



IBM System z10 Enterprise Class — The forward-thinking mainframe for the twenty-first century

Key prerequisites 3
Description 3
Product positioning 20
Statement of general direction 20



At a glance

The IBM System z10 EC is a world-class enterprise server designed to meet your business needs. The System z10 EC is built on the inherent strengths of the IBM System z platform and is designed to deliver new technologies and virtualization that provide improvements in price/performance for key new workloads. The System z10 EC further extends System z leadership in key capabilities with the delivery of expanded scalability for growth and large-scale consolidation, improved security and availability to reduce risk, and just-in-time capacity deployment, helping to respond to changing business requirements. The System z10 EC delivers:

- Improved total system capacity in a 64-way server, offering increased levels of performance and scalability to help enable new business growth.
- z10 quad-core 4.4 GHz processor chips that can help improve the execution of CPU-intensive workloads.
- Up to 1.5 terabytes of available real memory per server for growing application needs (with up to 1 TB real memory per LPAR).
- Increased scalability with 36 available subcapacity settings.
- Just-in-time deployment of capacity resources which can improve flexibility when making temporary or permanent changes. Activation can be further simplified and automated using z/OS® Capacity Provisioning (available on z/OS V1.9 with PTF and on z/OS V1.10, when available).
- New temporary capacity offering Capacity for Planned Event (CPE), a variation of Capacity Back Up (CBU). CPE can be used when capacity is unallocated, but available, and is needed for a short-term event.
- A new 16 GB fixed Hardware System Area (HSA) which is managed separately from customer memory. This fixed HSA is designed to improve availability by avoiding outages.
- Memory and books that are interconnected with a point-to-point symmetric multi processor (SMP) network running with an InfiniBand host bus bandwidth at 6 GBps designed to deliver improved performance.
- The new InfiniBand Coupling Links (planned to be available second quarter 2008) with a link data rate of 6 GBps, designed to provide a high-speed solution and increased distance (150 meters) compared to ICB-4 (10 meters).
- The new OSA-Express3 10 GbE LR (planned to be available second quarter 2008) with double the port density, increased throughput, and reduced latency.
- HiperSockets improvements with Multiple Write Facility for increased performance and Layer 2 support to host IP and non-IP workloads.

- Encryption accelerator provided on quad-core chip, which is designed to provide high-speed cryptography for protecting data in storage. CP Assist for Cryptographic Function (CPACF) offers more protection and security options with Advanced Encryption Standard (AES) 192 and 256 and stronger hash algorithm with Secure Hash Algorithm (SHA-512 and SHA-384).
- HiperDispatch for improved efficiencies between hardware and the z/OS operating system (z/OS 1.7 and above).
- Hardware decimal floating point unit on each core on the Processor Unit (PU), which can aid in decimal floating point calculations and is designed to deliver performance improvements and precision in execution.
- Large page support (1 megabyte pages).
- Up to 336 FICON™ Express4 channels.
- Fiber Quick Connect (FQC), a fiber harness integrated in the System z10 EC frame for a 'quick' connect to ESCON® and FICON LX channels.
- Support for IBM Systems Director Active Energy Manager (AEM) for Linux™ on System z for a single view of actual energy usage across multiple heterogeneous IBM platforms within the infrastructure. AEM V3.1 is a key component of IBM's Cool Blue portfolio within Project Big Green.¹
- The IBM System z9 Enterprise Class (z9 EC) and System z9 Business Class (z9 BC) servers are the last servers to support participation in the same Parallel Sysplex with IBM eServer zSeries 900 (z90), IBM eServer zSeries 800 (z800), and older System/390 Parallel Enterprise Server systems. This direction was announced in Hardware Announcement [ZG07-0286](#), dated April 18, 2007.

¹

This satisfies the direction announced in Software Announcement [ZP07-0479](#), dated November 13, 2007, for IBM Systems Director Active Energy Manager for POWER™, V3.1, which stated: Future System z servers plan to support the monitoring functions of IBM Systems Director Active Energy Manager.

Overview

The System z10 EC is a marriage of evolution and revolution, building on the inherent strengths of the System z™ platform, delivering new technologies and virtualization that are designed to offer improvements in price / performance for key workloads as well as enabling a new range of solutions. The z10 EC further extends the leadership of System z in key capabilities with the delivery of expanded scalability for growth and large-scale consolidation, availability to help reduce risk and improve flexibility to respond to changing business requirements, and improved security. The z10 EC is at the core of the enhanced System z platform that is designed to deliver technologies that business needs today along with a foundation to drive future business growth.

With a modular book design, the z10 EC E64 is designed to provide up to 1.7 times the total system capacity of the z9 EC Model S54 and up to three times the available memory of the z9 EC. Significant steps have been taken in the area of server availability in the z10 EC design. Preplanning requirements are minimized by delivering a fixed, reserved Hardware System Area (HSA) and new capabilities intended to allow you to seamlessly create logical partitions (LPARs), include logical subsystems, change logical processor definitions in an LPAR, and add cryptographic capabilities for an LPAR without a power-on reset.

z10 EC introduces just-in-time deployment of capacity resources designed to provide more flexibility to dynamically change capacity when business requirements change. You are no longer limited by one offering configuration; instead you can define one or more flexible configurations that can be used to solve multiple temporary situations. You can now have multiple configurations active at once and the configurations themselves are flexible so you can activate only what is needed from your defined configuration. As long as your total z10 EC infrastructure can support the maximums that are defined, they can be delivered. A significant change is the ability to add permanent capacity to the server when you are in a temporary state. The combination of these updates can change the way you think about on demand capacity.

New integrated clear-key encryption security features on z10 EC include support for a higher advanced encryption standard and more secure hashing algorithms. Performing these functions in hardware is designed to contribute to improved performance.

Integrated on the z10 EC processor unit is a Hardware Decimal Floating Point unit to accelerate decimal floating point transactions. This function is designed to markedly improve performance for decimal floating point operations which offer increased precision compared to binary floating point operations. This is expected to be particularly useful for the calculations involved in many

financial transactions.

New innovations on the z10 EC are designed to give needed capacity and memory along with the just-in-time management of resources. Advanced virtualization technologies aid in server consolidation, satisfying high I/O requests and dynamic provisioning of new servers.

IBM Global Financing can provide attractive low-rate financing for all new and upgraded z10 EC products, storage, software, and services. For more information, contact your local Global Financing sales representative or visit the Web site

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

Available worldwide for eligible customers acquiring products and services from IBM and IBM Business Partners.

Key prerequisites

Refer to the Hardware requirements and Software requirements sections of this announcement.

Planned availability dates

- Features and functions for the System z10 EC: February 26, 2008
- System z10 EC Models E12, E26, E40, E56, and E64: February 26, 2008
- z990 upgrades to System z10 EC: February 26, 2008
- z9 EC upgrades to System z10 EC: February 26, 2008
- MES features for Models E12, E26, E40, E56, and E64: May 26, 2008
- Model conversions for System z10 EC: May 26, 2008
- InfiniBand coupling links (#0163): Second quarter 2008
- System z9™ EC InfiniBand coupling links (#0167): Second quarter 2008
- System z9 BC InfiniBand coupling links (#0167): Second quarter 2008
- OSA-Express3 10 GbE LR (#3370): Second quarter 2008

Availability of programs with an encryption algorithm in France is subject to French government approval.

Description

The IBM System z10 Enterprise Class — A total systems approach to deliver leadership in enterprise computing: With a total systems approach designed to deploy innovative technologies, IBM System z introduces the z10 EC, supporting z/Architecture™, and offering the highest levels of reliability, availability, scalability, clustering, and virtualization. The z10 EC just-in-time deployment of capacity allows improved flexibility, administration, and the ability to enable changes as they happen. The expanded scalability on the z10 EC facilitates growth and large-scale consolidation. The z10 EC is designed to provide:

- Uniprocessor performance improvement up to 62% (based on LSPR mixed workload average)
- Non-uniprocessor performance improvement up to 50% (based on LSPR mixed workload average) for configurations with the same number of processors
- Up to 1.7 times the total system capacity of the z9 EC
- Up to 64 Processor Units (PUs) compared to a maximum of 54 on the z9 EC
- Up to 3 times as much total server available memory — up to 1.5 terabytes of total memory
- Up to 50% more subcapacity choices as compared to z9 EC
- Increased host base bandwidth using InfiniBand at 6 GBps
- Coupling with InfiniBand for improved distance and potential cost saving
- Performance improvements with HiperSockets Multiple Write Facility

- Improved Advanced Encryption Standard (AES) 192 and 256 and stronger hash algorithms with Secure Hash Algorithm (SHA) 384 and 512
- HiperDispatch for improved efficiencies between hardware and the z/OS operating system (z/OS 1.7 and above)
- Hardware Decimal Floating Point unit for improved numeric processing performance
- Reduction in the availability impact of preplanning requirements
 - Fixed Hardware System Area (HSA) designed so the maximum configuration capabilities can be exploited
 - Designed to reduce the number of planned Power-on-Resets
 - Designed to allow dynamic add/remove of a new logical partition (LPAR) to new or existing logical channel subsystem (LCSS)
- Open Systems Adapter-Express3 (OSA-Express3) 10 Gigabit Ethernet with double the port density and improved performance
- Up to 336 FICON channels
- Large page support (1 megabyte pages)
- Energy efficiency displays on System Activity Display (SAD) screens
- Just-in-time deployment of capacity for faster activation without dependency or referral to IBM
- Store System Information (STSI) change to support billing methodologies
- New temporary offering Capacity for Planned Event (CPE) available to manage system migrations, data center moves, maintenance activities, and similar situations
- Support for the IBM Systems Director Active Energy Manager (AEM) for Linux on System z

Model summary matrix

Model	PU _s	Memory	IB	I/O cages	CHPIDs
E12	1 to 12	16 to 352 GB	16	1 to 3	960
E26	1 to 26	16 to 752 GB	32	1 to 3	1024
E40	1 to 40	16 to 1136 GB	40	1 to 3	1024
E56	1 to 56	16 to 1520 GB	48	1 to 3	1024
E64	1 to 64	16 to 1520 GB	48	1 to 3	1024

Notes:

- Memory reserved for the fixed HSA is in addition to the purchased entitlement.
- The addition of the third and fourth books requires a reduction in the number of fanout cards plugged, to increase cooling around the MCM.
- Each LCSS supports up to 256 CHPIDs.

The performance advantage

IBM's Large Systems Performance Reference (LSPR) method is designed to provide comprehensive z/Architecture processor capacity ratios for different configurations of Central Processors (CPs) across a wide variety of system control programs and workload environments. For z10 EC, z/Architecture processor capacity indicator is defined with a (7XX) notation, where XX is the number of installed CPs.

In addition to the general information provided for z/OS 1.8, the LSPR also contains performance relationships for z/VM® and Linux operating environments.

Based on using an LSPR mixed workload, the performance of the z10 EC (2097) 701 is expected to be:

- Up to 1.62 times that of the z9 EC (2094) 701, and
- Up to 1.50 times the z9 EC for non-uniprocessor environments, assuming equal numbers of processors

Moving from a System z9 partition to an equivalently-sized System z10 partition, a z/VM workload will experience an ITR ratio that is somewhat related to the workload's instruction mix, MP factor, and level of storage over commitment. Workloads with higher levels of storage over commitment or higher MP factors are likely to experience lower than average z10 EC to z9 ITR scaling ratios. The range of likely ITR ratios is wider than the range has been for previous processor migrations.

The LSPR contains the Internal Throughput Rate Ratios (ITRRs) for the new z10 EC and the previous-generation zSeries® processor families based upon measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user may experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, and the workload processed. Therefore no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated. For more detailed performance information, consult the Large Systems Performance Reference (LSPR) available at

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

HiperDispatch

A z10 EC exclusive, HiperDispatch represents a cooperative effort between the z/OS operating system and the z10 EC hardware and is intended to provide improved efficiencies in both the hardware and the software in the following ways:

- Work may be dispatched across fewer logical processors therefore reducing the multi-processor (MP) effects and lowering the interference among multiple partitions.
- Specific z/OS tasks may be dispatched to a small subset of logical processors which Processor Resource/Systems Manager™ (PR/SM™) will tie to the same physical processors, thus improving the hardware cache re-use and locality of reference characteristics such as reducing the rate of cross-book communication.

Refer to the Software requirements section.

Hardware decimal floating point

Focused performance boost — hardware decimal floating point

Recognizing that speed and precision in numerical computing are essential, with the introduction of z10 EC each core on the PU has its own hardware decimal floating point unit, which is designed to improve performance of decimal floating point over that provided by System z9.

Decimal calculations are often used in financial applications and those done using other floating point facilities have typically been performed by software through the use of libraries. With a hardware decimal floating point unit, some of these calculations may be done directly and accelerated.

Software support for hardware decimal floating point on z10 EC is provided in several programming languages. Support is provided in Assembler Language in Release 4 or 5 of High Level Assembler. Decimal floating point data and instructions are also supported in Enterprise PL/1 V3.7 and resulting programs can be debugged by Debug Tool V8.1. Java™ applications, which make use of the BigDecimal Class Library, will automatically begin using the hardware decimal floating point instructions when running on a z10 EC. Support for decimal floating point data types is also provided in SQL as provided in DB2® Version 9. Refer to the Software requirements section.

Large page support for 1 megabyte pages

A change to the z/Architecture on z10 EC is designed to allow memory to be extended to support large (1 megabyte (MB)) pages. When large pages are used, in addition to the existing 4 KB page size, they are expected to reduce memory management overhead for exploiting applications.

Large page support is primarily of benefit for long running applications that are memory-access-intensive. Large page is not recommended for general use. Short lived processes with small working sets are normally not good candidates for large pages.

Large page support is exclusive to z10 EC and to z/OS. Refer to the Software requirements section.

Flexible memory

Flexible memory was first introduced on the z9 EC as part of the design changes and offerings to support enhanced book availability. Flexible memory provides the additional resources to maintain a constant level of memory when replacing a book. On z10 EC, the additional resources required for the flexible memory configurations are provided through the purchase of preplanned memory features (#1996) along with the purchase of your memory entitlement. In most cases, this implementation provides a lower-cost solution compared to z9 EC. Flexible memory configurations are available on Models E26, E40, E56, and E64 only and range from 32 GB to 1136 GB, model dependent. Contact your IBM representative to help you determine the appropriate configuration.

Cryptographic support for security-rich transactions

CP Assist for Cryptographic Function (CPACF): CPACF supports clear-key encryption. All CPACF functions can be invoked by problem state instructions defined by an extension of System z architecture. The function is activated using a no-charge enablement feature (#3863) and offers the following on every CPACF that is shared between two Processor Units (PUs) and designated as CPs and/or Integrated Facility for Linux (IFL):

- Data Encryption Standard (DES)
- Triple Data Encryption Standard (TDES)
- Advanced Encryption Standard (AES) for 128-bit keys
- Secure Hash Algorithm, SHA 1, SHA-224, and SHA-256
- Pseudo Random Number Generation (PRNG)

Enhancements to CP Assist for Cryptographic Function (CPACF): CPACF has been enhanced to include support of the following on CPs and IFLs:

- Advanced Encryption Standard (AES) for 192-bit keys and 256-bit keys
- SHA-384 and SHA-512 bit for message digest

SHA-1, SHA-256, and SHA-512 are shipped enabled and do not require the enablement feature.

Support for CPACF is also available using the Integrated Cryptographic Service Facility (ICSF). ICSF is a component of z/OS, and is designed to transparently use the available cryptographic functions, whether CPACF or Crypto Express2, to balance the workload and help address the bandwidth requirements of your applications.

The enhancements to CPACF are exclusive to the System z10 EC and supported by z/OS, z/VM, z/VSE®, and Linux on System z. Refer to the Software requirements section.

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

Crypto Express2 Coprocessor — for secure-key encrypted transactions (default) is:

- Designed to support security-rich cryptographic functions, use of secure-encrypted-key values, and User Defined Extensions (UDX)
- Designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification

Crypto Express2 Accelerator — for Secure Sockets Layer (SSL) acceleration:

- Is designed to support clear-key RSA operations
- Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol

Crypto Express2 features can be carried forward on an upgrade to the new System z10 EC, so users may continue to take advantage of the SSL performance and the configuration capability.

The configurable Crypto Express2 feature is supported by z/OS, z/VM, z/VSE, and Linux on System z. z/VSE offers support for clear-key SSL transactions only. Current versions of z/OS, z/VM, and Linux on System z offer support for both clear-key and secure-key operations.

Refer to the Software requirements section and also the Special features section of the Sales manual on the Web for further information.

<http://w3-3.ibm.com/sales/ssi/OIAccess.wss>

Additional cryptographic functions and features

Key management for remote loading of ATM and Point of Sale (POS) keys: The elimination of manual key entry is designed to reduce downtime due to key entry errors, service calls, and key management costs.

Improved key exchange with non-CCA cryptographic systems

New features added to IBM Common Cryptographic Architecture (CCA) are designed to enhance the ability to exchange keys between CCA systems, and systems that do not use control vectors by allowing the CCA system owner to define permitted types of key import and export while preventing uncontrolled key exchange that can open the system to an increased threat of attack.

These are supported by z/OS and by z/VM for guest exploitation. Refer to the Software requirements section.

Support for ISO 16609 CBC Mode T-DES Message Authentication (MAC) requirements

ISO 16609 CBC Mode T-DES MAC is accessible through ICSF function calls made in the PCI-X Cryptographic Adapter segment 3 Common Cryptographic Architecture (CCA) code.

This is supported by z/OS and by z/VM for guest exploitation. Refer to the Software requirements section.

Introducing support for RSA keys up to 4096 bits

The RSA services in the CCA API are extended to support RSA keys with modulus lengths up to 4096 bits. The services affected include key generation, RSA-based key management, digital signatures, and other functions related to these.

Refer to the ICSF Application Programmers Guide (SA22-7522) for additional details.

New crypto availability enhancement

Dynamically add crypto to a logical partition: Today, users can preplan the addition of Crypto Express2 features to a logical partition (LP) by using the Crypto page in the image profile to define the Cryptographic Candidate List, Cryptographic Online List, and Usage and Control Domain Indexes in advance of crypto hardware installation.

With the change to dynamically add crypto to a logical partition, changes to image profiles, to support Crypto Express2 features, are available without outage to the logical partition. Users can also dynamically delete or move Crypto Express2 features. Preplanning is no longer required.

This enhancement is supported by z/OS, z/VM for guest exploitation, and Linux on System z. Refer to the Software requirements section.

Continued support for TKE workstation and Smart Card Reader

TKE 5.2 workstation to enhance security and convenience: The Trusted Key Entry (TKE) workstation (#0839) and the TKE 5.2 level of Licensed Internal Code (#0857) are optional features on the System z10 EC. The TKE 5.2 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 three TKE workstations can be ordered.

The TKE workstation, feature #0839, is available on the System z10 EC, z9 EC, z9 BC, z990, and z890.

Refer also to the Special features section of the Sales manual on the Web for further information.

<http://w3-3.ibm.com/sales/ssi/OIAccess.wss>

Smart Card Reader: Support for an optional Smart Card Reader attached to the TKE 5.2 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).

TKE 5.2 Licensed Internal Code (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 5.2 LIC has limited the use of floppy diskettes to read-only. The TKE 5.2 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 optional TKE features are:

- TKE 5.2 LIC (#0857) and TKE workstation (#0839)
- TKE Smart Card Reader (#0887)
- TKE additional smart cards (#0888)

The Smart Card Reader, which can be attached to a TKE workstation with the 5.2 level of LIC, is available on the System z10 EC, z9 EC, z9 BC, z990, and z890.

System z10 EC cryptographic migration: Clients using a User Defined Extension (UDX) of the Common Cryptographic Architecture should contact their UDX provider for an application update before ordering a new System z10 EC machine, or planning to migrate or activate a UDX application to firmware driver 73 level.

- The Crypto Express2 feature is supported on the System z10 EC and can be carried forward on an upgrade to the System z10 EC.
- You must use TKE 5.2 workstations to control the System z10 EC.
- TKE 5.0 and 5.1 workstations (#0839 and #0859) may be used to control z9 EC, z9 BC, z890, and z990 servers.

FICON and FCP for connectivity to disk, tape, and printers

Extended distance FICON — Improved performance at extended distance: An enhancement to the industry standard FICON architecture (FC-SB-3) helps avoid degradation of performance at extended distances by implementing a new protocol for "persistent" Information Unit (IU) pacing. Control units that exploit the enhancement to the architecture can increase the pacing count (the number of IUs allowed to be in flight from channel to control unit). Extended distance FICON also allows the channel to "remember" the last pacing update for use on subsequent operations to help avoid degradation of performance at the start of each new operation.

Improved IU pacing can help to optimize the utilization of the link, for example help keep a 4 Gbps link fully utilized at 50 km, and allows channel extenders to work at any distance, with performance results similar to that experienced when using emulation.

The requirements for channel extension equipment are simplified with the increased number of commands in flight. This may benefit z/OS Global Mirror (Extended Remote Copy — XRC) applications as the channel extension kit is no longer required to simulate specific channel commands. Simplifying the channel extension requirements may help reduce the total cost of ownership of end-to-end solutions.

Extended distance FICON is transparent to operating systems and applies to all the FICON Express2 and FICON Express4 features carrying native FICON traffic (CHPID type FC). For exploitation, the control unit must support the new IU pacing protocol. The channel will default to current pacing values when operating with control units that cannot exploit extended distance FICON.

Exploitation of extended distance FICON is supported by IBM System Storage™ DS8000™ series Licensed Machine Code (LMC) level 5.3.1xx.xx (bundle version 63.1.xx.xx), or later.

Note: To support extended distance without performance degradation, the buffer credits in the FICON director must be set appropriately. The number of buffer credits required is dependent upon the link data rate (1 Gbps, 2 Gbps, or 4 Gbps), the maximum number of buffer credits supported by the FICON director or control unit, as well as application and workload characteristics. High bandwidth at extended distances is achievable only if enough buffer credits exist to support the link data rate.

FCP — Increased performance for small block sizes: The Fibre Channel Protocol (FCP) Licensed Internal Code has been modified to help provide increased I/O operations per second for small block sizes. With FICON Express4, there may be up to 57,000 I/O operations per second (all reads, all writes, or a mix of reads and writes), a 80% increase compared to System z9. These results are achieved in a laboratory environment using one channel configured as CHPID type FCP with no other processing occurring and do not represent actual field measurements. A significant increase in I/O operations per second for small block sizes can also be expected with FICON Express2.

This FCP performance improvement is transparent to operating systems and applies to all the FICON Express4 and FICON Express2 features when configured as CHPID type FCP, communicating with SCSI devices.

SCSI IPL now a base function: The SCSI Initial Program Load (IPL) enablement feature #9904, first introduced on z990 in October of 2003, is no longer required. The function is now delivered as a part of the server Licensed Internal Code. SCSI IPL allows an IPL of an operating system from an FCP-attached SCSI disk.

FICON Express4 — 1, 2, or 4 Gbps:

- Offers two unrepeated distance options (4 kilometer and 10 kilometer) when using single-mode fiber optic cabling
- Supports a 4 Gbps link data rate with auto-negotiation to 1 or 2 Gbps for synergy with existing switches, directors, and storage devices

The FICON Express4 features have two modes of operation designed for connectivity to servers, switches, directors, disks, tapes, and printers:

1. Native FICON and FICON Channel-to-Channel (CTC) traffic (CHPID type FC) in the z/OS, z/VM, z/VSE, z/TPF, TPF, and Linux on System z environments
2. Fibre Channel Protocol traffic (CHPID type FCP) in the z/VM, z/VSE, and Linux on System z environments

There are three FICON Express4 features from which to choose:

1. FICON Express4 10KM LX (#3321), with four channels per feature, is designed to support unrepeated distances up to 10 kilometers (6.2 miles) over single-mode fiber optic cabling. Use this feature when the unrepeated distance between devices is greater than 4 kilometers (km) or your link loss budget requirements exceed 2 dB end-to-end between devices.
2. FICON Express4 SX (#3322), with four channels per feature, is designed to carry traffic over multimode fiber optic cabling. Refer to the Standards section for the supported unrepeated distances.
3. FICON Express4 4KM LX (#3324), with four channels per feature, is designed to support unrepeated distances up to 4 kilometers (2.5 miles) over single-mode fiber optic cabling. This feature is designed to offer a cost-effective solution to satisfy the majority of your FICON/FCP single-mode fiber optic cabling distance requirements.

Note: The ANSI Fibre Channel Physical Interface (FC-PI-2) standard defines 10 kilometer (km) transceivers and 4 km transceivers when using 9 micron single-mode fiber optic cabling. IBM supports these FC-PI-2 variants.

IBM supports interoperability of 10 km transceivers with 4 km transceivers provided the unrepeated distance between a 10 km transceiver and a 4 km transceiver does not exceed 4 km (2.5 miles).

The FICON Express4 features have Small Form Factor Pluggable (SFP) optics to permit each channel to be individually serviced in the event of a fiber optic module failure. The traffic on the other channels on the same feature can continue to flow if a channel requires servicing.

All channels on a single FICON Express4 feature are of the same type — 4KM LX, 10KM LX, or SX. You may carry your current FICON Express2 and FICON Express features (#3319, #3320, #2319, #2320) forward from z990 or z9 EC.

Refer to the Software requirements section for operating system support for CHPID types FC and FCP.

FICON Express2 and FICON Express: Your 4-port FICON Express2 features (1 or 2 Gbps link data rate) can be carried forward to z10 EC. If you have 2-port FICON Express features (1 Gbps link data rate) you can also carry them forward to z10 EC. FICON Express LX (#2319) can be defined as CHPID type FCV (FICON bridge) to allow communication with ESCON control units using the ESCON Director Model 5 with the bridge feature. Migration to native FICON is encouraged. The ESCON Director Model 5 was withdrawn from marketing December 31, 2004.

Introducing Fiber Quick Connect for FICON LX environments

Fiber Quick Connect (FQC), an optional feature on z10 EC, is now being offered for all FICON LX (single-mode fiber) channels, in addition to the current support for ESCON (62.5 micron multimode fiber) channels. FQC is designed to significantly reduce the amount of time required

for on-site installation and setup of fiber optic cabling. FQC facilitates adds, moves, and changes of ESCON and FICON LX fiber optic cables in the data center, and may reduce fiber connection time by up to 80%.

FQC is for factory installation of Fiber Transport System (FTS) fiber harnesses for connection to channels in the I/O cage. FTS fiber harnesses enable connection to FTS direct-attach fiber trunk cables from IBM Global Technology Services.

FQC, coupled with FTS, is a solution designed to help minimize disruptions and to isolate fiber cabling activities away from the active system as much as possible.

IBM provides the direct-attach trunk cables, patch panels, and Central Patching Location (CPL) hardware, as well as the planning and installation required to complete the total structured connectivity solution. An ESCON example — four trunks, each with 72 fiber pairs, can displace up to 240 fiber optic jumper cables, the maximum quantity of ESCON channels in one I/O cage. This significantly reduces fiber optic jumper cable bulk.

At CPL panels you can select the connector to best meet your data center requirements. Small form factor connectors are available to help reduce the floor space required for patch panels.

CPL planning and layout is done prior to arrival of the server on-site using the default CHannel Path IDentifier (CHPID) placement report, and documentation is provided showing the CHPID layout and how the direct-attach harnesses are plugged.

Note: FQC supports all of the ESCON channels and all of the FICON LX channels in all of the I/O cages of the server.

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 for small, medium, and large enterprises.

- IBM Facilities Cabling Services — fiber transport system
- IBM IT Facilities Assessment, Design, and Construction Services — optimized airflow assessment for cabling

Planning and installation services for individual fiber optic cable connections are available. An assessment and planning for IBM Fiber Transport System (FTS) trunking components can also be performed.

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-Express) 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, 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 Optimized Airflow Assessment for Cabling option 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.

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

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

HiperSockets — "Network in a box"

HiperSockets Layer 2 support — For flexible and efficient data transfer for IP and non-IP workloads: Now, the HiperSockets internal networks on System z10 EC can support two transport modes: Layer 2 (Link Layer) as well as the current Layer 3 (Network or IP Layer). Traffic can be Internet Protocol (IP) Version 4 or Version 6 (IPv4, IPv6) or non-IP (AppleTalk, DECnet, IPX, NetBIOS, or SNA). HiperSockets devices are now protocol-independent and Layer 3 independent. Each HiperSockets device has its own Layer 2 Media Access Control (MAC) address, which is designed to allow the use of applications that depend on the existence of Layer 2 addresses such as Dynamic Host Configuration Protocol (DHCP) servers and firewalls.

Layer 2 support can help facilitate server consolidation. Complexity can be reduced, network configuration is simplified and intuitive, and LAN administrators can configure and maintain the mainframe environment the same as they do a non-mainframe environment.

With support of the new Layer 2 interface by HiperSockets, packet forwarding decisions are now based upon Layer 2 information, instead of Layer 3 information. The HiperSockets device can perform automatic MAC address generation to allow uniqueness within and across logical partitions (LPARs) and servers. MAC addresses can also be locally administered. The use of Group MAC addresses for multicast is supported as well as broadcasts to all other Layer 2 devices on the same HiperSockets network. Datagrams are delivered only between HiperSockets devices that are using the same transport mode (Layer 2 with Layer 2 and Layer 3 with Layer 3). A Layer 2 device cannot communicate directly with a Layer 3 device in another LPAR.

A HiperSockets device can filter inbound datagrams by Virtual Local Area Network identification (VLAN ID, IEEE 802.1q), the Ethernet destination MAC address, or both. Filtering can help reduce the amount of inbound traffic being processed by the operating system, helping to reduce CPU utilization.

Analogous to the respective Layer 3 functions, HiperSockets Layer 2 devices can be configured as primary or secondary connectors or multicast routers. This is designed to enable the creation of high-performance and high-availability Link Layer switches between the internal HiperSockets network and an external Ethernet or to connect the HiperSockets Layer 2 networks of different servers.

HiperSockets Layer 2 support is exclusive to System z10 EC, and supported by Linux on System z, and by z/VM for guest exploitation. Refer to the Software requirements section.

HiperSockets Multiple Write Facility for increased performance: HiperSockets performance has been enhanced to allow for the streaming of bulk data over a HiperSockets link between logical partitions (LPARs). The receiving LPAR can now process a much larger amount of data per I/O interrupt. This enhancement is transparent to the operating system in the receiving LPAR. HiperSockets Multiple Write Facility, with fewer I/O interrupts, is designed to reduce CPU utilization of the sending and receiving LPAR.

HiperSockets Multiple Write Facility is supported in the z/OS environment. Refer to the Software requirements section.

Local Area Network (LAN) connectivity

Introducing OSA-Express3 10 GbE LR — Designed to deliver increased throughput and reduced latency: OSA-Express3 10 Gigabit Ethernet (GbE) has been designed to increase throughput for standard frames (1492 byte) and jumbo frames (8992 byte) to help satisfy the bandwidth requirements of your applications. This increase in performance (compared to OSA-Express2 10 GbE) has been achieved through an enhancement to the architecture that supports direct host memory access by using a data router, eliminating store and forward delays.

Double the port density: The OSA-Express3 10 GbE has been designed with two PCI adapters, each with one port. Doubling the port density on a single feature helps to reduce the number of I/O slots required for high-speed connectivity to the Local Area Network (LAN). Each port continues to be defined as CHPID type OSD, supporting the Queued Direct Input/Output (QDIO) architecture for high-speed TCP/IP communication.

10 GbE cabling and connector: The OSA-Express3 10 GbE feature continues to be Long Reach (LR) supporting the 9 micron single mode fiber optic cabling environment. The connector is **new**; it is now the small form factor, LC Duplex connector. Previously the SC Duplex connector was supported. The LC Duplex connector is common with FICON, ISC-3, and OSA-Express2 Gigabit Ethernet LX and SX.

OSA-Express3 10 GbE LR (feature #3370) is exclusive to z10 EC and supports CHPID type OSD. It is supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z. Refer to the Software requirements section.

OSA-Express2

The OSA-Express2 family of LAN adapters includes:

- OSA-Express2 Gigabit Ethernet Long wavelength (GbE LX) and Short wavelength (GbE SX) for fiber optic connectivity to the LAN using:
 - QDIO (CHPID type OSD) for TCP/IP traffic when using Layer 3 and protocol-independent traffic when using Layer 2
 - OSA for NCP (CHPID type OSN) supporting LPAR-to-LPAR connectivity from operating systems that support Channel Data Link Control (CDLC) to IBM Communication Controller for Linux (CCL); supports Network Control Program (NCP) functions
- OSA-Express2 1000BASE-T Ethernet for Category 5 (copper) connectivity to the LAN using:
 - QDIO (CHPID type OSD)
 - Non-QDIO (CHPID type OSE) for SNA/APPN@/HPR and/or TCP/IP traffic
 - OSA-Integrated Console Controller (OSA-ICC) (CHPID type OSC) for emulation support for console session connections; TN3270E and local non-SNA DFT 3270 emulation
 - OSA for NCP (CHPID type OSN)

OSA-Express2 GbE (LX — #3364, SX — #3365) supports a link data rate of 1 Gigabit per second (Gbps) in each direction over 9 micron single-mode fiber (LX) or 50 or 62.5 micron multimode fiber (SX) with an LC Duplex connector.

OSA-Express2 1000BASE-T Ethernet (#3366) supports a link data rate of 10, 100, or 1000 Megabits per second (Mbps) auto-negotiated (target device must be set to auto-negotiate) and Category 5 Unshielded Twisted Pair (UTP) cabling with an RJ-45 connector.

Refer to the Software requirements section for operating system support for CHPID types, OSC, OSD, OSE, and OSN. Refer also to the Standards section.

Functions 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.
- 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).
- Checksum offload for IPv4 inbound and outbound packets — for calculating and validating the TCP/UDP and IP header checksums, reducing CPU cycles consumption (OSA performs the checksum calculations). Checksums are used to verify the contents of files when transmitted over a network.
- Multiple Image Facility (MIF) and spanned channels — for sharing OSA among logical channel subsystems.

InfiniBand coupling links for Parallel Sysplex®

The IBM System z10 EC introduces InfiniBand coupling link technology, designed to provide increased bandwidth at greater distances. At introduction, InfiniBand coupling links complement and do not replace Integrated Cluster Bus-4 (ICB-4) and InterSystem Channel-3 (ISC-3) available on z10 EC, z9 EC, z9 BC, z990, and z890 servers.

The IBM System z10 EC will support a 12x (12 lanes of fiber in each direction) InfiniBand-Double Data Rate (IB-DDR) coupling link which is designed to support a total interface link data rate of 6 gigabytes per second (GBps) in each direction. The maximum distance for this point-to-point link over fiber optic cabling is 150 meters (492 feet). This new InfiniBand coupling link provides improved performance over the current ISC-3 coupling link in data centers where servers are less than 150 meters apart.

A 12x InfiniBand-Single Data Rate (IB-SDR) coupling link is available on System z9 EC and System z9 BC servers configured as Internal Coupling Facilities (ICFs) only. This coupling link is designed to support a total interface link data rate of 3 gigabytes per second (GBps) in each direction. This new InfiniBand coupling link provides improved performance over the current ISC-3 coupling link in data centers where systems are less than 150 meters apart. When a System z10 EC server is connected to a System z9 server using point-to-point InfiniBand cabling, the link auto-negotiates to the highest common data rate of 3 GBps.

Other advantages of Parallel Sysplex using InfiniBand (PSIFB):

- InfiniBand coupling links also provide a new ability to define up to 16 CHPIDs on an HCA2-O fanout, allowing physical coupling links to be shared by multiple sysplexes. This also provides additional subchannels for Coupling Facility communication, improving scalability, and reducing contention in heavily utilized system configurations. It also allows for one CHPID to be directed to one CF, and another CHPID directed to another CF on the same target server, using the same port.
- Like other coupling links, external InfiniBand coupling links are also valid to pass time synchronization signals for Server Time Protocol (STP). Therefore the same coupling links can be used to exchange timekeeping information and Coupling Facility messages in a Parallel Sysplex.
- The IBM System z10 EC also takes advantage of InfiniBand as a higher-bandwidth replacement for the Self-Timed Interconnect (STI) I/O interface features found in prior System z servers.

The IBM System z10 EC will support up to 32 PSIFB links as compared to 16 PSIFB links on System z9 servers. For either System z10 EC or System z9 EC, there must be less than or equal to a total of 32 PSIFBs and ICB-4 links.

InfiniBand coupling links are CHPID type CIB.

Infiniband uses OM3 cables (50 micron multimode fiber rated at 2000 MHz-km).

The IBM System z9 Enterprise Class (z9 EC) and System z9 Business Class (z9 BC) servers are the last servers to support participation in the same Parallel Sysplex with IBM eServer zSeries 900 (z90), IBM eServer zSeries 800 (z800), and older System/390 Parallel Enterprise Server systems. This direction was announced in Hardware Announcement [ZG07-0286](#), dated April 18, 2007.

NTP Client support

The STP design has been enhanced to include support for a Simple Network Time Protocol (SNTP) client on the Support Element on z10 EC. With the STP feature enabled, you can initialize the time of an STP-only Coordinated Timing Network to the time provided by a Network Time Protocol (NTP) server, and maintain time accuracy. This allows an enterprise comprised of heterogeneous platforms to track to the same time source.

NTP Client support is also available on System z9 EC or System z9 BC servers that are at Driver 67L with the latest MCLs and have the STP Feature 1021 installed.

NTP Client support satisfies the direction stated in Hardware Announcement [ZG06-0814](#), dated October 10, 2006.

Additional information is available on the STP Web page

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

A Redpaper Server Time Protocol NTP Client support (REDP-4329) is available on the Redbooks™ Web site

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

Implementation Services for Parallel Sysplex DB2 data sharing

To assist with the assessment, planning, implementation, testing, and backup and recovery of a System z DB2 data sharing environment, IBM Global Technology Services announced and made available the IBM Implementation Services for Parallel Sysplex Middleware — DB2 data sharing on February 26, 2008.

This DB2 data sharing service is designed for clients who want to:

1. Enhance the availability of data
2. Enable applications to take full utilization of all servers' resources
3. Share application system resources to meet business goals
4. Manage multiple systems as a single system from a single point of control
5. Respond to unpredicted growth by quickly adding computing power to match business requirements without disruption
6. Build on the current investments in hardware, software, applications, and skills while potentially reducing computing costs

The offering consists of six selectable modules; each is a stand-alone module that can be individually acquired. The first module is an infrastructure assessment module, followed by five modules which address the following DB2 data sharing disciplines:

1. DB2 data sharing **planning**
2. DB2 data sharing **implementation**
3. **Adding** additional data sharing members
4. DB2 data sharing **testing**
5. DB2 data sharing **backup and recovery**

For more information on these services contact your IBM representative or refer to

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

Capacity on Demand

Capacity on Demand — Temporary Capacity: Just-in-time deployment of System z10 EC Capacity on Demand (CoD) is a radical departure from previous System z and zSeries servers. This new architecture allows:

- Up to four temporary records to be installed on the CEC and active at any given time
- Up to 200 temporary records to be staged on the SE
- Variability in the amount of resources that can be activated per record
- The ability to control and update records independent of each other
- Improved query functions to monitor the state of each record
- The ability to add capabilities to individual records concurrently, eliminating the need for constant ordering of new temporary records for different user scenarios
- Permanent LIC-CC upgrades to be performed while temporary resources are active

These capabilities allow you to access and manage processing capacity on a temporary basis, providing increased flexibility for on demand environments. The CoD offerings are built from a common Licensed Internal Code — Configuration Code (LIC-CC) record structure. These Temporary Entitlement Records (TERs) contain the information necessary to control which type of resource can be accessed and to what extent, how many times and for how long, and under what condition — test or real workload. Use of this information gives the different offerings their personality. Three temporary-capacity offerings will be available on February 26, 2008:

Capacity Back Up (CBU) — Temporary access to dormant processing units (PUs), intended to replace capacity lost within the enterprise due to a disaster. CP capacity or any and all specialty engine types (zIIP, zAAP, SAP, IFL, ICF) can be added up to what the physical hardware model can contain for up to 10 days for a test activation or 90 days for a true disaster recovery. Each CBU record comes with a default of five test activations. Additional test activations may be ordered in groups of 5 but a record can not contain more than 15 test activations. Each CBU

record provides the entitlement to these resources for a fixed period of time, after which the record is rendered useless. This time period can span from 1 to 5 years and is specified through ordering quantities of CBU years.

Capacity for Planned Events (CPE) — Temporary access to dormant PUs, intended to replace capacity lost within the enterprise due to a planned event such as a facility upgrade or system relocation. This is a new offering and is available only on the System z10 EC. CPE is similar to CBU in that it is intended to replace lost capacity; however, it differs in its scope and intent. Where CBU addresses disaster recovery scenarios that can take up to three months to remedy, CPE is intended for short-duration events lasting up to 3 days, maximum. Each CPE record, once activated, gives you access to all dormant PUs on the machine that can be configured in any combination of CP capacity or specialty engine types (zIIP, zAAP, SAP, IFL, ICF).

On/Off Capacity on Demand (On/Off CoD) — Temporary access to dormant PUs, intended to augment the existing capacity of a given system. On/Off CoD helps you contain workload spikes that may exceed permanent capacity such that Service Level Agreements cannot be met and business conditions do not justify a permanent upgrade. An On/Off CoD record allows you to temporarily add CP capacity or any and all specialty engine types (zIIP, zAAP, SAP, IFL, ICF) up to the following limits:

- The quantity of temporary CP capacity ordered is limited by the quantity of purchased CP capacity (permanently active plus unassigned).
- The quantity of temporary IFLs ordered is limited by quantity of purchased IFLs (permanently active plus unassigned).
- Temporary use of unassigned CP capacity or unassigned IFLs will not incur a hardware charge.
- The quantity of permanent zIIPs plus temporary zIIPs can not exceed the quantity of purchased (permanent plus unassigned) CPs plus temporary CPs and the quantity of temporary zIIPs can not exceed the quantity of permanent zIIPs.
- The quantity of permanent zAAPs plus temporary zAAPs can not exceed the quantity of purchased (permanent plus unassigned) CPs plus temporary CPs and the quantity of temporary zAAPs can not exceed the quantity of permanent zAAPs.
- The quantity of temporary ICFs ordered is limited by the quantity of permanent ICFs as long as the sum of permanent and temporary ICFs is less than or equal to 16.

Although the System z10 EC will allow up to four temporary records of any type to be installed, only one temporary On/Off CoD record may be active at any given time. An On/Off CoD record may be active while other temporary records are active.

Capacity provisioning — An installed On/Off CoD record is a necessary prerequisite for automated control of temporary capacity through z/OS MVST[™] Capacity Provisioning. z/OS MVS Capacity provisioning allows you to set up rules defining the circumstances under which additional capacity should be provisioned in order to fulfill a specific business need. The rules are based on criteria, such as: a specific application, the maximum additional capacity that should be activated, time and workload conditions. This support provides a fast response to capacity changes and ensures sufficient processing power will be available with the least possible delay even if workloads fluctuate. Refer to z/OS MVS Capacity Provisioning User's Guide (SA33-8299) for more information.

On/Off CoD Test — On/Off CoD allows for a no-charge test. No IBM charges are assessed for the test, including IBM charges associated with temporary hardware capacity, IBM software, or IBM maintenance. This test can be used to validate the processes to download, stage, install, activate, and deactivate On/Off CoD capacity nondisruptively. Each On/Off CoD-enabled server is entitled to only one no-charge test. This test may last up to a maximum duration of 24 hours commencing upon the activation of any capacity resources contained in the On/Off CoD record. Activation levels of capacity may change during the 24 hour test period. The On/Off CoD test automatically terminates at the end of the 24 hours period. In addition to validating the On/Off CoD function within your environment, you may choose to use this test as a training session for your personnel who are authorized to activate On/Off CoD.

Capacity on Demand — Permanent Capacity

Customer Initiated Upgrade (CIU) facility: When your business needs additional capacity quickly, Customer Initiated Upgrade (CIU) is designed to deliver it. CIU is designed to allow you to respond to sudden increased capacity requirements by requesting a System z10 EC PU and/or memory upgrade via the Web, using IBM Resource Link, and downloading and applying it to your System z10 EC server using your system's Remote Support connection. Further, with the

Express option on CIU, an upgrade may be made available for installation as fast as within a few hours after order submission.

Permanent upgrades: Orders (MESs) of all PU types and memory for System z10 EC servers that can be delivered by Licensed Internal Code, Control Code (LIC-CC) are eligible for CIU delivery. CIU upgrades may be performed up to the maximum available processor and memory resources on the installed server, as configured. While capacity upgrades to the server itself are concurrent, your software may not be able to take advantage of the increased capacity without performing an Initial Programming Load (IPL).

Increased flexibility with z/VM-mode partitions

System z10 EC provides for the definition of a z/VM-mode partition (LPAR) containing a mix of processor types including CPs and specialty processors IFLs, zIIPs, zAAPs, and ICFs. With the planned z/VM 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 provide an economical Java execution environment under z/OS on zAAPs, all in the same VM LPAR.

HMC system support

The new functions available on the Hardware Management Console (HMC) version 2.10.0 as described apply exclusively to z10 EC. However, the HMC version 2.10.0 will continue to support the systems as shown.

The 2.10.0 HMC will continue to support up to two 10/100 Mbps Ethernet LANs. Token Ring LAN is not supported. The 2.10.0 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.

Family	Machine type	Firmware driver	SE version
z10 EC	2097	73	2.10.0
z9 BC	2096	67	2.9.2
z9 EC	2094	67	2.9.2
z890	2086	55	1.8.2
z990	2084	55	1.8.2
z800	2066	3G	1.7.3
z900	2064	3G	1.7.3
9672 G6	9672/9674	26	1.6.2
9672 G5	9672/9674	26	1.6.2

TCP/IP Version 6 (IPv6)

HMC version 2.10.0 and Support Element (SE) version 2.10.0 can now communicate using TCP/IP Version 4 (IPv4), TCP/IP Version 6 (IPv6), or both. It is no longer necessary to assign a static IP address to an SE if it only needs to communicate with HMCs on the same subnet. An HMC and SE can use IPv6 link-local addresses to communicate with each other.

HMC/SE support is addressing the following requirements:

- The availability of addresses in the IPv4 address space is becoming increasingly scarce.
- The demand for IPv6 support is high in Asia/Pacific countries since many companies are deploying IPv6.
- The U.S. Department of Defense and other U.S. Government agencies are requiring IPv6 support for any products purchased after June 2008.

More information on the U.S. government requirements can be found at the following

<http://www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf>

http://www.whitehouse.gov/omb/egov/documents/IPv6_FAQs.pdf

Capacity On Demand

The HMC version 2.10.0 and SE version 2.10.0 will support new more flexible controls for temporary processor upgrades. You can temporarily add processors using the On/Off CoD (Capacity on Demand) (FC 9896), the CBU (Capacity BackUp) feature (FC 9910), or the CDP (Capacity Planned Event) (FC 9912). Highlights of the new more flexible controls for those features include the following.

- **Partial activation** — You can choose partial activation of resources up to the maximum you ordered.
- **On/Off CoD record reuse** — Each On/Off CoD record is initially active for 180 days. This record can be activated, deactivated, and reactivated many times while the record remains active. If increased capacity is needed for a longer period of time, or if you want to increase processor maximum capacities in the record, you can 'replenish' the record.
- **Permanent upgrade while temporary capacity is active** — You can add permanent processor or memory capacity while temporary On/Off CoD, CBU, or CPE records are active.
- **Multiple records can be simultaneously active** — Up to 4 records (On/Off CoD, CBU, and CPE) can be active at any given time. However, only one On/Off CoD record can be active at any given time.
- **Automatic deactivation** — When a record expires, the resource is automatically deactivated.
 - The record will not be deactivated if it means removing a dedicated processor or the last of a specific processor type.
 - Expiration warning messages will continue to be provided on the HMC prior to the date of expiration.

SNMP API (Simple Network Management Protocol Application Programming Interface) enhancements have also been made for the new Capacity On Demand features. More information can be found in the System z10 Enterprise Class Capacity On Demand User's Guide (SC28-6871).

Enhanced installation support for z/VM using the HMC

HMC version 2.10.0 along with SE version 2.10.0 on System z10 EC will now provide the ability to install Linux on System z in a z/VM virtual machine using the Hardware Management Console (HMC) DVD drive. This new function does not require a network connection between z/VM and the HMC, but instead, uses the existing communication path between the Support Element (SE) and the HMC.

Using the legacy support and the planned z/VM future 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 any network setup or a connection between an LPAR and the HMC.

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

Dynamic Enhancements

There are two new features which provide more dynamic capabilities without having to do preplanning.

- **Dynamic Add/Remove Cryptos**
 - The new "Change LPAR Cryptographic Controls" task allows dynamic changes without requiring a partition reactivation. Primary scenarios are as follows:
 - Dynamically add a Crypto to a partition for the first time
 - Dynamically add a Crypto to a partition already using Crypto
 - Dynamically remove Crypto from a partition
 - Notes:
 - Moving a Crypto from one partition to another is done via Remove/Add.
 - The above tasks don't require the Crypto hardware to be installed in the system.
 - This task can Change Running System and/or Save to Profiles.
 - A new Usage Domain Zeroize task will be provided to clear the appropriate partition Crypto keys for a given Usage Domain when removing a Crypto from a partition.
- **Dynamic Add Logical CPs without Preplanning**
 - Previously, the Image Profile defines the initial and reserved values for the different processor types for that partition. If those values weren't defined prior to partition activation/IPL, they could only be updated by reactivating that partition (including reIPL).

- The HMC/SE will now provide a task called Logical Processor Add which can:
 - Increase the "reserved" value for a given processor type (that is, GP, zAAP, zIIP, IFL, and so on)
 - Add a new processor type which is not in use yet for that partition
 - Increase the "initial" value for a given processor type
 - "Change Running System" and/or "Save to Profiles"
- Currently, exploitation of this support is limited to z/VM 5.3 with PTFs.

Enhanced Driver Maintenance (EDM)

There are several enhancements that have been made to the HMC/SE based on the feedback from the System z9 Enhanced Driver Maintenance field experience. Reliability, Availability, and Serviceability (RAS) enhancements were made. One example is a change to better handle intermittent customer network issues. EDM performance improvements will also be provided. Finally, new EDM user interface features were added to allow for customers and service personnel to better plan for the EDM. An example is a new option to check all licensed internal code change update EDM requirements. This option can be executed in advance of the EDM preload or activate.

Change Management

There were several enhancements made on the HMC/SE which provide more information for customers and service personnel as well as provide more flexibility.

The Query Channel/Crypto Configure Off/On Pending task will provide specific details on currently active Licensed Internal Code (LIC) change level and the levels which will be active after the Configure Off/On. In addition, the user will have the ability to determine which, if any, channels or Crypto Express2 features will require a configure off/on for a future LIC update process.

Customers and service personnel will be given the ability to redefine OSA-Express or Crypto Express2 LIC updates to be Configured Off/On if they desire the update to be done to one port or Crypto at a time rather than all at once for the same port/crypto type.

The System Information task has been updated to explicitly show any conditions where an LIC change update may not be truly active until an additional exception action is taken. This is generally an exception case that these conditions exist, but the information is now readily available on this one task.

Power/Thermal Monitoring

On System z9, IBM introduced power/thermal monitoring support with the HMC System Activity Display (SAD) task providing power consumption and air input temperature. On System z10, the HMC will now provide support for the Active Energy Manager (AEM) which will display power consumption/air input temperature as well as exhaust temperature. AEM will also provide some limited status/configuration information which might assist in explaining changes to the power consumption. AEM is exclusive to System z10.

Panel Wizards

Panel wizards were added to the HMC and SE in order to improve the user interface. The purpose of the wizards is to guide users through the panel options, provide recommended defaults where possible, and provide easier understanding of input and change of options. The following wizards were added. (Note that the existing tasks which the wizard provides the enhancement are still available).

- Manage User Wizard — provides a wizard for the following tasks:
 - User Profiles
 - Customize User Controls
 - Password Profiles
- Image Profile Wizard
 - Initial stage of a wizard for Customizing Image Activation Profiles Further enhancements are being investigated for the future.

z/VM Image Mode

On System z9, the supported Activation Image Profile Modes included the following. (Note that all of these modes have varying rules on what combination of processors and shared versus dedicated processors are allowed.)

- ESA/390 — Supports CPs, zAAPs, and zIIPs
- ESA/390 TPF — Support CPs
- Coupling Facility — Supports CPs and ICFs
- Linux only — Support CPs and IFLs

The HMC version 2.10.0 and SE version 2.10.0 will support an additional Activation Image Profile mode called **z/VM**. This image mode will support CPs, zAAPs, zIIPs, ICFs, and IFLs. It will allow all the varying rules and processor combinations in the above modes. The only requirement is that z/VM is the base operating system in that Image. This allows for easier Image Profile planning for whatever guest operating systems may run in that z/VM image. This also allows running different operating systems within that z/VM image for different purposes/processor requirements.

The key advantage of this support is for environments where customers need to use z/VM to host Linux and z/OS or z/VSE guests in the same "box," they will not have to artificially separate the management of those two environments if they do not want to. They can manage one z/VM image to host the entire collection of guests they want to deploy.

SNMP API Enhancements

In addition to the Capacity On Demand Simple Network Management Protocol Application Programming Interface (SNMP API) new features, the following SNMP API enhancements are also available:

- **Query Active Licensed Internal Code Change Levels API**
 - Returns Active Licensed Internal Code Change Levels.
 - Also returns if any exception conditions exists for Channel/Crypto Configure Off/On, Coupling Facility Control Code (CFCC) Reactivation, or Activation on next Power On Reset/System Activate.
- **Disabled Wait API Event**
 - Previously, SNMP Hardware Message Events had to be parsed for text of Hard Event, and there was no automation interface to obtain the Program Status Word (PSW).
 - This new SNMP Disabled Wait Event contains the PSW, Image Name, Partition ID, CPC Serial Number, and CPC Name, and will eliminate any need to parse text of Hardware Message Events.
- **Query PSW API**
 - New API support for obtaining PSW.
 - Only valid if Image is in not operating state.

CIM Automation APIs

The HMC will support Common Information Model (CIM) as an additional systems management API. The focus is on attribute query and operational management functions for System z: CPCs, Images, Activation Profiles. The goal is to provide similar functionality as the SNMP API. Some features (for example, indications (SNMP Trap equivalent), Capacity On Demand, processors) are not implemented in the CIM support yet.

CIM is defined by the Distributed Management Task Force

<http://www.dmtf.org>

The HMC object model extends the DMTF schema version 2.15. The Object Manager is OpenPegasus (V2.5.2).

<http://www.openpegasus.org>

Many toolkits exist to support client scripting. OpenPegasus comes with a C/C++ client toolkit. Standards Based LINUX Instrumentation for Manageability (SBLIM) Java Client includes other useful tools, including a Web-based class browser.

The IBM publication Common Information Model (CIM) Management Interface (SB10-7154) provides more information on System z10 CIM support.

Universal Lift Tool / Ladders

The Universal Lift Tool / Ladders feature (#3759) is designed to provide customers with enhanced system availability benefits by improving the service and upgrade times for larger, heavier devices. This feature includes a custom lift / lower mechanism that is specifically designed for use with System z10 frames, allowing these procedures to be accomplished quicker and with fewer people. It is recommended that one of these features be obtained for each customer account / datacenter.

Weight distribution plate

The weight distribution plate is designed to distribute the weight of a frame onto two floor panels in a raised-floor installation.

Certain configuration racks can weigh up to 2450 pounds per frame. The concentrated load on a caster or a leveling foot can be half of the total frame weight. For a multiple-system installation, it is possible that one floor panel could have two casters from two adjacent systems on it, and potentially induce a highly concentrated load onto a single floor panel. The weight distribution plate is designed to distribute the load over two floor panels and to eliminate a highly concentrated load on a single floor panel.

You are responsible for consulting with the floor tile manufacturer to determine the load rating of the floor tile and the pedestal structure supporting the floor tiles. Depending on the type of raised floors and the floor panels, additional panel supports (pedestals) may be required to restore or improve the structural integrity of the panel.

Note: Cable cutouts on a floor panel will significantly reduce the floor tile load rating (up to 50%).

Accessibility by people with disabilities

A U.S. 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

The IBM System z10 EC represents an evolution and a revolution of the IBM mainframe. With a modular design for affordable scalability and availability, the z10 EC 701 offers performance improvements of up to 1.62 times that of the z9 EC 701, and up to 1.50 times other z9 EC configurations for equivalent numbers of processors, up to 1.5 TB total memory, up to 12 PUs that can be defined for subcapacity use, and a new host bus interface using InfiniBand with a link data rate of 6 GBps. With a design for affordable scalability, z10 EC will continue to offer investment protection and improved price/performance with upgrades. Built on a foundation that improves recovery for unplanned outages and reduction of planned outages, the z10 EC goes further to offer a reduction in preplanning requirements by delivering and reserving a fixed Hardware System Area (HSA), and just-in-time deployment of resources that allows greater flexibility in defining and executing temporary capacity needs. The performance of z10 EC is designed to improve application performance, support more transactions, increase scalability, and assist in consolidation of workloads.

Statement of general direction

1x IB-DDR to be introduced to support extended distance for Coupling using InfiniBand: IBM intends to offer 1x InfiniBand at double data rate (1x IB-DDR), which is designed to complement and/or replace InterSystem Channel-3 (ISC-3) in a Parallel Sysplex environment. When 1x IB-DDR is introduced it will be designed to support one "lane" (one pair of fiber — a transmit and a receive) with a 5 gigabits per second (Gbps) link data rate at an unrepeated distance of 10 km using 9 micron single-mode fiber optic cabling with LC Duplex connectors.

The common fiber optic cabling environment is designed to facilitate a migration from ISC-3 (2 Gbps link data rate) to 1x IB-DDR (5 Gbps link data rate) in a Parallel Sysplex environment on

System z10 EC.

ICB-4 links to be phased out: IBM intends to not offer Integrated Cluster Bus-4 (ICB-4) links on future servers. IBM intends for System z10 to be the last server to support ICB-4 links.

Support of CHPID type FCV to be removed: IBM does not intend to offer configuration support for Channel Path Identifier (CHPID) type FCV on future servers. IBM intends for System z10 to be the last server to support FICON Express LX and CHPID type FCV. This feature is not orderable on the System z10 EC models. If the feature is installed at the time of an upgrade to the System z10 EC, it may be retained.

CHPID type FCV (Fibre Channel converted) is currently used with the FICON Express LX feature (#2319) to communicate with Enterprise Systems Connection (ESCON) control units using the FICON bridge card in the ESCON Director Model 5.

FICON support of 1 Gbps link data rates to be phased out: Future Fibre Connection (FICON) features, after FICON Express4, are not intended to support auto-negotiation to 1 Gbps link data rates.

The FICON Express4 features, supporting 1, 2, or 4 Gbps link data rates, will be the last features to support 1 Gbps link data rates.

System z FICON features to no longer support buffer credits for performance at extended distances: The FICON Express4 features are intended to be the last features to support extended distance without performance degradation. IBM intends to not offer FICON features with buffer credits for performance at extended distances. Future FICON features are intended to support distances up to 10 km without performance degradation. Extended distance solutions may include FICON directors or switches (for buffer credit provision) or Dense Wave Division Multiplexers (for buffer credit simulation).

System z9 to support 12x IB-SDR on z9 EC and z9 BC general purpose servers: Support of 12x InfiniBand at single data rate (12x IB-SDR) — a link data rate of 3 gigabytes per second (3 GBps) — is planned to be made available on z9 EC and z9 BC servers whether the Processor Units (PUs) are defined as Central Processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), or System Assist Processors (SAPs).

This expands the support for 12x IB-SDR on z9 EC and z9 BC beyond stand-alone Coupling Facilities (ICFs only).

z/VM LPAR enhancements: IBM intends to further enhance z/VM in a future release to exploit the System z10 EC support for a new logical partition (LPAR) mode "z/VM," exclusively for running z/VM. This new LPAR mode allows z/VM to utilize a wider variety of specialty processors in a single LPAR. For instance, in a z/VM mode LPAR, z/VM can manage Linux on System z guests running on IFL processors while also managing z/VSE and z/OS on central processors (CPs), and to offload z/OS system software overhead, such as DB2 workloads, on IBM System z9 and IBM System z10 Integrated Information Processors (zIIPs) and IBM System z9 and IBM System z10 Application Assist Processors (zAAPs).

Additional support for managing z/VM systems: IBM intends to further enhance z/VM in a future release to exploit the new Hardware Management Console (HMC) interface that allows the installation of Linux on System z into a z/VM virtual machine. Additionally, future support is planned for z/VM and the HMC to provide z/VM hypervisor-configuration tasks.

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.

Trademarks

System z, System z9, FICON, POWER, z/Architecture, PR/SM, Processor Resource/Systems Manager, System Storage, DS8000, Resource Link, Redbooks, and MVS are trademarks of International Business Machines Corporation in the United States or other countries or both.

z/OS, ESCON, z/VM, zSeries, DB2, z/VSE, APPN, Parallel Sysplex, and eServer are registered trademarks of International Business Machines Corporation in the United States or other countries or both.

Java is a trademark of Sun Microsystems, Inc.

Linux is a 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.

This announcement is provided for your information only. For additional information, contact your IBM representative, or visit the IBM worldwide contacts page at: <http://www.ibm.com/planetwide/>