex Timer Connectivity to Coupling Facilities)

As processor and Coupling Facility link technologies have improved over the years, the requirement for time synchronization tolerance between systems in a Parallel Sysplex environment has become ever more rigorous. In order to enable any exchange of timestamped information between systems in a sysplex involving the Coupling Facility to observe the correct time ordering, time stamps are now included in the message-transfer protocol between the systems and the Coupling Facility. Therefore, when a Coupling Facility is configured as an ICF on any z9 EC, the Coupling Facility will require connectivity to the same 9037 Sysplex Timer or Server Time Protocol (STP) configured Coordinated Timing Network (CTN) that the systems in its Parallel Sysplex cluster are using for time synchronization. If the ICF is on the same server as a member of its Parallel Sysplex environment, no additional connectivity is required, since the server already has connectivity to the Sysplex Timer.

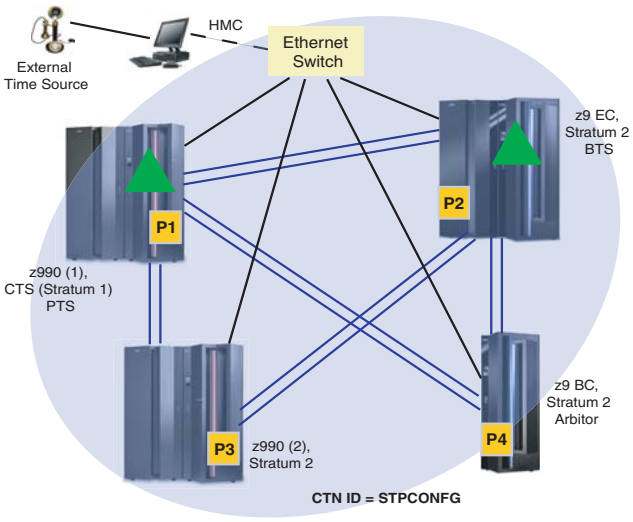

However, when an ICF is configured on any z9 EC which does not host any systems in the same Parallel Sysplex cluster, it is necessary to attach the server to the 9037 Sysplex Timer.

Server Time Protocol (STP)

In on demand business, two important objectives for survival are 1) systems designed to provide continuous availability, and 2) near transparent Disaster Recovery. Systems that are designed to provide continuous availability combine the characteristics of high availability and continuous operations to deliver high levels of service – targeted at 24x7. To attain these objectives, solutions such as GDPS are based on geographical clusters (such as Parallel Sysplex) and remote data mirroring across two or more data centers. An increasing number of enterprises are requiring that the geographical cluster or Parallel Sysplex be dispersed over distances of 100 km or larger to mitigate the risk that a single disaster could impact multiple data centers.

Server Time Protocol is a server-wide facility that is implemented in the Licensed Internal Code (LIC) of the IBM System z9, z990, and z890 servers and presents a single view of time to Processor Resource/Systems Manager (PR/SM). STP uses a message-based protocol in which timekeeping information is passed over externally defined Coupling Links between servers. The Coupling Links that can be used to transport STP messages include Inter System Channel-3 (ISC-3) links configured in peer mode, Integrated Cluster Bus-3 (ICB-3) links, and Integrated Cluster Bus-4 (ICB-4) links. These can be the same links that are already being used for Coupling Facility communication.

By using the same links to exchange time-keeping information and Coupling Facility messages in a Parallel Sysplex, STP can scale with distance. In other words, servers exchanging messages over ICB-3 and ICB-4 links can now meet more stringent synchronization requirements than servers exchanging messages over long ISC-3 links (distances up to 100 km), where the synchronization requirements are less stringent. This removes one of the restrictions of the current Sysplex Timer implementation.

The Server Timer Protocol feature is designed to simplify hardware configurations by:

- *Allowing clock synchronization for System z9, z990, and z890 servers and Coupling Facilities without requiring the Sysplex Timer and dedicated timer links. This helps reduce the need for separate hardware that needs to be ordered and maintained.*
- *Supporting a multi-site timing network of up to 100 km (62 miles) over fiber optic cabling, thus allowing a sysplex to span these distances. This overcomes the limitations of the Sysplex Timer to Sysplex Timer links being supported only up to 40 km.*
- *Potentially reducing the cross-site connectivity required for a multi-site Parallel Sysplex. Dedicated links are no longer required to transport timing information since STP and Coupling Facility messages may be transmitted over the same links.*

For more details, visit the STP Web site at: www-03.**ibm.com**/systems/z/pso/stp.html.

Parallel Sysplex Professional Services

IBM provides extensive services to assist customers in migrating their environments and applications to benefit from Parallel Sysplex clustering. A basic set of IBM services is designed to help address planning and early implementation requirements. These services can help you reduce the time and costs of planning a Parallel Sysplex environment and moving it into production.

IBM Global Services has a variety of IT and Geographically Dispersed Parallel Sysplex (GDPS) Services. Please consult: http://www-1.**ibm.com**/services/us/index.wss.

GDPS

GDPS is a multi-site or single-site end-to-end application availability solution that provides the capability to manage remote copy configuration and storage subsystems (including IBM TotalStorage®), to automate Parallel Sysplex operation tasks and perform failure recovery from a single point of control.

GDPS helps automate recovery procedures for planned and unplanned outages to provide near-continuous availability and disaster recovery capability.

For additional information visit: www-03.**ibm.com**/systems/z/gdps/.

Fiber Optic Cabling and System Connectivity

IBM Network Integration and Deployment Services for System z9 fiber cabling (System z9 fiber cabling services) enables businesses to choose the System z9 configuration that best matches their computing environment without having to worry about planning and implementing the fiber optic cabling. By teaming with IBM, businesses can receive a world-class solution for their System z9 fiber connectivity requirements, including consulting and project management, as well as the fiber optic jumper cables and installation to complete the System z9 integration.

System z9 fiber cabling now offers three options to address a solution for your fiber cable installation. Enterprise fiber cabling offers two additional options to help meet your structured (trunking) environment requirements.

System z9 fiber cabling:

- *Fiber optic jumper cabling package*
  will analyze your System z9 channel configuration and your existing fiber optic cabling to determine the appropriate fiber optic jumper cables required, then supply, label and install the fiber optic jumper cables and complete the installation with a detailed connection report.
- *Fiber optic jumper migration and reuse for a System z9 upgrade*
  will plan, organize, re-label, re-route and re-plug your existing fiber optic jumper cables for reuse with the upgraded System z9 server.
- *Fiber optic jumper cables and installation*
  will supply the fiber optic jumper cables you specify, then label and install the fiber optic jumper cables.

Enterprise fiber cabling options:

- *System z9 fiber optic trunk cabling package*
  will analyze your System z9 channel configuration and your existing fiber optic infrastructure to determine the appropriate fiber optic harnesses, fiber optic trunk cables and the fiber optic patch panel boxes required, then supply, label and install the fiber optic components to connect your new System z9 server to your existing structured fiber cabling infrastructure.
- *Enterprise fiber cabling package*
  will analyze your entire data center configuration and existing fiber optic infrastructure to determine the appropriate end-to-end enterprise solution for connectivity. This is a customized offering that includes trunk cables, zone cabinets, patch panels and direct attach harnesses for servers, directors and storage devices.

These tailored System z9 fiber cabling options use the same planning and implementation methodologies as IBM's customized enterprise fiber cabling services, but are focused on your System z9 fiber cabling needs.

Fiber Quick Connect (FQC): FQC, a System z9 configuration option for ESCON channels, helps reduce the cable bulk associated with the installation of potentially 240 (z800) to 256 (z900) to 420 (z9 EC and z990) ESCON channels in one I/O cage. Fiber harnesses, which are factory-installed, enable connection to IBM's Fiber Transport System (FTS) direct-attach fiber trunk cables. Each trunk can have up to 72 fiber pairs. Four trunks can displace the 240 to 256 fiber-optic cables.

In planning for System z9 servers, refer to *Planning for: Fiber Optic Links (ESCON, FICON, Coupling Links, and Open System Adapters)*, GA23-0367, and the Installation Manual Physical Planning (IMPP) manual. Refer to the services section of Resource Link for further details on the System z Fiber Cabling Service options and the Fiber Quick Connect configuration option.

Access Resource Link at **ibm.com**/servers/resourcelink.

System z9 EC Configuration Detail

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

All features that require I/O slots, and ICB-4 features, are included in the following table:

Feature	Minimum # of features	Maximum # of features	Maximum Connections	Increments per Feature	Purchase Increm.
ESCON, 16 port	0[1]	69	1024 channels	16 channels 1 reserved as a spare	4 channels
FICON Express4	0[1]	84	336 channels	4 channels	4 channels
FICON Express2*	0[1]	84	336 channels	4 channels	4 channels
FICON Express*	0[1]	60	120 channels	2 channels	2 channels
STI-3[3] ICB-3 link	0 0[1]	8 N/A	N/A 16 links[3]	2 outputs N/A	N/A 1 link
ICB-4	0[1]	N/A	16 links[3, 4]	N/A	1 link
ISC-3	0[1]	12	48 links[3]	4 links	1 link
OSA-Express2	0	24	48 ports	2 or 1 (10 GbE has 1)	2 ports/ 1 port
OSA-Express*	0	24	48 ports	2 ports	2 ports
Crypto Express2	0	8	16 PCI-X adapters	2 PCI-X adapters	2 PCI-X adapters[5]

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

2) Each STI-3 distribution card occupies one I/O slot (supports ICB-3s).

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

4) ICB-4s are not included in the maximum feature count for I/O slots but are included in the CHPID count.

5) Initial order of Crypto Express2 is 4 PCI-X adapters (two features). Each PCI-X adapter can be configured as a coprocessor or an accelerator.

** Available only when carried forward on an upgrade.*

Note: There is a maximum of 64 ESCON features/960 active channels and a maximum of 64 FICON features/256 channels Model S08.

Processor Unit Assignments

Model	Min. PU*	Max. PU	SAP Standard	Max Subcapacity Engines	Spares Standard
S08	1	8	2	8	2
S18	1	18	4	8	2
S28	1	28	6	8	2
S38	1	38	8	8	2
S54	1	54	8	8	2

PUs can be characterized as CPs, IFLs, ICFs, zAAPs, or Optional SAPs, up to a max number of PUs for the model

*Customer will be required to purchase at least one CP, IFL or ICF feature for any model.

Processor Memory

z9 EC Model	Minimum	Maximum
S08	16 GB	128 GB
S18	16 GB	256 GB
S28	16 GB	384 GB
S38	16 GB	512 GB
S54	16 GB	512 GB

Max 8 memory cards per z9 EC book. Memory cards 4 GB, 8 GB or 16 GB.

Channels

z9 EC Model	S08	S18	S28	S38	S54
ESCON Min	0	0	0	0	0
ESCON Max	960	1024	1024	1024	1024
FICON Min	0	0	0	0	0
FICON Express2 Max	256	336	336	336	336
FICON Express Max	120	120	120	120	120

ESCON and FICON Express2 increments in 4 channels; FICON Express increments in 2 channels.

Coupling Links

Links	IC	ICB-3*	ICB-4	ISC-3	Max Links
	0-32	0-16	0-16	0-48	Total External + Internal links = 64

*requires STI-3 card

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

Cryptographic Features

	Crypto Express2 Feature*
Minimum	0
Maximum	8

*Each feature has 2 PCI-X adapters; each adapter can be configured as a coprocessor or an accelerator.

OSA-Express2 and OSA-Express Features

	OSA-Express2 and OSA-Express Features*
Minimum	0
Maximum	24 (up to 48 ports; 10 GbE has 1 port)

*Any combination of OSA-Express2 features (1000BASE-T Ethernet, GbE LX, GbE SX, 10 GbE) and OSA-Express features (1000BASE-T Ethernet, GbE LX, GbE SX, Fast Ethernet).

z9 EC Frame and I/O Configuration Content: Planning for I/O

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

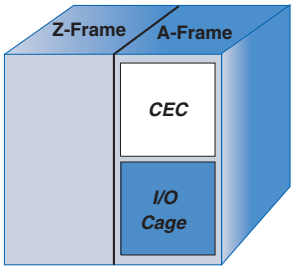

Single I/O cage

I/O Feature Type	Features	Maximum
ESCON	24	360 channels
FICON Express2/4	24	96 channels
FICON Express	24	48 channels
OSA-Express2	24	48 ports
OSA-Express	24	48 ports
Crypto Express2	8	16 adapters

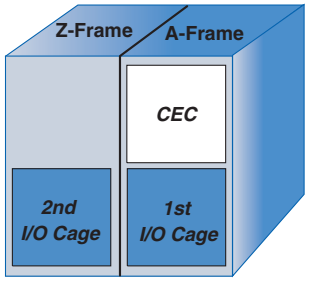

2 I/O cages

I/O Feature Type	Features	Maximum
ESCON	48	720 channels
FICON Express2/4	48	192 channels
FICON Express	48	96 channels
OSA-Express2	24	48 ports
OSA-Express	24	48 ports
Crypto Express2	8	16 adapters

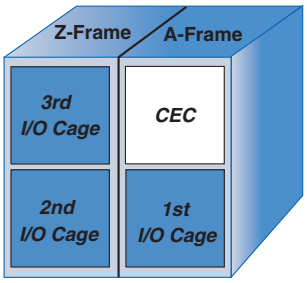

3 I/O cages

I/O Feature Type	Features	Maximum
ESCON	69	1024 channels
FICON Express2/4	84	336 channels
FICON Express	60	120 channels
OSA-Express2	24	48 ports
OSA-Express	24	48 ports
Crypto Express2	8	16 adapters

General Information:

- *ESCON configured in 4-port increments. Up to a maximum 69 cards, 1024 channels.*
- *OSA-Express2 can be Gigabit Ethernet (GbE), 1000BASE-T Ethernet or 10 GbE.*
- *OSA-Express can be Gigabit Ethernet (GbE), 1000BASE-T Ethernet or Fast Ethernet.*
- *If ICB-3 is required on the system, it will use up a single I/O slot for every 2 ICB-3 to accommodate the STI-3 card.*

Note: In the first and second I/O cage, the last domain in the I/O cage is normally used for ISC-3 and ICB-3 links. When the first 6 domains in an I/O cage are full, additional I/O cards will be installed in the next I/O cage. When all the first 6 domains in all I/O cages are full and no Coupling link or PSC cards are required, the last domain in the I/O cage will be used for other I/O cards making a total of 28 per cage.

System z9 EC Physical Characteristics

z9 EC Power/Heating/Cooling

System Power Consumption (kW)

Model / Config	1 I/O Cage	2 I/O Cage	3 I/O Cage
S08	6.3	9.2	12.1
S18	8.8	11.8	14.7
S28	10.9	13.9	16.9
S38 and S54	12.8	15.7	18.3

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

System Cooling (Air Flow Rate - CFM)

Config	m^3/min
1 I/O Cage	905
3 I/O Cage	1965

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

Heat Output (kBTU/hr)

Model / Config	1 I/O Cage	2 I/O Cage	3 I/O Cage
S08	21.5	31.4	41.3
S18	30	40.2	50.1
S28	37.2	47.4	57.6
S38 and S54	43.6	53.5	62.4

z9 EC Dimensions

	z9 EC
# of Frames	2 Frames IBF contained within 2 frames
Height (w/ covers) Width (w/ covers) Depth (w/ covers)	194.1 cm / 76.4 in (40 EIA) 157.7 cm / 62.1 in (each frame 30.2 in) 157.7 cm / 62.1 in
Height Reduction Width Reduction	178.5 cm / 70.3 in (38 EIA) None
Machine Area Service Clearance	2.49 sq. meters / 26.78 sq. feet 5.45 sq. meters / 58.69 sq. feet (IBF contained within the frame)

System z9 Operating System Support

The following table shows the operating systems' minimum levels that are required to support the System z9 and the new functions and features. Minimum level means that subsequent versions and releases also include the support.

In all cases, a PSP Bucket is required to provide the proper level of support.

	z/OS z/OS.e [1]	z/VM	Linux on System z9	z/VSE	z/TPF TPF [2]
Basic System z9 Support	1.4 [4] *	4.4 *	SLES 9 RHEL 4	3.1	1.1 4.1 [2]
LPAR Group Capacity	1.8 [4]	Not Supported	Not Supported	Not Supported	Not Supported
Hardware Decimal Floating Point	1.6 [4]	5.2 (Guest)	Not Supported	Not Supported	Not Supported
CFCC Level 15	1.6 [4]	5.2 (Guest)	Not Supported	Not Supported	
System-initiated CHPID Reconfiguration	1.9 [4]	Not Supported	Work in progress with LDP	Not Supported	Not Supported
60 Logical Partitions (30 for z9 BC)	1.4 [4] *	4.4 *	SLES 9 RHEL 4	3.1	1.1 4.1 [2]
63.75K Subchannels	1.4 [4] *	4.4 *	SLES 9, RHEL 4		
OSA-Express2 1000BASE-T Ethernet	1.4 [4] *	4.4 *	SLES 9 RHEL 4	3.1	1.1 4.1 PUT 13 [2]
MIDAW Facility	1.6	5.3 (Guest)	Not Supported		
CPACF Enhancements	1.6 [4]	4.4 *	SLES 9 [5], RHEL 4 [5]	3.1	
Crypto Express2 Exploitation	1.6 [4]	5.1	SLES 9, RHEL 4	3.1	
HiperSockets IPv6	1.7	5.2	Not Supported		
OSA-Express2 Large Send	1.6	5.1 (Guest)	SLES 9, RHEL 4		
OSA-Express2 CDLC Support	1.4 [4] *	5.1	SLES 9, RHEL 4	3.1	1.1
OSA Dynamic Idle	1.8 [4]	5.1 (Guest)	Not Supported	Not Supported	Not Supported
OSA Express2 Link Aggregation	TBC [4]	5.3	SLES 9, RHEL 4	Not Supported	Not Supported
OSA Layer 3 VMAC	1.8 [4]	5.1 (Guest)	Not Supported	Not Supported	Not Supported
OSA Express Network Traffic Analyzer	1.8 [4]	5.1 (Guest)	Not Supported	Not Supported	Not Supported
QDIO Diagnostic Sync	1.8 [4]	5.1 (Guest)	Not Supported	Not Supported	Not Supported
Multiple Subchannel Sets (MSS)	1.7	Not Supported	SLES 10, RHEL 5 [3]		
FICON Link Incident Report	1.7	4.4 *	Not Supported		
Multi-path IPL	1.6 [4]	Not Supported	Not Supported	Not Supported	Not Supported
Single System Image	1.6 up to 32	5.1 up to 24 5.3 up to 32	SLES 9 up to 32 RHEL 4 up to 32		1.1 up to 52
Enhanced Perf Assists for z/VM Guests	Not supported	5.2	SLES 10, RHEL 5 [3]		
N_Port ID Virtualization	Not supported	4.4* (Guest), 5.1	SLES 9, RHEL 5 [3]		3.1
FCP Program Directed re-IPL	Not supported	5.3 (Guest)	SLES 9, RHEL 5 [3]	4.1	
SubCapacity Systems	1.4 [4] *	4.4 *	IBM Software Group products are enabled [6]	3.1	1.1 4.1 [2]
zIIP Support	1.6	5.3 (Guest)	Not supported	Not supported	Not supported
Crypto Remote Key Loading	1.6 [4]	5.1	Not supported		
Crypto ISO 16609	1.6 [4]	5.1	Not supported		
FICON Express4 (CHPID type FC)	1.4 [4] *	4.4 *	SLES 9 RHEL 4	3.1	1.1 4.1PUT 16 [2]
FICON Express4 (CHPID type FCP)	Not supported	4.4 *	SLES 9, RHEL 4	3.1	

1) z/OS.e supported on z9 BC, z890, z800 only
2) indicates TPF
3) this function will be provided in future Linux on System z distribution releases/service updates
4) additional features, service or Web downloads required
5) IBM is working with LDPs on Kernel space exploitation[3]
6) Linux does not support it, the IBM Software Group products are enabled for it on all distributions
**) z/OS 1.4 and z/VM 4.4 are EOS.*
Please refer to the latest PSP bucket for latest PTFs for new functions/features. SLES = SUSE Linux Enterprise Server, RHEL = Red Hat Enterprise Linux

Coupling Facility - CF Level of Support

CF Level	Function	z800	z900	z890 / z990	z9 EC / z9 BC
15	Increasing the allowable tasks in the CF from 48 to 112				X
14	CFCC Dispatcher Enhancements		X	X	X
13	DB2 Castout Performance	X	X	X	X
12	z990 Compatibility 64-bit CFCC Addressability Message Time Ordering DB2 Performance SM Duplexing Support for zSeries	X X X X	X X X X X	X X X X X	X X X X X
11	z990 Compatibility SM Duplexing Support for 9672 G5/G6/R06				
10	z900 GA2 Level		X		
9	Intelligent Resource Director IC3 / ICB3 / ISC3 Peer Mode MQSeries® Shared Queues WLM Multi-System Enclaves	X X X X	X X X X	X X X X	X X X X
8	Dynamic ICF Expansion into shared ICF Pool Systems-Managed Rebuild	X X	X X	X X	X X
7	Shared ICF partitions on server models DB2 Delete Name Optimization	X X	X X	X X	X X

Note: G5/G6 and prior generation servers are not supported with System z9 for Coupling Facility or Parallel Sysplex levels.

System z9 and IBM eServer zSeries Features and Functions

The following table consists of a list of features (functions) that are supported on IBM System z9 and eServer zSeries servers.

Note: This table is not intended to include services, RPQs or specific quantities or measurements related to performance, memory size, bandwidth, HiperSockets, etc. The intention of this matrix is to provide a comparison of the standard and optional features (via a feature code) for the various servers.

I/O	System z9		zSeries			
	z9 BC	z9 EC	z990	z890	z900	z800
ESCON LED channels (17 MB)	O	O	O	O	O	O
Spare ESCON port(s)	S	S	S	S	S	S
FICON Express4 (4-Port)	O	O	–	–	–	–
FICON Express4 (2-Port)	O	O	–	–	–	–
FICON Express Channels (2 Gbps)	O	O	O	O	O	O
FICON CTC	S	S	S	S	S	S
MIDAW Facility	S	S	–	–	–	–
Multipath IPL	O	O	–	–	–	–
FICON Cascaded Director (system attached)	S	S	S	S	S	S
FICON FCP Support for z/VM, z/VSE and Linux (attach to SCSI devices)	O	O	O	O	O	O
FCP Program Directed re-IPL	O	O	–	–	–	–
FICON Link Incident Reporting	S	S	O	O	O	O
Request Node Identification Data (RNID)	S	S	–	–	–	–
FCP full-fabric connectivity to SCSI storage devices	O	O	O	O	O	O
FCP SCSI IPL	O	O	O	O	O	O
HiperSockets	S	S	S	S	S	S
HiperSockets support of IPv6	S	S	–	–	–	–
Open Systems Adapter-Express (OSA-Express):						
1000BASE-T Ethernet	O	O	O	O	–	–
Integrated ASCII ConsoleController (OSA-ICC)	O	O	O	O	–	–
OSA-Express2:						
Gigabit Ethernet	O	O	O	O	–	–
10 Gigabit Ethernet	O	O	O	O	–	–
1000BASE-T Ethernet	O	O	–	–	–	–
OSN (OSA for NCP)	O	O	–	–	–	–
OSN Dynamic LAN Idle	O	O	–	–	–	–
OSN Layer 3 VMAC	O	O	–	–	–	–
OSN Express2 Link Aggregation for z/VM	O	O	–	–	–	–
OSN Express2 Network Traffic Analyzer	O	O	–	–	–	–
QDIO Diagnostic Synchronization	O	O	–	–	–	–
Logical Channel Subsystems (LCSS): up to 256 channels per LCSS						
Up to two LCSSs	S	–	–	S	–	–
Up to four LCSSs	–	S	S	–	–	–
Multiple Subchannel Sets (MSS)	S	S	–	–	–	–
Spanned Channels (IC, HiperSockets, FICON, ICB-3, ICB-4, ISC-3, OSA)	S	S	S	S	–	–
System-initiated CHPID Reconfiguration	S	S	–	–	–	–

Cryptographic Features	System z9		zSeries			
	z9 BC	z9 EC	z990	z890	z900	z800
CP Assist Cryptographic Function (CPACF)	S	S	S	S	–	–
AES, PRNG, SHA-256*	S	S	–	–	–	–
Trusted Key Entry (TKE)	O	O	O	O	O	O
Crypto Express2 (2 Features)	O	O	O	O	–	–
Crypto Express2 (1Feature)	O		–	–	–	–

PR/SM	System z9		zSeries			
	z9 BC	z9 EC	z990	z890	z900	z800
LPAR Mode:	S	S	S	S	O	O
Up to 60 LPARs	–	S	–	–	–	–
Up to 30 LPARs	S	–	S	S	–	–
Up to 15 LPARs	–	–	–	–	S	S
EAL5 certification	S	S	S	S	S	S
LPAR Group Capacity Limit	O	O	–	–	–	–

Parallel Sysplex	System z9		zSeries			
	z9 BC	z9 EC	z990	z890	z900	z800
InterSystem Coupling-3 (ISC-3) Links	O	O	O	O	O	O
InterSystem Coupling-3 (ISC-3) Peer Mode (CFP)	O	O	O	O	O	O
Internal Coupling Channel-3 (IC-3)	S	S	S	S	S	S
Integrated Cluster Bus-3 (ICB-3) Links (1 GB) [1]	O	O	O	O	O	O
Integrated Cluster Bus-4 (ICB-4) Links (2 GB)	O	O	O	O	–	–
External Time Reference (ETR)	O	O	O	O	S	S
Server Time Protocol (STP)	O	O	O	O	–	–

Specialty Engines	System z9		zSeries			
	z9 BC	z9 EC	z990	z890	z900	z800
System z Application Assist Processor (zAAP)	O	O	O	O	–	–
System z9 Integrated Information Processor (zIIP)	O	O				
Integrated Facility for Linux (IFL)	O	O	O	O	O	O
Integrated Coupling Facility (ICF)	O	O	O	O	O	O

Parallel Sysplex	System z9		zSeries			
	z9 BC	z9 EC	z990	z890	z900	z800
On/Off Capacity on Demand (On/Off CoD)						
On/Off CoD with Extended Staging	S	S	S	S	–	–
On/Off CoD Test	S	S	S	S	–	–
Customer Initiated Upgrade	O	O	O	O	O	O
Capacity Backup (CBU) CP only	S	S	S	S	S	S
Capacity Backup (CBU) for IFLs, ICFs, zAAPs and zIIPs	S	S	–	–	–	–
Partial Memory Restart	S	S	S	S	S	–
Partial CP Restart	S	S	S	S	S	–
Partial I/O Restart	–	–	–	–	S	–
Failure Containment for the MBA	S	S	S	S	–	–
Dynamic Memory Sparing	S	S	S	S	S	S
Dynamic SAP Sparing	S	S	S	S	S	S
Enhanced Book Availability	–	S	–	–	–	–
Concurrent Book Add	–	S	S	–	–	–
Dual Power Feeds	S	S	S	S	S	S
Internal Battery Feature (IBF)	O	O	O	O	O	–
Concurrent PU Conversions	O	O	O	–	–	–
CFCC Enhanced Patch Apply	S	S	S	S	–	–

Key:
S = standard O = optional – = not supported

*1 While ICB-3 is supported for connection to z9 EC, z9 BC, z990, z890, z900, and z800, it is recommended that ICB-4 be used for connections to z9 EC, z9 BC, z990 and z890.
For a complete list of features refer to the System z9 Functional Matrix available from the Library area of Resource Link at: www.**ibm.com**/servers/resourcelink.

Endnote:

1) All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represents goals and objectives only.

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A.

Produced in the United States of America
04/07

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

IBM, IBM eServer, the IBM logo, the e-business logo, APPN, CICS, DB2, DB2 Connect, DB2 Universal Database, DFSMSdss, DRDA, ECKD, ESCON, FICON, Geographically Dispersed Parallel Sysplex, GDPS, HiperSockets, Hiperspace, IMS, Parallel Sysplex, PR/SM, Processor Resource/Systems Manager, RACF, Rational, Redbook, Resource Link, S/390, Sysplex Timer, System z, System z9, Tivoli, TotalStorage, Virtualization Engine, VM/ESA, WebSphere, z/Architecture, z/OS, z/VM, z/VSE, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation in the Unites States and other countries.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

ZSO03005-USEN-01