



IBM zEnterprise 196 enhancements deliver faster access to data

Table of contents

2 Overview	18 Product number
3 Key prerequisites	18 Publications
3 Planned availability date	20 Technical information
3 Description	28 Terms and conditions
16 Product positioning	29 Pricing
16 Statement of general direction	29 Announcement countries

At a glance

The world today is changing and people are more tech-savvy. Every aspect of life is benefiting from the interconnection of systems and the infusion of intelligence into those systems. Nothing is changing more than information technology and the opportunities for innovation and smarter computing models have never been greater. To take advantage of these opportunities, new systems need to communicate with each other, share data better, be easier to manage, and be more cost effective.

The IBM® zEnterprise™ System (zEnterprise) is a perfect fit in this world of smarter computing, being both the newest step in the evolution of System z® leadership and a premier solution for centrally managed enterprise cloud environments. zEnterprise is a true hybrid computing system comprised of virtualized heterogeneous resources integrated and managed as a single system, and optimized to your business objectives.

The zEnterprise includes a central processor complex (CPC), either the IBM zEnterprise 196 (z196) or the new IBM zEnterprise 114 (z114), the IBM zEnterprise BladeCenter® Extension (zBX) with its integrated optimizers and select IBM blades, and the management fabric that ties it all together, the IBM zEnterprise Unified Resource Manager. The enhancements announced today extend zEnterprise leadership, continuing to deliver a powerful, yet flexible server that goes beyond technology to drive business innovation.

Enhancements include:

Increased support of the zEnterprise System

- Enhancements to IBM zEnterprise Unified Resource Manager for continued support of a smarter business infrastructure

Security

- Changes to the Common Cryptographic Architecture, Crypto Express3, and Trusted Key Entry

Ease of use

- IPL from an alternate subchannel set

A new I/O infrastructure

- PCIe-based I/O infrastructure with more bandwidth - 8 GBps PCIe bus

- PCIe I/O drawer with 14% more capacity
- New form factor for port granularity
- Reduced power consumption

New features and functions for the storage area network

- A new form factor with FICON® Express8S; increased bandwidth and channel granularity for FICON, zHPF, and FCP
- Increased reliability for FCP channels with support of T10-DIF

New features and functions for networking

- A new form factor with OSA-Express4S; increased performance for inbound traffic and port granularity
- Improved scalability and performance with Inbound Workload Queuing for Enterprise Extender for z/OS®

New features and functions for Parallel Sysplex® environments for demanding data-sharing workloads

- Improved service times for 12x InfiniBand coupling links with a new generation of hardware - when using the 12x IFB3 protocol, synchronous service times are designed to be 40% faster than when using the 12x IFB protocol
- Increased connectivity for 1x InfiniBand coupling links
- Up to 32 subchannels (devices) per CHPID for 1x InfiniBand coupling links to help improve coupling throughput at extended distances with a new generation of hardware supporting four ports per feature
- Server Time Protocol (STP) recovery enhancement

Usability and security

- Enhancements for the Hardware Management Console (HMC) including increased security for file transfers

Overview

The zEnterprise gives you the freedom to bring mainframe and distributed technologies together in a system that can start to replace individual islands of computing, bringing applications closer to the data they need, and helping to reduce complexity while improving security. With enhancements to IBM zEnterprise Unified Resource Manager, we are continuing to provide a smarter business infrastructure.

In support of the heterogeneous technology available with the zEnterprise, today we are also announcing support for running the Linux® on System x® operating system on select IBM System x blades installed in the zBX. This is designed to help increase the application portfolio available for the zEnterprise. Refer to Hardware Announcement [ZG11-0200](#), dated July 12, 2011.

Security enhancements to the Common Cryptographic Architecture and Crypto Express3 coprocessor provide stronger algorithms for industry security standards and improved PIN protection - critical for automated teller machines (ATMs), debit cards, and point-of-sale (POS) devices. Introduction of the Trusted Key Entry 7.1 Licensed Internal Code helps keep in step with current recommendations for cryptographic strength such as new access control support, crypto module notebook support, and improvements to smart cards. Enhancements to the Hardware Management Console (HMC) are designed to ensure more security for File Transfer Protocol (FTP).

Since businesses need to deliver timely, integrated information to business leaders, support personnel, and customers 24x7, access to data and the network needs to be fast, secure, and dependable. The new z196 PCIe I/O subsystem improves robustness, resiliency, and performance, offers nondisruptive add and repair, and is

designed to help lower power requirements compared to the current I/O cage and I/O drawer.

An additional new function is being introduced to improve time coordination of zBX components to increase time accuracy to an NTP server. New features are being introduced for increased performance for 12x InfiniBand coupling links and more connectivity for 1x InfiniBand coupling links.

The new IBM zEnterprise 114 (z114) can upgrade to a zEnterprise z196 (z196) Model M15 (air-cooled model only). Refer to Hardware Announcement [ZG11-0207](#), dated July 12, 2011.

Key prerequisites

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

Planned availability date

- September 9, 2011
 - New build systems, MES orders, and new features and functions
- September 26, 2011
 - Unified Resource Manager functions
 - Manage suite (#0019) enhancements
 - Automate/Advanced Management Firmware Suite (#0020) enhancements
 - Manage Firmware System x Blade (#0042)
 - Advanced Management Firmware System x Blade (#0046)

Description

Enhancements to IBM zEnterprise Unified Resource Manager

IBM zEnterprise System continues to provide a smarter business infrastructure through augmentation of IBM zEnterprise Unified Resource Manager. Unified Resource Manager is designed to virtualize and manage all resources of a System z ensemble as a single pool of resources, integrating system and workload management across the multisystem, multitier, multiarchitecture environment.

Enhancements to the Manage suite (#0019) functions include:

- Hypervisor™ and virtual server life-cycle management to support the zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine). Hypervisor and virtual server life-cycle management was first introduced with Unified Resource Manager for PR/SM™, z/VM®, and PowerVM™.
- zEnterprise BladeCenter Extension (zBX) Model 002 integrated hypervisor for System x blades, which is shipped, deployed, and serviced as zEnterprise System licensed internal code (LIC) and is booted automatically on each System x blade at power-on-reset. Like the existing supported hypervisors (PR/SM, z/VM, and PowerVM), it is isolated on the intranode management network (INMN) and managed by Unified Resource Manager, providing a single point of control and supporting directed and dynamic virtual server provisioning.

Enhancements to the Automate/Advanced Management Firmware Suite (#0020) functions include:

- Representation of the physical and virtual resources that are used in the context of a deployed business function as a named workload for IBM System x blades. Workload representation enables focus on business goals rather than individual elements and better aligns resources with business needs.

- Ability to define a performance service level policy and to enable performance monitoring and reporting for IBM System x blades.

System x blades running Linux on System x are supported in the zBX, utilizing the zBX integrated hypervisor for IBM System x blades (using Kernel-based Virtual Machine), providing logical device integration between System z and System x blades for multitiered applications. System x blades are licensed separately and are enabled and managed as part of the ensemble by Unified Resource Manager.

For more information on IBM System x blades refer to Hardware Announcement [ZG11-0078](#), dated April 06, 2011, IBM BladeCenter HX5 is a scalable blade server designed to provide new levels of utilization, performance, and reliability for compute- and memory-intensive workloads.

Improved time coordination for zBX components

Network Time Protocol (NTP) clients, running on blades in zBX, can synchronize their time to the NTP server provided by the Support Element (SE) every hour. Therefore, it is important for the SE's clock to maintain time accuracy.

An enhancement has been made to improve the time accuracy of the SE's Battery Operated Clock (BOC) by synchronizing the SE's BOC to the server's Time-of-Day (TOD) clock every hour, instead of the previous synchronization which took place every 24 hours. This enhancement allows the SE's clock to maintain a time accuracy of 100 milliseconds to an NTP server configured as the External Time Source in an STP-only CTN.

In addition, this enhancement provides the capability for the components in zBX to maintain an approximate time accuracy of 100 milliseconds to an NTP server if they synchronize to the SE's NTP server at least once an hour. This enhancement is exclusive to z196 and z114.

For more information regarding IBM zEnterprise BladeCenter Extension (zBX) Model 002, refer to Hardware Announcement [ZG10-0263](#), dated July 22, 2010, IBM zEnterprise BladeCenter Extension (zBX).

Common Cryptographic Architecture and Crypto Express3 enhancements

To continue to satisfy customer requirements, new functions are now available for the Common Cryptographic Architecture (CCA) and the Crypto Express3 feature, when defined as a coprocessor, complementing the Central Processor Assist for Cryptographic Function (CPACF). The CPACF provides high-performance hardware encryption and decryption on every z196 or z114 processor unit defined as a central processor (CP) or Integrated Facility for Linux (IFL).

Crypto Express3 remains a tamper-sensing and tamper-responding programmable cryptographic feature. Each Crypto Express3 PCIe adapter contains dual processors that operate in parallel to support the Common Cryptographic Architecture with high reliability. The Crypto Express3 feature continues to reside in the I/O cage or I/O drawer of the zEnterprise System.

Common Cryptographic Architecture enhancements

The following enhancements have been added to the Common Cryptographic Architecture support which is used in the Crypto Express3 feature when it is configured as a coprocessor.

Expanded key support for AES algorithm: CCA currently supports the Advanced Encryption Standard (AES) algorithm to allow the use of AES keys to encrypt data. Expanded key support for AES adds a framework to support a much broader range of application areas, and lays the groundwork for future use of AES in areas where standards and customer applications are expected to evolve.

As stronger algorithms and longer keys become increasingly common, security requirements dictate that these keys must be wrapped using key encrypting keys

(KEKs) of sufficient strength. This feature adds support for AES key encrypting keys. These AES wrapping keys have adequate strength to protect other AES keys for transport or storage. The new AES key types use the variable-length key token. The supported key types are EXPORTER, IMPORTER, and for use in the encryption and decryption services, CIPHER. New application programming interfaces (APIs) have been added or modified to manage and use these new keys.

The following new or modified CCA API functions are also supported:

- Key Token Build2: Builds skeleton variable-length key tokens
- Key Generate2: Generates keys using random key data
- Key Part Import2: Creates keys from key part information
- Key Test2: Verifies the value of a key or key part
- Key Translate2
 - Translates a key: Changes the key encrypting key (KEK) used to wrap a key
 - Reformats a key: Converts keys between the legacy token format and the newer variable-length token format
- Symmetric Key Export: Modified to also export AES keys
- Symmetric Key Import2: Imports a key that has been wrapped in the new token format
- Secure Key Import2: Wraps key material under the master key or an AES KEK
- Restrict Key Attribute: Changes the attributes of a key token
- Key Token Parse2: Parses key attributes in the new key token
- Symmetric Algorithm Encipher and Symmetric Algorithm Decipher: Enhanced to encipher and decipher data using AES keys wrapped in the new variable-length token format

Expanded key support for AES algorithm and CCA API functions supported by the Crypto Express3 feature when defined as a coprocessor, is exclusive to z196 and z114, and is supported by z/OS and z/VM. Refer to the [Software requirements](#) section.

Enhanced ANSI TR-31 interoperable secure key exchange: ANSI TR-31 defines a method of cryptographically protecting Triple Data Encryption Standard (TDES) cryptographic keys and their associated usage attributes. The TR-31 method complies with the security requirements of the ANSI X9.24 Part 1 standard, although use of TR-31 is not required in order to comply with that standard. CCA has added functions that can be used to import and export CCA TDES keys in TR-31 formats. These functions are designed primarily as a secure method of wrapping TDES keys for improved and more secure key interchange between CCA and non-CCA devices and systems.

Enhanced ANSI TR-31 interoperable secure key exchange supported by the Crypto Express3 feature when defined as a coprocessor, is exclusive to z196 and z114, and is supported by z/OS and z/VM. Refer to the [Software requirements](#) section.

PIN block decimalization table protection: To help avoid a decimalization table attack to learn a personal identification number (PIN), a solution is now available in the CCA API to thwart this attack by protecting the decimalization table from manipulation. PINs are most often used for automated teller machines (ATMs) but are increasingly used at point-of sale, for debit and credit cards.

PIN block decimalization table protection supported by the Crypto Express3 feature when defined as a coprocessor, is exclusive to z196 and z114, and is supported by z/OS and z/VM. Refer to the [Software requirements](#) section.

PKA RSA OAEP with SHA-256 algorithm: RSA Encryption Scheme - Optimal Asymmetric Encryption Padding (RSA OAEP) is a public-key encryption scheme or method of encoding messages and data in combination with the RSA algorithm and a hash algorithm.

Currently, the Common Cryptographic Architecture and z/OS Integrated Cryptographic Service Facility (ICSF) provide key management services supporting the RSA OAEP method using the SHA-1 hash algorithm, as defined by the public key cryptographic standard (PKCS) #1 V2.0 standard. These services can be used to exchange AES or DES/TDES key values securely between financial institutions and systems. However, PKCS#1 V2.1 extends the OAEP method to include the use of the SHA-256 hashing algorithm to increase the strength of the key wrapping and unwrapping mechanism. The CCA key management services have been enhanced so that they can use RSA OAEP with SHA-256 in addition to RSA OAEP with SHA-1.

This provides support for PKCS that is mandated by some countries for interbank transactions and communication systems.

PKA RSA OAEP with SHA-256 algorithm is supported by the Crypto Express3 feature when defined as a coprocessor, is exclusive to z196 and z114, and is supported by z/OS and z/VM. Refer to the [Software requirements](#) section.

Additional Elliptic Curve Cryptography (ECC) functions: The Common Cryptographic Architecture has been extended to include the Elliptic Curve Diffie-Hellman (ECDH) algorithm.

Elliptic Curve Diffie-Hellman (ECDH) is a key agreement protocol that allows two parties, each having an elliptic curve public-private key pair, to establish a shared secret over an insecure channel. This shared secret may be used directly as a key, or to derive another key which can then be used to encrypt subsequent communications using a symmetric key cipher such as AES.

Enhancements include:

- Updated key management functions to support AES KEKs
 - Generating an ECC private key wrapped with an AES KEK
 - Importing and exporting an ECC private key wrapped with an AES KEK
- Support for ECDH with a new service

These additional Elliptic Curve Cryptographic functions are supported by the Crypto Express3 feature when defined as a coprocessor, are exclusive to z196 and z114, and are supported by z/OS and z/VM. Refer to the [Software requirements](#) section.

Trusted Key Entry (TKE) 7.1 Licensed Internal Code (LIC): The following functions are supported in the TKE 7.1 LIC:

- **New access control support for all TKE applications:** Every TKE application and the ability to create and manage crypto module and domain groups now require the TKE local cryptographic adapter profile to have explicit access to the TKE application or function the user wants to run. This was done to provide more control of what functions TKE users are allowed to perform.
- **New Migration Utility:** During a migration from a lower release of TKE to TKE 7.1 LIC, it will be necessary to add access control points to the existing roles. The new access control points can be added through the new Migrate Roles Utility or by manually updating each role through the Cryptographic Node Management Utility. The IBM-supplied roles created for TKE 7.1 LIC have all of the access control points needed to perform the functions they were permitted to use in TKE releases prior to TKE 7.1 LIC.
- **Single process for loading an entire key:** The TKE now has a wizard-like feature that takes users through the entire key loading procedure for a master or operational key. The feature preserves all of the existing separation of duties and authority requirements for clearing, loading key parts, and completing a key. The procedure saves time, by walking users through the key loading procedure. However, this feature does not reduce the number of people it takes to perform the key load procedure.
- **Single process for generating multiple key parts of the same type:** The TKE now has a wizard-like feature that allows a user to generate more than one key part at a time. The procedure saves time because the user has to start the

process only one time, and the TKE efficiently generates the desired number of key parts.

- **AES operational key support:** CCA V4.2 for the Crypto Express3 feature includes three new AES operational key types. From the TKE, users can load and manage the new AES EXPORTER, IMPORTER, and CIPHER operational keys from the TKE workstation crypto module notebook.
- **Decimalization table support:** CCA V4.2 for the Crypto Express3 feature includes support for 100 decimalization tables for each domain on a Crypto Express3 feature. From the TKE, users can manage the decimalization tables on the Crypto Express3 feature from the TKE workstation crypto module notebook. Users can manage the tables for a specific domain or manage the tables of a set of domains if they are using the TKE workstation Domain Grouping function.
- **Host cryptographic module status support:** From the TKE workstation crypto module notebook, users will be able to display the current status of the host cryptographic module that is being managed. If they view the Crypto Express3 feature module information from a crypto module group or a domain group, they will see only the status of the group's master module.
- **Display of active IDs on the TKE console:** A user can be logged onto the TKE workstation in privileged access mode. In addition, the user can be signed onto the TKE workstation's local cryptographic adapter. If a user is signed on in privileged access mode, that ID is shown on the TKE console. With this new support, both the privileged access mode ID and the TKE local cryptographic adapter ID will be displayed on the TKE console.
- **Increased number of key parts on smart card:** If a TKE smart card is initialized on a TKE workstation with a 7.1 level of LIC, it will be able to hold up to 50 key parts. Previously, TKE smart cards could hold only 10 key parts.
- **Use of ECDH to derive shared secret:** When the TKE workstation with a 7.1 level of LIC exchanges encrypted material with a Crypto Express3 at CCA level V4.2, Elliptic Curve Diffie-Hellman (ECDH) is used to derive the shared secret. This increases the strength of the transport key used to encrypt the material.

Trusted Key Entry 7.1 Licensed Internal Code is supported by the Crypto Express3 feature when defined as a coprocessor on z196, z114, z10™ EC, and z10 BC servers. TKE workstation (#0841) is required.

Ease of use enhancement - IPL from an alternate subchannel set

First came multiple subchannel sets (MSS) (two sets), which was designed to provide greater I/O device configuration capabilities for large enterprises. Two subchannel sets could be defined to each channel subsystem - the base addresses and aliases in subchannel set 0 and the parallel access volume aliases in subchannel set 1; 63.75K subchannels in set-0 and the addition of 64K-1 subchannels in set-1. MSS was delivered on System z9® in September of 2005 and supported by z/OS V1.7 and Linux on System z. This applied to the ESCON® and FICON protocols.

Then came a third subchannel set to extend the amount of addressable storage capacity - another 64k subchannels to help complement functions such as "large" or extended addressing volumes and HyperPAV and to help facilitate consistent device address definitions, simplifying addressing schemes for corresponding devices. This was delivered on zEnterprise 196 in September of 2010 and supported by z/OS V1.12 and Linux on System z. This applied to the ESCON, FICON, and zHPF protocols.

Now, we are introducing IPL from an alternate subchannel set to allow enterprises to IPL from subchannel set 1 or subchannel set 2 (z196 only), in addition to subchannel set 0. Devices used early during IPL processing can now be accessed using subchannel set 1 or subchannel set 2. This is intended to allow the users of Metro Mirror (PPRC) secondary devices defined using the same device number and a new device type in an alternate subchannel set to be used for IPL, IODF, and stand-alone dump volumes when needed.

IPL from an alternate subchannel set is exclusive to z196 and z114, applies to the FICON and zHPF protocols (CHPID type FC), and is supported by z/OS. Refer to

the [Software requirements](#) section. It is applicable to all of the supported FICON features. Note, the z114 supports up to two subchannel sets.

PCIe I/O drawer - increased capacity, granularity, bandwidth, and RAS

A new I/O drawer and new form factor I/O cards are being introduced to support a direct Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure with increased capacity, granularity, and infrastructure bandwidth, as well as increased reliability, availability, and serviceability. New hardware features are being introduced for the SAN and the LAN to support PCIe Gen2 - FICON Express8S and OSA-Express4S.

Capacity: With their compact design, two PCIe I/O drawers occupy the same space as one I/O cage while delivering 14% more capacity - up to 128 channels (64 I/O cards) are supported versus the 112 channels (28 I/O cards) offered with the I/O cage.

Granularity: With the PCIe I/O drawer we are introducing enhanced granularity features for the storage area network (SAN) and the local area network (LAN).

- For the FICON/zHPF/FCP storage area network: FICON Express8S for single mode fiber (LX) and multimode fiber (SX) environments with two channels per feature, and two channel path identifiers (CHPIDs).
- For the 10 Gigabit Ethernet and Gigabit Ethernet local area network: OSA-Express4S for single mode fiber (LX, LR) and multimode fiber (SX, LR) environments. The 10 GbE features have one port per feature and one CHPID. The GbE features have two ports per feature and one CHPID shared by the two ports.

Increased infrastructure bandwidth: The new I/O infrastructure introduces the industry-standard Peripheral Component Interconnect Express Generation 2 (PCIe Gen2) infrastructure from the processor book to the I/O cards, which now supports an 8 gigabytes per second (8 GBps) bus. This new infrastructure offers enhanced bandwidth in the connection to the processor book as well as the distribution buses contained in the drawer.

Reliability, Availability, Serviceability: The PCIe I/O drawer is designed with symmetrical, redundant cooling across all I/O cards and power supplies for improved RAS. Dual PCIe paths to the system processors from an I/O card provide enhanced resiliency. Also included are temperature monitoring of critical ASICs to optimize cooling and reliability.

FICON Express8S - a new generation for FICON , zHPF, and FCP

A new generation of features for the storage area network (SAN) is being introduced in support of the PCIe 8 GBps host bus and the PCIe I/O drawer. The new features for the multimode and single mode fiber optic cabling environments have path length reductions for High Performance FICON for System z (zHPF) and Fibre Channel Protocol (FCP), increased start I/Os, improved throughput for zHPF and FCP with the introduction of a hardware data router, and increased port granularity - two channels/ports per feature.

New design for increased performance for zHPF and FCP: FICON Express8S contains a new IBM ASIC which is designed to support the 8 GBps PCIe interface to the PCIe I/O drawer and increased start I/Os. In addition, a hardware data router has been added in support of the zHPF and FCP protocols for path length reduction and increased throughput. FICON Express8S supports a link data rate of 2, 4, or 8 Gbps auto-negotiated.

With these changes FICON Express8S, when supporting the zHPF or FCP protocols, has been designed to achieve full duplex line speed - 8 Gbps - in each direction.

The High Performance FICON for System z journey:

- In October of 2008 we made the initial zHPF announcement on System z10® with FICON Express4 and FICON Express2 channels.
 - Maximum of 31k zHPF IOs/sec, 2.2 times the FICON protocol
 - Single track limit for zHPF data transfers
- In July 2009 we introduced FICON Express8 channels on System z10.
 - Maximum of 52k zHPF IOs/sec
 - 64k byte limit for zHPF data transfers
- In July 2010 we introduced additional support on zEnterprise 196.
 - Extension to multi-tracks of zHPF data transfers
- Today with the new FICON Express8S channel exclusive to zEnterprise 196 and zEnterprise 114 we are introducing:
 - A hardware data router for more efficient zHPF data transfers.
 - The first FICON Express channel with hardware specifically designed to support zHPF - FICON Express8S. FICON Express8, FICON Express4, and FICON Express2 have a firmware-only zHPF implementation.

In laboratory measurements, using FICON Express8S in a z196 with the zHPF protocol and small data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum of 92,000 IOs/sec, compared to the maximum of 52,000 IOs/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 77% increase and applies to reads, writes, and a read/write mix.

In laboratory measurements, using FICON Express8S in a z196 with the zHPF protocol and a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 MB/sec (reads + writes) compared to a maximum of 770 MB/sec (reads + writes) achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase.

This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance 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.

Increased performance for the FCP protocol: A FICON Express8S feature, when defined as CHPID type FCP, conforms to the Fibre Channel Protocol (FCP) standard to support attachment of SCSI devices, to complement the classical storage attachment supported by FICON and zHPF channels.

In laboratory measurements, using FICON Express8S in a z196 with the FCP protocol for small data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum of 92,000 IOs/sec, compared to the maximum of 84,000 IOs/sec achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 10% increase and applies to reads, writes, and a read/write mix.

In laboratory measurements, using FICON Express8S in a z196 with the FCP protocol and an internal driver supporting the hardware data router, executing a mix of large sequential read and write data transfer I/O operations, FICON Express8S operating at 8 Gbps achieved a maximum throughput of 1600 MB/sec (reads + writes) compared to the maximum of 770 MB/sec (reads + writes) achieved with FICON Express8 operating at 8 Gbps. This represents approximately a 108% increase.

The FCP protocol is supported by z/VM, z/VSE™, and Linux on System z. Refer to the [Software requirements](#) section.

FCP channels to support T10-DIF for enhanced reliability

Recognizing that high reliability is important to maintaining the availability of business-critical applications, the System z Fibre Channel Protocol (FCP) has implemented support of the American National Standards Institute's (ANSI) T10 Data Integrity Field (DIF) standard. Data integrity protection fields are generated by the operating system and propagated through the storage area network (SAN). System z helps to provide added end-to-end data protection between the operating system and the storage device.

An extension to the standard, Data Integrity Extensions (DIX), provides checksum protection from the application layer through the host bus adapter (HBA), where cyclical redundancy checking (CRC) protection is implemented.

T10-DIF support by the FICON Express8S and FICON Express8 features, when defined as CHPID type FCP, is exclusive to z196 and z114. Exploitation of the T10-DIF standard requires support of the operating system and the storage device. Refer to the [Software requirements](#) section.

Increased port granularity: The FICON Express8S 10KM LX and SX features for single mode and multimode fiber optic cabling environments each now have two channels/ports per feature versus the four channels per feature for the FICON Express8 features. This design helps facilitate purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

The FICON Express8S features, supporting CHPID types FC (zHPF, FICON, channel-to-channel) and FCP, are exclusive to z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z. Refer to the [Software requirements](#) section.

OSA-Express4S - a new form factor of Ethernet for the LAN

A new generation of 10 Gigabit (10 GbE) and Gigabit Ethernet (GbE) OSA-Express features is being introduced in support of the PCIe 8 GBps host bus and the PCIe I/O drawer. These new features retain the robust performance characteristics of the OSA-Express3 GbE and 10 GbE features while providing increased port granularity.

Port granularity for increased flexibility: The OSA-Express4S features have half the number of ports per feature compared to OSA-Express3. This design facilitates purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

OSA-Express4S 10 GbE with one port per feature: The OSA-Express4S 10 GbE features have one CHPID per feature and one port associated with the CHPID. A feature can be configured to support CHPID type OSD or OSX.

- CHPID type OSD continues to support the Queued Direct Input/Output (QDIO) architecture, TCP/IP traffic when Layer 3 (uses IP address), protocol-independent when Layer 2 (uses MAC address).
- CHPID type OSX continues to support connectivity and access control to intraensemble data network (IEDN) from z196 or z114 to zBX (OSA-Express 10 GbE features only)

OSA-Express4S 10 Gigabit Ethernet Long Reach (LR) is available for single mode fiber environments. OSA-Express4S 10 Gigabit Ethernet Short Reach (SR) is available for multimode fiber environments.

OSA-Express4S GbE with two ports per feature: The OSA-Express4S GbE features have one CHPID per feature (CHPID type OSD) and two ports associated with the CHPID.

Note: CHPID type OSN (OSA-Express for NCP) is not being offered for the OSA-Express4S GbE features.

OSA-Express4S GbE long wavelength (LX) is available for single mode fiber environments. OSA-Express4S GbE short wavelength (SX) is available for multimode fiber environments.

The OSA-Express4S 10 Gigabit Ethernet and Gigabit Ethernet features are exclusive to z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z. Refer to the [Software requirements](#) section.

Checksum offload for IPv6 packets is now available for z/OS environments and for z/VM guest exploitation: When the checksum function is offloaded from the host, CPU cycles are reduced, improving performance. With the introduction of OSA-Express4S, the checksum offload function is now performed for IPv6 packets as well as IPv4 packets, whether the traffic goes out to the local area network (LAN), comes in from the LAN, or flows logical partition-to-logical partition through OSA-Express4S.

Checksum offload provides the capability of calculating the Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Protocol (IP) header checksums for Internet Protocol Version 4 (IPv4) packets and now, IPv6 packets.

When checksum offload was introduced in May of 2003, it was limited to IPv4 packets. Checksum offload for IPv6 packets is exclusive to OSA-Express4S features (CHPID types OSD and OSX) on z196 and z114. It is supported by z/OS and z/VM for guest exploitation. Refer to the [Software requirements](#) section. Checksum offload for IPv4 packets is currently available for all in-service releases of z/OS and Linux on System z.

Checksum offload for LPAR-to-LPAR traffic in the z/OS environment is included in the OSA-Express4S design for both IPv4 and IPv6 packets. Refer to the [Software requirements](#) section.

Large send for IPv6 packets is now available for z/OS environments: Large send (also referred to as TCP segmentation offload) is designed to improve performance by offloading **outbound** TCP segmentation processing from the host to an OSA-Express4S feature by employing a more efficient memory transfer into OSA-Express4S.

Large send support for IPv6 packets applies to the OSA-Express4S features (CHPID type OSD and OSX), and is exclusive to z196 and z114. It is supported by z/OS. Refer to the [Software requirements](#) section. Large send for IPv4 packets is currently available for all in service releases of z/OS, Linux on System z, and z/VM for guest exploitation.

Note: Large send is not supported for LPAR-to-LPAR packets.

Design change yields performance benefit for OSA-Express4S 10 Gigabit Ethernet inbound traffic

Performance using jumbo frames: In laboratory measurements, using an OSA-Express4S 10 Gigabit Ethernet (10 GbE) feature in a z196 defined as CHPID type OSD with an inbound-to-the-host streams workload operating at 10 Gbps, we achieved a maximum user-payload throughput of 1,180 megabytes per second (MBps) compared to a maximum of 680 MBps achieved with an OSA-Express3 10 GbE feature on a z196. This represents approximately a 70% increase for jumbo frames (8000 byte frames).

With a mixed-direction streams workload in the same jumbo frames environment, using an OSA-Express4S 10 GbE feature, we achieved a maximum user-payload throughput of 2,080 MBps with an OSA-Express4S 10 GbE feature compared to a maximum of 1,240 MBps on an OSA-Express3 10 GbE feature on z 196. This represents approximately a 70% increase for jumbo frames.

Performance using standard frames: In laboratory measurements, using an OSA-Express4S 10 GbE feature in a z196 defined as CHPID type OSD with an inbound-to-the-host streams workload operating at 10 Gbps, we achieved a

maximum user-payload throughput of 1,120 MBps compared to a maximum of 615 MBps achieved with an OSA-Express3 10 GbE feature on a z196. This represents approximately an 80% increase for standard frames (1492 byte frames).

With a mixed-direction streams workload in the same standard frames environment, using an OSA-Express4S 10 GbE feature, we achieved a maximum user-payload throughput of 1,680 MBps with an OSA-Express4S 10 GbE compared to a maximum of 1,180 MBps with an OSA-Express3 10 GbE feature on a z196. This represents approximately a 40% increase for standard frames.

OSA-Express4S 10 GbE performance was measured in a controlled environment using IBM Application Workload Modeler (AWM). The actual throughput or performance that any user may experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the network options and configuration, and the workload processed. One MBps represents 1,048,576 bytes per second.

Access AWM information at

<http://www.ibm.com/software/network/awm/index.html>

OSA-Express4S Gigabit Ethernet is already capable of line speed for jumbo frames and standard frames.

Inbound workload queuing for Enterprise Extender - for improved scalability and performance: Inbound workload queuing (IWQ) for the OSA-Express features has been enhanced to differentiate and separate inbound Enterprise Extender traffic to a new input queue. The Enterprise Extender separation and processing associated with the Enterprise Extender input queue provides improved scalability and performance for Enterprise Extender.

With each input queue representing a unique type of workload, each having unique service and processing requirements, the IWQ function allows z/OS to use appropriate processing resources for each input queue. This approach allows multiple concurrent z/OS processing threads to process each unique input queue to avoid traditional resource contention. In a heavily mixed workload environment, this "off the wire" network traffic separation provided by OSA-Express IWQ reduces the conventional z/OS processing required to identify and separate unique workloads.

Inbound workload queuing for Enterprise Extender is supported by the OSA-Express4S and OSA-Express3 features when defined as CHPID type OSD or OSX. It is exclusive to z196 and z114, and is supported by z/OS and by z/VM for guest exploitation. Refer to the [Software requirements](#) section.

For more information on inbound workload queuing, refer to Hardware Announcement [ZG10-0249](#), dated July 22, 2010, IBM zEnterprise System -- A new dimension in computing.

Server Time Protocol (STP) enhancement

Recovery enhancement: The new generation of host channel adapters (HCA3-O or HCA3-O LR), introduced for coupling, have been designed to send a reliable unambiguous "going away signal" (GOSIG) to indicate that the server on which the HCA3 is running is about to enter a failed (check stopped) state. When the GOSIG sent by the Current Time Server (CTS) in an STP-only Coordinated Timing Network (CTN) is received by the Backup Time Server (BTS), the BTS can safely take over as the CTS without relying on the previous recovery methods of Offline Signal (OLS) in a two-server CTN or the Arbiter in a CTN with three or more servers.

This enhancement is exclusive to z196 and z114 and is available only if you have an HCA3-O or HCA3-O LR on the CTS communicating with an HCA3-O or HCA3-O LR on the BTS. Note that the already available STP recovery design is still available for the cases when a GOSIG is not received or for other failures besides a server failure.

A new generation of 12x and 1x InfiniBand coupling links

Support for demanding data-sharing workloads: Simplify your Parallel Sysplex connectivity requirements and do more with less, with a new generation of host channel adapters (HCA3s) for coupling - the HCA3-O fanout for 12x InfiniBand (12x IFB) with improved service times using the 12x InfiniBand3 (12x IFB3) protocol and the HCA3-O LR fanout for 1x InfiniBand (1x IFB) with four ports of connectivity and optional additional subchannels for extended-distance solutions.

- 12x InfiniBand coupling links support a link data rate of 6 gigabytes per second (GBps) at distances up to 150 meters (492 feet).
- 1x IFB coupling links support a link data rate of 5 gigabits per second (Gbps) server-to-server or 2.5 or 5 Gbps when using dense wavelength division multiplexers (DWDMs). The 1x IFB coupling links support distances up to 10 km (6.2 miles) unrepeated. Greater distances are possible with an RPQ or when using qualified DWDM solutions.
- Both 12x and 1x InfiniBand coupling links also allow you to share physical links by defining multiple logical links (channel path identifiers - CHPIDs) - up to 16 CHPIDs across two ports for 12x IFB, and up to 16 CHPIDs across four ports for 1x IFB.

Note: The 12x InfiniBand HCA3-O fanout can connect to 12x InfiniBand HCA2-O fanout on a z196, z114, or z10. An HCA3-O fanout cannot connect to an HCA1-O fanout on a z9™.

Two protocols - 12x IFB and 12x IFB3 - for the new generation of 12x InfiniBand coupling links: There are now two protocols supported by the HCA3-O for 12x IFB feature - 12x IFB and 12x IFB3.

1. **12x IFB3 protocol:** When HCA3-Os are communicating with HCA3-Os and have been defined with four or fewer CHPIDs per port, the 12x IFB3 protocol is utilized.
2. **12x IFB protocol:** If more than four CHPIDs are defined per HCA3-O port, or HCA3-O features are communicating with HCA2-O features on zEnterprise or System z10 servers, links will run with the 12x IFB protocol.

Improved service times when using the 12x IFB3 protocol: The HCA3-O feature supporting 12x InfiniBand coupling links has been designed to deliver improved services times. When no more than four CHPIDs are defined per HCA3-O port, the 12x IFB3 protocol is used. When using the 12x IFB3 protocol, synchronous service times are designed to be 40% faster than when using the 12x IFB protocol.

Improved physical connectivity with 1x InfiniBand coupling links: The HCA3-O long reach (LR) fanout for 1x InfiniBand coupling links has been designed with four ports per feature to satisfy requirements for more physical connectivity. The added connectivity will be helpful as clients migrate from InterSystem Channel-3 (ISC-3) to 1x InfiniBand coupling links.

Up to 32 subchannels (devices) per CHPID for 1x InfiniBand coupling links: To improve link utilization and coupling throughput at extended distances between the coupling facility (CF) and the operating system or between CFs separated by extended distances, we are now optionally supporting 32 subchannels (devices) per CHPID, versus the current 7 subchannels per CHPID. The additional subchannel definition is available whether using an HCA3-O LR or HCA2-O LR feature for 1x IFB. This increase is designed to help eliminate the need to add CHPIDs or physical links in order to help improve coupling throughput at extended distances.

The HCA3-O feature (#0171) for 12x IFB and the HCA3-O LR feature (#0170) for 1x IFB are exclusive to z196 and z114, utilize CHPID type CIB, and are supported in the z/OS environment. Refer to the [Software requirements](#) section.

Enhancements for the Hardware Management Console/Support Element

Enhancements for usability, and an update that strengthens secure File Transfer Protocol support, have been added to the Hardware Management Console/Support Element (HMC/SE) to help satisfy customer and service requirements.

Systems and Technology Group (STG) Lab Services

In support of this announcement, STG Lab Services has developed enablement services for the IBM zEnterprise BladeCenter Extension (zBX). Whether the environment is z/OS, z/VM, or Linux on System z, Lab Services can help to set up your systems and train your IT staff. Lab Services will enable the ensemble, install and configure the necessary software components, and provide hardware connectivity services that will help to demonstrate the capabilities and functions of this new technology, as well as help you to quickly realize the value and investment that you have made in the new hybrid system.

The four services available are:

IBM Smart Analytics Optimizer Enablement Services

A three- to four-week service offering that provides services to assist current DB2® for z/OS clients with the planning and installation activities for the integration of the optimizer into their data center as well as configuring the hardware components necessary to enable a basic functional IBM Smart Analytics Optimizer environment and IBM System Storage® unit.

zEnterprise Ensemble Enablement JumpStart Assistance for zBX Blades

A three- to four-week service offering that is targeted for implementations involving zBX Blades and is designed to help accelerate the implementation of an ensemble. The service will set up the z/OS and zBX infrastructure, use a test application to demonstrate the functions of the ensemble environment, and provide education to the client with a focus on Unified Resource Manager "Manage" functions.

zEnterprise Ensemble Enablement JumpStart Assistance for z/VM

A three- to four-week service offering that is targeted for implementations that will use the Unified Resource Manager to manage Linux virtual servers under the z/VM Hypervisor. The service will set up the z/VM and Unified Resource Manager infrastructure, use a test application to show the functions of the ensemble environment, and provide education to the client with a focus on Unified Resource Manager "Manage" functions.

zEnterprise Ensemble Enablement JumpStart Assistance for DataPower® XI50z Blades

A two-week service offering that is targeted for implementations involving the configuration and enablement of the DataPower XI50z within a zBX environment. An education workshop is also provided that will assist you to leverage the new blades quickly and assist in the training of staff on new functions available on the Hardware Management Console (HMC) to support the XI50z and zEnterprise environments.

Other z196- and z114-related service offerings from STG Lab Services

STG Lab Services provides other service offerings in support of IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114). A sampling of these services that might be applicable include:

- z/OS Infrastructure Solutions, such as
 - Parallel Sysplex
 - Hardware, Performance, and DFSMSHsm™ Health Inspections
- Middleware Solutions, such as

- DB2, CICS®, and WebSphere® Services
- z/OS and Linux on System z Security Services
- Linux on System z Services and Training Kits
- Systems Agenda Consulting
- Cross-Platform Consulting Services
 - Data Center Services
 - IBM High Availability Center of Competency briefings and education on high availability best practices
 - IBM Systems Technical Training

For more information visit

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

Contact your IBM representative, your IBM Business Partner, or the IBM zEnterprise Enablement Services team (systemz@us.ibm.com).

IBM zEnterprise Setup and Migration Services

IBM Global Technology Services (GTS) has released the IBM Implementation Services for System z - IBM zEnterprise setup and migration. The services are designed to assist you as you plan for, configure, implement, and migrate to a zEnterprise System.

The IBM zEnterprise Setup and Migration Services are comprised of a four-day standard service with three optionally selectable activities that can be acquired separately to address specific installation requirements.

The four-day standard service includes:

- Onsite or remote planning session. IBM will review your overall plan, define milestones, identify potential migration scenarios, and assess your readiness for installation.
- Onsite or remote review and updating of hardware, network, and storage provisioning.

IBM will:

- Review and update your I/O configuration (IODF)
- Assist you, as needed, with the OSA-Express implementation and ensure the console definitions are valid
- Review and update your VTAM® definitions
- Review and update your TCP/IP parameters settings
- Review the operating system (z/OS, z/VM, z/VSE, and Linux on System z) settings to ensure they are appropriate for the zEnterprise System
- Review the procedures for migrating the logical partition (LPAR) configurations to the new Hardware Management Console (HMC) and assist as needed
- Identify and order the operating system (z/OS, z/VM, and z/VSE) and subsystem (DB2, IMS™, and CICS) maintenance
- One day onsite within 120 days of contract signing to facilitate deployment.

The following optional activities can be acquired individually, using the recommended number of the onsite daily assist feature:

- **Two weeks of services to upgrade z/OS (within coexistence policy) to the current version and release** - Applies skilled IBM resources using best practices to migrate back-level z/OS versions (within coexistence policy) to work with the zEnterprise System. The contract can be expanded to include other migration activities such as: subsystems, operational migration, Independent Software Vendor (ISV), user-exits migration, testing, and deployment.

- **Two weeks of services for sysplex clients exploiting Parallel Sysplex InfiniBand using InfiniBand (PSIFB) with the zEnterprise System** - Assists you in planning, implementing, and migrating your ICB and ISC-3 configurations to PSIFB.
- **Two weeks of services to assist on other installation and Implementation requirements** - Assists you with infrastructure work specific to your requirements.

For more information about the services, contact your Business Partner or IBM representative, or send an email to

express@us.ibm.com

Accessibility by people with disabilities

A US Section 508 Voluntary Product Accessibility Template (VPAT) containing details on accessibility compliance can be requested at

http://www.ibm.com/able/product_accessibility/index.html

Product positioning

Updates to the zEnterprise 196 and the zEnterprise System are designed to bring you a whole new set of applications, new levels of scalability, security, and reliability, as well as new means to incorporate service management into your data center.

IBM zEnterprise continues to provide a smarter business infrastructure through augmentation of IBM zEnterprise Unified Resource Manager. Unified Resource Manager is designed to virtualize and manage all resources of a System z ensemble as a single pool of resources, integrating system and workload management across the multisystem, multitier, multiarchitecture environment.

Support of select IBM System x blades in the zBX allows the zEnterprise to access a whole new application portfolio. Front-end applications that need access to centralized data serving would be a good fit for running on the blades, as well as applications that are a front end to core CICS or IMS transaction processing such as IBM WebSphere.

Security enhancements to the Common Cryptographic Architecture and the Crypto Express3 coprocessor of the z196 as well as a new level of Trusted Key Entry Licensed Internal Code provide enhanced hardware encryption and decryption.

The new industry-standard PCIe I/O subsystem in the z196 improves robustness and resiliency, and combined with High Performance FICON for System z (zHPF) offers excellent performance improvements. The new PCIe I/O drawer is designed to deliver nondisruptive add and repair and lower power requirements compared to the I/O cage or I/O drawer. Increased port granularity is also available with the new features being offered for the PCIe I/O drawer.

For demanding data-sharing workloads, enhancements to Parallel Sysplex 12x InfiniBand (12x IFB) coupling links have been designed to provide 40% faster synchronous service times (12x IFB3 protocol) than the current 12x IFB protocol. And, you now have more physical connectivity for 1x InfiniBand coupling links with double the number of ports. In addition, to improve link utilization and coupling throughput at extended distances, there are now up to 32 subchannels (devices) per CHPID for 1x InfiniBand coupling links.

Statement of general direction

Application Program Interfaces (APIs) for Unified Resource Manager: IBM intends to offer APIs for IBM zEnterprise Unified Resource Manager. These APIs are designed to provide access to the same underlying functions that support the

Unified Resource Manager user interface and can be exploited to enable discovery, monitoring, and provisioning use cases.

IBM intends to enhance the Tivoli Integrated Service Management for System z portfolio to take advantage of the additional zEnterprise ensemble monitoring and management information provided by the Unified Resource Manager APIs.

Dynamic discovery of storage resources by Unified Resource Manager:

IBM intends to offer dynamic discovery of storage resources by Unified Resource Manager. A server administrator will be able to trigger discovery of additional storage resources through the user interface of Unified Resource Manager.

HiperSockets™ Completion Queue: IBM plans to support transferring HiperSockets messages asynchronously, in addition to the current synchronous manner on z196 and z114. This could be especially helpful in burst situations. The Completion Queue function is designed to allow HiperSockets to transfer data synchronously if possible and asynchronously if necessary, thus combining ultra-low latency with more tolerance for traffic peaks. HiperSockets Completion Queue is planned to be supported in the z/VM and z/VSE environments in a future deliverable.

HiperSockets integration with the IEDN: Within a zEnterprise environment, it is planned for HiperSockets to be integrated with the intraensemble data network (IEDN), extending the reach of the HiperSockets network outside of the central processor complex (CPC) to the entire ensemble, appearing as a single Layer 2 network. HiperSockets integration with the IEDN is planned to be supported in z/OS V1.13 and z/VM in a future deliverable.

The IBM zEnterprise 196 and the IBM zEnterprise 114 are the last System z servers to support ESCON channels. IBM plans to not offer ESCON channels as an orderable feature on System z servers that follow the z196 (machine type 2817) and z114 (machine type 2818). In addition, ESCON channels cannot be carried forward on an upgrade to such follow-on servers. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to features 2323 and 2324.

System z customers should continue to eliminate ESCON channels from the mainframe wherever possible. Alternate solutions are available for connectivity to ESCON devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to FICON. This offering is designed to help customers to simplify and manage a single physical and operational environment - FICON channels on the mainframe with continued connectivity to ESCON devices.

This is a restatement of general direction already published. Refer to Hardware Announcement [ZG11-0019](#), dated February 15, 2011.

The IBM zEnterprise 196 and the zEnterprise 114 are the last System z servers to support the Power Sequence Controller (PSC) feature. IBM intends to not offer support for the PSC (feature 6501) on future System z servers after the z196 (machine type 2817) and z114 (machine type 2818). PSC features cannot be ordered and cannot be carried forward on an upgrade to such a follow-on server.

This is a revision to the PSC statement of general direction published July 22, 2010. Refer to Hardware Announcement [ZG10-0249](#), dated July 22, 2010, IBM zEnterprise System -- A new dimension in computing. At that time we indicated the PSC feature could not be ordered on future servers. We are further clarifying that PSC features will also not be supported as carry forward on an upgrade.

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of

any future features or functionality described for our products remains at our sole discretion.

Product number

Description	Machine type	Model	Feature
zEnterprise 196	2817	M15 M32 M49 M66 M80	
Manage FW System x Blade			0042
Adv Mgmt FW System x Blade			0046
PCIe fanout			0169
HCA3-0 LR fanout for 1x IFB			0170
HCA3-0 fanout for 12x IFB			0171
PCIe interconnect			0400
OSA-Express4S GbE LX			0404
OSA-Express4S GbE SX			0405
OSA-Express4S 10 GbE LR			0406
OSA-Express4S 10 GbE SR			0407
FICON Express8S 10KM LX			0409
FICON Express8S SX			0410
TKE 7.1 LIC			0867
PCIe I/O drawer-A			4003
PCIe I/O drawer-w			4006

Description	Machine type	Model	Feature
System z10 EC	2097	E12 E26 E40 E56 E64	

TKE 7.1 LIC 0867

Description	Machine type	Model	Feature
System z10 BC	2098	E10	
TKE 7.1 LIC			0867

Publications

The following publications have been updated and are available now in the *Library* section of Resource Link®.

Title	Order number
z196 Installation Manual for Physical Planning	GC28-6897
z196 System Overview	SA22-1086
Functional Matrix	ZSW0-1335
IOCP User's Guide	SB10-7037
PR/SM Planning Guide	SB10-7155
CHPID Mapping Tool User's Guide	GC28-6900

The following publications are shipped with the product. They have been updated and are available in the *Library* section of Resource Link.

Title	Order number
z196 Safety Inspection	GC28-6894
Systems Safety Notices	G229-9054

System z Statement of Limited Warranty	GC28-6883
z196 Installation Manual	GC28-6890
z196 Service Guide	GC28-6892
Service Guide for TKE Workstations (Version 7.0)	GC28-6901
Systems Environmental Notices and User Guide	Z125-5823

The following publications have been updated and will be available at planned availability in the *Library* section of Resource Link.

Title	Order number
System z Application Programming Interfaces	SB10-7030
System z Application Programming Interfaces for Java™	API-JAVA
System z CIM Management Interface	SB10-7154
Capacity On Demand User's Guide	SC28-2605
Service Guide for HMCs and SES	GC28-6861
Hardware Management Console Operations Guide (V2.11.1)	SC28-6905
Support Element Operations Guide (V2.11.1)	SC28-6096
Stand-Alone IOCP User's Guide	SB10-7152
Coupling Links I/O Interface Physical Layer	SA23-0395
Fibre Channel Connection (FICON) I/O Physical Layer	SA24-7172
ESCON and FICON CTC Reference	SB10-7034
Maintenance Information for Fiber Optic Links	SY27-2597
Planning for Fiber Optic Links	GA23-0367
Fiber Optic Cleaning Procedures	SY27-2604
Introduction to Ensembles	GC27-2609
Ensemble Planning and Configuring Guide	GC27-2608
Ensemble Performance Management Guide	GC27-2607
HMC Operations Guide for Ensembles (V2.11.1)	SC27-2615
z196 Parts Catalog	GC28-6898
OSA-Express Customer Guide and Reference	SA22-7935

Publications can be obtained at Resource Link by accessing the following website

<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.

The following Redbooks® have been updated and are available now.

Title	Order number
zEnterprise Technical Introduction	SG24-7832
zEnterprise Technical Guide	SG24-7833
zEnterprise Unified Resource Manager	SG24-7921
System z Connectivity Handbook	SG24-5444

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

Specified operating environment

Physical specifications

Approximate weight for the zEnterprise 196 (z196) machine type 2817:

	New Build Minimum System Model M15 One I/O Cage	New Build Maximum System Model M80 Max # of I/O Cages
System with IBF Feature		
- kg	1448.0	2489.0
- lb	3258.0	5600.0
- kg (O/H IO cable exit)	1532.4	2573.0
- lb (O/H IO cable exit)	3448.0	5790.0
System without IBF Feature		
- kg	1248.0	2191.0
- lb	2807.0	4930.0
- kg (O/H IO cable exit)	1332.4	2278.0
- lb (O/H IO cable exit)	2997.0	5125.0

"O/H IO cable exit" refers to Top Exit I/O cabling.

Standards

The FICON Express8S features are designed to conform to the following standards:

- Fibre Channel Single-Byte Command Code Sets-3 (FC-SB-3)
The INCITS 374-2003, Information Technology - Fibre Channel Single-Byte Command Code Sets-3 (FC-SB-3) standard describes the channel mapping protocol associated with the Single-Byte Command Code Sets.

This standard was developed by Task Group T11 of Accredited Standards Committee INCITS during 2002-2003. The standards approval process started in 2002.

SB-3 is a mapping protocol, referred to as an FC-4. An FC-4 is a mapping protocol that maps a particular Upper Level Protocol (ULP) instance to Fibre Channel. The SB-3 ULP is based on the SB-2 ULP, which is based on the Single-Byte Command Code Set.
- Fibre Channel - Framing and Signaling (FC-FS)
The INCITS 373, Information Technology - Fibre Channel - Framing and Signaling (FC-FS) standard provides a general transport vehicle for Upper Level Protocols (ULPs) (for example, Small Computer System Interface (SCSI) command sets, Internet Protocol (IP), and others).
- FCP - Fibre Channel Protocol

The ANSI INCITS Project 1144D "Fibre Channel Protocol for SCSI, Second Version (FCP-2)" is the FCP analogy to FC-SB-3 and describes the mapping of SCSI onto Fibre Channel.

- Fibre Channel-Physical Interface-4 (FC-PI-4) revision 8.00 to INCITS

The FICON Express8S features have the following characteristics:

FICON Express8S 10KM LX (#0409)

- Data rate: 2, 4, or 8 Gbps
- Defined as: CHPID types FC (FICON, zHPF, CTC) and FCP (communication with SCSI devices)
- Connector type: LC Duplex
- Port count: Two channels per feature
- Cable type: Single mode fiber optic cabling (9 micron)

FICON Express8S SX (#0410)

- Data rate: 2, 4, or 8 Gbps
- Defined as: CHPID types FC (FICON, zHPF, CTC) and FCP (communication with SCSI devices)
- Connector type: LC Duplex
- Port count: Two channels per feature
- Cable type: Multimode fiber optic cabling (50 or 62.5 micron)

For details regarding the unrepeated distances for FICON Express8S 10KM LX and FICON Express8S SX refer to *System z Planning for Fiber Optic Links (GA23-0367)* available on zEnterprise 196 at planned availability in the *Library* section of Resource Link

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

The OSA-Express4S 10 Gigabit Ethernet and Gigabit Ethernet features are designed to conform to the following standards:

- Gigabit Ethernet (1000BASE-SX, 1000BASE-LX) IEEE 802.3ac IEEE 802.1q IEEE 802.3x IEEE 802.3z DIX Version 2
- 10 Gigabit Ethernet LR (10GBASE-LR) and SR (10GBASE-SR) IEEE 802.3ae IEEE 802.1q IEEE 802.3x - flow control DIX Version 2

The OSA-Express4S features have the following characteristics:

OSA-Express4S Gigabit Ethernet LX (long wavelength) (#0404)

- Data rate: 1000 Mbps (1 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 LX ports per feature
- Cable type: Single mode fiber optic cabling (9 micron)
- Unrepeated distance: 5 km (3.1 miles)

OSA-Express4S Gigabit Ethernet SX (short wavelength) (#0405)

- Data rate: 1000 Mbps (1 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 SX ports
- Cable type: Multimode fiber optic cabling (50 or 62.5 micron)
- Unrepeated distance:
 - With 50 micron fiber at 500 MHz-km: 550 meters (1804 feet)
 - With 62.5 micron fiber at 200 MHz-km: 275 meters (902 feet)
 - With 62.5 micron fiber at 160 MHz-km: 220 meters (722 feet)

OSA-Express4S 10 Gigabit Ethernet LR (long reach) (#0406)

- Data rate: 10 gigabits per second (10 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID types OSD and OSX
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: One LR port
- Cable type: Single mode fiber optic cabling (9 micron)
- Unrepeated distance: 10 km (6.2 miles)

OSA-Express4S 10 Gigabit Ethernet SR (short reach) (#0407)

- Data rate: 10 gigabits per second (10 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID types OSD and OSX
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: One SR port
- 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)

Operating environment

Hardware requirements

The hardware requirements for the machine type, models, features, and functions contained in this announcement are identified.

Machine Change Levels (MCLs) are required. Descriptions of the MCLs are available now on Resource Link.

Access Resource Link at

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

MCLs are designed to be applied concurrently. Contact IBM service personnel for further information.

Required Driver Level and MCLs to be made available at planned availability:

HMC (V2.11.1) plus MCLs
Support Element (V2.11.1) plus MCLs

You should review the PSP buckets for minimum Machine Change Levels (MCLs) and software PTF levels before IPLing operating systems.

Peripheral hardware and device attachments

IBM devices previously attached to IBM System z10, System z9, and zSeries® servers are supported for attachment to zEnterprise 196 channels, unless otherwise noted. The subject I/O devices must meet ESCON or FICON 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 zEnterprise 196 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.

Information on switches and directors qualified for IBM System z FICON and FCP channels can be found in the *Library* section of Resource Link

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

Software requirements

Listed are the operating systems and the minimum versions and releases supported by zEnterprise 196, its functions, and its features. Select the releases appropriate to your operating system environments.

Note: Refer to the z/OS, z/VM, z/VSE subsets of the 2817DEVICE Preventive Service Planning (PSP) bucket prior to installing a zEnterprise 196.

Regarding support of z/OS, the notes below may apply:

(1) z/OS V1.10 supports zEnterprise 196. However, z/OS V1.10 support will be withdrawn September 30, 2011. After that date, the IBM z/OS Lifecycle Extension for z/OS V1.10 (5656-A01) is required for z/OS V1.10 support with the zEnterprise 196. Talk to your IBM representative for details. For the complete list of software support, see the PSP buckets. For more information on the IBM Lifecycle Extension for z/OS V1.10, refer to Software Announcement [ZP11-0006](#), dated February 15, 2011.

(2) z/OS V1.9 support was withdrawn September 30, 2010. However, with the z/OS Lifecycle Extension (5646-A01), z/OS V1.9 supports the zEnterprise 196. Talk to your IBM representative for details. No exploitation of new zEnterprise System functions is available with z/OS V1.9. Certain functions and features of the zEnterprise 196 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 V1.9, refer to Software Announcement [ZP10-0028](#), dated May 11, 2010.

(3) z/OS V1.8 support was withdrawn September 30, 2009. However, with the z/OS Lifecycle Extension (5638-A01), z/OS V1.8 supports the zEnterprise 196. Talk to your IBM representative for details. No exploitation of new zEnterprise 196 functions is available with z/OS V1.8. Certain functions and features of the zEnterprise 196 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 V1.8, refer to Software Announcement [ZS10-0086](#), dated June 15, 2010.

zEnterprise 196, for the contents of this announcement, requires at a minimum:

- z/VM V5.4 with PTFs.
- z/VSE V4.2 and V4.3 with PTFs.

All other operating system versions and releases remain unchanged from original planned availability of z196. Refer to Hardware Announcement [ZG10-0249](#), dated July 22, 2010, IBM zEnterprise System -- A new dimension in computing.

Expanded support for AES algorithm, Enhanced ANSI TR-31 Secure Key Exchange, PIN block decimalization table protection, and Additional Elliptic Curve Cryptography (ECC) functions for the Crypto Express3 feature (when defined as a coprocessor) require at a minimum:

- z/OS V1.11, z/OS V1.12, or z/OS V1.13 with the Cryptographic Support for z/OS V1R11-V1R13 web deliverable (planned to be available September 9, 2011). This may be obtained at <http://www.ibm.com/systems/z/os/zos/downloads/>
- z/VM V5.4 with PTFs for guest exploitation.
- Linux on System z:
 - Support is planned for a future release of Common Cryptographic Architecture Support Program for Linux on System z.

PKA RSA OAEP with SHA-256 algorithm for the Crypto Express3 feature (when defined as a coprocessor) requires one of the following at a minimum:

- z/OS V1.13 with a PTF planned to be available September 9, 2011.
- z/OS V1.11, z/OS V1.12, or z/OS V1.13 with the Cryptographic Support for z/OS V1R11-V1R13 web deliverable planned to be available September 9, 2011.
- z/OS V1.10 (1), z/OS V1.11, or z/OS V1.12 with the Cryptographic Support for z/OS V1R10-V1R12 web deliverable with a PTF planned to be available September 9, 2011.
- z/VM 5.4 for guest exploitation.

The Cryptographic web deliverables are available at

<http://www.ibm.com/systems/z/os/zos/downloads/>

IPL from an alternate subchannel set requires at a minimum:

- z/OS V1.13.
- z/OS V1.11 or V1.12 with PTFs.

FICON Express8S (CHPID type FC) when utilizing FICON or Channel-To-Channel (CTC), requires at a minimum:

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - Novell SUSE Linux Enterprise Server (SLES) 10 and SLES 11.
 - Red Hat Enterprise Linux (RHEL) 5 and RHEL 6.

FICON Express8S (CHPID type FC) for support of zHPF single-track operations requires at a minimum:

- z/OS V1.11.
- z/OS V1.8 (3), V1.9 (2), or V1.10 (1) with PTFs.
- Linux on System z distributions:
 - SLES 11 SP1.
 - RHEL 6.0.

FICON Express8S (CHPID type FC) for support of zHPF multitrack operations requires at a minimum:

- z/OS V1.11 with PTFs.
- z/OS V1.9 (2) or V1.10 (1) with PTFs.
- Linux on System z distributions:
 - RHEL 6.1.
 - IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

FICON Express8S (CHPID type FCP) for support of SCSI devices requires at a minimum:

- z/VM V5.4 with PTFs.
- z/VSE V4.2.
- Linux on System z distributions:
 - SLES 10 and SLES 11.
 - RHEL 5 and RHEL 6.

FICON Express8S (CHPID type FCP) support of hardware data router requires at a minimum:

- Linux on System z distributions:
 - IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

T10-DIF support by the FICON Express8S and FICON Express8 features when defined as CHPID type FCP requires at a minimum:

- z/VM 5.4 with PTFs for guest exploitation.
- Linux on System z distributions:
 - IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

OSA-Express4S GbE LX (#0404) and GbE SX (#0405) require at minimum:

CHPID type OSD with exploitation of two ports per CHPID:

- z/OS V1.10 (1).
- z/OS V1.9 (2) with PTFs.
- z/OS V1.8 (3) with PTFs.
- z/VM V6.1.
- z/VM V5.4 with PTFs.
- z/VSE V4.2.
- z/TPF V1.1 PUT 4 with PTFs.
- Linux on System z distributions:
 - SLES 10 SP2 and SLES 11.
 - RHEL 5.2 and RHEL 6.0.

CHPID type OSD without maximum port exploitation (one port on the PCIe adapter is available for use):

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.

- Linux on System z distributions:
 - SLES 10 and SLES 11.
 - RHEL 5 and RHEL 6.

OSA-Express4S 10 GbE LR (#0406) and 10 GbE SR (#0407) require at a minimum:

CHPID type OSD:

- z/OS V1.8 (3).
- z/VM V5.4.
- z/VSE V4.2.
- z/TPF V1.1.
- Linux on System z distributions:
 - SLES 10 and SLES 11.
 - RHEL 5 and RHEL 6.

CHPID type OSX for access control to the intraensemble data network (IEDN) from z196 to Unified Resource Manager functions:

- z/OS V1.12 with PTFs.
- z/OS V1.11 with PTFs.
- z/OS V1.10 (1) with PTFs.
- z/VM V6.1 with PTFs.
- z/VM V5.4 with PTFs to define, modify, and delete OSX CHPID types when z/VM is the controlling LPAR for dynamic I/O.
- z/VSE V5.1 (when available).
- z/TPF V1.1 PUT 4 with PTFs.
- Linux on System z distributions:
 - SLES 10 SP4, SLES 11 SP1 (maintenance update).
 - RHEL 5.6 and RHEL 6.0.

Inbound workload queuing for Enterprise Extender for the OSA-Express4S and OSA-Express3 features (CHPID types OSD or OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 with PTFs for guest exploitation.

OSA-Express4S checksum offload for IPv6 packets (CHPID types OSD and OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 for guest exploitation.

OSA-Express4S checksum offload for LPAR-to-LPAR traffic for IPv4 and IPv6 packets (CHPID types OSD and OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 for guest exploitation.

OSA-Express4S Large Send for IPv6 packets (CHPID types OSD and OSX) requires at a minimum:

- z/OS V1.13.
- z/VM V5.4 for guest exploitation.

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 cabling for the FICON Express8S and OSA-Express4S features is a customer responsibility.

Cabling responsibilities

Fiber optic cables, cable planning, labeling, and placement are all customer responsibilities for new installations and upgrades. 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 - for new installations and upgrades:

- All fiber optic cable planning.
- All purchasing of correct fiber optic cables.
- 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 z196 installation if the above cabling tasks are not accomplished as required.

Fiber Quick Connect (FQC), a fiber harness integrated in the z196 frame for "quick" connect, is offered as a feature on the z196 for connection to ESCON and FICON LX channels.

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

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

Security, auditability, and control

The zEnterprise 196 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.

Global Technology Services

Contact your IBM representative for the list of selected services available in your country, either as standard or customized offerings, for the efficient installation, implementation, and/or integration of this product.

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

Field installable feature

Yes

Warranty period

One year.

An IBM part or feature installed during the initial installation of an IBM machine is subject to a full warranty effective on the date of installation of the machine. An IBM part or feature that replaces a previously installed part or feature assumes the remainder of the warranty period for the replaced part or feature. An IBM part or feature added to a machine without replacing a previously installed part or feature is subject to a full warranty effective on its date of installation. Unless specified otherwise, the warranty period, type of warranty service, and service level of a part or feature are the same as those for the machine in which it is installed.

Customer setup

No

IBM Licensed Internal Code (LIC) is licensed for use by a customer on a Specific Machine, designated by serial number, under the terms and conditions of the IBM License Agreement for Machine Code, to enable a Specific Machine to function in accordance with its Specifications, and only for the capacity authorized by IBM and acquired by the customer.

You can obtain the agreement at

http://www.ibm.com/systems/support/machine_warranties/machine_code.html

or by contacting your IBM representative.

Pricing

For all local charges, contact your IBM representative.

Announcement countries

All European, Middle Eastern, and African countries.

Trademarks

zEnterprise, Hypervisor, PR/SM, PowerVM, z10, z/VSE, z9, DFSMSHsm, IMS, HiperSockets and Electronic Service Agent are trademarks of IBM Corporation in the United States, other countries, or both.

IBM, System z, BladeCenter, FICON, z/OS, Parallel Sysplex, System x, z/VM, System z9, ESCON, System z10, DB2, System Storage, DataPower, CICS, WebSphere, VTAM, Resource Link, Redbooks and zSeries 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.

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/>