IBM System z10 Business Class - The smart choice for your business. z can do IT better

Table of contents

At a glance

The IBM® System z10 BC is a world-class enterprise server built on the inherent strengths of the IBM System z® platform. It is designed to deliver new technologies and virtualization that provide improvements in price/performance for key new workloads. The System z10 BC further extends System z leadership in key capabilities with the delivery of granular growth options, business-class consolidation, improved security and availability to reduce risk, and just-in-time capacity deployment helping to respond to changing business requirements. Whether you want to deploy new applications quickly, grow your business without growing IT costs, or consolidate your infrastructure for reduced complexity, look no further - z Can Do IT. The System z10 BC delivers:

- The IBM z10 Enterprise Quad Core processor chip running at 3.5 GHz, designed to help improve CPU intensive workloads.
- A single model E10 offering increased granularity and scalability with 130 available capacity settings.
- Up to a 5-way general purpose processor and up to 5 additional Specialty Engine processors or up to a 10-way IFL or ICF server for increased levels of performance and scalability to help enable new business growth.
- 50% price reduction on Specialty Engines for System z10 BC, now $47,500*

- Integrated Encryption 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 algorithms with Secure Hash Algorithm (SHA-384 and SHA-512). Support for Longer Personal Account Numbers for stronger data protection on Crypto Express2. Trusted Key Entry Licensed Internal Code 5.3 enhancement to support Advanced Encryption Standard (AES) encryption algorithm, audit logging, and an infrastructure for payment card industry data security standard (PCIDSS).
- Integrated 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.
- Up to 120 GB of available real memory per server for growing application needs (with up to 248 GB of real memory planned in June 2009). Also a new 8GB fixed Hardware System Area (HSA) which is managed separately from customer memory. This fixed HSA is designed to improve availability by avoiding outages.
- 62% price reduction on Memory for System z10 BC or EC when purchased after October 21, 2008, and with a Specialty Engine[1] for new workloads, now $2,250 per GB[2].
- Plan ahead memory that allows for nondisruptive memory increases.
- 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). Additionally, increased flexibility with the ability for more temporary offerings installed on the CPC and ways to acquire capacity backup.
- 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.
- Production workload may now be executed on a CBU Upgrade during a CBU Test provided that certain contract terms are in effect with IBM.
- InfiniBand host bus bandwidth at 6 GBps designed to deliver improved performance.
- The new InfiniBand Coupling Links 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).
- Long reach 1x InfiniBand coupling links - an alternative to ISC-3 facilitating coupling link consolidation.
- Coupling Facility Control Code Level 16 - to help deliver faster service time for CF Duplexing, and improvements to the efficiency of workload distribution when using shared queues in the Coupling Facility.
- Time accuracy, availability, and system management improvements with new STP enhancements.
- Improved access to data with High Performance FICON® for System z (zHPF) on both FICON Express4 and FICON Express2. Additionally, enhanced problem determination, analysis, and manageability of the storage area network (SAN) by providing registration information to the fabric on the name server for both FICON and FCP.
- FCP - increased performance for small block sizes.
- SCSI Initial Program Load (IPL) - now a base function.
- Platform and name server registration in FICON channel.
- Extended-distance FICON - helps avoid degradation of performance at extended distances.
- Increased performance for Local Area Network connectivity with new OSA-Express3 I/O features providing double the port density, increased throughput, and reduced latency. OSA-Express3 10 GbE Long Reach (LR) and Short Reach (SR), OSA Express3 GbE 4-port LX and SX, OSA-Express3-2P GbE SX, OSA-Express3 1000BASE-T 4-port card and OSA-Express3-2P 1000BASE-T.
- HiperSockets™ improvements with Multiple Write Facility for increased performance and Layer 2 support to host IP and non-IP workloads.

- 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.
- Non-Raised floor support.
- Balanced Power Plan Ahead feature to minimize and balance currents on all three inputs of your three phase line cord.

[1] Prices are stated in U.S. currency for the U.S. and may vary by country. Specialty Engines do not include Internal Coupling Facilities (ICFs).

[2] Limited to 16 GB per Specialty Engine.

Overview

The IBM System z10 Business Class (z10 BC) delivers innovative technologies for small and medium enterprises that give you a whole new world of capabilities to run modern applications. Ideally suited as the cornerstone of your new enterprise data center, this competitively priced server delivers unparalleled qualities of service to help manage growth and reduce cost and risk in your business.

The z10 BC further extends the leadership of System z by delivering expanded granularity and optimized scalability for growth, enriched virtualization technology for consolidation of distributed workloads, improved availability and security to help increase business resiliency, and just-in-time management of resources. The z10 BC is at the core of the enhanced System z platform and is the new face of System z.

The z10 BC Model E10 is designed to provide up to 1.5 times the total system capacity for general purpose processing, nearly two times the available memory, and over 40% more configurable processors than the z9™ BC Model S07. In June 2009, memory options up to 248 GB will be available providing nearly four times the available memory of the z9 BC Model S07.

For those customers with distributed servers trying to reduce complexity of operations and operating costs, the z10 BC facilitates consolidation of dozens to hundreds of individual distributed servers into virtual images on one z10 BC server. z10 BC delivers improvements in capacity, memory, I/O infrastructure, and virtualization technology that you need in one small footprint.

In the area of server availability, enhancements have been engineered into the z10 BC to help eliminate unwanted down time. For example, preplanning requirements are minimized by delivering a fixed, reserved Hardware System Area (HSA) that enables dynamic creation of logical partitions, including logical channel subsystems, subchannel sets, and devices, using dynamic I/O without preplanning. Additionally, new capabilities are intended to allow you to dynamically change logical processor definitions and cryptographic co-processor definitions for a logical partition without requiring the logical partition to be deactivated and re-activated.

Further improvement to availability and flexibility is achieved with just-in-time deployment of capacity resources designed to dynamically change capacity when business requirements change. You are no longer limited by one offering configuration; instead one or more flexible configurations can be defined that can be used to solve multiple temporary situations. You can choose from multiple configurations and the configurations themselves are flexible so you can activate only what is needed from your defined configuration. Another significant change is the ability to add permanent capacity to the server when you are in a temporary state. These z10 BC enhancements are designed to allow you to take advantage of the technology helping to provide on-demand capacity more effectively. There are new terms governing System z Capacity Back Up (CBU) now available which allow customers to execute production workload on a CBU Upgrade during a CBU Test.

IBM continues the long history of providing integrated technologies to optimize a variety of workloads. Specialty engines have been available to help users expand

the use of the mainframe for new workloads while helping to lower the cost of ownership. The z10 BC processor unit now delivers an integrated 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.

Additionally, integrated clear-key encryption security features on z10 BC 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 in a security-rich environment.

High speed connectivity and high bandwidth out to the data and the network are critical in achieving high levels of transaction throughput and enabling resources inside and outside the server to maximize application requirements. The z10 BC has a new host bus interface with a link data rate of 6 GB using the industry-standard InfiniBand protocol to help satisfy requirements for coupling (ICF and server-to-server connectivity), cryptography (Crypto Express2 with secure coprocessors and SSL transactions), I/O (ESCON®, FICON, or FCP), and LAN (new OSA-Express3 Gigabit, 10 Gigabit and 1000BASE-T Ethernet features). New High Performance FICON for System z (zHPF) also brings new levels of performance when accessing data on zHPF enabled storage devices such as the IBM System Storage™ DS8000™.

IBM Global Financing can provide attractive low rate financing for all new and upgraded z10 BC 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

IBM Global Financing is 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

- October 28, 2008
  - System z10 BC Model E10
  - z10 BC features and functions
  - z890 upgrades to z10 BC
- November 26, 2008
  - z9 BC upgrades to z10 BC
- January 28, 2009
  - z10 BC model capacity conversions
  - z10 BC MES features
  - z10 BC E10 to z10 EC Model E12
  - OSA-Express QDIO data connection isolation for the z/VM® environment
  - Non-raised floor support
- June 30, 2009
  - Memory size maximum of 248 GB

Note: On/Off Capacity on Demand is available starting October 28, 2008.

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

Description

The IBM System z10 Business 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 BC, supporting z/Architecture®, and offering the highest levels of reliability, availability, scalability, clustering, and virtualization. The z10 BC just-in-time deployment of capacity allows improved flexibility and administration, and the ability to enable changes as they happen. The expanded scalability on the z10 BC facilitates growth and large-scale consolidation. The z10 BC is designed to provide:

- Uniprocessor performance up to 1.4 times the uniprocessor performance of the z9 BC S07 Z01 (based on LSPR mixed workload average)
- Up to 1.5 times the total system capacity for general purpose processing, of the z9 BC
- Up to 12 Processor Units (PUs) including SAPs, as compared to a maximum of 8 on the z9 BC (including SAPs)
- Up to 1.9 times as much total server available memory as a z9 BC - up to 120 gigabytes of total memory
- Up to 3.8 times as much total server available memory as a z9 BC by June 30, 2009 - up to 248 gigabytes of total memory
- Up to 78% more subcapacity choices as compared to z9 BC
- Increased host base bandwidth using InfiniBand at 6 GBps
- Hardware support for HiperDispatch
- Hardware Decimal Floating Point unit for improved numeric processing performance
- Large page support (1 megabyte pages)
- Up to 128 FICON channels
- High Performance FICON for System z (zHPF), providing improvement in performance and RAS on both FICON Express4 and FICON Express2 features
- Platform and name server registration in FICON channel
- Extended-distance FICON to help avoid degradation of performance at extended distances
- FCP - increased performance for small block sizes
- SCSI Initial Program Load (IPL) - now a base function
- Performance improvements with HiperSockets Multiple Write Facility
- 12x DDR Coupling with InfiniBand for improved distance compared to ICB-4 links and potential cost saving by ISC-3 link consolidation
- 1x DDR Coupling over InfiniBand links supporting 10 km unrepeated distance
- STP time accuracy, availability, and system management improvements
- Improved Advanced Encryption Standard (AES) 192 and 256 and stronger hash algorithms with Secure Hash Algorithm (SHA) 384 and 512
- 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 a new or existing logical channel subsystem (LCSS)
- Open Systems Adapter-Express3 (OSA-Express3) 10 Gigabit Ethernet with double the port density and improved performance
- 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
- Temporary offering Capacity for Planned Event (CPE) available to manage system migrations, data center moves, maintenance activities, and similar situations
- Improved performance management with Capacity Provisioning
- Plan ahead memory that allows for nondisruptive memory increases
- Support for the IBM Systems Director Active Energy Manager (AEM) for Linux on System z

Model summary matrix

Model	PUs	Memory	IFB	I/O drawers	CHPIDs
E10	1 to 10	4 to 120 GB	0 to 12	0 to 4	512

Notes:

- The total maximum number of PUs is 12 when you include SAPs.
- Memory reserved for the fixed HSA is in addition to the purchased entitlement.
- Each LCSS supports up to 256 CHPIDs.
- Memory size maximum to increase from 120 GB to 248 GB on June 30, 2009.
- Single phase line cords support up to a maximum of two I/O drawers using line cords (#8991), (#8990), (#8991), or (#8999). For more than two I/O drawers, it is necessary to use three phase line cords (#8983), (#8984), (#8986), (#8987), (#8988), or (#8998).

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 BC, z/Architecture processor subcapacity indicator is defined with a (A0x-Z0x) notation, where x is the number of installed CPs, from one to five. There are a total of 26 subcapacity levels, designated by the letters A through Z.

In addition to the general information provided for z/OS V1.9, 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 BC (2098) Z01 is expected to be up to 1.4 times that of the z9 BC (2096) S07 Z01.

Moving from a System z9® partition to an equivalently sized System z10 BC 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 BC 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 BC 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/lspr/

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

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 BC 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 BC 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/I 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 BC. 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 BC is designed to allow memory to be extended to support large (1 megabyte (MB)) pages. Use of large pages can improve CPU utilization 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 System z10 running either z/OS or Linux on System z. Refer to the Software requirements section.

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 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 to support secure and clear-key RSA operations
- The tamper-responding hardware and lower-level firmware layers are validated to U.S. Government FIPS 140-2 standard: *Security Requirements for Cryptographic Modules* at Level 4.

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 BC, 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 operations 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://www.ibm.com/common/ssi/index.wss

Crypto Express2-1P

An option of one PCI-X adapter per feature, in addition to the current two PCI-X adapters per feature, is being offered for the z10 BC to help satisfy small and midrange security requirements while maintaining high performance.

The Crypto Express2-1P feature, with one PCI-X adapter, can continue to be defined as either a Coprocessor or an Accelerator. A minimum of two features must be ordered.

Additional cryptographic functions and features with Crypto Express2 and Crypto Express2-1P

Key management: Added 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: Added 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: 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.

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.

Cryptographic enhancements to Crypto Express2 and Crypto Express2-1P

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, z/VSE, and Linux on System z. Refer to the Software requirements section.

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

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

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

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

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

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

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

TKE 5.3 workstation: The Trusted Key Entry (TKE) workstation (#0839) and the TKE 5.3 level of Licensed Internal Code (#0854) are optional features on the System z10 BC. The TKE 5.3 Licensed Internal Code (LIC) is loaded on the TKE workstation prior to shipment. The TKE workstation offers security-rich local and remote key management, providing authorized persons a method of operational and master key entry, identification, exchange, separation, and update. The TKE workstation supports connectivity to an Ethernet Local Area Network (LAN) operating at 10 or 100 Mbps. Up to ten TKE workstations can be ordered.

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

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

The TKE workstation (#0839) and TKE 5.3 LIC (#0854) are available on the z10 BC, z10 EC, z9 EC, and z9 BC.

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

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

Smart Card Reader - new feature: Support for an optional Smart Card Reader (#0855) attached to the TKE 5.3 workstation allows for the use of smart cards that contain an embedded microprocessor and associated memory for data storage. Access to and the use of confidential data on the smart cards is protected by a user-defined Personal Identification Number (PIN).

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

The Smart Card Reader, attached to a TKE workstation with the 5.3 level of LIC will support System z10 BC, z10 EC, z9 EC, and z9 BC. However, TKE workstations with 5.0, 5.1 and 5.2 LIC must be upgraded to TKE 5.3 LIC.

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

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

- The Crypto Express2 feature is supported on the z9 BC and can be carried forward on an upgrade to the System z10 BC.
- You may continue to use TKE workstations with 5.3 licensed internal code to control the System z10 BC.
- 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

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

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

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

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

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

The FICON Express4 and FICON Express2 features will support both the existing FICON protocol and the zHPF protocol concurrently in the server Licensed Internal Code. High performance FICON is supported by z/OS for DB2, VSAM, PDSE, and zFS applications. Refer to the Software requirements section. zHPF applies to all FICON Express4 and FICON Express2 features (CHPID type FC) and is exclusive to System z10. Exploitation is required by the control unit.

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

Platform and name server registration in FICON channel

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

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

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

It will register:

- Platforms:
  - Worldwide node name (node name for the platform - same for all channels)
  - Platform type (host computer)
  - Platform name (includes vendor ID, product ID, and vendor-specific data from the node descriptor)
- Channels:
  - Worldwide port name (WWPN)
  - Node port identification (N_PORT ID)
  - FC-4 types supported (always 0x1B and additionally 0x1C if any Channel-to-Channel (CTC) control units are defined on that channel)
  - Classes of service support by the channel

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

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

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 can 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 5.4.1.xx.xx (bundle version 64.1.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), an 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 that support FCP, 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

Choose the FICON Express4 features that best meet your business requirements

To meet the demands of your Storage Area Network (SAN), provide granularity, facilitate redundant paths, and satisfy your infrastructure requirements, there are five features from which to choose.

```
Feature                    Feat   Infrastructure     Ports per
                                                     feature

FICON Express4 10KM LX     3321   Single mode fiber  4
FICON Express4 4KM LX      3324   Single mode fiber  4
FICON Express4-2C 4KM LX   3323   Single mode fiber  2

FICON Express4 SX          3322   Multimode fiber    4
FICON Express4-2C SX       3318   Multimode fiber    2
```

Choose the features that best meet your granularity, fiber optic cabling, and unrepeated distance requirements. If you have a requirement for:

- Two FICON channels, select feature #3318 (SX) or #3323 (LX)
- A maximum of four FICON channels, you may choose to order two FICON Express4-2C features, with each of the features in a separate I/O domain for high availability
- A maximum of six FICON channels, you may choose to order one FICON Express4 four-channel feature and one FICON Express4-2C feature
- A mix of SX (multimode) and LX (single mode) fiber optic cabling, you may choose to order the FICON Express4-2C 4KM LX feature to satisfy your single mode fiber optic cabling requirements, and order the FICON Express4 SX four-channel feature for your multimode fiber optic cabling requirements
- Eight or more channels - order the FICON Express4-2C feature only if connectivity to FICON control units cannot be spread over two I/O domains for high availability using only FICON Express4 four-channel features

Note: A 4KM LX transceiver is designed to interoperate with a 10KM LX transceiver.

Refer to the Standards section for the characteristics of each of the features.

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 to System z10 BC.

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

FICON Express2 and FICON Express: Your current FICON Express2 features (1 or 2 Gbps link data rate) can be carried forward to z10 BC. If you have FICON Express features (1 Gbps link data rate) you can also carry them forward to z10 BC. 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.

Fiber Quick Connect for FICON LX environments

Fiber Quick Connect (FQC), an optional feature on z10 BC, is 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 drawer. 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 drawer. 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 IDdentifier (CHPID) placement report, and documentation is provided showing the CHPID layout and how the direct-attach harnesses are plugged.

FQC supports all of the ESCON channels and all of the FICON LX channels in the I/O drawer 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 BC, 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 BC, 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 BC 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 nonmainframe 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 BC, 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 - a new generation

OSA-Express3 - a new family of LAN adapters: The third generation of Open Systems Adapter-Express (OSA-Express3) features have been introduced to help reduce latency and overhead, deliver double the port density of OSA-Express2, and provide increased throughput.

The OSA-Express3 features support the following environments:

```
CHPID     OSA-Express3
type        features               Purpose/Traffic


OSC³      1000BASE-T   OSA-Integrated Console Controller (OSA-ICC)
                       TN3270E, non-SNA DFT to IPL CPCs and LPARs
                       Operating system console operations


OSD³      1000BASE-T   Queued Direct Input/Output (QDIO)
          GbE          TCP/IP traffic when Layer 3
          10 GbE       Protocol-independent when Layer 2


OSE³      1000BASE-T   Non-QDIO, SNA/APPN/HPR and/or TCP/IP passthru


OSN       1000BASE-T   OSA-Express for NCP
          GbE          Supports channel data link control (CDLC)
                       LPAR-to-LPAR communication exclusively; no
                       external communication
```

[3] Note that software PTFs or a new release may be required (depending on CHPID type) to support all ports.

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

```
Feature                  Feat   Infrastructure   Ports per
                                                  feature

OSA-Express3 GbE LX      3362   Single mode fiber   4
OSA-Express3 10 GbE LR   3370   Single mode fiber   2

OSA-Express3 GbE SX      3363   Multimode fiber     4
OSA-Express3 10 GbE SR   3371   Multimode fiber     2
OSA-Express3-2P GbE SX   3373   Multimode fiber     2

OSA-Express3 1000BASE-T  3367   Copper              4
OSA-Express3-2P 1000BASE-T 3369 Copper              2
```

Refer to the Standards section for the characteristics of each of the features.

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

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

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

Port density or granularity: The OSA-Express3 features have Peripheral Component Interconnect Express (PCI-E) adapters. The previous table identifies whether the feature has 2 or 4 ports for LAN connectivity. Select the density that best meets your business requirements. 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.

Note: The two port features (OSA-Express3-2P GbE SX, and OSA-Express3-2P 1000BASE-T) are exclusive to the z10 BC.

10 GbE cabling and connector: The OSA-Express3 10 GbE features support Long Reach (LR) using 9 micron single mode fiber optic cabling and Short Reach (SR) using 50 or 62.5 micron multimode fiber optic cabling. The connector is new; it is now the small form factor, LC Duplex connector. Previously the SC Duplex connector was supported for LR. The LC Duplex connector is common with FICON, ISC-3, and OSA-Express2 Gigabit Ethernet LX and SX.

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

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

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

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

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

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

OSA-Express3 Ethernet features - Summary of benefits: OSA-Express3 10 GbE LR (single mode fiber), 10 GbE SR (multimode fiber), GbE LX (single mode fiber), GbE SX (multimode fiber), and 1000BASE-T (copper) are designed for use in high-speed enterprise backbones, for local area network connectivity between campuses, to connect server farms to z10, and to consolidate file servers onto z10. With reduced latency, improved throughput, and up to 96 ports of LAN connectivity, (when all are 4-port features, 24 features per server), you can "do more with less."

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

- Reduced latency (up to 45% reduction) and increased throughput (up to 4x) for applications

- More physical connectivity to service the network and fewer required resources:
  - Fewer CHPIDs to define and manage
  - Reduction in the number of required I/O slots
  - Possible reduction in the number of I/O drawers
  - Double the port density of OSA-Express2
  - A solution to the requirement for more than 48 LAN ports (now up to 96 ports)

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

OSA-Express2 availability: OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet continue to be available for ordering, for a limited time, if you are not yet in a position to migrate to the latest release of the operating system for exploitation of two ports per PCI-E adapter and if you are not resource-constrained.

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

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

OSA-Express QDIO data connection isolation for the z/VM environment

Multi-tier security zones are fast becoming the network configuration standard for new workloads. Therefore, it is essential for workloads (servers and clients) hosted in a virtualized environment (shared resources) to be protected from intrusion or exposure of data and processes from other workloads.

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

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

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

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

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

Coupling connectivity for Parallel Sysplex

Introducing long reach InfiniBand coupling links

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

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

Depending on the capability of the attached DWDM, the link data rate will automatically be set to either SDR or DDR.

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

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

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

Five coupling link options: The z10 BC supports Internal Coupling channels (ICs), Integrated Cluster Bus-4 (ICB-4), InterSystem Channel-3 (ISC-3) (peer mode), and 12x and 1x InfiniBand (IFB) links for communication in a Parallel Sysplex environment.

1. Internal Coupling Channels (ICs) can be used for internal communication between Coupling Facilities (CFs) defined in LPARs and z/OS images on the same server.
2. Integrated Cluster Bus-4 (ICB-4) links are for short distances. ICB-4 links use 10 meter (33 feet) copper cables, of which 3 meters (10 feet) is used for internal routing and strain relief. ICB-4 is used to connect z10 BC-to-z10 BC, z10 EC, z9 EC, z9 BC, z990, and z890. Note: If connecting to a z9 BC or a z10 BC with ICB-4, those servers cannot be installed with the non-raised floor feature. Also, if the z10 BC is ordered with the non-raised floor feature, ICB-4 cannot be ordered.

3. InterSystem Channel-3 (ISC-3) supports communication over unrepeated distances of up to 10 km (6.2 miles) using 9 micron single mode fiber optic cables and even greater distances with System z qualified optical networking solutions. ISC-3s are supported exclusively in peer mode (CHPID type CFP).

4. 12x InfiniBand coupling links (12x IB-SDR or 12x IB-DDR) offer an alternative to ISC-3 in the data center and facilitate coupling link consolidation; physical links can be shared by multiple systems or CF images on a single system. The 12x IB links support distances up to 150 meters (492 feet) using industry-standard OM3 50 micron fiber optic cables.

5. Long Reach 1x InfiniBand coupling links (1x IB-SDR or 1x IB-DDR) are an alternative to ISC-3 and offer greater distances with support for point-to-point unrepeated connections of up to 10 km (6.2 miles) using 9 micron single mode fiber optic cables. Greater distances can be supported with System z qualified optical networking solutions. Long reach 1x InfiniBand coupling links support the same sharing capability as the 12x InfiniBand version allowing one physical link to be shared across multiple CF images on a system.

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

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

The whitepaper is available at

http://www.ibm.com/systems/z/advantages/pso/whitepaper.html

Coupling Facility Control Code Level 16

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

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

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

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

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

Access the tool at

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

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

Implementation Services for Parallel Sysplex

IBM Implementation Services for Parallel Sysplex CICS® and WAS Enablement

IBM Implementation Services for Parallel Sysplex Middleware - CICS enablement consists of five fixed-price and fixed-scope selectable modules:

1. CICS application review
2. z/OS CICS infrastructure review (module 1 is a prerequisite for this module)
3. CICS implementation (module 2 is a prerequisite for this module)
4. CICS application migration
5. CICS health check

IBM Implementation Services for Parallel Sysplex Middleware - WebSphere Application Server enablement consists of three fixed-price and fixed-scope selectable modules:

1. WebSphere Application Server network deployment planning and design
2. WebSphere Application Server network deployment implementation (module 1 is a prerequisite for this module)
3. WebSphere Application Server health check

For a detailed description of this service, refer to Services Announcement ZS08-0146, dated June 24, 2008.

IBM 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

Server Time Protocol (STP)

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

Server Time Protocol enhancements

The following Server Time Protocol (STP) enhancements are available on the z10 EC, z10 BC, z9 EC, and z10 BC. The prerequisites are that you install STP feature #1021 and that the latest MCLs are installed for the applicable driver.

NTP client support: This enhancement addresses the requirements of customers who need to provide the same accurate time across heterogeneous platforms in an enterprise.

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

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

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

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

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

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

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

This avoids a manual reconfiguration of the BTS to be the CTS, if the PTS/CTS is not able to access its ETS. In an ETR network when the primary Sysplex Timer is not able to access the ETS device, the secondary Sysplex Timer takes over the role of the primary - a recovery action not always accepted by some customers. The STP design provides continuous availability of ETS while maintaining the special roles of PTS and BTS assigned by the customer.

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

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

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

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

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

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

Note that this enhancement is also available on the z990 and z890 servers.

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

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

Additional information is available on the STP Web page

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

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

  http://www.redbooks.ibm.com/
- Server Time Protocol: Planning Guide, SG24-7280
- Server Time Protocol: Implementation Guide, SG24-7281

Capacity on Demand

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

- Up to eight temporary records to be installed on the CPC 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 Licensed Internal Code - Configuration Code (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 made available on October 28, 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. Presently 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.

CBU Tests: Customers may now execute production workload during a CBU test provided that a) an amount of System z production workload capacity equivalent to the CBU upgrade is shut down or otherwise made unusable by the customer for the duration of the test, and b) the appropriate contracts are in place. All new CBU contract documents contain these new CBU Test terms. Existing CBU customers will need to execute IBM Customer Agreement Amendment for IBM System z Capacity Backup Upgrade Tests, form number Z125-8145.

Capacity for Planned Event (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. 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.
- The quantity of temporary SAPs ordered is limited by the quantity of permanent SAPs as long as the sum of permanent and temporary SAPs is less than or equal to 32.

Although the System z10 BC will allow up to eight 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.

Management of temporary capacity through On/Off CoD is further enhanced through the introduction of resource tokens. For CP capacity, a resource token represents an amount of processing capacity that will result in 1 MSU of SW cost for 1 day - an MSU-day. For specialty engines, a resource token represents activation of 1 engine of that type for 1 day - an IFL-day, a zIIP-day, or a zAAP-day. The different resource tokens are contained in separate pools within the On/Off CoD record. The customer, via the Resource Link ordering process, determines how many tokens go into each pool. Once On/Off CoD resources are activated, tokens will be decremented from their pools every 24 hours. The amount decremented is based on the highest activation level for that engine type during the previous 24 hours.

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

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

Capacity provisioning: An installed On/Off CoD record is a necessary prerequisite for automated control of temporary capacity through z/OS Capacity Provisioning. z/OS 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, and 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. See *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-hour 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.

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 Capacity On Demand User's Guide* (SC28-6871).

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 BC PU and/or memory upgrade via the Web, using IBM Resource Link, and downloading and applying it to your System z10 BC 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 BC servers that can be delivered by Licensed Internal Code - Configuration 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).

Plan ahead memory

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

The preplanned memory feature is offered in 4 gigabyte (GB) increments. The quantity assigned by the configuration tool is the number of 4 GB blocks necessary to increase the physical memory from that required for the "starting" logical memory to the physical memory required for the "target" logical configuration. Activation of any preplanned memory requires the purchase of preplanned memory activation features. One preplanned memory activation feature (#1992) is required for each preplanned memory feature (#1991). You now have the flexibility to activate memory to any logical size offered between the starting and target size.

Increased flexibility with z/VM-mode partitions

System z10 BC allows you to define a z/VM-mode partition (LPAR) containing a mix of processor types including CPs and specialty engines - IFLs, zIIPs, zAAPs, and ICFs. With z/VM V5.4 support, this new capability increases flexibility and simplifies systems management by allowing z/VM to manage guests to operate Linux on System z on IFLs, to operate z/VSE and z/OS on CPs, to offload z/OS system software overhead, such as DB2 workloads, on zIIPs, and to offer an economical Java execution environment under z/OS on zAAPs, all in the same VM LPAR.

HMC system support

The new functions available on the Hardware Management Console (HMC) version 2.10.1 as described apply exclusively to z10 BC. However, the HMC version 2.10.1 will continue to support the systems as shown.

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

Family	Machine type	Firmware driver	SE version
z10 BC	2098	76	2.10.1
z10 EC	2097	76	2.10.1
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

Internet Protocol, Version 6 (IPv6)

HMC version 2.10.1 and Support Element (SE) version 2.10.1 can now communicate using IP Version 4 (IPv4), 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

http://www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf
http://www.whitehouse.gov/omb/egov/documents/IPv6_FAQs.pdf

HMC/SE Console Messenger: On systems prior to z9, the remote browser capability was limited to Platform Independent Remote Console (PIRC), with a very small subset of functionality. Full functionality via Desktop On-Call (DTOC) was limited to one user at a time; it was slow, and was rarely used.

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

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

HMC z/VM tower systems management enhancements: Building upon the previous z/VM Systems Management support from the Hardware Management Console (HMC), which offered management support for already defined virtual resources, new HMC capabilities are being made available allowing selected virtual

resources to be defined. In addition, further enhancements have been made for managing defined virtual resources.

Enhancements are designed to deliver out-of-the-box integrated graphical user interface-based (GUI-based) management of selected parts of z/VM. This is especially targeted to deliver ease-of-use for enterprises new to System z. You can more seamlessly perform hardware and selected operating system management using the HMC Web browser-based user interface.

Support for HMC z/VM tower systems management enhancements is exclusive to z/VM 5.4 and the IBM System z10.

Enhanced installation support for z/VM using the HMC: HMC version 2.10.1 along with Support Element (SE) version 2.10.1 on z10 BC and corresponding z/VM 5.4 support will now give you the ability to install Linux on System z in a z/VM virtual machine using the HMC DVD drive. This new function does not require an external network connection between z/VM and the HMC, but instead, uses the existing communication path between the HMC and SE.

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

Using the legacy support and the z/VM 5.4 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 external network setup or a connection between an LPAR and the HMC.

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

Support for the enhanced installation support for z/VM using the HMC is exclusive to z/VM 5.4 and the IBM System z10.

Dynamic enhancement: The following feature is available without requiring preplanning.

- **Dynamic Add Logical CPs without Preplanning**
  - Previously, the Image Profile defined the initial and reserved values for the different processor types for that partition. If those values were not defined prior to partition activation/IPL, they could only be updated by reactivating that partition (including reIPL).
  - The HMC/SE now offers a task called Logical Processor Add which can:
    - -- Increase the "reserved" value for a given processor type (for example, CP, zAAP, zIIP, IFL)
    - -- Add a new processor type which is not in use yet for that partition
    - -- Increase the "initial" value for a given processor type
    - -- Change Running System and/or Save to Profiles

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

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

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-Express3 and OSA-Express2 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 or Crypto type.

The System Information task has been updated to explicitly show any conditions where a 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 are still available with the enhancement.)

- 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 - Supports CPs
- Coupling Facility - Supports CPs and ICFs
- Linux only - Supports CPs and IFLs

System z10 supports 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 or processor requirements. The key advantage of this support

is this: for environments where users need to use z/VM 5.4 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. z/VM 5.4 availability satisfies the statement of direction made in Hardware Announcement ZG08-0289, dated February 26, 2008.

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 the contents of the PSW
  - Only valid if the image is in not operating state

CIM automation APIs: The HMC will support Common Information Model (CIM) as an additional systems management API with functionality similar to the SNMP API. The capabilities include attribute query, operational management functions for System z, CPCs, images, Activation Profiles, indications (SNMP Trap equivalent), Capacity on Demand, and processors. 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

The HMC also conforms to additional DTMF profiles related to Virtual System, System Virtualization, and Software inventory. 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.

Up to 30 logical partitions: The z10 BC supports 30 logical partitions (LPARs) and provides the ability to define up to two Logical Channel Subsystems (LCSS). Each LCSS is capable of supporting up to 256 CHPID definitions and 15 logical partitions. With Processor Resource/Systems Manager™ (PR/SM™) and Multiple Image Facility (MIF), you can share ESCON and FICON channels, coupling channels, HiperSocket CHPIDs, and OSA ports across LPARs. All except ESCON channels can span to LPARs defined in different Logical Channel Subsystems.

Support of up to 30 LPARs is supported by z/OS and z/OS.e, z/VSE, z/VM, z/TPF, TPF, and Linux on System z9. Refer to the Hardware requirements and Software requirements sections of this announcement.

HiperDispatch

A System z10 exclusive, HiperDispatch represents a cooperative effort between the z/OS operating system and Processor Resource/Systems Manager (PR/SM) on System z10 hardware.

- Work may be dispatched across fewer logical processors, thereby reducing the multi-processor (MP) effects and potentially lowering the interference among multiple partitions.
- Specific z/OS tasks may be dispatched to a small subset of logical processors. 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.

The cooperation between the z10 hardware and the z/OS operating system to increase efficiency will provide minimal, if any, benefit on the z10 BC due to the limited number of processors and therefore lower MP effects inherent in the z10 BC design.

Refer to

http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/Techdocs

Search on the keyword HIPERDISPATCH for more specific information related to HiperDispatch. Refer to the Software requirements section.

LPAR dynamic PU reassignment

System configuration has been enhanced to optimize the CPU-to-book allocation of physical processors (physical units - PUs) dynamically. The initial allocation of customer-usable PUs to physical books can change dynamically to better suit the actual logical partition configurations that are used on the server. Swapping of specialty engines and Central Processors (CPs - general purpose processors) with each other can now occur, as the system attempts to "pack" logical partition configurations into physical configurations that span the least number of books. The effect of this can be observed in dedicated as well as shared partitions that utilize HiperDispatch. The effect of Dynamic PU reassignment will provide minimal, if any, benefit on the z10 BC due to the limited number of processors and the hardware infrastructure inherent in the z10 BC design.

Universal Lift Tool / Ladders

The Universal Lift Tool / Ladders feature (#3759) is designed to provide users 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.

IBM Lifecycle Extension for z/OS V1.7

z/OS V1.7 support was withdrawn September 30, 2008. The Lifecycle Extension for z/OS V1.7 (5637-A01) makes fee-based corrective service for z/OS V1.7 available through September 2010. With the Lifecycle Extension, z/OS V1.7 supports the z10 BC server. Certain functions and features of the z10 BC require later releases of z/OS. For the complete list of software support, see the PSP buckets and the Software requirements section.

Non-raised floor environment

An IBM System z10 Business Class (z10 BC) feature may be ordered to allow use of the z10 BC in a non-raised floor environment. This capability may help ease the cost of entry into the z10 BC; a raised floor may not be necessary for some infrastructures.

The non-raised floor z10 BC implementation is designed to meet all electromagnetic compatibility standards. Feature #7998 must be ordered if the z10 BC is to be used in a non-raised floor environment. A Bolt-down kit (#7992) is also available for use with a non-raised floor z10 BC, providing frame stabilization and bolt-down hardware to help secure a frame to a non-raised floor. The Bolt-down kit (#7992) may be ordered for initial box or MES starting January 28, 2009.

Balanced Power Plan Ahead

Phase currents are minimized when they are balanced among the three input phases. Balanced Power Plan Ahead (#3002) is designed to allow you to order the full complement of bulk power regulators (BPRs) on any configuration, to help ensure that the configuration will be in a balanced power environment. The addition of BPRs on an already installed System z10 BC will be disruptive.

You must have a three phase line cord (#8983), (#8984), (#8986), (#8987), (#8988), or (#8998) when using Balanced Power Plan Ahead (#3002).

Bolt-down kits

Bolt-down kits are available for use with z10 BC for the purpose of physically securing your system in place. The need for such a kit is determined by installation particulars, but may be used to help insure that the equipment stays in place in the event of some type of shock or seismic event.

Three varieties of Bolt-down kits are available:

- (#7990) - Bolt-Down Kit, High-Raised Floor 2098 - This feature provides frame stabilization and bolt-down hardware to help secure a frame to a concrete floor beneath a 11.75- to 16.0-inch (298- to 405-mm) raised floor.
- (#7991) - Bolt-Down Kit, Low-Raised Floor 2098 - This feature provides frame stabilization and bolt-down hardware to help secure a frame to a concrete floor beneath a 9.25- to 11.75-inch (235- to 298-mm) raised floor.
- (#7992) - Bolt-Down Kit, Non-Raised Floor 2098 - This feature provides frame stabilization and bolt-down hardware to help secure a frame to a non-raised floor.

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 future runs on System z. IBM's new System z10 BC delivers a new face for midrange enterprise computing that gives you a whole new world of capabilities to run modern applications. This competitively priced server delivers unparalleled qualities of service to help manage growth and reduce cost and risk in your business, making it ideally suited as the cornerstone of your new enterprise data center.

With a midrange focus, System z10 BC delivers peace of mind when it comes to advanced availability and affordability through a low entry point with very granular scalability, offering 130 different capacities to grow as you do. The System z10 BC is designed to provide performance improvements of up to 1.5 times the total system capacity for general purpose processing of the z9 BC Model S07 and with nearly two times the available memory of the z9 BC.

If you are not currently utilizing System z today, you are missing out on capabilities the z10 BC servers provide: greater than 10,000 secure Web transactions per second for your business, up to 120 GB of memory, 2x the previous generation, (expanding to 248 GB of memory by June 2009), and up to 10 customizable Processor Units (PUs) which deliver the kind of computing horsepower you need for varying workloads. Whether your requirements are to run Online Transaction Processing, Data Serving, Batch Processing, Web Serving, Application Development, or all at the same time, the System z10 BC supports 5 different operating systems for unmatched flexibility. Industry-leading virtualization lets you do it all at the same time with resource sharing for further cost savings. And a new host bus interface uses InfiniBand with a link data rate of 6 GBps, enough to support the full capacity and processing power of the new IBM System z10 BC.

With a design for affordable scalability, System z10 BC will continue to offer investment protection and improved price/performance with upgrades. For example, if you have an IFL specialty engine running z/VM and/or Linux on a System z9 BC, an upgrade to z10 BC will provide up to a 1.4 times improvement in processing capacity at no additional cost in most cases - true investment protection.

As part of our commitment to deliver ongoing price/performance improvement (founded in the Mainframe Charter) and to help increase the economic value to our clients, we are taking actions to reduce the cost of deploying and growing new workloads on System z. Our commitment is to continually assess our client needs and industry conditions and to make changes as required. Our goal is to assure we continue to provide highly competitive alternatives for new workloads being deployed on System z.

Built on a foundation that improves recovery for unplanned outages and reduction of planned outages, the z10 BC 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. If you need more capacity for a short period, with a little preplanning, you just turn it on when you need it. The performance of z10 BC is designed to improve application performance, support more transactions, increase scalability, offer more flexibility, and assist in consolidation of workloads.

Whether you are an existing customer or a new customer looking for better solutions to improve and leverage your company's IT investments, the new face of System z makes the System z10 BC ideally suited as the cornerstone of your new enterprise data center.

Statement of general direction

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

(zAAPs), System z9 Integrated Information Processors (zIIPs), or System Assist Processors (SAPs).

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

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

Product number

For a list of machine types, models, and features, see the Web version of Hardware Announcement ZG08-0806, dated October 21, 2008.

Education support

Visit the following Web site for additional information

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

Contact your IBM representative for course information.

Publications

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

```
        Title                                          Order number

z10 BC System Overview                                 SA22-1085
z10 BC Installation Manual - Physical Planning (IMPP)  GC28-6875
PR/SM Planning Guide                                   SB10-7153
```

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

```
     Title                                      Order number

z10 BC Installation Manual                      GC28-6874
z10 BC Service Guide                            GC28-6878
z10 BC Safety Inspection Guide                  GC28-6877
System Safety Notices                           G229-9054
```

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

```
        Title                                          Order number

Application Programming Interfaces for Java            API-JAVA
Application Programming Interfaces                     SB10-7030
Capacity on Demand User's Guide                        SC28-6871
Agreement for Licensed Machine Code                    SC28-6872
CHPID Mapping Tool User's Guide                        GC28-6825
Common Information Model (CIM) Management Interface    SB10-7154
Coupling Links I/O Interface Physical Layer            SA23-0395
ESCON  and FICON CTC Reference                         SB10-7034
ESCON I/O Interface Physical Layer                     SA23-0394
FICON  I/O Interface Physical Layer                    SA24-7172
Hardware Management Console Operations Guide (V2.10.1) SC28-6873
IOCP User's Guide                                      SB10-7037
Maintenance Information for Fiber Optic Links          SY27-2597
z10 BC Parts Catalog                                   GC28-6876
Planning for Fiber Optic Links                         GA23-0367
SCSI IPL - Machine Loader Messages                     SC28-6839
Service Guide for HMCs and SEs                         GC28-6861
Service Guide for Trusted Key Entry Workstations       GC28-6862
Standalone IOCP User's Guide                           SB10-7152
Support Element Operations Guide (Version 2.10.1)      SC28-6879
TKE PCIX Workstation User's Guide                      SA23-2211
System z Functional Matrix                             ZSW0-1335
OSA-Express Customer's Guide                           SA22-7935
OSA-ICC User's Guide                                   SA22-7990
```

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

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

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

The following Redbook publications are available now:

```
        Title                                       Order number

System z10 Business Class Technical Overview        SG24-7632
z10 BC Technical Guide                              SG24-7516
System z Connectivity Handbook                      SG24-5444
```

The following Redbook publications will provide additional information, once they become available:

```
        Title                                       Order number

System z10 Capacity on Demand                       SG24-7504
Getting Started with InfiniBand on System z10       SG24-7539
 and System z9
Configuration Setup                                 SG24-7571
Server Time Protocol: Planning Guide                SG24-7280
Server Time Protocol: Implementation Guide          SG24-7281
OSA-Express Implementation Guide                    SG24-5948
```

For other IBM Redbooks publications, refer to

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

Services

Global Technology Services

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

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

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

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

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

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

For details on education offerings related to specific products, visit

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

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

Technical information

Specified operating environment

Physical specifications

```
Dimensions:
                     Depth     Width     Height

System with All Covers
  - Inches            71.0      30.9      79.26
```

```
          – Centimeter              185.4       78.5      201.32

System with Covers and Reduction
   – Inches                          71.0       30.9       70.3
   – Centimeter                     185.4       78.5      178.5

Frame on Casters with Packaging (Domestic)
   – Inches                          51.4       32.4      79.76
   – Centimeter                     130.6       82.2     202.58

Frame With Packaging (ARBO Crate)
   – Inches                          51.5       36.5       87.6
   – Centimeter                     130.8       92.7      222.5

Approximate weight:


                              New Build
                               Minimum
                               System
                               Model E10



System with IBF Feature
   – kg                       952.5
   – lb                       2100
System without IBF Feature
   – kg                       857.3
   – lb                       1890
```

To assure installability and serviceability in non-IBM industry-standard racks, review the installation planning information for any product-specific installation requirements.

Standards

The FICON Express4 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).
- 4 KM LX Variant Draft. The 4 KM LX variant may be found in draft document FC-PI-2, Rev. 9.
- 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.

The FICON Express4 features have the following characteristics:

FICON Express4 4KM LX (4 kilometer long wavelength) (#3324) and FICON Express4-2C 4KM LX (#3323)

- Data rate: 1, 2 or 4 Gbps auto-negotiated
- Operating mode: Full duplex
- Defined as: CHPID types FC or FCP (per channel)
- Connector type: LC Duplex

- Channel count: Four LX channels (#3324) or two LX channels (#3323) per feature; a 4KM LX transceiver is designed to interoperate with a 10KM LX transceiver
- Cable type: Single mode fiber optic cabling (9 micron)
- Unrepeated distance at 4 Gbps: 4 km (2.48 miles)
- Unrepeated distance if using MCP cables at 1 Gbps: 550 meters (1804 feet)

FICON Express4 10KM LX (10 kilometer long wavelength) (#3321)

- Data rate: 1, 2 or 4 Gbps auto-negotiated
- Operating mode: Full duplex
- Defined as: CHPID types FC or FCP (per channel)
- Connector type: LC Duplex
- Channel count: Four LX channels per feature; a 10KM LX transceiver is designed to interoperate with a 4KM LX transceiver
- Cable type: Single mode fiber optic cabling (9 micron)
- Unrepeated distance at 4 Gbps: 10 km (6.2 miles)
- Unrepeated distance if using MCP cables at 1 Gbps: 550 meters (1804 feet)

FICON Express4 SX (short wavelength) (#3322) and FICON Express4-2C SX (#3318)

- Data rate: 1, 2, or 4 Gbps auto-negotiated
- Operating mode: Full duplex
- Defined as: CHPID types FC or FCP (per channel)
- Connector type: LC Duplex
- Channel count: Four SX channels (#3322) or two SX channels (#3318) per feature
- Cable type: Multimode fiber optic cabling (50 or 62.5 micron)
- Unrepeated distance:
  - 1 Gbps with 50 micron fiber at 2000 MHz-km: 860 meters (2822 feet)
  - 2 Gbps with 50 micron fiber at 2000 MHz-km: 500 meters (1640 feet)
  - 4 Gbps with 50 micron fiber at 2000 MHz-km: 270 meters (886 feet)
  - 1 Gbps with 50 micron fiber at 500 MHz-km: 500 meters (1640 feet)
  - 2 Gbps with 50 micron fiber at 500 MHz-km: 300 meters (984 feet)
  - 4 Gbps with 50 micron fiber at 500 MHz-km: 150 meters (492 feet)
  - 1 Gbps with 62.5 micron fiber at 200 MHz-km: 300 meters (984 feet)
  - 2 Gbps with 62.5 micron fiber at 200 MHz-km: 150 meters (492 feet)
  - 4 Gbps with 62.5 micron fiber at 200 MHz-km: 70 meters (230 feet)
  - 1 Gbps with 62.5 micron fiber at 160 MHz-km: 250 meters (820 feet)
  - 2 Gbps with 62.5 micron fiber at 160 MHz-km: 120 meters (394 feet)
  - 4 Gbps with 62.5 micron fiber at 160 MHz-km: 55 meters (180 feet)

The OSA-Express3 and OSA-Express2 features are designed to conform to the following standards:

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

The OSA-Express3 and OSA-Express2 features have the following characteristics:

OSA-Express3 Gigabit Ethernet LX (long wavelength) (#3362) and OSA-Express2 Gigabit Ethernet LX (long wavelength) (#3364)

- Data rate: 1000 Mbps (1 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID type OSD or OSN; OSA-Express3 two ports per PCIe adapter, two PCIe adapters per feature
- Defined as: CHPID type OSD or OSN; OSA-Express2 one port per PCIe adapter, two PCIe adapters per feature
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: Four LX ports per feature (#3362)
- Port count: Two LX ports per feature (#3364)
- Cable type: Single mode fiber optic cabling (9 micron); accommodates reuse of existing multimode fiber (50 or 62.5 micron) when used with a pair of mode conditioning patch (MCP) cables
- Unrepeated distance: 5 km (3.1 miles)
- Unrepeated distance if using MCP cables: 550 meters (1804 feet)

OSA-Express3 Gigabit Ethernet SX (short wavelength) (#3363), OSA-Express3-2P Gigabit Ethernet SX (short wavelength) (#3373), and OSA-Express2 Gigabit Ethernet SX (short wavelength) (#3365)

- Data rate: 1000 Mbps (1 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID type OSD or OSN; two ports per PCIe adapter, one and two PCIe adapters per feature
- Defined as: CHPID type OSD or OSN
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: Four SX ports/two PCIe adapters per feature (#3363)
- Port count: Two SX ports/one PCIe adapter per feature (#3373)
- Port count: Two SX ports/two PCIe adapters per feature (#3365)
- 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-Express3 1000BASE-T Ethernet (#3367), OSA-Express3-2P 1000BASE-T Ethernet (#3369), and OSA-Express2 1000BASE-T Ethernet (#3366)

- Data rate: 10, 100, or 1000 Mbps
- Operating modes: Autonegotiate, half duplex, full duplex
- Defined as: CHPID types OSC, OSD, OSE, OSN;
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: RJ-45
- Port count: Four 1000BASE-T ports/two PCIe adapters per feature (#3367)
- Port count: Two 1000BASE-T ports/one PCIe adapter per feature (#3369)
- Port count: Two 1000BASE-T ports/two PCIe adapters per feature (#3366)

- Cable type: EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cable with a maximum length of 100 meters (328 feet)

OSA-Express3 10 Gigabit Ethernet LR (long reach) (#3370)

- Data rate: 10 Gigabits per second (10 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID type OSD
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex (#3370)
- Port count: Two LR ports/two PCIe adapters per feature (#3370)
- Cable type: Single mode fiber optic cabling (9 micron)
- Unrepeated distance: 10 km (6.2 miles)

OSA-Express3 10 Gigabit Ethernet SR (short reach) (#3371)

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

Operating environment

- Temperature: 10 to 32 degrees C (50 to 89 degrees F) for all models up to 900 meters; maximum ambient reduces 1 degree C per 300 meters above 900 meters
- Relative humidity: 8 to 80%
- Wet bulb (Caloric Value): 23 degrees C (73 degrees F) operating mode
- Maximum dew point: 17 degrees C (62.6 degrees F) operating mode
- Electrical power:

  7.3 kVA (typically 0.999 PF at 200 V)
  7.35 kVA (typically 0.99 PF at 380 V)
  7.4 kVA (typically 0.98 PF at 48 0V)

  Note: The above kVA is for a maximum configuration in a warm room (system inlet temperature > 28 deg C/82.4 deg F). Typical configurations in normal environment will average 4 kVA. Exact values for specific configurations will be available using the Power Estimation Tool for this system.

- Capacity of exhaust: 2440 cubic meters / hour (1435 CFM)
- Noise level:

  Declared A-Weighted Sound Power Level, LWAd(B) = 7.2
  Declared A-Weighted Sound Pressure Level, LpAm(dB) = 54
- Leakage and starting current: 105 mA / 135 A (~10 ms)

Hardware requirements

You should review the PSP buckets for minimum MCL and software PTF levels before IPLing operating systems.

The hardware requirements for the System z10 BC and its features and functions are identified below.

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

Access Resource Link at

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

Peripheral hardware and device attachments: IBM devices previously attached to IBM System z9 and zSeries servers are supported for attachment to System z10 BC 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 z10 BC supports devices as described above, IBM does not commit to provide support or service for an IBM device that has reached its End of Service effective date as announced by IBM.

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

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

Software requirements

Listed are the operating systems and the minimum versions and releases supported by z10 BC, 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 2098DEVICE Preventive Service Planning (PSP) bucket prior to installing a z10 BC.

System z10 BC requires at a minimum:

- z/OS V1.7 , with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs. Note that the IBM zIIP Support for z/OS and z/OS.e V1R6/R7 Web deliverable is required to be installed for HiperDispatch (a zIIP processor is not required).

  Note: z/OS V1.7 support was withdrawn September 30, 2008. The Lifecycle Extension for z/OS V1.7 (5637-A01) makes fee-based corrective service for z/OS V1.7 available through September 2010. With the Lifecycle Extension, z/OS V1.7 supports the z10 BC server. Certain functions and features of the z10 BC require later releases of z/OS.
- z/OS V1.8, V1.9, or V1.10 with PTFs.
- z/OS.e V1.8, with PTFs.
- z/VM.
  - CFCC Level 16 Guest Exploitation: z/VM V5.2, V5.3 with PTFs, and V5.4.
  - QDIO Data Connection Isolation : z/VM V5.3 and V5.4 with PTFs.
  - Enhanced installation support for z/VM using the HMC: z/VM V5.4.
  - z/VM Mode Partitions: z/VM V5.4.
  - HCD Support: z/VM V5.2, V5.3, and V5.4 with PTFs.
  - IOCP Support: z/VM V5.2. V5.3, and V5.4 with PTFs.
- z/VSE V3.1 with PTFs, V4.1 with PTFs, or z/VSE V4.2.
- z/TPF V1.1 is required to support 64 engines per z/TPF LPAR.
- TPF V4.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.

– Red Hat RHEL 4 and RHEL 5.

z/VM mode partitions requires at a minimum:

- z/VM V5.4.

Installing Linux from the HMC requires at a minimum:

- z/VM V5.4.

Dynamic Add Logical CPs requires at a minimum:

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

HCA2-O fanout (#0163) supporting InfiniBand coupling links (12x IB-SDR on z9 and 12x IB-DDR on z10) at 150 meters (492 feet) on z10 BC and z10 EC requires at a minimum:

- z/OS V1.7, with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01), and PTFs.
- z/OS V1.8, or z/OS V1.9, or z/OS V1.10 with PTFs.
- z/OS.e V1.8 with PTFs.
- z/VM V5.3 to define, modify, and delete an InfiniBand coupling link, CHPID type CIB, when z/VM is the controlling LPAR for dynamic I/O.

HCA2-O LR fanout (#0168) supporting InfiniBand coupling links (1x IB-SDR or 1x IB-DDR) at an unrepeated distance of 10 km (62 miles) requires at a minimum:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and PTFs.
- z/OS V1.8, or z/OS V1.9, or z/OS V1.10 with PTFs.
- z/OS.e V1.8 with PTFs.
- z/VM V5.3 to define, modify, and delete an InfiniBand coupling link, CHPID type CIB, when z/VM is the controlling LPAR for dynamic I/O.

Coupling Facility Control Code Level 16 on z10 BC requires at a minimum for exploitation of new features:

- z/OS V1.7, with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs.
- z/OS V1.8, V1.9, or V1.10 with PTFs.
- z/OS.e V1.8 with PTFs.
- z/VM V5.2 and V5.3 with PTFs, and V5.4 for guest virtual coupling exploitation.

Hardware Decimal Floating Point on System z10 BC requires at a minimum:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7 and PTFs (for High Level Assembler support).
- z/OS V1.8 with PTFs (for High Level Assembler, Language Environment®, DBX, and CDA RTLE support).
- z/OS.e V1.8 with PTFs (for High Level Assembler, Language Environment, DBX, and CDA RTLE support).
- z/OS V1.9 with PTFs for full support, for C/C++.
- (Optionally) IBM 64-bit SDK for z/OS, Java Technology Edition, V6.0.0 SR1.
- z/VM V5.3.

Capacity provisioning on System z10 BC requires at a minimum:

- z/OS V1.9 or z/OS V1.10 with PTFs (see *z/OS MVS Capacity Provisioning User's Guid*e (SA33-8299) for z/OS functions that must be enabled).

- Linux on System z distributions:
  - Novell SUSE SLES 10 SP2.
  - IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases.

Large Page support (1 megabyte pages) on System z10 BC requires at a minimum:

- z/OS V1.9 or z/OS V1.10 with PTFs.
- Linux on System z distributions:
  - Novell SUSE SLES 9 SP2.
  - Red Hat RHEL 5.2.

CP Assist for Cryptographic Function (CPACF) (#3863) on the System z10 BC requires at a minimum:

- z/OS.
  - z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and either the Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 Web deliverable (no longer available), the Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7 Web deliverable (no longer available), or the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
  - z/OS V1.8.
  - z/OS.e V1.8.
- z/VM V5.2 for guest exploitation.
- z/VSE V3.1 and IBM TCP/IP for VSE/ESA™ V1.5.0 with PTFs.
- z/TPF V1.1.
- TPF V4.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4.3 and RHEL 5.

Enhancements to CP Assist for Cryptographic Function (CPACF) on the System z10 BC requires at a minimum:

- z/OS.
  - z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and either the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
  - z/OS V1.8 or z/OS V1.9 with either the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
  - z/OS.e V1.8 with either the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
  - z/OS V1.10.
- z/VM V5.2 for guest exploitation.
- z/VSE V4.1 and IBM TCP/IP for VSE/ESA V1.5.0 with PTFs.
- Linux on System z distributions:
  - Novell SUSE SLES 10 SP2.
  - Red Hat RHEL 5.2.

Configurable Crypto Express2 and Crypto Express2-1P on the System z10 BC requires at a minimum:

- z/OS.
  - z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and the Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/V1R7 Web deliverable (no longer available), or with the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
  - z/OS V1.8.
  - z/OS.e V1.8.
- z/VM V5.2 for guest exploitation.
- zVSE V3.1 and IBM TCP/IP for VSE/ESA V1.5.0 with PTFs.
- z/TPF V1.1 (acceleration mode only).
- Linux on System z distributions:
  - Novell SUSE SLES 9 SP3 and SLES 10.
  - Red Hat RHEL 4.4 and RHEL 5.1.

Note: z/VSE supports clear-key operations only. Linux on System z and z/VM V5.2, and later, support clear- and secure-key operations.

Note: The Cryptographic Support Web deliverables may be obtained at

http://www-03.ibm.com/systems/z/os/zos/downloads/

Key management for remote loading of ATM and Point of Sale (POS) keys on System z10 BC requires at a minimum:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and the Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/V1R7 Web deliverable (no longer available), or with the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2 for guest exploitation.

Improved Key Exchange with Non-CCA Cryptographic systems on System z10 BC requires at a minimum:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/VM V5.2 for guest exploitation.

Support for ISO 16609 CBC Mode T-DES Message Authentication (MAC) requirements on System z10 BC requires at a minimum:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) and the Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/V1R7 Web deliverable (no longer available), or with the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2 for guest exploitation.

Support for RSA keys up to 4096 bits in length on System z10 BC requires at a minimum:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01), with the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS V1.8 or z/OS V1.9 with either the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS.e V1.8 with either the Cryptographic Support for z/OS V1R7-V1R9 and z/OS.e V1R7-V1R8 Web deliverable, or the Cryptographic Support for z/OS V1R8-V1R10 and z/OS.e V1R8 Web deliverable planned to be available November 21, 2008.
- z/OS V1.10.
- z/VM V5.2 for guest exploitation.

Dynamically Add Crypto to Logical Partition on System z10 BC requires at a minimum:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2 for guest exploitation.
- z/VSE V4.2.
- Linux on System z distributions:
  - Novell SUSE SLES 10 SP1.
  - Red Hat RHEL 5.1.

Secure Key AES on System z10 BC requires at a minimum:

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

Updates to Crypto Facility Query (CFQ) Function on System z10 BC requires at a minimum:

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

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

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

High Performance FICON for System z (zHPF) (CHPID type FC) requires at a minimum:

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

FICON Express4 (CHPID type FC), including Channel-To-Channel (CTC), on z10 BC requires at a minimum:

- z/OS V1.8.
- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1.
- TPF V4.1 at PUT 16.
- z/TPF 1.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

FICON Express4 (CHPID type FCP) for support of SCSI disks on z10 BC requires at a minimum:

- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

HiperSockets Layer 2 support on the z10 BC requires at a minimum:

- z/VM V5.2 for guest exploitation.
- Linux on System z distributions:
  - Novell SUSE SLES 10 SP2.
  - Red Hat RHEL 5.2.

HiperSockets Multiple Write Facility on the z10 BC requires at a minimum:

- z/OS V1.9 with PTFs.
- z/OS V1.10.

OSA-Express3 GbE LX (#3362) and GbE SX (#3363 and #3373) on z10 BC require at minimum:

Supporting CHPID types OSD with **exploitation of four ports per feature** on #3362 and #3363 and **exploitation of two ports per feature** on #3373

- z/OS V1.8 or z/OS V1.9 with PTFs.
- z/OS.e V1.8 with PTFs.
- z/OS V1.10.
- z/VM V5.2 with PTFs.
- z/VSE V4.1 with PTFs.
- z/TPF 1.1 PUT 4 with APARs.
- Linux on System z distributions - for four ports per feature on #3362 and #3363:

- Novell SUSE SLES 10 SP2.
- Red Hat RHEL 5.2.
- Linux on System z distributions - for two ports per feature on #3373:
    - Novell SUSE SLES 9 and SLES 10.
    - Red Hat RHEL 4 and RHEL 5.

Supporting CHPID types OSD with exploitation of **two** ports per feature on #3362 and #3363 and exploitation of **one** port per feature on #3373.

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- TPF V4.1 at PUT 13 with PTF.
- z/TPF 1.1.
- Linux on System z distributions:
    - Novell SUSE SLES 9 and SLES 10.
    - Red Hat RHEL 4 and RHEL 5.

Supporting CHPID type OSN in support of OSA-Express for NCP:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01) with PTFs.
- z/OS V1.8, z/OS.e V1.8, or V1.9 with PTFs.
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- z/TPF 1.1 PUT 4 with APARs.
- Linux on System z distributions:
    - Novell SUSE SLES 9 SP2 and SLES 10.
    - Red Hat RHEL 4.3 and RHEL 5.

OSA-Express3 1000BASE-T (#3367 and #3369) on z10 BC requires at minimum:

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

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

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

For CHPID type OSD and **exploitation of four ports per feature (#3367) and two ports per feature (#3369)**:

- z/OS V1.8 or z/OS V1.9 with PTFs.
- z/OS.e V1.8 with PTFs.
- z/OS V1.10.
- z/VM V5.2 with PTFs.
- z/VSE V4.1 with PTFs.
- z/TPF 1.1 PUT 4 with APARs.

- Linux on System z distributions - for four ports per feature (#3367):
  - Novell SUSE SLES 10 SP2.
  - Red Hat RHEL 5.2.
- Linux on System z distributions - for two ports per feature (#3369):
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

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

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- z/TPF V4.1 PUT 13 with PTF.
- z/TPF V1.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

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

- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1.

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

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- z/TPF 4.1.
- z/TPF 1.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 SP2 and SLES 10.
  - Red Hat RHEL 4.3 and RHEL 5.

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

Supporting CHPID type OSD and two ports per feature:

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01).
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1 with PTFs.
- z/TPF V4.1 PUT 13 with PTF.
- z/TPF V1.1.

- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

OSA-Express QDIO data connection isolation on System z10 and System z9 (CHPID type OSD) requires at a minimum: z/VM V5.3 with PTFs (planned to be available in fourth quarter 2008).

OSA-Express3 10 GbE LR (#3370) and 10 GbE SR (#3371) on z10 BC require at a minimum:

Supporting CHPID type OSD and two ports per feature:

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

OSA-Express2 GbE LX (#3364) and GbE SX (#3365) on z10 BC require at a minimum:

For CHPID type OSD:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1.
- TPF V4.1 at PUT 13 with PTF.
- z/TPF 1.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

For CHPID type OSN in support of OSA-Express for NCP:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7 and PTFs.
- z/OS V1.8 with PTFs.
- z/OS.e V1.8 with PTFs.
- z/VM V5.2.
- z/VSE 3.1 with PTFs.
- TPF 4.1.
- z/TPF 1.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 SP2 and SLES 10.
  - Red Hat RHEL 4.3 and RHEL 5.

OSA-Express2 1000BASE-T Ethernet (#3366) on z10 BC requires at a minimum:

For CHPID type OSC:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1.

For CHPID type OSD:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1.
- TPF V4.1 at PUT 13 with PTF.
- z/TPF 1.1.
- Linux on System z distributions:
  - Novell SUSE SLES 9 and SLES 10.
  - Red Hat RHEL 4 and RHEL 5.

For CHPID type OSE:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE V3.1.

For CHPID type OSN in support of OSA-Express for NCP:

- z/OS V1.7 with IBM Lifecycle Extension for z/OS V1.7 and PTFs.
- z/OS V1.8.
- z/OS.e V1.8.
- z/VM V5.2.
- z/VSE 3.1 with PTFs.
- TPF 4.1.
- z/TPF 1.1
- Linux on System z distributions:
  - Novell SUSE SLES 9 SP2 and SLES 10.
  - Red Hat RHEL 4.3 and RHEL 5.

Planning information

Customer responsibilities

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

Cable orders

Fiber optic cable orders

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

Three phase line cords required for more than two I/O drawers

Single phase line cords support up to a maximum of two I/O drawers using line cords (#8991), (#8990), (#8991), or (#8999). For more than two I/O drawers, it is necessary to use three phase line cords (#8983), (#8984), (#8986), (#8987), (#8988), or (#8998).

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

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

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

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

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

The services are packaged as follows:

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

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

See the Planning Information, Cabling Responsibilities section of this announcement for more information or contact IBM Global Technology Services for details.

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

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

Cabling responsibilities

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

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

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

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

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

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

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

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

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

Installability

The average installation time for a z10 BC is approximately 16 installer hours. This does not include planning hours. This assumes the Pre-Installation Configuration Service, a full System Assurance Product Review, and implementation of the cable services have been performed. See your IBM representative for details on these services.

Security, auditability, and control

The z10 BC 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

Warranty period

One year

Warranty service

The specified level of maintenance service may not be available in all worldwide locations. Additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country- and location-specific information.

IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well lit, and suitable for the purpose.

Warranty service:

The specified level of maintenance service may not be available in all worldwide locations. Additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country and location specific information.

IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well lit, and suitable for the purpose.

The following service is available as warranty for your machine type.

- IBM Onsite Repair, Same Business Day 6 hours average Onsite response Time, 24 hours per day, 7 days a week

Warranty service upgrades

The specified level of maintenance service may not be available in all worldwide locations. Additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country and location specific information.

IBM On-site Service:IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well lit, and suitable for the purpose.

Warranty service upgrades

The specified level of maintenance service may not be available in all worldwide locations. Additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country and location specific information.

IBM On-site Service:IBM will repair the failing machine at your location and verify its operation. You must provide a suitable working area to allow disassembly and reassembly of the IBM machine. The area must be clean, well lit, and suitable for the purpose.

There are no warranty service upgrades.

Usage plan machine

No

IBM hourly service rate classification

A

Maintenance service offerings

This machine is eligible under Terms and Conditions of the IBM ServiceSuite™ (SSU), the IBM Enterprise Service Agreement (ESA) or the IBM Maintenance Agreement. Consult your IBM representative for details.

The maintenance service offerings are as follows :

- IBM Onsite Repair, Same Business Day Onsite Response Time, 9 hours per day, Monday through Friday excluding holidays, Latest Call Registration 12:00.

- ESA and SSU customers will receive 2 hours coverage extensions for no charge: IBM Onsite Repair, Same Business Day Onsite Response Time, 11 hours per day, Monday through Friday excluding holidays, Latest Call Registration 12:00.
- IBM Onsite Repair, Same Business Day Onsite Response Time, 18 hours per day, Monday through Saturday excluding holidays, Latest Call Registration 18:00.
- IBM Onsite Repair, Same Business Day 6 hours average Onsite Response Time, 24 hours per day, 7 days a week.

Committed Services (CS): For service options with a committed level of service or any other special service option, please contact your IBM representative. Refer to the following European documents:

- Announcement Letter ZS03-0150 for IBM Customer Agreement (ICA)
- Announcement Letter ZS04-0135 for Enterprise Agreement Contract
- Announcement Letter ZS98-0118 for ServiceSuite Contract
- Hardware Maintenance Operational Guides and Service Level Code Description Table available at

    http://www-5.ibm.com/services/europe/maintenance/

When a type of service involves the exchange of a machine part, the replacement may not be new, but will be in good working order.

Field-installable features

Yes

Model conversions

Yes

Machine installation

Installation is performed by IBM. IBM will install the machine in accordance with the IBM installation procedures for the machine. Outside the U.S., contact the local IBM office.

Graduated program license charges apply

No

Licensed internal code

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 Agreement for Licensed Internal Code, to enable a specific machine to function in accordance with its specifications, and only for the capacity authorized by IBM and which the customer has acquired. You can obtain the agreement at the Web site below or by contacting your IBM representative.

    http://www.ibm.com/servers/support/machine_warranties/
    licensed_internal_code.html

Specific Machine LIC Type Model: 2098-E10

Europe Business Partner terms and conditions

Category: The products are added to the discount categories A and B.

Exhibit: The products are added to the System z hardware product exhibit.

Pricing

For all local charges, contact your IBM representative.

IBM Global Financing

IBM Global Financing offers competitive financing to credit-qualified customers to assist them in acquiring IT solutions. Offerings include financing for IT acquisition, including hardware, software, and services, from both IBM and other manufacturers or vendors. Offerings (for all customer segments: small, medium, and large enterprise), rates, terms, and availability can vary by country. Contact your local IBM Global Financing organization or visit

http://www.ibm.com/financing

IBM Global Financing offerings are provided through IBM Credit LLC in the United States, and other IBM subsidiaries and divisions worldwide to qualified commercial and government customers. Rates are based on a customer's credit rating, financing terms, offering type, equipment type, and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension, or withdrawal without notice.

Announcement countries

All European, Middle Eastern, and African countries.

Trademarks

HiperSockets, IBM Systems Director Active Energy Manager, Cool Blue, z9, System Storage, DS8000, Resource Link, z/VSE, IMS, MVS, Processor Resource/Systems Manager, PR/SM, VSE/ESA, Electronic Service Agent and ServiceSuite are trademarks of IBM Corporation in the United States, other countries, or both.

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

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

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

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

Terms of use

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

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

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

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

Corrections

(Corrected on October 27, 2008)

The planned availability date for z9 BC upgrades to z10 BC was revised, and other small corrections were made.