

IBM System z10 Enterprise Class - The future runs on System z10, the future begins today

Table of contents

3 Key prerequisites	18 Publications
3 Planned availability dates	19 Services
4 Description	19 Technical information
16 Product positioning	25 IBM Electronic Services
16 Statement of general direction	25 Terms and conditions
17 Product number	26 Prices
17 Education support	26 AP distribution

At a glance



The future runs on System z® - The future begins today.

The IBM® System z10 Enterprise Class (z10 EC) server is the next step in the evolution of the large-scale System z mainframe, fulfilling our promise to deliver technology improvements in reliability, availability, serviceability, security, scalability, virtualization, and energy efficiency. These core competencies make it an excellent choice to be the cornerstone of the IT infrastructure with innovative technologies to stay competitive. Controlling costs, improving provisioning speed, and providing data center security and resiliency help create an environment to run modern applications side-by-side with traditional mission-critical workloads.

Enhancements today include:

- Improved access to data with High Performance FICON® for System z (zHPF) on both FICON Express4 and FICON Express2
- Enhanced problem determination, analysis, and manageability of the storage area network (SAN) by providing registration information to the fabric on the name server for both FICON and FCP
- Increased performance for Local Area Network connectivity with introduction of Open Systems Adapter-Express3 (OSA-Express3) 1000BASE-T Ethernet

- A Short Reach version of OSA-Express3 10 Gigabit Ethernet with two ports per feature for infrastructures with multimode fiber optic cabling
- OSA-Express3 Gigabit Ethernet and 1000BASE-T for the Network Control Program, providing Channel Data Link Control protocol support between the z10 EC and IBM Communications Controller for Linux® on System z (CCL) allowing systems administrators to configure, manage, and operate their CCL Network Control Programs in the same manner as their ESCON-attached 374x NCPs
- Protection from network intrusion with OSA-Express QDIO data connection isolation for z/VM® virtualized environments on System z10 and System z9®
- Long Reach 1x InfiniBand coupling links - an alternative to ISC-3 facilitating coupling link consolidation
- Coupling Facility Control Code Level 16 - to help deliver faster service time for CF Duplexing, and improvements to the efficiency of workload distribution when using shared queues in the Coupling Facility
- Updates to Server Time Protocol for enhanced time accuracy, availability, and systems management
- Support for Longer Personal Account Numbers for stronger data protection on Crypto Express2
- Trusted Key Entry Licensed Internal Code 5.3 enhancement to support Advanced Encryption Standard (AES) encryption algorithm, audit logging, and an infrastructure for payment card industry data security standard (PCIDSS)
- Increased flexibility for just-in-time offerings with ability for more temporary offerings installed on the CPC and ways to acquire capacity backup
- Plan Ahead Memory for nondisruptive memory upgrades
- 62% price reduction on memory for z10 EC when purchased after October 21, 2008, and with a specialty engine for new workloads, now \$2,250 per GB¹
- Ability to allow production workload to be executed on a CBU Upgrade during a CBU test provided that certain contract terms are in effect with IBM

¹ Limited to 16 GB per specialty engine

Overview

The IBM System z10 Enterprise Class (z10 EC) server is designed to meet the challenges of today's business world and to be the cornerstone of an evolutionary new model for efficient IT delivery called the new enterprise data center. This model helps reset the economics of IT and can dramatically improve operational efficiency, security, and responsiveness - to help keep a business competitive.

The z10 EC, with its advanced combination of reliability, availability, serviceability, security, scalability, and virtualization, delivers the technology that can help define this framework for the future. The z10 EC delivers improvements to performance, capacity, and memory which can help enterprises grow their existing businesses while providing a cost-effective infrastructure for large-scale consolidation.

Today's announcements extend the z10 EC leadership with improved access to data and the network; tighter security with longer Personal Account Numbers for stronger protection of data; enhancements for improved performance when connecting to the network; increased flexibility in defining your options to handle backup requirements; and enhanced time accuracy to an external time source.

Any successful business needs to be able to deliver timely, integrated information to business leaders, support personnel, and customers on a 24x7 basis. This means that access to data needs to be fast, secure, and dependable. Enhancements made to z/Architecture® and the FICON interface architecture with the new High Performance FICON for System z (zHPF) are optimized for online transaction processing (OLTP) workloads. The FICON Express4 and FICON Express2 features support the native FICON protocol and the new zHPF protocol simultaneously.

The System z10 was introduced with a new connectivity option for LANs - Open Systems Adapter-Express3 (OSA-Express3). Today we are announcing two new OSA-Express3 Ethernet

features: 10 Gigabit Ethernet Short Reach complements the Long Reach feature and supports infrastructures with multimode fiber optic cabling while the 1000BASE-T Ethernet feature supports OSA-Integrated Console Controller (OSC), TCP/IP Layer 3 and Layer 2 (OSD), non-QDIO TCP/IP and SNA/APPN/HPR (OSE), and OSA-Express for NCP (OSN). Both OSA-Express3 and OSA-Express2 have protection from network intrusion with QDIO data connection isolation for z/VM virtualized environments.

An IT system needs to be available and protected every day. The z10 EC offers availability enhancements which include faster service time for CF Duplexing, updates to Server Time Protocol (STP) for enhanced time accuracy to an External Time Source, and support for heterogeneous platforms in an enterprise to track to the same time source. Security enhancements to the Crypto Express2 feature deliver new support for 13-, 14-, 15-, 16-, 17-, 18-, and 19-digit Personal Account Numbers for stronger protection of data.

The z10 EC has a new architectural approach for temporary offerings that have the potential to change the thinking about on demand capacity. The z10 EC can have one or more flexible configuration definitions that can be available to solve multiple temporary situations and multiple capacity configurations that can be active at once. This means that On/Off Capacity on Demand (CoD) can be active and up to seven other offerings can be active simultaneously. Tokens are available that can be purchased for On/Off CoD either before or after execution. The System z10 offers Plan Ahead Memory, giving users the option of planning for nondisruptive memory upgrades.

Updates to the z10 EC are designed to help improve IT today, outline a compelling case for the future running on System z, and lock in the z10 EC as the cornerstone of the new enterprise data center.

There are new terms governing System z Capacity Back Up (CBU) which allow you to execute production workload on a CBU Upgrade during a CBU test.

Key prerequisites

Refer to the [Hardware requirements](#) and [Software requirements](#) sections of this announcement.

Planned availability dates

October 28, 2008

- z10 EC features and functions
- TKE 5.3 LIC (#0854) on z10 EC, z9™ EC, z9 BC
- TKE additional smart cards (#0884) on z10 EC, z9 EC, z9 BC
- TKE smart card reader (#0885) on z10 EC, z9 EC, z9 BC
- STP enhancements on z10 EC

January 28, 2009

Functions:

- OSA-Express QDIO data connection isolation for the z/VM environment on System z10 and System z9
- STP enhancements on z9 EC and z9 BC
- STP configuration and time information saved across Power-on-Resets or power outages on z990 and z890

MES features:

- HCA2-O LR fanout card for Long Reach 1x InfiniBand (#0168) on z10 EC
- OSA-Express3 1000BASE-T (#3367) and 10 GbE SR (#3371) on z10 EC
- Capacity on Demand features (#9917 - #9929) on z10 EC

- TKE 5.3 LIC (#0854) on z10 EC, z9 EC, z9 BC
- TKE additional smart cards (#0884) on z10 EC, z9 EC, z9 BC
- TKE smart card reader (#0885) on z10 EC, z9 EC, z9 BC

First half 2009

- HMC feature #0090 on z10 EC, z9 EC, z9 BC

Description

LPAR dynamic PU reassignment

System configuration has been enhanced to optimize the CPU-to-book allocation of physical processors (physical units - PUs) dynamically. The initial allocation of customer-usable PUs to physical books can change dynamically to better suit the actual logical partition configurations that are used on the server.

Swapping of specialty engines and Central Processors (CPs - general purpose processors) with each other can now occur, as the system attempts to "pack" logical partition configurations into physical configurations that span the least number of books. The effect of this can be observed in dedicated as well as shared partitions that utilize HiperDispatch.

LPAR dynamic PU reassignment is exclusive to System z10 and is transparent to operating systems.

CPU measurement facility

The CPU measurement facility is a hardware facility which consists of counters and samples. The facility provides a means to collect runtime data for software performance tuning. The detailed architecture information for this facility can be found in the System z10 *Library* section of Resource Link™. Search for *CPU-Measurement Facility* at

<http://www.ibm.com/servers/resourcelink>

CPU measurement facility is exclusive to System z10 and is supported by z/OS®. Refer to the [Software requirements](#) section.

Cryptographic enhancements for security-rich transactions

Secure Key AES: The Advanced Encryption Standard (AES) is a National Institute of Standards and Technology specification for the encryption of electronic data. It is expected to become the accepted means of encrypting digital information, including financial, telecommunications, and government data. AES is the symmetric algorithm of choice, instead of Data Encryption Standard (DES) or Triple-DES, for the encryption and decryption of data. The AES encryption algorithm is supported with secure (encrypted) keys of 128, 192, and 256 bits.

The secure key approach, similar to what is supported today for DES and TDES, offers the ability to keep the encryption keys protected at all times, including the ability to import and export AES keys, using RSA public key technology.

Support for AES encryption algorithm includes the master key management functions required to load or generate AES master keys, update those keys, and re-encipher key tokens under a new master key.

Secure key AES is exclusive to System z10 and is supported by z/OS and z/VM for guest exploitation. Refer to the [Software requirements](#) section.

Support for 13- through 19-digit Personal Account Numbers: Credit card companies sometimes perform card security code computations based on Personal Account Number (PAN) data. Currently, Integrated Cryptographic Service Facility (ICSF) callable services CSNBCSV (VISA CVV Service Verify) and CSNBMSG (VISA CVV Service Generate) are used to verify and to generate a VISA Card Verification Value (CVV) or a MasterCard Verification Code (CVC).

The ICSF callable services currently support 13-, 16-, and 19-digit PANs. To deliver additional flexibility, new keywords PAN-14, PAN-15, PAN-17, and PAN-18 are implemented in the rule array for both CSNBCSV and CSNBCSG to indicate that the PAN data is comprised of 14, 15, 17, or 18 PAN digits, respectively.

Support for 13- through 19-digit PANs is exclusive to System z10 and is offered by z/OS and z/VM for guest exploitation. Refer to the [Software requirements](#) section.

Trusted Key Entry 5.3 LIC: The Trusted Key Entry (TKE) workstation and the TKE 5.3 level of Licensed Internal Code (#0854) are optional features on the z10 EC. The TKE 5.3 Licensed Internal Code (LIC) is loaded on the TKE workstation prior to shipment. The TKE workstation offers security-rich local and remote key management, providing authorized persons a method of operational and master key entry, identification, exchange, separation, and update. The TKE workstation supports connectivity to an Ethernet Local Area Network (LAN) operating at 10 or 100 Mbps. Up to ten TKE workstations can be ordered.

Enhancement with TKE 5.3 LIC: The TKE 5.3 level of LIC includes support for the AES encryption algorithm, adds 256-bit master keys, and includes the master key management functions required to load or generate AES master keys to cryptographic coprocessors in the host.

Also included is an embedded screen capture utility to permit users to create and to transfer TKE master key entry instructions to diskette or DVD. Under "Service Management" a "Manage Print Screen Files" utility is available to all users.

TKE 5.3 LIC has added the capability to store key parts on DVD-RAMs and continues to support the ability to store key parts on paper, or optionally on a smart card. TKE 5.3 LIC has limited the use of floppy diskettes to read-only. The TKE 5.3 LIC can remotely control host cryptographic coprocessors using a password-protected authority signature key pair either in a binary file or on a smart card.

The TKE workstation and TKE 5.3 LIC (#0854) are available on the z10 EC, z10 BC, z9 EC, and z9 BC.

Refer also to the *Special features* section of the *Sales Manual* on the Web for further information

<http://www.ibm.com/common/ssi/index.wss>

Smart Card Reader - new feature: Support for an optional Smart Card Reader (#0885) attached to the TKE 5.3 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 Smart Card Reader, attached to a TKE workstation with the 5.3 level of LIC, supports z10 EC, z10 BC, z9 EC, and z9 BC. However, TKE workstations with 5.0, 5.1, and 5.2 LIC must be upgraded to TKE 5.3 LIC.

TKE additional smart cards - new feature: You have the capability to order Java-based blank smart cards (#0884) which offer a highly efficient cryptographic and data management application built-in to read-only memory for storage of keys, certificates, passwords, applications, and data. The TKE blank smart cards are compliant with FIPS 140-2 Level 2. When you place an order for a quantity of one, you are shipped 10 smart cards.

Capacity on Demand

Changes have been made to enhance the Capacity on Demand (CoD) experience for System z10 EC customers:

- The number of temporary records that can be installed on the Central Processor Complex (CPC) has increased from four to eight.
- Resource tokens are now available for On/Off CoD.

Management of temporary capacity through On/Off CoD is further enhanced through the introduction of resource tokens. For Central Processor (CP) capacity, a resource token

represents an amount of processing capacity that will result in one Million Service Unit (MSU) of software cost for one day - an MSU-day. For specialty engines, a resource token represents activation of one engine of that type for one day - an IFL-day, a zIIP-day, or a zAAP-day. The different resource tokens are contained in separate pools within the On/Off CoD record.

Using the Resource Link ordering process, you determine how many tokens go into each pool. Once On/Off CoD resources are activated, tokens will be decremented from their pools every 24 hours. The amount decremented is based on the highest activation level for that engine type during the previous 24 hours.

Resource tokens are intended to help you bound the hardware costs associated with using On/Off CoD. The use of resource tokens is optional and they are available on either a prepaid or post-paid basis. When prepaid, you are billed for the total amount of resource tokens contained within the On/Off CoD record. When post-paid, the total billing against the On/Off CoD record is limited by the total amount of resource tokens contained within the record.

Resource Link offers an ordering wizard to help you determine how many tokens you need to purchase for different activation scenarios. Resource tokens within an On/Off CoD record may also be replenished. For more information on the use and ordering of resource tokens, refer to the *Capacity on Demand Users Guide* (SC28-6871).

Plan Ahead Memory: Future memory upgrades can now be preplanned to be nondisruptive. The preplanned memory feature will add the necessary physical memory required to support target memory sizes. The granularity of physical memory in the System z10 design is more closely associated with the granularity of logical, entitled memory, leaving little room for growth. If you anticipate an increase in memory requirements, a "target" logical memory size can now be specified in the configuration tool along with a "starting" logical memory size. The configuration tool will then calculate the physical memory required to satisfy this target memory. Should additional physical memory be required, it will be fulfilled with the currently available preplanned memory features.

The preplanned memory feature is offered in 16 gigabyte (GB) increments. The quantity assigned by the configuration tool is the number of 16 GB blocks necessary to increase the physical memory from that required for the "starting" logical memory to the physical memory required for the "target" logical configuration. Activation of any preplanned memory requires the purchase of a preplanned memory activation feature. One preplanned memory activation feature (#1997) is required for each preplanned memory feature (#1996). You now have the flexibility to activate memory to any logical size offered between the starting and target size.

Plan ahead memory is exclusive to System z10 and is transparent to operating systems.

Increased flexibility with z/VM-mode partitions

System z10 EC allows you to define a z/VM-mode partition (LPAR) containing a mix of processor types including Central Processor (CP) and specialty engines - Integrated Facility for Linux (IFL), System z Integrated Information Processor (zIIP), System z Application Assist Processor (zAAP), and Internal Coupling Facility (ICF). With z/VM V5.4 support, this new capability increases flexibility and simplifies systems management by allowing z/VM to manage guests to operate Linux on System z on IFLs, to operate z/VSE™ and z/OS on CPs, to offload z/OS system software overhead, such as DB2® workloads on zIIPs, and to offer an economical Java™ execution environment under z/OS on zAAPs, all in the same z/VM LPAR.

HMC system support

The new functions available on the Hardware Management Console (HMC) version 2.10.1 apply exclusively to System z10. However, the HMC version 2.10.1 will continue to support System z9, zSeries®, and S/390® G5/G6 servers.

The 2.10.1 HMC will continue to support up to two 10 Mbps or 100 Mbps Ethernet LANs. A Token Ring LAN is not supported. The 2.10.1 HMC applications have been updated to support HMC hardware without a diskette drive. DVD-RAM, CD-ROM, and/or USB flash memory drive media will be used.

HMC/SE console messenger: On servers prior to System z9, the remote browser capability was limited to Platform Independent Remote Console (PIRC), with a very small subset of

functionality. Full functionality using Desktop-On-Call (DTC) was limited to one user at a time and was slow, so it was rarely used.

With System z9, full functionality to multiple users was delivered with a fast Web browser solution. You liked this, but requested the ability to communicate to other remote users.

There is now a new console messenger task that offers basic messaging capabilities to allow system operators or administrators to coordinate their activities. The new task may be invoked directly, or using a new option in *Users and Tasks*. This capability is available for HMC and SE local and remote users permitting interactive plain-text communication between two users and also allowing a user to broadcast a plain-text message to all users. This feature is a limited messenger application and does not interact with other messengers.

HMC z/VM tower systems management enhancements: Building upon the previous z/VM systems management support from the Hardware Management Console (HMC), which offered management support for already defined virtual resources, new HMC capabilities are being made available allowing selected virtual resources to be defined. In addition, further enhancements have been made for managing defined virtual resources.

Enhancements are designed to deliver out-of-the-box integrated graphical user interface-based (GUI-based) management of selected parts of z/VM. This is especially targeted to deliver ease-of-use for enterprises new to System z.

This helps to avoid the purchase and installation of additional hardware or software, which may include complicated setup procedures. You can more seamlessly perform hardware and selected operating system management using the HMC Web browser-based user interface.

Enhanced installation support for z/VM using the HMC: HMC version 2.10.1, along with Support Element (SE) version 2.10.1 on z10 EC, now gives you the ability to install Linux on System z in a z/VM virtual machine using the HMC DVD drive. This new function does not require an external network connection between z/VM and the HMC, but instead uses the existing communication path between the HMC and the SE.

Note: This support is intended for environments that have no alternative, such as a LAN-based server, for serving the DVD contents for Linux installations. The elapsed time for installation using the HMC DVD drive can be an order of magnitude, or more, longer than the elapsed time for LAN-based alternatives.

Using the current support and the z/VM support, z/VM can be installed in an LPAR and both z/VM and Linux on System z can be installed in a virtual machine from the HMC DVD drive without requiring an external network setup or a connection between an LPAR and the HMC.

This addresses security concerns and additional configuration efforts using the only other previous solution of the external network connection from the HMC to the z/VM image.

Enhanced installation support using the HMC is exclusive to System z10 and is supported by z/VM. Refer to the *Software requirements* section.

Dynamic add logical CPs without preplanning

- Previously, the Image Profile defined the initial and reserved values for the different processor types for that partition. If those values were not defined prior to partition activation/IPL, they could only be updated by reactivating that partition (including reIPL).
- HMC/SE now offers a task called Logical Processor Add which concurrently can:
 - Increase the "reserved" value for a given processor type (that is, CP, zAAP, zIIP, IFL)
 - Add a new processor type which is not in use yet for that partition
 - Increase the "initial" value for a given processor type

Dynamic add of logical CPs without preplanning is exclusive to System z10 and is supported by z/OS and z/VM. Refer to the [Software requirements](#) section.

Enhanced Driver Maintenance (EDM): There are several reliability, availability, and serviceability (RAS) enhancements that have been made to the HMC/SE based on the feedback from the System z9 Enhanced Driver Maintenance field experience.

- Change to better handle intermittent network issues
- Performance improvements
- New EDM user interface features to allow for you and service personnel to better plan for the EDM
- A new option to check all licensed internal code which can be executed in advance of the EDM preload or activate

FICON enhancements for the storage area network

High Performance FICON for System z (zHPF) - improvement in performance and RAS:

Enhancements have been made to the z/Architecture and the FICON interface architecture to deliver optimizations for online transaction processing (OLTP) workloads. When exploited by the FICON channel, the z/OS operating system, and the control unit, zHPF is designed to help reduce overhead and improve performance.

Additionally, the changes to the architecture offer end-to-end system enhancements to improve reliability, availability, and serviceability (RAS).

zHPF channel programs can be exploited by OLTP I/O workloads - DB2, VSAM, PDSE, and zFS - which transfer small blocks of fixed size data (4K blocks). zHPF implementation by the DS8000™ is exclusively for I/Os that transfer less than a single track of data.

The maximum number of I/Os is designed to be improved up to 100% for small data transfers that can exploit zHPF. Realistic production workloads with a mix of data transfer sizes can see up to 30 to 70% of FICON I/Os utilizing zHPF resulting in up to a 10 to 30% savings in channel utilization. Sequential I/Os transferring less than a single track size (for example, 12x4k bytes/IO) can also benefit.

The FICON Express4 and FICON Express2 features support both the existing FICON protocol and the zHPF protocol concurrently in the server Licensed Internal Code.

High performance FICON is supported by z/OS for DB2, VSAM, PDSE, and zFS applications. Refer to the [Software requirements](#) section. zHPF applies to all FICON Express4 and FICON Express2 features (CHPID type FC) and is exclusive to System z10. Exploitation is required by the control unit.

IBM System Storage™ DS8000 Release 4.1 delivers new capabilities to support High Performance FICON for System z, which can improve FICON I/O throughput on a DS8000 port by up to 100%. The DS8000 series Licensed Machine Code (LMC) level 5.4.1xx.xx (bundle version 64.1.xx.xx), or later, is required.

Platform and name server registration in FICON channel: The FICON channel now delivers the same information to the fabric as is commonly offered by open systems, registering with the name server in the attached FICON directors. With this information, your storage area network (SAN) can be more easily and efficiently managed, enhancing your ability to perform problem determination and analysis.

Registration allows other nodes and/or SAN managers to query the name server to determine what is connected to the fabric and what protocols are supported (FICON, FCP), and to gain information about the System z10 using the attributes that are registered (see following).

The FICON channel is now designed to perform registration with the Fibre Channel's Management Service and Directory Service.

It will register:

- Platforms:
 - Worldwide node name (node name for the platform - same for all channels)
 - Platform type (host computer)
 - Platform name (includes vendor ID, product ID, and vendor-specific data from the node descriptor)
- Channels:
 - Worldwide port name (WWPN)

- Node port identification (N_PORT ID)
- FC-4 types supported (always 0x1B and additionally 0x1C if any Channel-to-Ch (CTC) control units are defined on that channel)
- Classes of service support by the channel

Platform registration is a service defined in the Fibre Channel - Generic Services 4 (FC-GS-4) standard (INCITS (ANSI) T11 group).

Platform and name server registration applies to all of the FICON Express4, FICON Express2, and FICON Express features (CHPID type FC). This support is exclusive to System z10 and is transparent to operating systems.

Local Area Network (LAN) connectivity - a new generation

Additions to the OSA-Express3 family of LAN adapters - 1000BASE-T and 10 GbE SR: The third generation of Open Systems Adapter-Express (OSA-Express3) features were introduced in May 2008 to help reduce latency and overhead, deliver double the port density of OSA-Express2, and deliver increased throughput. At that time we introduced OSA-Express3 10 Gigabit Ethernet LR and Gigabit Ethernet LX for infrastructures with single mode fiber optic cabling and Gigabit Ethernet SX for infrastructures with multimode fiber optic cabling.

The new additions to the OSA-Express3 family of LAN adapters are 1000BASE-T Ethernet, with four ports per feature, for copper cabling infrastructures (Category 5 and Category 6) and a Short Reach (SR) version of 10 Gigabit Ethernet (10 GbE SR), with two ports per feature, to satisfy the requirements of infrastructures with multimode fiber optic cabling.

OSA-Express3 1000BASE-T Ethernet: Each port on the OSA-Express3 1000BASE-T Ethernet four-port feature supports attachment to either a 10BASE-T (10 Mbps), 100BASE-TX (100 Mbps), or 1000BASE-T (1000 Mbps or 1 Gbps) Ethernet Local Area Network (LAN). The feature supports auto-negotiation and automatically adjusts to 10, 100, or 1000 Mbps, depending upon the LAN. OSA-Express3 1000BASE-T Ethernet can be defined as a spanned channel and can be shared among LPARs within and across LCSSs. When configured at 1 Gbps, the 1000BASE-T Ethernet feature operates in full duplex mode only and supports jumbo frames when in QDIO mode (CHPID type OSD).

The OSA-Express3 1000BASE-T Ethernet feature supports the following environments:

CHPID type	Purpose/Traffic
OSC	OSA-Integrated Console Controller (OSA-ICC) TN3270E, non-SNA DFT to IPL CPCS and LPARS operating system console operations
OSD	Queued Direct Input/Output (QDIO) TCP/IP traffic when Layer 3 Protocol-independent when Layer 2
OSE	Non-QDIO, SNA/APPN/HPR and TCP/IP passthru (LCS)
OSN	OSA-Express for Network Control Program (NCP) Supports channel data link control (CDLC) protocol LPAR-to-LPAR communication exclusively; no external communication

For CHPID type OSC note that one port per PCI-E adapter is available for use. CHPID type OSC does not recognize the second port on a PCI-E adapter. Thus, if both CHPIDs on an OSA-Express3 1000BASE-T Ethernet feature are defined as CHPID type OSC, then two of the four ports are recognized.

OSA-Express3 10 GbE SR: The OSA-Express3 10 Gigabit Ethernet (GbE) Short Reach (SR) feature, with two ports per feature, is designed to support attachment to a 10 Gigabits per second (Gbps) Ethernet Local Area Network (LAN) or Ethernet switch capable of 10 Gbps. The connector on the switch must be LC Duplex (or a conversion kit will be required) and the transceiver must be Short Reach (SR), supporting a 50 or 62.5 micron multimode fiber optic infrastructure. The sending and receiving transceivers must be the same (SR-to-SR).

OSA-Express3 10 GbE SR supports CHPID type OSD exclusively. It can be defined as a spanned channel and can be shared among LPARs within and across LCSSs.

CHPID type	Purpose/Traffic
OSD	Queued Direct Input/Output (QDIO) TCP/IP traffic when Layer 3 Protocol-independent when Layer 2

Choose the OSA-Express3 features that best meet your business requirements: To meet the demands of your applications, offer granularity, facilitate redundant paths, and satisfy your infrastructure requirements, there are five features from which to choose. In the 10 GbE environment, Short Reach (SR) is being offered for the first time.

Feature	Feat	Infrastructure	Ports per feature
OSA-Express3 GbE LX	3362	Single mode fiber	4
OSA-Express3 10 GbE LR	3370	Single mode fiber	2
OSA-Express3 GbE SX	3363	Multimode fiber	4
OSA-Express3 10 GbE SR	3371	Multimode fiber	2
OSA-Express3 1000BASE-T	3367	Copper	4

Refer to the [Standards](#) section for the characteristics of each of the features.

OSA-Express3 for reduced latency and improved throughput: Like the OSA-Express3 10 Gigabit Ethernet LR and Gigabit Ethernet LX and SX features delivered in May 2008, OSA-Express3 1000BASE-T Ethernet and 10 Gigabit Ethernet SR help reduce latency and improve throughput with introduction of an Ethernet hardware data router; what was previously done in firmware (packet construction, inspection, and routing) is now performed in hardware. With the Ethernet data router, there is now **direct memory access**; packets flow directly from host memory to the LAN without firmware intervention. OSA-Express3 is also designed to help reduce the round-trip networking time between systems. Up to a 45% reduction in latency at the TCP/IP application layer has been measured.

The OSA-Express3 features are also designed to improve throughput for standard frames (1492 byte) and jumbo frames (8992 byte) to help satisfy the bandwidth requirements of your applications. Up to a 4x improvement has been measured (compared to OSA-Express2).

The above statements are based on OSA-Express3 performance measurements performed in a laboratory environment on a System z10 and do not represent actual field measurements. Results can vary.

OSA-Express3 Ethernet features - Summary of benefits: OSA-Express3 10 GbE LR (single mode fiber), 10 GbE SR (multimode fiber), GbE LX (single mode fiber), GbE SX (multimode fiber), and 1000BASE-T (copper) are designed for use in high-speed enterprise backbones, for local area network connectivity between campuses, to connect server farms to System z10, and to consolidate file servers onto System z10. With reduced latency, improved throughput, and up to 96 ports of LAN connectivity, you can "do more with less."

For the operating system to recognize all four ports on an OSA-Express3 Gigabit Ethernet or 1000BASE-T Ethernet feature, when using QDIO and CHPID type OSD, a new release or PTF is required. If software updates are not applied, only two of the four ports can be exploited and used by the operating system.

The key benefits of OSA-Express3 compared to OSA-Express2 are:

- Reduced latency (up to 45% reduction) and increased throughput (up to 4x) for applications (software changes are not required to exploit the two ports on OSA-Express3 10 GbE or to exploit two of the four ports on OSA-Express3 GbE SX, GbE LX, or 1000BASE-T)
- More physical connectivity to service the network and fewer required resources:

- Fewer CHPIDs to define and manage
- Reduction in the number of required I/O slots
- Potential reduction in the number of I/O cages
- Double the port density of OSA-Express2
- Solution to the requirement for more than 48 LAN ports (now up to 96 ports)

The OSA-Express3 features are exclusive to System z10. Refer to the [Software requirements](#) section for the operating systems supported by each channel path identifier (CHPID) type.

OSA-Express3 support for OSA-Express for NCP: OSA-Express for Network Control Program (NCP), channel path identifier (CHPID) type OSN, is now available for use with the OSA-Express3 GbE features as well as the OSA-Express3 1000BASE-T Ethernet feature.

OSA-Express for NCP, supporting the channel data link control (CDLC) protocol, delivers connectivity between System z operating systems and IBM Communication Controller for Linux (CCL). CCL allows you to keep your business data and applications on the mainframe operating systems while moving NCP functions to Linux on System z.

CCL delivers a foundation to help enterprises simplify their network infrastructure while supporting traditional Systems Network Architecture (SNA) functions such as SNA Network Interconnect (SNI).

Communication Controller for Linux on System z (program number 5724-J38) is the solution for companies that want to help improve network availability by replacing token-ring networks and ESCON® channels with an Ethernet network and integrated LAN adapters on System z10, OSA-Express3 GbE or 1000BASE-T and OSA-Express2 GbE or 1000BASE-T.

Note: For CHPID type OSN, none of the ports are used for external communication. OSA-Express for NCP does not use ports. All communication is LPAR-to-LPAR. Thus, if both CHPIDs on an OSA-Express3 GbE or 1000BASE-T Ethernet feature are defined as CHPID type OSN, then none of the four ports are used for external communication.

OSA-Express for NCP is supported in the z/OS, z/VM, z/VSE, TPF, z/TPF, and Linux on System z environments. Refer to the [Software requirements](#) section.

Historical summary: Functions that continue to be supported by OSA-Express3 and OSA-Express2

- Queued Direct Input/Output (QDIO) - uses memory queues and a signaling protocol to directly exchange data between the OSA microprocessor and the network software for high-speed communication.
 - QDIO Layer 2 (Link layer) - for IP (IPv4, IPv6) or non-IP (AppleTalk, DECnet, IPX, NetBIOS, or SNA) workloads. Using this mode the Open Systems Adapter (OSA) is protocol-independent and Layer-3 independent. Packet forwarding decisions are based upon the Medium Access Control (MAC) address.
 - QDIO Layer 3 (Network or IP layer) - for IP workloads. Packet forwarding decisions are based upon the IP address. All guests share OSA's MAC address.
- Jumbo frames in QDIO mode (8992 byte frame size) when operating at 1 Gbps (fiber or copper) and 10 Gbps (fiber).
- 640 TCP/IP stacks per CHPID - for hosting more images.
- Large send for IPv4 packets - for TCP/IP traffic and CPU efficiency, offloading the TCP segmentation processing from the host TCP/IP stack to the OSA-Express feature.
- Concurrent LIC update - to help minimize the disruption of network traffic during an update; when properly configured, designed to avoid a configuration off or on (applies to CHPID types OSD and OSN).
- Multiple Image Facility (MIF) and spanned channels - for sharing OSA among logical channel subsystems.

OSA-Express QDIO data connection isolation for the z/VM environment on System z10 and System z9: Multi-tier security zones are fast becoming the network configuration standard for new workloads. Therefore, it is essential for workloads (servers/clients) hosted in

a virtualized environment (shared resources) to be protected from intrusion or exposure of data and processes from other workloads.

With Queued Direct Input/Output (QDIO) data connection isolation you:

- Have the ability to adhere to security and HIPAA-security guidelines and regulations for network isolation between the operating system instances sharing physical network connectivity.
- Can establish security zone boundaries that have been defined by your network administrators.
- Have a mechanism to isolate a QDIO data connection (on an OSA port), ensuring all internal OSA routing between the isolated QDIO data connections and all other sharing QDIO data connections are disabled. In this state, only external communications to and from the isolated QDIO data connection are allowed. If you choose to deploy an external firewall to control the access between hosts on an isolated virtual switch and sharing LPARs then an external firewall needs to be configured and each individual host or LPAR must have a route added to their TCP/IP stack to forward local traffic to the firewall.

Internal "routing" can be disabled on a per QDIO connection basis. This support does not affect the ability to share an OSA-Express port. Sharing occurs as it does today, but the ability to communicate between sharing QDIO data connections may be restricted through the use of this support. You decide whether an operating system's or z/VM's Virtual Switch OSA-Express QDIO connection is to be *nonisolated* (default) or *isolated*.

Note: QDIO data connection isolation applies to the device statement defined at the operating system level. While an OSA-Express CHPID may be shared by an operating system, the data device is not shared.

QDIO data connection isolation applies to the z/VM environment and to all of the OSA-Express3 and OSA-Express2 features (CHPID type OSD) on System z10 and to the OSA-Express2 features on System z9. Refer to the [Software requirements](#) section.

InfiniBand coupling links for Parallel Sysplex

Introducing Long Reach 1x InfiniBand coupling links

Now, InfiniBand can be used for Parallel Sysplex® coupling and STP communication at unrepeated distances up to 10 km (6.2 miles) and greater distances when attached to qualified optical networking solutions. InfiniBand coupling links supporting extended distance is referred to as Long Reach 1x (one pair of fiber) InfiniBand.

- Long Reach 1x InfiniBand coupling links support single data rate (SDR) at 2.5 gigabits per second (Gbps) when connected to a DWDM capable of SDR (1x IB-SDR).
- Long Reach 1x InfiniBand coupling links support double data rate (DDR) at 5 Gbps when connected to a DWDM capable of DDR (1x IB-DDR).

The link data rate will auto-negotiate from SDR to DDR depending upon the capability of the attached equipment.

Long Reach 1x InfiniBand coupling links utilize the Host Channel Adapter2 Optical Long Reach fanout card (HCA2-O LR #0168). Like the 12x InfiniBand coupling link feature (HCA2-O fanout card #0163), the HCA2-O LR fanout card can also be used to exchange timekeeping messages for Server Time Protocol (STP).

This environment supports use of 9 micron single mode fiber optic cables with LC Duplex connectors, the same fiber optic cable you have been using with InterSystem Channel-3 (ISC-3).

There is no change to the Channel Path Identifier (CHPID). It remains CHPID type CIB whether 12x IB-SDR or DDR or 1x IB-SDR or DDR. HCA2-O LR fanout cards supporting Long Reach 1x InfiniBand are exclusive to System z10 and are supported by z/OS and by z/VM to define, modify, and delete an InfiniBand coupling link, CHPID type CIB, when z/VM is the controlling LPAR for dynamic I/O. Refer to the [Software requirements](#) section.

You now have five coupling link options for communication in a Parallel Sysplex environment:

1. Internal Coupling Channels (ICs) can be used for internal communication between Coupling Facilities (CFs) defined in LPARs and z/OS images on the same server.
2. Integrated Cluster Bus-4 (ICB-4) is for short distances. ICB-4 links use 10 meter (33 feet) copper cables, of which 3 meters (10 feet) is used for internal routing and strain relief. ICB-4 is used to connect z10 EC-to-z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890. **Note:** If connecting to a z10 BC or a z9 BC with ICB-4, those servers cannot be installed with the nonraised floor feature. Also, if the z10 BC is ordered with the nonraised floor feature, ICB-4 cannot be ordered.
3. 12x InfiniBand coupling links (12x IB-SDR or 12x IB-DDR) offer an alternative to ISC-3 in the data center and facilitate coupling link consolidation. Physical links can be shared by multiple operating system images or Coupling Facility images on a single system. The 12x InfiniBand links support distances up to 150 meters (492 feet) using industry-standard OM3 50 micron multimode fiber optic cables.
4. Long Reach 1x InfiniBand coupling links (1x IB-SDR or 1x IB-DDR) are an alternative to ISC-3 and offer greater distances with support for point-to-point unrepeated distances up to 10 km (6.2 miles) using 9 micron single mode fiber optic cables. Greater distances can be supported with System z-qualified optical networking solutions. Long Reach 1x InfiniBand coupling links support the same sharing capabilities as the 12x InfiniBand version, allowing one physical link to be shared by multiple operating system images or Coupling Facility images on a single system.
5. InterSystem Channel-3 (ISC-3) supports communication at unrepeated distances up to 10 km (6.2 miles) using 9 micron single mode fiber optic cables and greater distances with System z-qualified optical networking solutions. ISC-3s are supported exclusively in peer mode (CHPID type CFP).

Note: The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload. Specifically, with 12x InfiniBand coupling links, while the link data rate is higher than that of ICB, the service times of coupling operations are greater, and the actual throughput is less.

Refer to the *Coupling Facility Configuration Options* whitepaper for a more specific explanation of when to continue using the current ICB or ISC-3 technology versus migrating to InfiniBand coupling links.

The whitepaper is available at

<http://www.ibm.com/systems/z/advantages/ps0/whitepaper.html>

Coupling Facility Control Code Level 16

Improved service time with Coupling Facility Duplexing enhancements: Prior to Coupling Facility Control Code (CFCC) Level 16, System-Managed Coupling Facility (CF) Structure Duplexing required two duplexing protocol exchanges to occur synchronously during processing of each duplexed structure request. CFCC Level 16 allows one of these protocol exchanges to complete asynchronously. This allows faster duplexed request service time, with more benefits when the Coupling Facilities are further apart, such as in a multi-site Parallel Sysplex.

List notification improvements: Prior to CFCC Level 16, when a shared queue (subsidiary list) changed state from empty to nonempty, the CF would notify ALL active connectors. The first one to respond would process the new message, but when the others tried to do the same, they would find nothing, incurring additional overhead.

CFCC Level 16 can help improve the efficiency of coupling communications for IMS™ Shared Queue and WebSphere® MQ Shared Queue environments. The Coupling Facility notifies only one connector in a sequential fashion. If the shared queue is processed within a fixed period of time, the other connectors do not need to be notified, saving the cost of the false scheduling. If the shared queue is not read within the time limit, then the other connectors are notified as they were prior to CFCC Level 16.

When migrating to new CF levels, lock, list, and cache structure sizes might need to be increased to support new function. For example, when you upgrade from CFCC Level 15 to Level 16 the required size of the structure might increase. This adjustment can have an impact when the system allocates structures or copies structures from one coupling facility to another at different CF levels.

The coupling facility structure sizer tool can size structures for you and takes into account the amount of space needed for the current CFCC levels.

Access the tool at

<http://www.ibm.com/servers/eserver/zseries/cfsizer/>

CFCC Level 16 is exclusive to System z10 and is supported by z/OS and z/VM for guest exploitation. Refer to the [Software requirements](#) section.

Server Time Protocol (STP) enhancements

STP is a message-based protocol in which timekeeping information is transmitted between servers over externally defined coupling links. ICB-4, ISC-3, and InfiniBand coupling links can be used to transport STP messages.

The following STP enhancements are available on System z10 and System z9 servers. STP feature (#1021) and the latest Machine Change Levels are required.

Enhanced Network Time Protocol (NTP) client support: This enhancement addresses the requirements for those who need to provide the same accurate time across heterogeneous platforms in an enterprise.

The STP design has been enhanced to include support for a Simple Network Time Protocol (SNTP) client on the Support Element. By configuring an NTP server as the STP External Time Source (ETS), the time of an STP-only Coordinated Timing Network (CTN) can track to the time provided by the NTP server, and maintain a time accuracy of 100 milliseconds.

Note: NTP client support has been available since October 2007.

Enhanced accuracy to an External Time Source: The time accuracy of an STP-only CTN has been improved by adding the capability to configure an NTP server that has a pulse per second (PPS) output signal as the ETS device. This type of ETS device is available worldwide from several vendors that provide network timing solutions.

STP has been designed to track to the highly stable, accurate PPS signal from the NTP server, and maintain an accuracy of 10 microseconds as measured at the PPS input of the System z server. A number of variables such as accuracy of the NTP server to its time source (GPS, radio signals for example), and cable used to connect the PPS signal will determine the ultimate accuracy of STP to Coordinated Universal Time (UTC).

In comparison, the IBM Sysplex Timer® is designed to maintain an accuracy of 100 microseconds when attached to an ETS with a PPS output. If STP is configured to use a dial-out time service or an NTP server without PPS, it is designed to provide a time accuracy of 100 milliseconds to the ETS device.

For this enhancement, the NTP output of the NTP server has to be connected to the Support Element (SE) LAN, and the PPS output of the same NTP server has to be connected to the PPS input provided on the External Time Reference (ETR) feature of the System z10 or System z9 server.

Continuous availability of NTP servers used as External Time Source: Improved External Time Source (ETS) availability can now be provided if you configure different NTP servers for the Preferred Time Server (PTS) and the Backup Time Server (BTS). Only the PTS or the BTS can be the Current Time Server (CTS) in an STP-only CTN. Prior to this enhancement, only the CTS calculated the time adjustments necessary to maintain time accuracy. With this enhancement, if the PTS/CTS cannot access the NTP server or the pulse per second (PPS) signal from the NTP server, the BTS, if configured to a different NTP server, may be able to calculate the adjustment required and propagate it to the PTS/CTS. The PTS/CTS in turn will perform the necessary time adjustment steering.

This avoids a manual reconfiguration of the BTS to be the CTS, if the PTS/CTS is not able to access its ETS. In an ETR network when the primary Sysplex Timer is not able to access the ETS device, the secondary Sysplex Timer takes over the role of the primary - a recovery action

not always accepted by some environments. The STP design provides continuous availability of ETS while maintaining the special roles of PTS and BTS assigned by the enterprise.

The improvement is available when the ETS is configured as an NTP server or an NTP server using PPS.

NTP server on Hardware Management Console (HMC): Improved security can be obtained by providing NTP server support on the HMC. If an NTP server (with or without PPS) is configured as the ETS device for STP, it needs to be attached directly to the Support Element (SE) LAN. The SE LAN is considered by many users to be a private dedicated LAN to be kept as isolated as possible from the intranet or Internet.

Since the HMC is normally attached to the SE LAN, providing an NTP server capability on the HMC addresses the potential security concerns most users may have for attaching NTP servers to the SE LAN. The HMC, using a separate LAN connection, can access an NTP server available either on the intranet or Internet for its time source. Note that when using the HMC as the NTP server, there is no pulse per second capability available. Therefore, you should not configure the ETS to be an NTP server using PPS.

Enhanced STP recovery when Internal Battery Feature is in use: Improved availability can be obtained when power has failed for a single server (PTS/CTS), or when there is a site power outage in a multisite configuration where the PTS/CTS is installed (the site with the BTS is a different site not affected by the power outage).

If an Internal Battery Feature (IBF) is installed on your System z server, STP now has the capability of receiving notification that customer power has failed and that the IBF is engaged. When STP receives this notification from a server that has the role of the PTS/CTS, STP can automatically reassign the role of the CTS to the BTS, thus automating the recovery action and improving availability.

STP configuration and time information saved across Power-on-Resets (POR) or power outages: This enhancement delivers system management improvements by saving the STP configuration across PORs and power failures for a single server STP-only CTN. Previously, if there was a POR of the server or the server experienced a power outage, the time and assignment of the PTS and CTS roles would have to be reinitialized. You will no longer need to reinitialize the time or reassign the role of PTS/CTS across POR or power outage events.

Note: This enhancement is also available on the z990 and z890 servers, in addition to System z10 and System z9 servers.

Application Programming Interface (API) to automate STP CTN reconfiguration: The concept of "a pair and a spare" has been around since the original Sysplex Couple Data Sets (CDSs). If the primary CDS becomes unavailable, the backup CDS would take over. Many sites have had automation routines bring a new backup CDS online to avoid a single point of failure. This idea is being extended to STP. With this enhancement, if the PTS fails and the BTS takes over as CTS, an API is now available on the HMC so you can automate the reassignment of the PTS, BTS, and Arbiter roles. This can improve availability by avoiding a single point of failure after the BTS has taken over as the CTS.

Prior to this enhancement, the PTS, BTS, and Arbiter roles had to be reassigned manually using the System (Sysplex) Time task on the HMC. For additional details on the API, refer to *System z Application Programming Interfaces* (SB10-7030-11).

Additional information is available on the STP Web page

<http://www.ibm.com/systems/z/pso/stp.html>

Additional information is also available from the following Redbooks® available at the Web site

<http://www.redbooks.ibm.com/>

- Server Time Protocol Planning Guide, SG24-7280
- Server Time Protocol Implementation Guide, SG24-7281

Accessibility by people with disabilities

Product positioning

The future begins today - The z10 EC was introduced in February 2008 and represents an evolution and a revolution in design. At the heart of the server is the 4.4 GHz quad core processor chip which targets high frequency. This new processor chip allows expanded scalability, and when combined with the larger memory capability of the z10 EC, sub-capacity options, a new host bus interface, and the premier virtualization of z/VM, the z10 EC is perfectly positioned to offer greater growth to enable large-scale consolidation.

Built on a foundation that improves recovery for unplanned outages and reduction in planned outages, the z10 EC can also help reduce preplanning requirements by reserving a fixed Hardware System Area (HSA). On demand capacity gets a boost with the new just-in-time deployment of resources that can help to rapidly respond to changing business requirements. Plan Ahead Memory offers the option to plan for nondisruptive memory upgrades.

IBM introduced the z10 EC as the cornerstone of the new enterprise data center. An evolutionary model for the delivery of efficient IT, the new enterprise data center is designed to deliver new economics through the simplification of IT, provide rapid deployment of services with improved manageability, and enable tight alignment of business and IT to support innovation. The technology evolution of System z that delivers the z10 EC allows it to be uniquely positioned to hold the role as the cornerstone in this new enterprise data center. Its core strengths - reliability, availability, serviceability, security, scalability, and virtualization - are key enabling technologies that underpin the z10 EC in this role.

Enhancements today strengthen the z10 EC's role in the future. High Performance FICON for System z (zHPF) can help you connect to data faster, enhancements to Crypto Express2 can help provide stronger protection of data, and OSA-Express3 improvements can help provide access to data in the network. Server Time Protocol (STP) delivers enhanced time accuracy to an External Time Source and offers support for heterogeneous platforms in an enterprise to track to the same time source. Responding to the on demand needs of a 24x7 global economy is improved with just-in-time capacity on the server.

Delivering technologies that businesses need today for running traditional workloads, implementing modern applications, and consolidating distributed servers for better economics, the z10 EC is the server best qualified to protect and support future business growth.

Statement of general direction

System z9 to support 12x IB-SDR on z9 EC and z9 BC general purpose servers: IBM intends to expand the support for connecting a System z10 with 12x InfiniBand Single Data Rate (12x IB-SDR) to z9 EC and z9 BC beyond dedicated Coupling Facilities (ICFs only), to include support for System z9 general purpose servers. A System z10 will then be able to connect with 12x InfiniBand coupling links to a System z9 server whether the System z9 has Processor Units (PUs) that are defined as Central Processors (CPs), Integrated Facilities for Linux (IFLs), System z9 Application Assist Processors (zAAPs), System z9 Integrated Information Processors (zIIPs), or System Assist Processors (SAPs).

Note: Support of 12x InfiniBand Single Data Rate coupling links on System z9 for attachment to System z10 are for migration purposes only. IBM does not intend to support 12x InfiniBand coupling links to connect two System z9 servers. The new Long Reach 1x InfiniBand coupling links are exclusive to System z10 and will not be offered on System z9.

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.

Product number

Description	Machine type	Model	Feature	Init/MES
System z10 EC	2097	E12 E26 E40 E56 E64		
HMC w/Dual EN (see Note 1)			0090	Both
HCA2-0 LR Fanout Card			0168	Both
TKE 5.3 LIC			0854	Both
TKE addl smart cards			0884	Both
TKE Smart Card Reader			0885	Both
OSA-Express3 1000BASE-T			3367	Both
OSA-Express3 10 GbE SR			3371	Both
LC Dup 5ft harness			7958	Both
1 MSU-day			9917	MES
100 MSU-days			9918	MES
10000 MSU-days			9919	MES
1 IFL-day			9920	MES
100 IFL-day			9921	MES
1 ICF-day			9922	MES
100 ICF-day			9923	MES
1 zIIP-day			9924	MES
100 zIIP-day			9925	MES
1 zAAP-day			9926	MES
100 zAAP-day			9927	MES
1 SAP-day			9928	MES
100 SAP-day			9929	MES

Description	Machine type	Model	Feature	Init/MES
System z9 EC	2094	S08 S18 S28 S38 S54		
HMC w/Dual EN (see Note 1)			0090	Both
TKE 5.3 LIC			0854	Both
TKE addl smart cards			0884	Both
TKE Smart Card Reader			0885	Both

Description	Machine type	Model	Feature	Init/MES
System z9 BC	2096	R07 S07		
HMC w/Dual EN (see Note 1)			0090	Both
TKE 5.3 LIC			0854	Both
TKE addl smart cards			0884	Both
TKE Smart Card Reader			0885	Both

If field installed on a purchased machine, parts removed or replaced become the property of IBM and must be returned.

Note 1: Feature #0090 to be available on z10 EC, z9 EC, and z9 BC first half 2009.

Education support

Visit the following Web site for additional information

<http://www.ibm.com/training/us>

Contact your IBM representative for course information.

Publications

The following publications are available in the *Library* section of Resource Link:

Title	Order number
z10 EC System Overview	SA22-1084
z10 EC Installation Manual - Physical Planning (IMPP)	GC28-6865
z10 EC PR/SM™ Planning Guide	SB10-7153

The following publications are shipped with the product and available in the *Library* section of Resource Link:

Title	Order number
z10 EC Installation Manual	GC28-6864
z10 EC Service Guide	GC28-6866
z10 EC Safety Inspection Guide	GC28-6870
System Safety Notices	G229-9054

The following publications will be available at planned availability in the *Library* section of Resource Link:

Title	Order number
Agreement for Licensed Machine Code	SC28-6872
Application Programming Interfaces for Java	API-JAVA
Application Programming Interfaces	SB10-7030
Capacity on Demand User's Guide	SC28-6871
CHPID Mapping Tool User's Guide	GC28-6825
Common Information Model (CIM) Management Interface	SB10-7154
Hardware Management Console Operations Guide (V2.10.1)	SC28-6873
IOCP User's Guide	SB10-7037
SCSI IPL - Machine Loader Messages	SC28-6839
Service Guide for HMCs and SEs	GC28-6861
Service Guide for Trusted Key Entry Workstations	GC28-6862
Standalone IOCP User's Guide	SB10-7152
Statement of Limited warranty	GC28-6863
Support Element Operations Guide (Version 2.10.1)	SC28-6879
System z Functional Matrix	ZSW0-1335
OSA-Express Customer's Guide	SA22-7935
OSA-ICC User's Guide	SA22-7990

Publications for System z10 Enterprise Class can be obtained at Resource Link by accessing the following Web site

<http://www.ibm.com/servers/resourcelink>

Using the instructions on the Resource Link panels, obtain a user ID and password. Resource Link has been designed for easy access and navigation.

Updates to the following Redbook publications are available now:

Title	Order number
Server Time Protocol: Planning Guide	SG24-7280
Server Time Protocol: Implementation Guide	SG24-7281
Getting Started with InfiniBand on System z10 and System z9	SG24-7539

System z10 Enterprise Class Technical Introduction	SG24-7515
System z10 Enterprise Class Technical Guide	SG24-7516
System z Connectivity Handbook	SG24-5444
System z10 Enterprise Class Capacity on Demand	SG24-7504
OSA-Express Implementation Guide	SG24-5948

For other IBM Redbooks publications, refer to

<http://www.redbooks.ibm.com/>

Services

Global Technology Services

IBM services include business consulting, outsourcing, hosting services, applications, and other technology management.

These services help you learn about, plan, install, manage, or optimize your IT infrastructure to be an On Demand Business. They can help you integrate your high-speed networks, storage systems, application servers, wireless protocols, and an array of platforms, middleware, and communications software for IBM and many non-IBM offerings. IBM is your one-stop shop for IT support needs.

For details on available services, contact your IBM representative or visit

<http://www.ibm.com/services/>

For details on available IBM Business Continuity and Recovery Services, contact your IBM representative or visit

<http://www.ibm.com/services/continuity>

For details on education offerings related to specific products, visit

<http://www.ibm.com/services/learning/index.html>

Select your country, and then select the product as the category.

Technical information

EMC conformance

- Taiwan BSMI CNS13438 (Taiwan EMC Standard)
- ACA C-Tick (Australia and New Zealand)

Specified operating environment

Standards

The OSA-Express3 1000BASE-T Ethernet and 10 Gigabit Ethernet Short Reach features are designed to conform to the following standards:

- Ethernet (10BASE-T)
 - IEEE 802.2 Logical Link Control Protocol
 - IEEE 802.3 CSMA/CD Protocol
 - ISO/IEC 8802-3
 - DIX Version 2
- Ethernet (100BASE-TX)
 - IEEE 802.3u CSMA/CD Protocol
- Ethernet (1000BASE-T)
 - IEEE 802.1p
 - IEEE 802.1q

- IEEE 802.3ab
- IEEE 802.3ac
- IEEE 802.3u
- IEEE 802.3x
- 10 Gigabit Ethernet SR (10GBASE-SR)
 - IEEE 802.3ae
 - IEEE 802.1q
 - IEEE 802.3x - flow control
 - DIX Version 2

The OSA-Express3 features have the following characteristics:

OSA-Express3 1000BASE-T Ethernet (#3367)

- Data rate: 10, 100, or 1000 Mbps
- Operating modes: Auto-negotiate, half duplex, full duplex
- Defined as: CHPID types OSC, OSD, OSE, OSN
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: RJ-45
- Port count: Four 1000BASE-T ports per feature
- Cable type: EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cable with a maximum length of 100 meters (328 feet)

OSA-Express3 10 Gigabit Ethernet SR (#3371)

- Data rate: 10 Gigabits per second (10 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID type OSD
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: Two Short Reach (SR) ports per feature
- Cable type: Multimode fiber optic cabling (50 or 62.5 micron)
- Unrepeated distance:
 - With 50 micron fiber at 2000 MHz-km: 300 meters (984 feet)
 - With 50 micron fiber at 500 MHz-km: 82 meters (269 feet)
 - With 62.5 micron fiber at 200 MHz-km: 33 meters (108 feet)

Hardware requirements

You should review the PSP buckets for minimum Machine Change Levels (MCLs) and software PTF levels before IPLing operating systems. To support new functions and features, MCLs are required.

Descriptions of the MCLs are available now through Resource Link

Access Resource Link at

<http://www.ibm.com/servers/resourcelink>

Peripheral hardware and device attachments: IBM devices previously attached to IBM System z9 and zSeries servers are supported for attachment to System z10 EC channels, unless otherwise noted. The subject I/O devices must meet ESCON or FICON/FCP architecture requirements to be supported. I/O devices that meet OEMI architecture requirements are supported only using an external converter. Prerequisite Engineering Change Levels may be required. For further detail, contact IBM service personnel. While the z10 EC supports devices as described above, IBM does not commit to provide support or service for an IBM device that has reached its End of Service effective date as announced by IBM.

Note: IBM cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions regarding the capabilities of non-IBM products should be addressed to the suppliers of those products. For a list of the current supported FICON devices, refer to the following Web site

<http://www.ibm.com/systems/z/connectivity/>

Software requirements

Listed are the operating system minimum versions and releases. Select the releases appropriate to your operating system environments.

Note: Refer to the z/OS, z/VM, z/VSE subsets of the 2097DEVICE Preventive Service Planning (PSP) bucket prior to installing a z10 EC.

IBM Lifecycle Extension for z/OS V1.7: z/OS V1.7 support was withdrawn September 30, 2008. The Lifecycle Extension for z/OS V1.7 (5637-A01) makes fee-based corrective service for z/OS V1.7 available through September 2010. With the Lifecycle Extension, z/OS V1.7 supports the z10 EC server. Certain functions and features of the z10 EC require later releases of z/OS. For the complete list of software support, see the PSP buckets. For more information on the IBM Lifecycle Extension for z/OS V1.7 Refer to Software Announcement [AP08-0274](#), dated August 12, 2008.

CPU measurement facility on z10 EC requires at a minimum:

- z/OS V1.8, V1.9, or V1.10 with PTFs.

Secure Key AES on z10 EC requires at a minimum:

- z/OS V1.8, z/OS V1.9, or z/OS V1.10 with the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/VM V5.2 for guest exploitation.

Support for 13- through 19-digit Personal Account Numbers on z10 EC requires at a minimum:

- z/OS V1.8, z/OS V1.9, or z/OS V1.10 with the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/VM V5.2 for guest exploitation.

Enhanced installation support for z/VM using the HMC on z10 EC requires at a minimum:

- z/VM V5.4.

Dynamic add logical CPs without preplanning on z10 EC requires at a minimum:

- z/OS V1.10.
- z/VM V5.4.
- z/VM V5.3 with PTFs.

High Performance FICON (zHPF) support for CHPID type FC on z10 EC requires at a minimum:

- z/OS V1.8, V1.9, or V1.10 with PTFs.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs.

OSA-Express3 1000BASE-T (#3367) on z10 EC requires at minimum:

For CHPID type OSC supporting TN3270E and non-SNA DFT:

Note: One port per PCI-E adapter is available for use. CHPID type OSC does not recognize the second port on a PCI-E adapter.

- z/OS V1.8.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/VM V5.2.
- z/VSE V3.1.
- TPF 4.1 and z/TPF 1.1.

For CHPID type OSD and **exploitation of four ports per feature:**

- z/OS V1.10.
- z/OS V1.8 or V1.9 with PTFs.
- z/VM V5.4.
- z/VM V5.2 or V5.3 with PTFs.
- z/VSE V4.1 with PTFs.
- z/TPF 1.1 PUT 4 with APARs.
- Linux on System z distributions:
 - Novell SUSE SLES 10 SP2.
 - Red Hat RHEL 5.2.

For CHPID type OSD and use of one port per PCI-E adapter, two of the four ports per feature:

- z/OS V1.8.
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- TPF V4.1 at PUT 13 with PTF.
- z/TPF 1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 9 and SLES 10.
 - Red Hat RHEL 4 and RHEL 5.

For CHPID type OSE and support of 4 or 2 ports per feature:

- z/OS V1.8.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/VM V5.2.
- z/VSE V3.1.

For CHPID type OSN:

Note: CHPID type OSN does not use ports. All communication is LPAR-to-LPAR.

- z/OS V1.8.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/VM V5.2.
- z/VSE 3.1 with PTFs.
- TPF 4.1.
- z/TPF 1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 9 SP2 and SLES 10.

- Red Hat RHEL 4.3 and RHEL 5.

OSA-Express3 10 GbE SR (#3371) on z10 EC requires at a minimum:

Supporting CHPID type OSD and two ports per feature:

- z/OS V1.8.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- TPF 4.1 at PUT 13 with PTFs.
- z/TPF 1.1.
- Linux on System z distributions:
 - Novell SUSE SLES 9 and SLES 10.
 - Red Hat RHEL 4 and RHEL 5.

OSA-Express QDIO data connection isolation on z10 EC supporting OSA-Express3 and OSA-Express2 and z9 EC supporting OSA-Express2 (CHPID type OSD) requires at a minimum:

- z/VM V5.3 with PTFs.

HCA2-O LR fanout (#0168) supporting Long Reach 1x InfiniBand coupling links on z10 EC requires at a minimum:

- z/OS V1.8, V1.9, or V1.10 with PTFs.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs.
- z/VM 5.3 to define, modify, and delete an InfiniBand coupling link, CHPID type CIB, when z/VM is the controlling LPAR for dynamic I/O.

Coupling Facility Control Code Level 16 on z10 EC requires at a minimum:

- z/OS V1.8, V1.9, or V1.10 with PTFs.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs.
- z/VM V5.2 for guest exploitation with PTFs.

Planning information

Customer responsibilities

Information on customer responsibilities for site preparation can be found in the Library section of Resource Link at

<http://www.ibm.com/servers/resourcelink>

Cable orders

Fiber optic cable orders

Fiber optic cables for the z10 EC, z10 BC, z9 EC, z990, z9 BC, and z890 are available from IBM Site and Facilities Services.

IBM Site and Facilities Services has a comprehensive set of scalable solutions to address IBM cabling requirements, from product-level to enterprise-level. The IBM Facilities Cabling Services - fiber transport system and the IBM IT Facilities Assessment, Design, and Construction Services - optimized airflow assessment for cabling, offered by IBM Site and Facilities Services provides services for small, medium, and large enterprises:

- Assessment and planning for IBM Fiber Transport System (FTS) trunking components
- Planning and installation services for individual fiber optic connections

IBM Global Technology Services has the expertise and personnel available to effectively plan and deploy the appropriate cabling with the future in mind. These services may include assessment, planning, consultation, cable selection, installation, and documentation, depending upon the services selected.

These services are designed to be right-sized for your products or the end-to-end enterprise, and to take into consideration the requirements for all of the protocols and media types supported on the System z10 EC, System z9, and zSeries (for example, ESCON, FICON, Coupling Links, OSA) whether the focus is the data center, the Storage Area Network (SAN), the Local Area Network (LAN), or the end-to-end enterprise.

IBM Site and Facilities Services are designed to deliver convenient, packaged services to help reduce the complexity of planning, ordering, and installing fiber optic cables. The appropriate fiber cabling is selected based upon the product requirements and the installed fiber plant.

The services are packaged as follows:

Under **IBM Facilities Cabling Services** there is the option to provide IBM Fiber Transport System (FTS) trunking commodities (fiber optic trunk cables, fiber harnesses, panel-mount boxes) for connecting to the z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890. IBM can reduce the cable clutter and cable bulk under the floor. An analysis of the channel configuration and any existing fiber optic cabling is performed to determine the required FTS trunking commodities. IBM can also help organize the entire enterprise. This option includes enterprise planning, new cables, fiber optic trunking commodities, installation, and documentation.

Under IBM IT Facilities Assessment, Design, and Construction Services there is the option to provide the Optimized Airflow Assessment for Cabling to provide you with a comprehensive review of your existing data center cabling infrastructure. This service provides an expert analysis of the overall cabling design required to help improve data center airflow for optimized cooling, and to facilitate operational efficiency through simplified change management.

See the [Cable orders](#) section of this announcement for more information or contact IBM Global Technology Services for details.

Refer to the services section of Resource Link for further details. Access Resource Link at

<http://www.ibm.com/servers/resourcelink>

Cabling responsibilities

Fiber optic cables, cable planning, labeling, and placement are all customer responsibilities for new installations and upgrades. Fiber optic conversion kits and Mode Conditioning Patch (MCP) cables are not orderable as features on a z10 EC. Installation Planning Representatives (IPRs) and System Service Representatives (SSRs) will not perform the fiber optic cabling tasks without a services contract.

The following tasks are required to be performed by the customer prior to machine installation:

- All fiber optic cable planning.
- All purchasing of correct fiber optic cables.
- All installation of any required Mode Conditioning Patch (MCP) cables.
- All installation of any required Conversion Kits.
- All routing of fiber optic cables to correct floor cutouts for proper installation to server.
 - Use the Physical Channel Identifier (PCHID) report or the report from the Channel Path Identifier (CHPID) Mapping Tool to accurately route all cables.
- All labeling of fiber optic cables with PCHID numbers for proper installation to server.
 - Use the PCHID report or the report from the CHPID Mapping Tool to accurately label all cables.

Additional service charges may be incurred during the server installation if the above cabling tasks are not accomplished as required.

Fiber Quick Connect (FQC), a fiber harness integrated in the z10 EC frame for "quick" connect, is offered as a feature on the z10 EC for connection to ESCON and FICON LX channels.

Cables for ICB links continue to be available as features. Refer to the *Special features* section of the *Sales Manual* on the Web for a list of these features and cables for ICB links.

<http://www.ibm.com/common/ssi/OIX.wss>

For further details also refer to the *Installation Manual for Physical Planning* (IMPP), available on Resource Link.

Note: IBM Site and Facilities Services can satisfy your fiber optic as well as your copper cabling requirements.

Security, auditability, and control

The z10 EC uses the security and auditability features and functions of host hardware, host software, and application software.

The customer is responsible for evaluation, selection, and implementation of security features, administrative procedures, and appropriate controls in application systems and communications facilities.

IBM Electronic Services

IBM has transformed its delivery of hardware and software support services to help you achieve higher system availability. Electronic Services is a Web-enabled solution that offers an exclusive, no-additional-charge enhancement to the service and support available for IBM servers. These services are designed to provide the opportunity for greater system availability with faster problem resolution and preemptive monitoring. Electronic Services comprises two separate, but complementary, elements: Electronic Services news page and Electronic Services Agent.

The Electronic Services news page is a single Internet entry point that replaces the multiple entry points traditionally used to access IBM Internet services and support. The news page enables you to gain easier access to IBM resources for assistance in resolving technical problems.

The Electronic Service Agent™ is no-additional-charge software that resides on your server. It monitors events and transmits system inventory information to IBM on a periodic, client-defined timetable. The Electronic Service Agent automatically reports hardware problems to IBM. Early knowledge about potential problems enables IBM to deliver proactive service that may result in higher system availability and performance. In addition, information collected through the Service Agent is made available to IBM service support representatives when they help answer your questions or diagnose problems. Installation and use of IBM Electronic Service Agent for problem reporting enables IBM to provide better support and service for your IBM server.

To learn how Electronic Services can work for you, visit

<http://www.ibm.com/support/electronic>

Terms and conditions

MES discount applicable

No

Field installable feature

Yes

Warranty period

One year

This feature or features assume the same warranty or maintenance terms as the machine in which it is installed for the full warranty or maintenance period announced for such machine.

Customer setup

No

Machine code

Same license terms and conditions as base machine

Prices

For all local charges, contact your IBM representative.

AP distribution

Country/Region	Announced
AP IOT	
ASEAN*	Yes
India/South Asia**	Yes
Australia	Yes
People's Republic of China	Yes
Hong Kong S.A.R of the PRC	Yes
Macao S.A.R of the PRC	Yes
Taiwan	Yes
Korea	Yes
New Zealand	Yes
Japan IOT	
Japan	Yes

* Brunei Darussalam, Indonesia, Cambodia, Lao People's Democratic Republic, Malaysia, Philippines, Singapore, Thailand, and Vietnam

** Bangladesh, Bhutan, India, Sri Lanka, Maldives, Nepal, and Afghanistan

Trademarks

z9, Resource Link, z/VSE, DS8000, System Storage, IMS, PR/SM and Electronic Service Agent are trademarks of IBM Corporation in the United States, other countries, or both.

System z, IBM, FICON, z/VM, System z9, z/Architecture, z/OS, DB2, zSeries, S/390, ESCON, Parallel Sysplex, WebSphere, Sysplex Timer and Redbooks are registered trademarks of IBM Corporation in the United States, other countries, or both.

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

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

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

Terms of use

IBM products and services which are announced and available in your country can be ordered under the applicable standard agreements, terms, conditions, and prices in effect at the time. IBM reserves the right to modify or withdraw this announcement at any time without notice. This announcement is provided for your information only. Reference to other products in this announcement does not necessarily imply those products are announced, or intend to be announced, in your country. Additional terms of use are located at:

<http://www.ibm.com/legal/us/en/>

For the most current information regarding IBM products, consult your IBM representative or reseller, or visit the IBM worldwide contacts page

<http://www.ibm.com/planetwide/>