



IBM System z10 Enterprise Class: Helping to meet global 24x7 demands for information services with improvements for Internet access and coupling

Table of contents

2 Key prerequisites	10 Publications
2 Planned availability dates	11 Services
3 Description	12 Technical information
8 Product positioning	16 IBM Electronic Services
10 Product number	16 Terms and conditions
10 Education support	17 Prices

At a glance

In today's global economy, businesses have requirements to provide users with information services that are fast, always available, secure, and dependable. The cornerstone of your enterprise data center, the System z10 EC, delivers new technologies that scale, resource management and virtualization capabilities, along with System z™ leadership providing availability, security, and reliability.

The OSA-Express3 features are designed to help meet the demands of your business by offering double the physical connectivity of OSA-Express2 while reducing latency and improving throughput with the introduction of a hardware data router (moving functions from firmware to hardware), a new microprocessor, and a PCI Express bus.

The System z10 and System z9® servers now offer InfiniBand coupling links for a clustered environment, to facilitate communication and time synchronization at increased distance within the data center.

Available on System z10 EC:

- OSA-Express3 Gigabit Ethernet SX and LX with four ports per feature
- OSA-Express3 10 Gigabit Ethernet LR, announced February 26, 2008, with two ports per feature
- InfiniBand coupling links, announced February 26, 2008, supporting:
 - Sharing of physical links with multiple CHPIDs per port
 - Point-to-point connectivity up to 150 meters (492 feet) using high-bandwidth, industry-standard, fiber optic cables
- Balanced Power Plan Ahead to help ensure that the configuration will be in a balanced power environment

Available on System z9 EC dedicated coupling facilities (CFs) and z9™ BC Model S07 dedicated CFs communicating with System z10 EC:

- InfiniBand coupling links, announced February 26, 2008, supporting:
 - Sharing of physical links with multiple CHPIDs per port
 - Point-to-point connectivity up to 150 meters (492 feet) using high-bandwidth, industry-standard, fiber optic cables

For ordering, contact your IBM representative, an IBM Business Partner, or IBM Americas Call Centers at 800-IBM-CALL (Reference: YE001).

Accessing intranets and the Internet with a new generation: OSA-Express3

The OSA-Express3 features (Gigabit Ethernet SX and LX, 10 Gigabit Ethernet LR) provide connectivity to clients and servers using 1 Gbps or 10 Gbps local area networks. They are System z10 EC integrated hardware features installed in an I/O cage, making them integral components of the server I/O subsystem. With the Open Systems Adapter-Express (OSA-Express) functions delivered since June 1999, you have the connectivity, bandwidth, availability, reliability, and recovery that you have come to expect from the mainframe. This new generation is designed to deliver reduced latency, improved throughput, and double the port density.

A new technology for Parallel Sysplex® # InfiniBand coupling links

Clustering technology is designed to bring the strengths of parallel processing to business-critical applications. A Parallel Sysplex cluster can consist of up to 32 z/OS® images, coupled to one or more coupling facilities, using specialized channels (coupling links) for communication and exchanging messages for time synchronization (Server Time Protocol). The coupling facilities, at the heart of the cluster, enable record-level read/write data sharing among the images in a cluster. When configured properly, a cluster can provide users with near-continuous application availability for planned and unplanned outages.

InfiniBand coupling links provide an additional option for your Parallel Sysplex cluster on System z10 and System z9. When used in the data center, InfiniBand coupling links are designed to complement but may not replace Integrated Cluster Bus-4 (ICB-4) and InterSystem Channel-3 (ISC-3) links.

With InfiniBand coupling links, the design allows up to 16 CHPIDs to be defined across two ports on an HCA-O fanout; physical coupling links can be shared by multiple sysplexes. For example, this capability allows for one CHPID to be directed to one coupling facility and a second CHPID to be directed to a separate coupling facility on the same target server using the same port. With the MBA fanout, supporting ICB-4, there are two CHPIDs, one per port.

OSA-Express3 10 GbE LR (#3370), HCA2-O fanout for InfiniBand coupling links on z10 EC (#0163), and HCA1-O fanout for InfiniBand coupling links on z9 EC and z9 BC Model S07 (#0167) were announced in Hardware Announcement [108-154](#), dated February 26, 2008.

Key prerequisites

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

Planned availability dates

May 6, 2008, for:

- z10 EC Site Tool Kit (#9968) on System z10 EC new build models and MES features

May 30, 2008, for:

- OSA-Express3 features (#3362, #3363, #3370) on System z10 EC
- HCA2-O fanout for InfiniBand coupling links (#0163) on System z10 EC
- HCA1-O fanout for InfiniBand coupling links (#0167) on System z9 EC dedicated CFs and z9 BC Model S07 dedicated CFs
- Balanced Power Plan Ahead (#3001) on System z10 EC new build models
- Balanced Power Plan Ahead (#3001) on System z10 EC upgrades from z9 EC
- Balanced Power Plan Ahead (#3001) on System z10 EC upgrades from z990

August 29, 2008, for:

- Feature MESs that include Balanced Power Plan Ahead (#3001) on System z10 EC models
- Feature MES adds to System z10 EC models which have Balanced Power Plan Ahead (#3001) already installed

Description

Accessing intranets and the Internet with a new generation: OSA-Express3

OSA-Express3 is designed for high-speed communication in the enterprise backbone or between campuses, to connect server farms, or to consolidate file servers onto z10 EC. The workload can be Internet Protocol-based (IP) or non-IP-based. OSA-Express3 is an integrated hardware feature optimized for mainframe local area network (LAN) connectivity.

OSA-Express3: Why a new design? What makes it faster?

New hardware infrastructure

The OSA-Express3 design includes:

- New microprocessor -# The IBM application-specific integrated circuit (ASIC) has been redesigned and includes a hardware data router.
- New PCI adapter, PCI Express (PCI-E).

With this new hardware infrastructure, the OSA-Express3 features are designed to deliver reduced latency and improved throughput for standard (1492 byte) and jumbo frames (8992 byte).

Direct Memory Access (DMA)

OSA-Express3 and the operating systems share a common storage area for memory-to-memory communication, reducing system overhead and improving performance. There are no read or write channel programs for data exchange. For write processing, no I/O interrupts have to be handled. For read processing, the number of I/O interrupts is minimized.

Hardware data router

With OSA-Express3, much of what was previously done in firmware (packet construction, inspection, and routing) is now performed in hardware. This allows packets to flow directly from host memory to the LAN without firmware intervention.

With the hardware data router, the "store and forward" technique is no longer used, which enables true direct memory access, a direct host memory-to-LAN flow, returning CPU cycles for application use.

The family of OSA-Express3 features include:

- Gigabit Ethernet short wavelength (GbE SX) used with multimode fiber optic cabling
- Gigabit Ethernet long wavelength (GbE LX) used with single mode fiber optic cabling
- 10 Gigabit Ethernet long reach (10 GbE LR) used with single mode fiber optic cabling

OSA-Express3 is designed to outperform OSA-Express2

OSA-Express3: Up to a 45% reduction in latency

OSA-Express3 is designed to reduce the minimum round-trip networking time between z10 EC systems. With OSA-Express3 10 Gigabit Ethernet, the minimum round-trip time at the TCP/IP application layer is 40% less than with OSA-Express2 10 Gigabit Ethernet. With OSA-Express3 Gigabit Ethernet, it is 45% less than with OSA-Express2 Gigabit Ethernet.

OSA-Express3 10 GbE throughput: Up to a 4x improvement

OSA-Express3 is also designed to improve throughput. When handling a mix of inbound and outbound TCP streams traffic, OSA-Express3 10 Gigabit Ethernet delivers 1 gigabyte per second (1 GBps, or 1,000,000,000 bytes per second) when the packet maximum transmission unit (MTU) size is 1492 bytes. When the MTU size is 8992-byte, throughput is 1.1 GBps. These improvements range approximately from 3x to 4x that of OSA-Express2 10 Gigabit Ethernet measured on a z9 EC system.

In a laboratory environment, with a highly efficient traffic generator, OSA-Express3 10 Gigabit Ethernet reaches 90% of Ethernet line speed when sending 1506-byte frames outbound to the network. When sending 4048-byte frames outbound, line speed is effectively reached. The maximum one-way speed of the link is 1.25 GBps.

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

Double the port density on all OSA-Express3 features -# up to 96 ports

The OSA-Express3 features have double the port density of their predecessors, OSA-Express2.

- OSA-Express3 10 Gigabit Ethernet LR has two ports per feature.
- OSA-Express3 Gigabit Ethernet SX and LX each have four ports per feature.

When all ports are OSA-Express3, there can be up to 96 ports, versus the maximum of 48 ports when all ports are OSA-Express2. With more ports per feature, the number of required I/O slots and I/O resources can be reduced.

Note: The maximum of 96 ports can be achieved only when there are 24 OSA-Express3 Gigabit Ethernet features, with 4 ports per feature.

Four-port exploitation on OSA-Express3 GbE SX and LX

For the operating system to recognize all four ports on an OSA-Express3 Gigabit Ethernet feature, a new release and/or PTF is required. If software updates are not applied, only two of the four ports will be "visible" to the operating system.

Activating all four ports on an OSA-Express3 feature provides you with more physical connectivity to service the network and reduces the number of required resources (I/O slots, I/O cages, fewer CHPIDs to define and manage).

Four-port exploitation is supported by z/OS, z/VM®, z/VSE™, z/TPF, and Linux® on System z. Refer to the [Software requirements](#) section for availability dates and required operating system levels for this exploitation.

OSA-Express3 10 GbE has a new, small form factor, connector

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

QDIO architecture continues to be optimized for high-speed communication

As the capacity of the mainframe has increased and the amount of data that is processed has grown, it has been necessary to continue to optimize the architecture to support the demand for reduced overhead and latency and improved throughput. An enhanced version of the I/O architecture, the Queued Direct Input/Output Architecture (QDIO), is designed to "handle the data." Since its introduction in 1999, it too has been enhanced to keep pace with application bandwidth demands.

QDIO is a highly efficient data transfer mechanism that is designed to dramatically reduce system overhead and improve throughput by using data queues in main memory and a signaling

protocol (a signal adapter instruction referred to as a SIGA). The SIGA is processed by the System Assist Processor (SAP) and passes a pointer to the data, which already resides in main memory. QDIO is the interface between the operating system and the OSA-Express3 hardware.

QDIO and Layer 3 versus Layer 2 -# two transport modes

- Layer 3 transport is the Internet Protocol (IP) or network layer and is used for IP traffic.
- Layer 2 transport is the Media Access Control (MAC) or link layer and is used for IP and non-IP traffic.

Layer 2 transport mode: When would it be used?

If you have an environment with an abundance of Linux images in a guest LAN environment, or you need to define router guests to provide the connection between these guest LANs and the OSA-Express3 features, then using the Layer 2 transport mode may be the solution.

If you have Internetwork Packet Exchange (IPX), NetBIOS, and SNA protocols, in addition to Internet Protocol Version 4 (IPv4) and IPv6, use of Layer 2 could provide "protocol independence."

The OSA-Express3 features have the capability to perform like Layer 2 type devices, providing the capability of being protocol- or Layer-3-independent (that is, not IP-only). With the Layer 2 interface, packet forwarding decisions are based upon Link Layer (Layer 2) information, instead of Network Layer (Layer 3) information. Each operating system attached to the Layer 2 interface uses its own MAC address. This means the traffic can be IPX, NetBIOS, SNA, IPv4, or IPv6.

An OSA-Express3 feature can filter inbound datagrams by Virtual Local Area Network identification (VLAN ID, IEEE 802.1q), and/or the Ethernet destination MAC address. Filtering can reduce the amount of inbound traffic being processed by the operating system, reducing CPU utilization.

Layer 2 transport mode is supported by z/VM and Linux on System z.

Layer 3 Virtual MAC for z/OS

To help simplify configuration definition and management, as well as facilitating load balancing when an LPAR is sharing the same OSA-Express MAC address with another LPAR, each operating system instance can have its own unique "logical" or "virtual" MAC (VMAC) address. All IP addresses associated with a TCP/IP stack are accessible using their own VMAC address, instead of sharing the MAC address of an OSA port. This applies to Layer 3 mode and to an OSA-Express3 or OSA-Express2 port shared among Logical Channel Subsystems. This support is designed to:

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

VMACs are currently available for Layer 2 mode in the z/VM and Linux on System z environments. Layer 3 VMAC is applicable to the OSA-Express3 and OSA-Express2 features when configured as CHPID type OSD (QDIO). It is supported by z/OS V1.8 with PTFs and z/VM V5.2 for guest exploitation.

Internet Protocol (IP) assist functions: OSA-Express3 QDIO assists in IP processing and offloads TCP/IP stack functions. This is a "highlights" list and may not include all IP assist functions.

- Address Resolution Protocol (ARP) processing -# mapping of IP addresses to Media Access Control (MAC) addresses
- Multicast support -# for sending data to multiple recipients
- Broadcast support -# simultaneously transmitting data to more than one destination; messages are transmitted to all stations in a network (for example, a warning message from a system operator)
- Building MAC and Logical Link Control (LLC) headers
- Packet filtering -# TCP/IP stack sees only IP datagrams
- Checksum offload for IPv4 packets -# calculates the TCP/UDP and IP header checksums, verifies the correctness of files, and reduces host CPU cycles
- Large send for IPv4 packets -# host sends 64 KB blocks to Open Systems Adapter for packet processing, returning CPU cycles for application use

OSA-Express3 10 GbE LR (#3370), GbE SX (#3363), and GbE LX (#3362) are exclusive to z10 EC and support CHPID type OSD. They are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z. Refer to the [Software requirements](#) section.

Preview # OSA-Express3 support for OSA-Express for NCP

OSA-Express for Network Control Program (NCP), channel path identifier (CHPID) type OSN, will be made available for use with the OSA-Express3 GbE features.

With the introduction of OSA-Express for NCP, supporting the channel data link control (CDLC) protocol, it became possible to provide 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 provides 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 z, OSA-Express2 or OSA-Express3.

OSA-Express for NCP is currently supported on the OSA-Express2 GbE features in the z/OS, z/VM, z/VSE, TPF, z/TPF, and Linux on System z environments.

A new technology for Parallel Sysplex # InfiniBand coupling links

InfiniBand, a new generation of coupling links, was announced February 26, 2008. InfiniBand coupling links (IFBs) are now being made available for System z10 EC environments and for System z9 EC and z9 BC Model S07 dedicated coupling facilities (CFs).

InfiniBand coupling links # 150 meters

Coupling facilities can now be separated by up to 150 meters (492 feet). The fiber optic cable is a high-bandwidth, industry-standard OM3 50 micron multimode fiber rated at 2000 MHz-km with a Multi-fiber Push-On (MPO) connector (the same connector that has been used with Fiber Quick Connect # a fiber optic harness on System z # for many years). It is a data center "proven" connector.

If ISC-3 links are used within the data center, InfiniBand coupling links may be an alternative.

Now 12 "lanes" for Parallel Sysplex traffic

InfiniBand coupling links use fiber optic cabling containing 12 pairs (12x) of fiber compared to one pair (1x) of fiber used with ISC-3 fiber optic cabling.

- InfiniBand coupling links support double data rate (DDR) when a z10 EC is communicating with another z10 EC.

- InfiniBand coupling links support single data rate (SDR) when a z10 EC is communicating with a z9 EC dedicated CF or z9 BC Model S07 dedicated CF.
- There are 12 "lanes" (12x) with 12 fibers for sending traffic and 12 fibers for receiving traffic (24 fibers).
- When a link negotiates to 12x DDR, each "lane" supports a data rate of 5 Gbps (0.5 GBps) for a total link data rate of 6 GBps.
- When a link negotiates to 12x SDR, each "lane" supports a data rate of 2.5 Gbps (0.25 GBps) for a total link data rate of 3 GBps.

When the InfiniBand coupling link is z10 EC-to-z10 EC, the link auto-negotiates to 6 GBps. A z10 EC system auto-negotiates to 3 GBps when connected to a z9 EC or z9 BC dedicated coupling facility.

Note: The InfiniBand link data rate of 6 GBps or 3 GBps does 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. With InfiniBand coupling links, while the link data rate may be higher than that of ICB, the service times of coupling operations are greater, and the actual throughput may be less than with ICB links.

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

The whitepaper is available at

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

New hardware in support of InfiniBand coupling links

A new infrastructure was created to support an InfiniBand coupling link environment. Host channel adapter optical (HCA-O) fanouts are being introduced for z10 EC and for z9 EC and z9 BC dedicated coupling facilities. The HCA-O fanouts, with two ports per fanout, reside on the front of each processor book. The fiber optic cables are plugged directly into the front of the HCA-O fanouts.

- HCA2-O fanout for z10 EC
- HCA1-O fanout for z9 EC and z9 BC Model S07 dedicated coupling facilities

New definitions for InfiniBand coupling links

There is a new physical definition to associate with a channel path identifier, an adapter identification. Unlike channels installed in an I/O cage, which are identified by a physical channel path identifier (PCHID) number related to their physical location, HCA-O fanouts and ports are identified by an adapter identification (AID) value which is determined by its physical location. The AID must be used to assign a CHPID to the fanout in the hardware configuration definition. The CHPID assignment is done by associating the CHPID to an AID and port.

The AID assigned to a fanout can be found in the PCHID report provided for each new server or for upgrades on z10 EC, z9 EC, and z9 BC.

There is also a new CHPID type CIB (coupling using InfiniBand). CHPID type CIB is common for z10 EC, z9 EC, and z9 BC.

On z10 EC, z9 EC, and z9 BC the design allows up to 16 CHPIDs to be defined across the two ports on each HCA-O fanout. This can reduce the number of coupling links; physical coupling links to be shared by multiple sysplexes. For example, this capability allows for one CHPID to be directed to one coupling facility and a second CHPID to be directed to a separate coupling facility on the same target server, using the same port. With the MBA fanout, supporting ICB-4, there are two CHPIDs, one per port.

The 16 CHPIDs can be defined in any combination across the two ports. There is no specific limit per port, as long as 16 CHPIDs per HCA-O fanout is not exceeded.

An increased number of CHPIDs per physical link can help to facilitate consolidation of ISC-3 links onto InfiniBand coupling links.

Additional coupling link options

- Internal coupling channel (IC) for internal communication between CFs defined in LPARs, and z/OS images on the same server. ICs operate at memory speed.
- Integrated Cluster Bus-4 (ICB-4) for coupling facility communication over short distances, using 10 meter (33 feet) copper cables. ICB-4 supports a link data rate of 2 GBps.
- InterSystem Channel-3 (ISC-3) for coupling facility communication at unrepeated distances up to 10 km (6.2 miles) using single mode fiber optic cables. ISC-3 supports a link data rate of 2 gigabits per second (2 Gbps).

InfiniBand coupling links can also be used to exchange timekeeping messages for Server Time Protocol (STP).

Refer to the online Sales Manual on the Web for minimums and maximums for IC, ICB-4, IFB, and ISC-3

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

You can choose the coupling links that best suit your business needs: IC, ICB, IFB, or ISC-3.

InfiniBand coupling links are supported on z10 EC as well as z9 EC dedicated coupling facilities and z9 BC Model S07 dedicated coupling facilities. Operating system support is provided by z/OS and z/VM for dynamic I/O. Refer to the [Software requirements](#) section.

Balanced Power Plan Ahead

Phase currents are minimized when they are balanced among the three input phases. Balanced Power Plan Ahead (#3001) 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 EC will be disruptive. The configuration tool will determine if additional internal battery features (#3211) are required.

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

Section 508 of the U.S. Rehabilitation Act

System z10 Enterprise Class servers are capable on delivery, when used in accordance with IBM's associated documentation, of satisfying the applicable requirements of Section 508 of the Rehabilitation Act of 1973, 29 U.S.C. Section 794d, as implemented by 36 C.F.R. Part 1194, provided that any Assistive Technology used with the product properly interoperates with it.

Product positioning

Today we are announcing:

- A new generation of Gigabit Ethernet (GbE) features, with double the port density of the OSA-Express2 GbE features. These features provide connectivity to multimode and single mode fiber optic networks at a link data rate of 1 gigabit per second (1 Gbps). Better yet, the new GbE features are designed to help reduce latency and to improve throughput for standard frames (1492 byte) and jumbo frames (8992 byte).
- The availability date for the OSA-Express3 10 Gigabit Ethernet long reach (10 GbE LR) feature, which has double the port density of the OSA-Express2 10 GbE LR feature, along with reduced latency and improved throughput.

- The availability date for InfiniBand coupling links (IFBs) on z10 EC as well as z9 EC dedicated coupling facilities (CFs) and z9 BC Model S07 dedicated CFs.

OSA-Express3

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

For the operating system to recognize all four ports on an OSA-Express3 Gigabit Ethernet feature, a new release and/or PTF is required. If software updates are not applied, only two of the four ports will be "visible" to the operating system.

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

- Reduced latency (up to 45% reduction) and increased throughput (up to 4x) for applications (software changes are not required to exploit the two ports on OSA-Express3 10 GbE or to exploit two of the four ports on OSA-Express3 GbE SX or LX)
- More physical connectivity to service the network and less resources:
 - Fewer CHPIDs to define and manage
 - Reduction in the number of required I/O slots
 - May reduce the number of I/O cages
 - Double the port density of OSA-Express2
 - Satisfies requirement for more than 48 LAN ports (now up to 96 ports)

OSA-Express2 GbE continues 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 four ports per feature and if you are not resource-constrained.

Note: OSA-Express for NCP supporting the channel data link control (CDLC) protocol and providing connectivity between System z operating systems and IBM Communication Controller for Linux (CCL), is currently only supported on OSA-Express2. Refer to the Preview in the [Description](#) section.

When to select InfiniBand coupling links

InfiniBand coupling links may be an ideal choice for consolidating multiple ISC-3 links within a data center at distances up to 150 meters. In addition, you now have the capability to define multiple CHPIDs on the ports of an HCA-O fanout.

While InfiniBand coupling links have flexible sharing capability and provide increased distance using industry-standard fiber optic cabling, compared to Integrated Cluster Bus (ICB), InfiniBand coupling links may not be appropriate for all workloads. ICB-3 (available on z9 EC, z9 BC, z990, z890) and ICB-4 (available on z10 EC, z9 EC, z9 BC, z990, z890) may still be the preferred choice for workloads that drive a high rate of synchronous coupling facility operations.

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

The Coupling Facility Configuration Options whitepaper is available at

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

Refer also to Getting Started with InfiniBand on System z10 and System z9 (SG24-7539) at

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

Balanced Power Plan Ahead

Phase currents are minimized when they are balanced among the three input phases. Balanced Power Plan Ahead (#3001) allows you to order the full complement of bulk power regulators on any configuration, to help ensure that the configuration will be in a balanced power environment.

Product number

Description	Machine type	Model	Feature
System z10 EC	2097	E12 E26 E40 E56 E64	
OSA-Express3 GbE LX			3362
OSA-Express3 GbE SX			3363
z10 EC Site Tool Kit			9968
Balanced Power Plan Ahead			3001

Business Partner information

If you are a Direct Reseller - System Reseller acquiring products from IBM, you may link directly to Business Partner information for this announcement. A PartnerWorld ID and password are required (use IBM ID).

<https://www.ibm.com/partnerworld/mem/sla.jsp?num=108-296>

Education support

Visit the following Web site for additional information

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

Call IBM IT Education Services at 800-IBM-TEACH (426-8322) for catalogs, schedules, and enrollments.

Publications

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

Title	Order number
z10 EC System Overview	SA22-1084
z10 EC Installation Manual - Physical Planning	GC28-6865
z9 EC System Overview	SA22-6833
z9 BC System Overview	SA22-1083
z9 EC Installation Manual - Physical Planning	GC28-6844
z9 BC Installation Manual - Physical Planning	GC28-6855
System z9 PR/SM™ Planning Guide	SB10-7041
Open Systems Adapter-Express Customer Guide and Reference	SA22-7935

The following publications are shipped with the product and will be available in the Library section of Resource Link at planned availability:

Title	Order number
z10 EC Installation Manual	GC28-6864
z10 EC Service Guide	GC28-6866

Services

Global Technology Services

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

Developed in partnership with Systems and Technology Group, GTS Integrated Communications Services has made available **Network Integration Services for data center networks - private optical networking** to help provide design and implementation services for implementing private optical networking connectivity between data centers. This service is a key component of an end-to-end solution and is a holistic approach to enabling the availability of business data and applications across mirrored and backup data centers -# bringing together System z using GDPS®, other server platforms, and System Storage™ using their respective protocols -# all over an optical networking transport.

System z has announced OSA-Express3, the newest generation of integrated LAN adapters, exclusively available on System z10. When LAN connectivity is extended beyond the borders of a single data center, it is often incorporated into the connectivity design of Parallel Sysplex and GDPS installations. The design of optical networking fabrics, both within and between data centers, may be impacted with the introduction of OSA-Express3 10 Gigabit Ethernet.

To help alleviate potential problems associated with improper planning of the data center infrastructure, and to offer IBM customers a qualified end-to-end solution, GTS is also jointly announcing enhancements to its private optical networking service offerings. These include providing technical focal points for the entire network solution, not just a subset, improving the design of linkages between servers, storage, switches, and WDM fabric connections, and offering additional preinstallation connectivity planning support. Combined with enhanced connectivity from System z, this offering enables customers to fully exploit the potential of the new enterprise data center.

For details on the Network Integration Services for data center networks - private optical networking, refer to Services Announcement [608-029](#), dated May 06, 2008.

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 GDPS, contact your IBM representative or visit

<http://www.ibm.com/systems/z/advantages/gdps/index.html>

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.

Specified operating environment

Standards

The OSA-Express3 Gigabit Ethernet SX and LX features are designed to conform to the following standards:

- Gigabit Ethernet (1000BASE-SX, 1000BASE-LX)
 - IEEE 802.3ac
 - IEEE 802.1q
 - IEEE 802.3x
 - IEEE 802.3z
 - DIX Version 2

The OSA-Express3 10 Gigabit Ethernet Long Reach feature is designed to conform to the following standards:

- 10 Gigabit Ethernet LR (10GBASE-LR)
 - IEEE 802.3ae
 - IEEE 802.1q
 - IEEE 802.3x - flow control
 - DIX Version 2

The OSA-Express3 features have the following characteristics:

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

- Data rate: 1000 Mbps (1 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID type OSD, two CHPIDs, one CHPID for each PCI-E adapter (one CHPID shared by the two ports on a PCI-E adapter)
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: **Four** LX ports per feature
- 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)

- Data rate: 1000 Mbps (1 Gbps)
- Operating mode: Full duplex
- Defined as: CHPID type OSD, two CHPIDs, one CHPID for each PCI-E adapter (one CHPID shared by the two ports on a PCI-E adapter)
- Frame size: DIX V2: 1492 bytes; for jumbo frame 8992 bytes
- Connector type: LC Duplex
- Port count: **Four** SX ports per feature
- 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 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
- Port count: **Two** LR ports per feature
- Cable type: Single mode fiber optic cabling (9 micron)
- Unrepeated distance: 10 km (6.2 miles)

Operating environment

Hardware requirements

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>

Software requirements

Listed are the operating systems and the minimum versions and releases supported by System z10 EC, z9 EC, and z9 BC. Select the releases appropriate to your operating system environments.

Note: Refer to the z/OS, z/VM, z/VSE subsets of the 2097DEVICE Preventive Service Planning (PSP) bucket prior to installing a z10 EC or adding InfiniBand coupling links or OSA-Express3 features to a previously installed z10 EC.

Note: Refer to the z/OS, z/VM, z/VSE subsets of the 2094DEVICE and 2096DEVICE Preventive Service Planning (PSP) bucket prior to installing a z9 EC or z9 BC, or adding InfiniBand coupling links or OSA-Express3 features to a previously installed z9 EC or z9 BC.

There may be a minimum Machine Change Level (MCL) required by operating systems. It is critical that all bucket information is reviewed and adhered to prior to IPLing the operating system.

HCA2-O fanouts (z10 EC) and HCA1-O fanouts (z9 EC, z9 BC Model S07) supporting InfiniBand coupling links require at a minimum:

- z/OS V1.7 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.

OSA-Express3 Gigabit Ethernet LX (#3362) and SX (#3363) on the z10 EC require at a minimum:

Supporting CHPID type OSD and exploitation of four ports per feature:

- z/OS V1.9 with PTFs (planned to be available third quarter 2008).
- z/VM V5.2 with PTFs (planned to be available July 2008).
- z/VSE V4.1 with PTFs (planned to be available July 2008).
- z/TPF 1.1 PUT 4 with APARs (planned to be available June 2008).

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

Supporting CHPID type OSD and **two** of the four ports per feature:

- z/OS V1.7.
- z/VM V5.2.
- z/VSE V3.1 with PTF.
- 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.

OSA-Express3 10 Gigabit Ethernet LR (#3370) on the z10 EC requires at a minimum:

Supporting CHPID type OSD and two ports per feature:

- z/OS V1.7.
- z/VM V5.2.
- z/VSE V3.1 with PTF.
- 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.

Planning information

Customer responsibilities

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

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

Fiber optic cable orders

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

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

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

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

These services are designed to be right-sized for your products or the end-to-end enterprise, and to take into consideration the requirements for all of the protocols and media types supported

on the System z10 EC, System z9, and zSeries® (for example, ESCON, FICON, ISC-3, OSA-Express) whether the focus is the data center, the Storage Area Network (SAN), the Local Area Network (LAN), or the end-to-end enterprise.

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

The services are packaged as follows:

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

Under **IBM IT Facilities Assessment, Design, and Construction Services** there is the 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 **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 EC. Installation Planning Representatives (IPRs) and System Service Representatives (SSRs) will not perform the fiber optic cabling tasks without a services contract.

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

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

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

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

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

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

Security, auditability, and control

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

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

IBM Electronic Services

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

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

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

To learn how Electronic Services can work for you, visit

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

Terms and conditions

MES discount applicable

No

Field installable feature

Yes

Warranty period

One year.

Features assume the same warranty or maintenance terms as the machine in which they are installed for the full warranty or maintenance period announced for such machine.

Customer setup

No

Machine code

Same license terms and conditions as base machine

Prices

Prices are available upon request.

Description	Mach type	Mod	Feat	Init/ MES
System z10 EC	2097	E12 E26 E40 E56 E64		
OSA-Express3 GbE LX*			3362	Both
OSA-Express3 GbE SX*			3363	Both
z10 EC Site Tool Kit*			9968	Both
Balanced Power Plan Ahead*			3001	Both

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

Trademarks

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

System z9, Parallel Sysplex, z/OS, IBM, ESCON, FICON, PartnerWorld, OpenPower, zSeries, z/VM, WebSphere, GDPS, ibm.com and System/390 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.

Microsoft and Windows are registered trademarks of Microsoft Corporation 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. 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/us/>