
**IBM System z
Introduction
February 2008**

**IBM System z10
Enterprise Class (z10 EC)**

Frequently Asked Questions

Worldwide



Table of Contents

Announcement Overview	4
IBM Destination z	9
z10 EC Hardware	12
z10 EC Upgradeability	19
Availability Enhancements	20
IFL, zAAP and zIIP specialty engines	22
z10 EC Pricing	30
Performance	34
Networking Enhancements	39
I/O Enhancements	41
Capacity on Demand Offerings	43
Cryptographic Enhancements	50
Parallel Sysplex Enhancements	59
IBM Systems Director Active Energy Manager (AEM) on Linux for System z V3.1	61
System z and Cell Broadband Engine Architecture	62
GDPS Enhancements	65
IBM Global Financing (IGF)	66
Statements of Direction	71

Announcement Overview

Question:

What is the new System z™ product that was announced on February 26, 2008?

Answer:

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

The z10 EC delivers:

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

offers more protection and security options with Enhanced Encryption Standard (AES) 192 and 256 and stronger hash algorithm with Secure Hash Algorithm (SHA-512 and SHA-384).

- HiperDispatch for improved efficiencies between hardware and the z/OS operating system (z/OS V1.7 and above).
- Hardware Decimal Floating Point unit on each core on the Processor Unit (PU), which can aid in decimal floating point calculations and is designed to deliver performance improvements and precision in execution.
- Large page support (1 MB pages).
- Up to 336 FICON® Express4 channels.
- Fiber Quick Connect (FQC), a fiber harness integrated in the z10 EC frame for a 'quick' connect to ESCON® and FICON LX channels.
- Support for the IBM System Director Active Energy Manager (AEM) for Linux® on System z V3.1 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.

Question:

What part does the z10 EC play in a New Enterprise Data Center?

Answer:

Clearly the strengths of the mainframe - particularly its flexibility and responsiveness - make the z10 EC a great fit for the future as the cornerstone for the New Enterprise Data Center. The z10 is highly virtualized and secure, green by design (selected by IBM for "Project Big Green"), is policy driven with automated service management and provisioning, and has the new 'just-in-time' provisioning for resources.

Question:

Who would be interested in the z10 EC?

Answer:

The z10 EC has several audiences that will be interested in its announcements. Existing System z customers who are looking for massive scalability for their data center but want to minimize floor space will find the high performance, energy efficient, scalable z10 EC to be a perfect fit. Customers who want the financial and environmental benefits available by doing large scale consolidation will find the z10 EC, along with z/VM® and Linux for System z, will hit the mark for their needs. Customers who want a system that can be the cornerstone of a New Enterprise Data Center will find the z10 EC to be an excellent match with its intelligent policy driven / workload management capabilities, its security rich environment, its data server abilities and its superior virtualization technologies.

Question:

How do I order the z10 EC?

Answer:

Contact your IBM System z Representative or IBM Business Partner.

Question:

What are the planned availability dates of the z10 EC announced on February 26, 2008?

Answer:

For the z10 EC:

- Features and functions for the z10 EC: February 26, 2008
 - z10 EC Models E12, E26, E40, E56 and E64: February 26, 2008
 - IBM eServer™ zSeries® 990 (z990) upgrades to z10 EC: February 26, 2008
 - System z9® Enterprise Class (z9™ EC) upgrades to z10 EC: February 26, 2008
 - MES features for Models E12, E26, E40, E56 and E64: May 26, 2008
 - Model conversions for z10 EC: May 26, 2008
 - InfiniBand coupling links (#0163): Second quarter, 2008
 - System z9 EC InfiniBand coupling links (#0167): Second quarter, 2008
 - System z9 BC InfiniBand coupling links (#0167): Second quarter, 2008
 - OSA-Express3 10 GbE LR (#3370): Second quarter, 2008
-
-

Question:

Will IBM continue to market the IBM System z9 Enterprise Class (z9 EC) and IBM System z9 Business Class (z9 BC)?

Answer:

Yes, both the z9 EC and the z9 BC will continue to be manufactured and sold.

Question:

Will IBM offer a midmarket System z10 EC product?

Answer:

IBM does not comment on future plans. The IBM System z9 BC will continue to be our midmarket program offering.

Question:

What software is supported on the z10 EC?

Answer:

Listed are the operating systems and the minimum versions and releases supported by z10 EC, 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 2097 DEVICE Preventive Planning (PSP) bucket prior to installing a z10 EC.

The z10 EC requires at a minimum:

- z/OS V1.7, V1.8, V1.9 with PTFs (z/OS V1.7 requires the IBM zIIP Support for z/OS V1.7 Web deliverable to be installed to enable HiperDispatch)
- z/VM V5.2 and V5.3. Guest exploitation of the z10 EC at the level of System z9 functionality with the PTFs for APARs VM64180 and VM64242
- z/VSE V3.1, V4.1 Compatibility Support with PTFs

- 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
-

Question:

What levels of z/OS are supported with the z10 EC?

Answer:

z/OS V1.7 and higher, including z/OS V1.10 when available, are supported on the z10 EC (some maintenance required). In addition, z/OS R1.7 requires the zIIP Web Deliverable to enable HiperDispatch.

Note: z/OS V1.7 goes out of service September 2008. You should consider ordering and migrating to z/OS V1.9 as soon as possible to stay within IBM coexistence and migration guidelines.

Information on migration is available at:

- February 2008 z/OS Hot Topics article
www.ibm.com/systems/z/os/zos/bkserv/hot_topics.html
 - IBM Education Assistant modules
Publib.boulder.
ibm.com
[/infocenter/ieduasst/stgv1r0/index.jsp?topic=/com.ibm.iea.zos/zos/1.9/InstallationAndMigration.html](http://infocenter/ieduasst/stgv1r0/index.jsp?topic=/com.ibm.iea.zos/zos/1.9/InstallationAndMigration.html)
 - Upcoming Migration Conference calls (planned - March 2008)
 - Information at www.ibm.com/systems/z/os/zos/index.html
 - Migration Checker
www.ibm.com/systems/z/os/zos/downloads/#mchecker
-

Question:

Where can I find the FAQs on the z10 EC?

Answer:

The URL: www.ibm.com/systems/z/faq/ that carried this FAQ list.

Note: This URL will also have the z9 BC FAQs.

Question:

Where can I find the data sheets for the System z10 EC?

Answer:

The URLs are as follows: **z10 EC Data Sheet:**

ibm.com/common/ssi/fcgi-bin/ssialias?infotype=pm&subtype=sp&pdf=yes&appname=STG_ZS_USEN&htmlfid=ZSD03005USEN

z9 BC Data sheet

ibm.com/common/ssi/fcgi-bin/ssialias?infotype=pm&subtype=sp&pdf=yes&appname=STG_ZS_USEN&htmlfid=ZSD01858USEN

Question:

Where can I find the FAQs for the Operating Systems that support the z10 EC?

Answer:

The URL: www.ibm.com/systems/z/faq/ #os has FAQs for z/VM, z/VSE, z/OS and Linux on System z.

IBM Destination z

Question:

What is IBM Destination z?

Answer:

IBM Destination z names the community of resellers, ISVs, SI's, Academic Institutions and Client groups engaged with IBM System z. These organizations have joined to actively participate in: continuously improving the mainframe user experience, innovating on the mainframe as the system for the future, collaborating around best practices and ideas, and jointly offering customers a "one-stop shop." Announced WW on June 21, 2007, IBM Destination z has grown its membership WW and established foundations in key developing regions, including China and India.

Question:

What does IBM Destination z provide the IT community of: current customers, prospects, press and industry analysts, and academics?

Answer:

IBM Destination z provides many real and potential benefits to the IT community. It first serves to remind the marketplace that System z has a broad, growing, and rich ecosystem of institutions dedicated to mainframe innovation, value, and community. IBM Destination z also provides a tighter coupling between the large and growing number of worldwide academic institutions who are fortifying their respective curricula with mainframe courses, mastery standards, and certification. That, in turn, ensures a "tens of thousands" number of students trained to support mainframes and enable IT to provide unprecedented value. IBM Destination z also provides industry influencers with incontrovertible proof of the growing mainframe support and affirmation. Analysts cite IBM Destination z as a prime example that the mainframe is alive and well, and worthy of investment by institutions wanting a flexible, controllable, cost-effective, secure, and energy efficient infrastructure.

Question:

Who belongs to IBM Destination z?

Answer:

Community members of IBM Destination z vary from "small" ISVs to large organizations – and many points in between these. Those who belong actively contribute knowledge, case studies, references, and points of view on the mainframe. This widely spread; diverse input assures vitality, relevance, and modernity to how the mainframe environment matches to the primary, global IT challenges. Universities WW belong, as do significant User communities, some with thousands of individual members representing a few thousand organizations. Simply, IBM Destination z seems highly representative.

Question:

How is IBM Destination z tied into the Academic Initiative (AI)?

Answer:

IBM Destination z ties into the IBM System z Academic Initiative by providing access through the Web (and other means) to latest mainframe information. IBM Destination z members help align standards, content, and curriculum ideas with the colleges and universities, WW, who belong to the Academic Initiative. IBM Destination z also integrates with Academic Initiative's Student Opportunity System, a beneficial approach to connecting Mainframe-trained talented students with organizations hiring those types of employees. The Academic Initiative and IBM Destination z also link with various faculty members to tap their expertise in mainframe disciplines, and provide incentives for more research, better student performance, and more mainframe-related courses. It is increasingly the case that Academic Initiative institutions take the "next step" in mainframe support, and join IBM Destination z. Very important to note: over 47,000 students have been mainframe-trained worldwide in recent years. These significant numbers help assure adequate-plus talent to make owning a mainframe an exceptionally prudent decision. For more information and a list of IBM Destination z academic members: <http://www-03.ibm.com/systems/z/destinationz/academia.html>.

Question:

What is the process for joining IBM Destination z?

Answer:

Joining IBM Destination z is considered straightforward. There are enrollment forms that ask basic information about an organization interested in community membership. IBM Destination z wants to be sure those joining have mainframe expertise either developing software, implementing or integrating mainframe solutions, and selling IBM mainframe hardware, software, services and solutions. Business Partner applicants must be a PartnerWorld[®] member in good standing.

The applicant submits the form that is reviewed against defined criteria to ensure eligibility. The enrollment is confirmed when;

- A. Eligibility is confirmed.
 - B. The applicant organization provides a public reference or proof of promoting growth on the IBM System z platform
 - C. IBM Destination z will formally announce the new enrollee, and include their logo and reference on the IBM Destination z Web site with the intent of triggering collaboration and partnership with other IBM Destination z members.
-
-

Question:

What can my organization expect when/if we enroll?

Answer:

Your organization, whether a Business Partner, University or Client Group can expect a welcome, immediate inclusion in the community, access to portals and resources containing vital information, in depth facts and technology, connectivity to other sites, co-marketing opportunities with other IBM Destination z members and with IBM, and a strong voice in determining the attributes of the mainframe as its worldwide uptake increases. You can also expect to provide, as a Business Partner member, a yearly public reference. Additionally, we carry your company's logo on our IBM Destination z home page, and we allow you to "fly" IBM

IBM Destination z's emblem on your organization's Web site. Most interesting is the level of idea exchange, technical support, and open discussion through various Blogs and other Web-initiated exchanges on topics that further the utility, cost-effectiveness, and innovation in mainframe operating environments. You will also be included in early disclosure of new products, features and solutions as business practices allow.

Question:

Who is the contact point for joining and other activities?

Answer:

- From a Worldwide standpoint, the leader for IBM Destination z or Resellers interested in joining is Marc Smith, whose e-mail is: smarc@us.ibm.com
 - Specifically for System Integrators interested in joining, the contact point is Travis Smith, whose e-mail is: ttsmith@us.ibm.com
 - The other contact point for ISVs, especially those in Europe, wanting to join is Dinka Kalinic: dkalinic@us.ibm.com
 - For Academic Institutions interested in joining IBM Destination, the contact point is Michael Todd: todd@us.ibm.com
-

Question:

What is the progress in mainframe skills growth?

Answer:

- The IBM Academic Initiative has assisted over 400 schools world wide with IBM mainframe courses and resources. There is a world wide list of schools and contacts who are teaching System z. See: ibm.com/university/systemz - Participating Schools.
- Educators have reported that over 44,000 students have received mainframe education.
- Over 7000 students from 982 schools world wide have competed in IBM's Student Mainframe Contests.

If you have specific questions about z skills, send an e-mail to zSkills@us.ibm.com

z10 EC Hardware

Question:

Tell me about this new processor chip on the z10 EC:

Answer:

The IBM z10™ Enterprise Quad Core processor chip is the engine for the next generation of IBM System z mainframe systems. It features a new microprocessor core which is a giant leap in operating frequency, from 1.7 GHz on System z9 to 4.4 GHz on the System z10. The z10 EC chip is made up of four microprocessor cores, each with a private 3 MB cache and its own Hardware Decimal Floating Point unit. There are two coprocessor units on the z10 EC chip, each of which implements cryptographic and data compression functions, and each shared by two of the four cores. The chip also includes a memory controller (MC), I/O bus controller and a switch which connects all four cores to a shared interface with the SMP (symmetrical multi processing) hub chip and its shared cache.

The rich shared cache structure of the new chip is optimized for the enterprise data serving workloads which are the heart of the mainframe, and the comprehensive high frequency design helps the new chip yield more performance for CPU-intensive applications which are a growing part of System z new workloads.

Question:

What about the z10 EC makes it so good for consolidation of server farms?

Answer:

Scale is the first big factor making the z10 EC great for consolidation of distributed x86 servers. The new 4.4 GHz chip means more processing power per chip, along with the increase in number of available processing units per z10 EC. And the z10 EC 64-way offers 1.7x more total server capacity than a System z9 EC 54-way. The z10 EC also delivers up to 3x (1.5 terabytes) the memory of the z9 EC (1.5 TB per server, 1 TB per LPAR).

You can also save real dollars when consolidating due to less software costs and less floor space requirements. Using z/VM, the gold standard of virtualization software, you have automation capabilities, provisioning, management and the most sophisticated and complete hypervisor available. And Linux on System z has over 1,900 applications available from over 400 vendors.

z/OS on System z10 servers have the potential for consolidation opportunities too. z/OS V1.9 (and V1.10, when available) supports up to 64 engines (zIIPs, zAAPs, CPs) per LPAR. In addition, z/OS V1.8 and later releases support up to 4 TB of real memory; on the z10 EC, you can have up to 1 TB of real memory per LPAR.

Question:

It looks like you are using the same modular design you had on the z990 and the z9 EC – what is the significance of this design?

Answer:

The modular design creates the opportunity to address the ever increasing costs related to building systems with increasing capacities. The flexible 'book' design means that you can start with exactly the configuration you need, and expansion in the family can be easy and almost always concurrent (upgrading from any model to the enhanced capacity Model E64 will require a planned outage).

The modular book design also allows for a reduction in planned and unplanned outages by offering concurrent repair, replace and upgrade functions for processors, memory and I/O.

Question:

Is the z10 EC still a superscalar processor?

Answer:

Yes. A scalar processor is a processor that is based on a single issue architecture, which means that only a single instruction is executed at a time. A superscalar processor allows concurrent execution of instructions by adding additional resources onto the microprocessor to achieve more parallelism by creating multiple, longer pipelines, each working on their own set of instructions. Software compilers and interpreters can be designed to work with superscalar processor implementations. In the case of the z10 EC, the C++ compiler and Java™ Virtual Machine for z/OS exploit microprocessor superscalar implementation. The intent is to improve the performance advantage for workloads such as WebSphere® and Java applications.

Question:

What is the machine type of the z10 EC?

Answer:

The machine type is 2097.

Question:

What z10 EC models were announced?

Answer:

The following models were announced - note that the last two digits of the model indicate the maximum number of processor units (PUs) available for purchase on the model:

- A z10 EC E12 model can be a 1-way through 12-way - which means there are 12 processor units (PUs) contained on one book.
- A z10 EC E26 model can be a 1-way through 26-way (26 PUs) contained on two books.
- A z10 EC E40 model can be a 1-way through 40-way (40 PUs) contained on three books.
- A z10 EC E56 model can be a 1-way through 56-way (56 PUs) contained on four books.
- The enhanced capacity z10 EC E64 model can be a 1-way through 64-way (64 PUs) contained on four books.

The PUs can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z10 Application Assist Processors (zAAPs), System z10 Integrated Information Processors (zIIPs), additional System Assist Processors (SAPs), Internal Coupling Facilities (ICFs) or used as additional spares.

Only twelve subcapacity processors can be active on the server (and it doesn't matter which model you have). When more than twelve CPs have been purchased on servers that have more than one book, a selection can be made to activate only 12 or fewer subcapacity features. This means that the new subcapacity settings are available on any of the models as long as they are configured (not the same as purchased) with twelve or fewer general purpose processors.

Question:

Does the z10 EC offer more available subcapacity processors than the z9 EC?

Answer:

Yes up to 50% more subcapacity choices as compared to the z9 EC. The z10 EC can have up to 12 subcapacity Central Processors (CPs) per server while the z9 EC accommodates up to 8 subcapacity CPs. These can reside in multiple books.

Question:

If the total memory available on the z10 EC is 1.5 TB, and I get 16 GB of memory for my HSA, how much memory is available per book?

Answer:

The z10 EC server supports up to 1.5 TB of real memory per server, and 1.0 TB per LPAR, but the actual maximum physical memory sizes are related to the number of books in the system. Each book may contain a maximum of 384 GB of physical memory and in a 4 book system up to 1520 GB (1.5 TB) of physical memory can be purchased. This is equal to 4 books times 384 GB minus 16 GB reserved for the HSA.

Memory sizes in each book do not have to be similar; different books may contain different amounts of memory.

The minimum initial amount of memory that can be ordered is 16 GB for all models.

Question:

How does the z10 EC report processor information such as model number and CPU ID information?

Answer:

There are two instructions used to obtain processor information. Store System Information Instruction (STSI) and Store CPU ID instruction (STIDP). In response to customer requests, the STSI instruction is modified on the z10 EC, to give more useful and detailed information about the base configuration as well as information about possible concurrent upgrade options. The change should make it easier for customers to resolve billing situations where Independent Software Vendor (ISV) products are in use.

The STSI instruction has been updated to return the Model Capacity Identifier for the permanent configuration, as well as the Model Capacity Identifier for any temporary capacity such as On/Off CoD or CBU capacity.

The Store CPU ID (STIDP) provides information about the processor type, serial number, and logical partition identifier.

Question

What is different about the Model E64?

Answer:

The Model E64 is an enhanced capacity model, which contains a different configuration of MCMs. The z10 EC is fully populated with four books and 64 orderable PUs. You can customize the machine to be a 1 to 64-way. Like the other four book model, the E56, the E64 can be ordered with a minimum of 16 GB of memory up to a maximum of 1.5 TB.

Upgrade from any other model of the z10 EC to a Model E64 will require a planned outage.

Question:

How many spare processing units are on the z10 EC?

Answer:

There are two spare processing units on the z10 EC. These spares can be shared across the books. The z10 EC offers core (engine) level sparing.

Question:

Can any of these spare processing units be used for other purposes?

Answer:

No. These spare processing units are only used to provide failover in the remote event of a processor failure: general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z10 Application Assist Processors (zAAPs), System z10 Integrated Information Processors (zIIPs) or Internal Coupling Facility (ICFs).

Question:

How many System Assist Processors (SAPs) are on the z10 EC?

Answer:

The answer depends on the model. The number of SAPs provided to the customer as standard are as follows:

- The z10 EC Model E12 has three.
- The z10 EC Model E26 has six.
- The z10 EC Model E40 has nine.
- The z10 EC Model E56 has ten.
- The z10 EC Model E64 has eleven.

Additionally, the customer can acquire more SAPs from among the pool of available processing units within the model.

Question:

What is new about the Hardware System Area (HSA) on the z10 EC?

Answer:

On the z10 EC, a 16 GB area is set aside for the HSA and will be managed independently of purchased memory. The fixed size of the HSA eliminates planning for future expansion of HSA since HCD/IOCP will always reserve:

- Four Channel Subsystems (CSSs)
 - Fifteen logical partitions in each CSS for a total of 60 Logical Partitions
 - Subchannel set 0 with 63.75 K devices in each CSS
 - Subchannel set 1 with 64 K devices in each CSS.
-
-

Question:

What is meant by 'just-in-time' deployment of capacity on the z10 EC?

Answer:

Customers have asked us to ease the administrative tasks required to activate temporary capacity. There is a need to be able to access and implement our temporary offerings (CBU and On/Off CoD) much more quickly than we've been able to in the past, in some cases to automate it so the server itself can determine that capacity may be required or no longer required, and handle the changes automatically. Significant rework has been done in the provisioning architecture on the z10 EC to make your ability to handle capacity changes much more efficiently and 'just-in-time' to meet your needs.

Question:

What is the new Hardware Decimal Floating Point unit on the z10 EC processor chip and what types of applications will take advantage of it?

Answer:

Speed and precision in numerical computing are important for all our customers. Each core on the PU of the z10 EC has its own hardware decimal floating point unit, designed to improve performance of decimal floating point over that provided by the 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.

Question:

What is HiperDispatch?

Answer:

Please refer to the Performance Section of this document.

Question:

What is the large page support announced for the z10 EC?

Answer:

Page frames are allocated with a 4K size. The z10 EC additionally supports a larger page frame size of 1 MB (1 MB). It is expected that long running, memory access intensive applications will benefit from large page frames. It should be noted that large pages are treated as fixed pages and are never paged out.

Question:

What is the Universal Lift Tool / Ladders feature (#3759)?

Answer:

The Universal Lift Tool / Ladders feature includes a custom lift / lower mechanism that is specifically designed for use with the z10 EC frames. It is designed to provide customers with enhanced system availability benefits by improving the service and upgrade times for larger, heavier devices.

It is recommended that one of these features be obtained for each customer account / data center.

Question:

Can I buy a z10 EC that has only IFL or ICF processors without including a general purpose processor (CP)?

Answer:

Yes. Similar to the z9 EC, you can order only IFLs or ICFs in a z10 EC, using a software model of 700 with 1 to 64 IFLs or a maximum of 16 ICFs.

Question:

Will the z10 EC support Token-Ring on the Hardware Management Console (HMC)?

Answer:

No.

Question:

What is the Frame Bolt-Down kit used for?

Answer:

A Bolt-Down Kit for raised floor and non-Raised floor environments provides frame stabilization floor tie-down hardware to help secure a System z frame to a concrete floor beneath a 9- to 13-inch or 12- to a 22-inch raised floor.

These are designed to help secure the frames and their contents from damage when exposed to vibrations and shocks such as those in a seismic event. The frame tie downs are intended for securing a System z frame weighing less than 3600 lbs per frame. These tie downs are designed to help secure the frame on a raised floor.

These kits are provided on AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Question:

Are there different kits available?

Answer:

Yes, the following kits are available. For z10 EC, you need a quantity of two Bolt-Down kits.

(#7994) - Bolt-Down Kit, High-Raised Floor 2097 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 (298mm to 405mm) raised floor.

(#7993) - Bolt-Down Kit, Low-Raised Floor 2097 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 (235mm to 298mm) raised floor.

Question:

Will the z10 EC offer the Power Monitor and Power Estimation capabilities available on the System z9 servers?

Answer:

Yes. The tool is designed to help monitor the power consumption and temperature of the z10 EC or System z9 server. The System Activity Display (SAD) on the Hardware Management Console (HMC) can display the current total power consumption in watts and Btu/hour and can also display the input temperature.

When doing data center planning, or server upgrade planning, there is a tool on IBM Resource Link™ that can be used to estimate the anticipated power consumption of a particular machine model and its associated configuration. A user will input the machine model, memory, and I/O configuration and the tool will output an estimate of the power requirements needed for the system.

z10 EC Upgradeability

Question:

What are the upgrade paths from a z9 EC or a z990 to the z10 EC?

Answer:

All models of the z9 EC and the z990 can upgrade directly to a z10 EC.

Question:

Is there an upgrade path from a z9 BC to the z10 EC?

Answer:

No.

Question:

Can I upgrade any z10 EC to a Model E64?

Answer:

Yes, but it is important to note that upgrading other z10 EC models to the E64 will be disruptive due to the difference in MCMs on the E64.

Question:

Will I be able to upgrade from a full capacity z10 EC to a subcapacity z10 EC?

Answer:

Yes. Each of first twelve general purpose processors on the z10 EC can be divided into one full capacity and three sub-units. This creates a 12 by 4 matrix of settings. As long as upgrades are positive capacity growth, you can move around anywhere within the matrix when adding capacity. When your number of general purpose processors exceeds twelve, then all of the general purpose processors must be full capacity.

Question:

Can I have subcapacity processors if I have more than one book?

Answer:

Yes. But remember that only twelve subcapacity processors can be active on the server (and it doesn't matter which model you have). When more than twelve CPs have been purchased on servers that have more than one book, a selection can be made to activate only 12 or fewer subcapacity features. This means that the subcapacity settings are available on any of the models as long as they are configured (not the same as purchased) with twelve or fewer general purpose processors.

Availability Enhancements

Question:

What new availability enhancements are on the z10 EC?

Answer:

Our availability focus in System z is to prevent failures from occurring in the first place. As our hardware products have matured, we continue to provide enhancements to avoid scheduled and unscheduled outages. For the z10 EC, updates that focused on these types of outages include concurrent firmware fixes, concurrent driver upgrades, concurrent parts replacement and concurrent hardware upgrades.

The z10 EC put more focused attention on avoidance of planned outages. These include the new 'just-in-time' deployment of capacity, the design to eliminate unnecessary CBU passwords, enhanced driver maintenance and the design to eliminate a logical partition deactivation/activate/IPL. Having the new 16 GB defined HSA allows for changes such as the dynamic add of features and functions to the server without the need for Power-On-Reset.

Question:

What does it mean when IBM says the z10 EC will help 'eliminate' preplanned outages for hardware changes?

Answer:

Preplanning requirements have been reduced for the z10 EC primarily with the introduction of the fixed size reserved 16 GB HSA. By having enough memory available that is not carved from user purchased memory, a full I/O configuration can be created that takes into account configuration possibilities that might happen in the future. It is important to have these changes already defined, because with the z10 EC, it is now possible to dynamically:

- Add a logical partition
- Add a Logical Channel Subsystem (LCSS)
- Add a Subchannel set
- Add a logical CP to a logical partition
- Remove a logical CP from a logical partition
- Add a cryptographic coprocessor
- Remove a cryptographic coprocessor
- Enable I/O connections
- Swap processor types

Question:

Will the z10 EC offer the flexible memory option?

Answer:

Yes. Flexible memory was introduced on the z9 EC as part of design changes and offerings to support enhanced book availability. Flexible memory provides the additional resources to maintain a constant level of memory when replacing a book. On z10 EC, the additional resources required for the flexible memory configurations are provided through the purchase of preplanned memory features (#1996) along with the purchase of your memory entitlement. Flexible memory configurations are available on the E26, E40, E56 and E64 only and range from 32 GB to 1136 GB, depending on the model.

Question:

What is the cryptography availability enhancement designed into the z10 EC?

Answer:

With the cryptography availability enhancement to dynamically add cryptographic features to logical partitions you can make changes to image profiles to support Crypto Express2 features without outage to the logical partition. You can also dynamically delete or move Crypto Express2 features. And, all without preplanning.

Question:

Remind me about Enhanced Book Availability and Redundant I/O Interconnect?

Answer:

Introduced on the System z9, Enhanced Book Availability allows a single book in a multi book server to be concurrently removed and reinstalled for an upgrade or a repair. Removing a book would mean that connectivity to the I/O connected to that book is lost - but to prevent connectivity loss, the Redundant I/O Interconnect feature allows you to maintain full connection to critical devices when a book is removed.

Question:

What is AutoIPL in the new z/OS v1.10?

Answer:

AutoIPL support will provide the capability to request that the system automatically IPL standalone dump, z/OS, or both, when a disabled wait state is requested by a system component. In a Parallel Sysplex® environment, the Sysplex Failure Manager (SFM) policy can result in actions that load disabled wait states on systems to be partitioned out of the sysplex, which can also trigger AutoIPL processing. AutoIPL capability is intended to help you achieve faster failure data capture and recovery after system failures.

IFL, zAAP and zIIP specialty engines

Question:

Why the funny naming for zIIPs and zAAPs?

Answer:

The official name for the zAAP on the z10 EC is the System z10 Application Assist Processor. A zAAP is supported on the z10 EC, z9 EC, z9 BC, z990 and z890.

The official name for the zIIP on the z10 EC is the System z10 Integrated Information Processor. The zIIP on the System z9 servers is System z9 Integrated Information Processor. The zIIP is only supported on the z10 EC and System z9 servers.

Question:

What new uses of the Specialty Engines are introduced with the February 26, 2008 announcement?

Answer:

There are several new exploiters for System z specialty engines.

z/OS Global Mirror (formerly Extended Remote Copy XRC) is enabled for the zIIP. Specifically, z/OS DFSMS™ allows most System Data Mover (SDM) processing associated with z/OS Global Mirror to be eligible to run on the zIIP

Also, z/OS V1.10 is planned to enable additional XML processing to be made eligible for the zIIP and zAAP specialty processors. IBM middleware (such as DB2® 9 and Enterprise COBOL V4.1) and other products can benefit from this new functionality in addition to taking advantage of the z/OS XML System Services capabilities. These enhancements are expected to help improve the price performance of XML processing on z/OS and ultimately may help facilitate the decision to develop more XML-based applications on z/OS.

IBM plans to enhance the XML Toolkit for z/OS (5655-J51) so eligible workloads can use z/OS XML System Services. This allows eligible XML Toolkit processing for non-validating parse requests to exploit the zAAP. This function is planned to be available on the XML Toolkit for z/OS V1.9 with SPE. XML Toolkit support for processing validating parse requests using z/OS XML System Services, with the appropriate offload of eligible work to zAAP is planned at a future date. Delivery of this function satisfies the Statements of Direction in Hardware Announcement 107-190, dated April 18, 2007 and Software Announcement 207-175 dated August 7, 2007.

Question:

What is z/OS Global Mirror and what is the value of zIIP enabling it?

Answer:

z/OS Global Mirror (previously known as Extended Remote Copy or XRC) provides a long-distance remote copy solution across two sites for open systems and zSeries data using asynchronous technology. It uses System Data Mover (SDM) to perform the data movement and management of the consistency of the data. The z/OS Global Mirror maintains a copy of the data asynchronously at a remote location, and can be implemented over unlimited distances. It is a combined hardware and software solution that offers data integrity and data availability and can be used as part of business continuance solutions, for workload movement, and for data

migration. Although the main z/OS Global Mirror implementation consists of host resident software, special z/OS Global Mirror support is required in the DS8000™ that attaches the z/OS Global Mirror primary volumes.

With zIIP assisted z/OS Global Mirror, the zIIP essentially becomes a z/OS data mirroring engine that can provide better price performance and improved utilization of resources at the recovery site. Most DFSMS system data mover (SDM) processing is eligible to be redirected to a zIIP processor, which can help lower server utilization at the recovery site, or create server "white space" to be used for other projects.

Question

What is required for zIIP assisted z/OS Global Mirror?

Answer:

The capability is available with:

- z/OS V1.10 (when available), or z/OS V1.9 and V1.8 with PTF
- IBM System Storage™ DS8000, or any storage controller supporting z/OS Global Mirror
- Any z10 EC, z9 EC, or z9 BC with a zIIP.

An updated white paper is planned for z/OS Global Mirror functional improvements, including more detail on zIIP assisted zGM.

Question:

Can I order subcapacity specialty engines?

Answer:

No. All specialty engines on the z10 EC run at full capacity.

Question:

What is the value of the Integrated Facility for Linux (IFL)?

Answer:

The attractively priced IFL processor enables you to purchase additional processing capacity exclusively for Linux workloads, without affecting the MSU rating of the IBM System z model designation. This means that an IFL will not increase charges for System z software running on general purpose processors in the server.

Question:

How does Linux benefit from running on the z10 EC?

Answer:

The increased power of the z10 EC processor combined with other architectural advances like the increased number of Logical Partitions (LPARs) and the improved internal I/O throughput gives the z10 EC the ability to run increased workloads.

Since the price of the Integrated Facility for Linux (IFL) for this solution package is similar, whether you buy it for the z10 EC, z9 EC, z990, or z900, and the price for Linux software is typically processor-based. This can mean customers may be able to either support larger numbers of users for a similar software price, or provide greater performance for the same number of users at a similar price.

So customers get a price/performance improvement for Linux workload with each successive generation of mainframe technology. Actual improvements are dependent on the environment.

Question:

Is the functionality of the IFL the same as a general purpose processor?

Answer:

Yes, the IFL has the same functionality as a general purpose processor on a System z. It supports On/Off Capacity on Demand (On/Off CoD), Capacity Upgrade on Demand for the nondisruptive addition of one or more IFLs, and Capacity BackUp (CBU) for emergency situations. HiperSockets can be used for communication between Linux systems, or Linux and other operating systems on the same System z platform.

Question:

Can I buy a System z10 EC server that is made up entirely of IFL processors and no general purpose processors?

Answer:

Yes. A dedicated System z10 EC Linux server is available, meaning that all processors are IFLs on this System z10 EC server. Please see your IBM or Business Partner representative for more information on it.

Question:

What is a System z10 Application Assist Processor (zAAP)? What is the benefit to z/OS customers?

Answer:

zAAPs help enable customers to strategically integrate new application technologies such as their Java technology-based Web applications today and XML-based data interchange services with their core business database environment by providing a more cost-effective, specialized z/OS application Java execution environment.

Question:

How do Java-based applications benefit from the zAAP?

Answer:

zAAPs can enable customers to run Java Web applications next to mission-critical data for integrated, security-rich and efficient application and database serving. Execution of these new applications within the same z/OS LPAR as their associated database subsystems can help simplify server infrastructures and improve operational efficiencies by reducing the number of TCP/IP programming stacks, firewalls, physical interconnections and their associated processing latencies that might otherwise be required when the application servers and their database servers are deployed on separate physical server platforms.

By executing the Java cycles on a zAAP, you can reduce the demands and capacity requirements on general purpose CPs which may then be available for reallocation to other IBM mainframe workloads.

Furthermore, zAAPs allow customers to purchase additional processing power exclusively for Java workload execution without affecting the total MSU rating or machine model designation, as zAAPs do not carry a rated capacity. Consequently, IBM does not impose software charges on zAAP capacity. Additional IBM software charges will apply only when additional general purpose CP capacity is used.

Moreover, zAAPs may have the effect of reducing charges for subcapacity eligible IBM software products by lowering the rolling 4-hour average MSUs for LPARs with assigned zAAPs.

Best of all, the IBM JVM processing cycles can be executed on the configured zAAPs with no anticipated modifications to the Java application(s).

Question:

What are the hardware and software prerequisites of the IBM zIIP?

Answer:

The z10 EC will ship with the appropriate driver code to support a zIIP should you order one. Operating system support, which enables the redirecting of eligible workloads to the zIIP, is provided by z/OS V1.7 via a Web download. zIIP exploitation may be exploited by ISV products as well, see your ISV for details. Additional zIIP exploitation is planned by z/OS Communications Server for portions of IPsec network encryption processing (available with z/OS V1.8 and PTFs, and native in z/OS V1.9 when available).

Question:

How does data serving (specifically with regard to DB2 for z/OS V8 and DB2 9) work with zIIPs?

Answer:

Using a zIIP can help free up capacity on the general purpose processor. When the zIIP is available, DB2 Universal Database™ (DB2 UDB) for z/OS V8 (and later) will exploit the zIIP by sending eligible work to z/OS 1.7 (or later) that can be directed to the zIIP.

The types of eligible DB2 UDB for z/OS V8 workloads executing in SRB mode, which all or a portion of can be sent to the zIIP, are:

1. Network Connected Application Types – An application (running on UNIX®, Linux, Intel®, Linux on System z10 EC (or z/OS) may access a DB2 UDB for z/OS V8 database that is hosted on a z10 EC. Eligible work that can be directed to the zIIP are portions of those requests made from the application server, to the host, through SQL calls over a DRDA® via

TCP/IP connection. Examples of workloads that may be running on the server connected over DRDA via TCP/IP to the z10 EC may include BI, ERP or CRM application serving.

2. Data Warehousing Application Types – Requests that utilize DB2 UDB for z/OS V8 for long running queries, including complex star schema parallel queries, may have portions of these SQL requests directed to the zIIP when DB2 gives z/OS the necessary information. These queries are typical in data warehousing implementations. The addition of select long running parallel queries may provide more opportunities for DB2 customers to optimize their environment for data warehousing while leveraging the unique qualities of service provided by z10 EC and DB2.
 3. Utility Function Types – A portion of DB2 utility functions that are used to maintain index maintenance structures (LOAD, REORG and REBUILD INDEX) can be redirected to the zIIP.
-

Question:

Why was DB2 DRDA processing chosen for zIIP eligible workload?

Answer:

Database workloads such as CICS®, IMS™, stored procedures and batch have become increasingly efficient and cost effective for the mainframe. Today, customers are looking to further leverage their data on the mainframe and are turning to the mainframe more and more frequently for new application workloads. These application workloads, such as Enterprise Resource Planning (ERP), Customer Relationship Management (CRM) and Business Intelligence (BI) often use DB2 as a database server.

Added system costs for network access to DB2 results in hardware and software charges that are substantially higher for remote applications, when compared to the same workload running as a local batch application. The zIIP is aimed at helping customers more cost effectively run these database serving applications on the mainframe. Our objective with the zIIP is to help bring the costs of network access to DB2 more closely in line with the costs of running similar workloads under CICS, IMS or batch on the mainframe.

Question:

If my application is a business intelligence application that is also network attached (using DRDA over TCP/IP), does this mean that my workload can leverage the eligible redirect for star schema parallel processing and DRDA over TCP/IP?

Answer:

Yes, you can benefit from both eligible redirects. If a star schema parallel query comes in remotely via DRDA over TCP/IP, a portion of the workload coming across the DRDA over TCP/IP connection can be redirected to the IBM zIIP, as well as the portion of the star schema parallel query processing that is redirected.

Question:

What is z/OS XML System Services?

Answer:

z/OS XML System Services (z/OS XML) is a system level XML parser that is integrated with the base z/OS operating system (available with z/OS V1.8 (and V1.7 with PTF)). It is intended for use by system components, middleware and applications that need a simple, efficient, non-validating XML parsing solution.

XML plays a vital role in the development and adoption of SOA (Service-Oriented Architecture) solutions because in many cases, the messages flowing between the SOA services are XML. As a result, having XML processing on z/OS and making portions of XML processing eligible for the System z specialty engines is a logical direction for the adoption of XML-based technologies on z/OS.

Question:

What z/OS XML System Services processing is eligible for the zAAP and zIIP processors?

Answer:

z/OS XML System Services (z/OS XML) processing eligibility for zAAPs and zIIPs will be rolled out over time. The following table summarizes the XML workload, the exploiter, its availability and other requirements.

Workload	Examples	Available	Redirect	Requirements
z/OS XML System Services , non-validating parsing, – executing in TCB mode z/OS 1.9, or z/OS 1.8 (with OA20308), or z/OS 1.7 (w/ OA16303, OA20308)	- local applications inserting/ saving XML data, and XML table loads on DB2 9	Yes (with z/OS V1.9 GA)	100% of z/OS XML System Services parsing - eligible for zAAP	DB2 9 New Function Mode C API for z/OS XML System Services available with z/OS 1.9, and z/OS 1.7 & 1.8 with APAR OA18713
	- any SW using z/OS XML System Services parsing in TCB - select XML Toolkit for z/OS V1.9 parsing workloads	Soon (with XML Toolkit SPE)	same	XML Toolkit for z/OS (5655-J51) V1.9 with SPE (#tbd)
	- Enterprise COBOL V4.1, using XMLPARSE option	Yes (with COBOL 4.1 GA)	same	Enterprise COBOL V4.1 (5655-J51) z/OS V1.7- V1.9 with OA22777 (SW Announce 207-339).
z/OS XML System Services, non validating parsing, executing in enclave SRB mode	- DB2 9 inserting/ saving XML data using DRDA via TCP/IP	Yes (with GA of z/OS V1.8)	Same % as the zIIP-eligible work (DRDA)	z/OS 1.8, or z/OS 1.7 w/ OA16303 DB2 9 New Function Mode
	- any SW (including DB2 9) using z/OS XML System Services in enclave SRB mode	Soon (GA of z/OS V1.10)	100% of z/OS XML System Services parsing eligible for zIIP	z/OS 1.9 and 1.8 (both with APAR #OA22035) DB2 9 NFM
z/OS XML System Services with validating parsing, both enclave SRB and TCB modes.	- any SW using z/OS XML System Services validating parsing	Soon (GA z/OS V1.10)	100% of z/OS XML System Services validating parsing eligible for zAAP (TCB) or zIIP (enclave SRB)	z/OS V1.10 *
	- select XML Toolkit for z/OS V1.9 workloads	SOD*	TBD	TBD
Java-based XML parsing	- applications using Java-based XML parser in IBM SDK - any SW performing XML parsing/ processing in Java	Yes (with availability of zAAP)	100% of Java-based XML parsing eligible for zAAP	Any z/OS, System z processor with zAAP support.

Question:

Once a zIIP is installed, how is workload on it measured?

Answer:

Once a zIIP is installed [along with z/OS 1.7 (with PTF) or later, and DB2 UDB for z/OS, V8 (with PTF)], monitoring the zIIP activity will be similar to monitoring zAAP activity. Specifically:

- Set up WLM policy with Service Class(es) for SUBSYSTEM TYPE=DDF
 - RMF™ Monitor 1 Type 70 Record will monitor overall zIIP activity:
 - Logical processor busy as seen by z/OS is reported
 - Physical processor busy as seen by LPAR is reported
 - RMF Monitor 1 Type 72 Record will show more detail:
 - The amount of time spent executing on zIIP processors is reported
 - Usage and Delay sample counts for zIIP eligible work is reported. In addition, DB2 accounting trace records can provide information on the zIIP. IBM Tivoli® OMEGAMON® XE for DB2 Performance Expert on z/OS, DB2 Performance Expert or IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS can be used to monitor the zIIP information.
 - In addition to RMF - SMF 30 and 70 records will be generated.
-
-

Question:

If Linux on System z10 is running in an LPAR on the z10 EC, accessing DB2 UDB for z/OS V8, can portions of the eligible work be redirected to the zIIP?

Answer:

Yes. Whether the customer is using HiperSockets or some other network to access the DB2 database, the Linux application can take advantage of redirection of processing to the zIIP.

Question:

Will the z/OS support for WLM goal mode management for zAAP processors be extended to the zIIP specialty engines?

Answer:

Yes. WLM will manage zIIPs the same way it already manages CPs and zAAPs.

Question:

What does z/VM 5.3 provide for zAAP and zIIP support?

Answer:

z/VM 5.3 provides new guest support for IBM System z10 Application Assist Processors (zAAPs) and System z10 Integrated Information Processors (zIIPs) including:

- Simulation support - z/VM can create virtual specialty processors for virtual machines by dispatching the virtual processors on real Central Processors (CPs) allowing users to assess the operational and CPU utilization implications of configuring a z/OS system with zIIP or zAAP processors without requiring the real specialty processor hardware. zIIPs can be simulated only on z10 EC or System z9 (z9 EC and z9 BC) servers. zAAPs can be simulated only on z10 EC, z9 EC, z9 BC, z990, and z890 servers.

- Virtualization support - z/VM can create virtual specialty processors for virtual machines by dispatching the virtual processors on corresponding real specialty processors of the same type and may help improve your total cost of ownership by allowing available zAAP and zIIP capacity not being used by z/OS LPARs to be allocated to a z/VM LPAR hosting z/OS guests running Java and DB2 workloads.
-

Question:

Where can I get more information on the zIIP?

Answer:

For more information on the zIIP, please look at the Web site: www.ibm.com/systems/z/ziip.

z10 EC Pricing

Question:

Is hardware pricing for z10 EC similar to pricing on z9 EC?

Answer:

Yes, hardware pricing for z10 EC is very similar to z9 EC in approach. As in the past, each customer order will carry a unique price driven by the customer's specific configuration needs and desires.

Question:

How will diagonal upgrades from z9 EC (or z990) to z10 EC be priced?

Answer:

The platform's value proposition of net pricing diagonal upgrades (upgrades from z9 EC or z990) will not change. As in the past, microprocessor and memory feature conversions will be the key net priced items.

Question:

What is the memory pricing on the z10 EC?

Answer:

System z10 EC memory pricing will be priced as low as \$6K per GB in the United States. Pricing may vary in other countries.

Question:

How will maintenance be priced on z10 EC?

Answer:

Monthly 24x7 ESA list prices will exist and be applied against the z10 EC feature codes similar to z9 EC or z990.

Question:

What is the cost of the IFL, the zAAP and the zIIP?

Answer:

The price of the IFL, the zAAP and the zIIP will be US \$125,000 in the United States. Pricing may vary in other countries.

Question:

Will there be a 10% software tech dividend on the z10 EC similar to what was offered on the z9 EC and the z990 before it? Will it be cumulative if I upgrade, for example if I'm two generations back do I get both pricing initiatives?

Answer:

Yes, the approach for pricing software on the z10 EC is similar to pricing on the z9 EC. The announced MSUs for software pricing purposes on the z10 EC are on average 10% less than those for z9 EC. With this, in many cases, software savings can be realized on the z10 EC versus the z9 EC. And yes, the announced MSUs for software pricing purposes on the z10 EC are on average 19% less than those for z990. (The z9 EC offered 10% savings in software pricing MSUs when coming from a z9 EC and the z10 EC has an additional 10% savings for a total average of 19% over the z990).

Question:

How does a customer pay for software on the z10 EC server?

Answer:

Customers may pay for Monthly License Charge software under either Workload License Charges (WLC) or Parallel Sysplex License Charges (PSLC). Full-Capacity WLC is available on z10 EC with the z/OS operating system. Subcapacity WLC is available on a z10 EC that has the z/OS operating system and has discontinued both use and licensing of the OS/390[®] operating system, subject to applicable terms and conditions. Once qualified, customers may determine when they wish to adopt WLC. PSLC is available on z10 EC in all environments, including a standalone (non-coupled) z10 EC, subject to applicable terms and conditions.

IBM System z New Application License Charges (zNALC) for z/OS, announced January 9, 2007 offers a reduced price for z/OS operating system on LPARs where customers are running eligible 'new qualified workloads' such as Java language business applications running under WebSphere Application Server, Lotus[®] Domino[®], SAP, PeopleSoft, and Siebel. The zNALC offering continues the IBM commitment to subcapacity pricing, allowing customers with qualified new workload to obtain a reduced price for z/OS based on the size of the LPAR(s) executing new workload (assuming all applicable terms and conditions are met).

Customers using z/VSE V3.1 products running on a z10 EC server will qualify for the same announced pricing terms available on z9 EC. Extended License Charge (ELC) applies for servers over 80 MSUs.

Midrange Workload License Charge (MWLC) for z/VSE announced January 9, 2007. MWLC is a monthly license charge price metric on the z10 EC that applies to z/VSE V4 and 12 key VSE middleware programs such as CICS TS for VSE, DB2 Server for VSE, and ACF/VTAM[®] for VSE. MWLC is available on z10 EC servers with z/VSE 4.

z/VM 5.1 and later, and its support charges are priced based on the number of processors according to International Program License Agreement (IPLA) and Subscription and Support (S&S) terms and conditions.

Question:

What is a "software model"?

Answer:

Like the previous System z mainframe servers, the machine type and model of the z10 EC does not indicate the machine's capacity. Instead, you must know the number of installed general purpose processors (CPs) to determine the machine's Model-Capacity Identifier for software licensing/charging purposes. The software model follows the nomenclature nxx, where the n indicates the subcapacity size (4, 5, or 6, and the original '7' will indicate a full capacity processor), and where xx indicates the number of installed general purpose processors. You may research MSU ratings per software model on the Web at ibm.com/zseries/library/swpriceinfo.

Question:

Are there purchase restrictions limiting the number of zIIPs I can order?

Answer:

A customer may order zIIPs up to the number of permanently purchased general purpose processors (CPs) on a given z10 EC model. This requirement is at a server level, so a customer could have an LPAR with more zIIPs than general purpose processors (CPs), as long as there are enough general purpose processors (CPs) in the entire server to meet the one-for-one requirement.

Question:

If I currently have one general purpose processor (CP) and one zAAP, can I order one zIIP without needing to order another general purpose processor (CP) to meet the restriction listed above?

Answer:

Yes. One zAAP and one zIIP can share one general purpose processor (CP) to satisfy the ordering restrictions.

Question:

What impact could the IBM zIIP have on IBM software charges?

Answer:

IBM does not impose software charges on zIIP capacity. Additional IBM software charges will apply when general purpose processor (CP) capacity is used. The amount of general purpose processor savings will vary based on the amount of workload executed by the zIIP, among other factors.

Question:

What impact could the IBM zIIP have on non IBM software charges?

Answer:

Customers are encouraged to contact their specific ISVs/USVs directly to determine if or how their charges will be affected.

Question:

What is DB2 for z/OS Value Unit Edition?

Answer:

On February 26th, IBM will be announcing a new offering, DB2 for z/OS Value Unit Edition (VUE). This offering provides the same robust DB2 for z/OS data server at a One-Time Charge price. It can only be used for net new workloads that run on a zNALC LPAR or system. It is not intended to replace existing MLC workload, but instead to offer customers alternatives for net new workload on System z and DB2. DB2 for z/OS VUE gives customers the ability to grow their workload without increasing their existing MLC stack. This offering continues to strengthen the role of System z and DB2 as the cornerstone of new applications like SOA, dynamic data warehousing, operational BI, ERP and WebSphere. For further information contact your IBM Software Sales Rep.

DB2 9 for z/OS VUE is planned to GA on February 29,2008.

DB2 8 for z/OS VUE is planned to GA on March 26, 2008.

Question:

Where can I get more information on IBM software charges?

Answer:

www.ibm.com/servers/eserver/zseries/swprice/

Performance

Question:

What is the basis of the Large System Performance Ratios (LSPR)?

Answer:

The LSPR ratios reflect the range of performance between prior zSeries servers and the z10 EC as measured using a wide variety of application benchmarks. The latest release of LSPR contains a number of updates to reflect the continuing evolution of zSeries customer applications and configurations. First, the workload suite has changed: the Java based batch workload (CB-J) is replaced with a more customer-like Java based batch workload (ODE-B). Second, all workloads were moved to z/OS 1.8 and more recent levels of subsystem and compiler software. Third, the new HiperDispatch feature (described later below) is turned on for all z10 EC z/OS LSPR measurements. LSPR continues to include both single-image z/OS and multi-image z/OS in separate tables. For details on the workloads and performance ratios, please reference ibm.com/servers/eserver/zSeries/lSpr.

Question:

Why are there two tables in LSPR?

Answer:

The LSPR was enhanced to include performance ratios reflecting both “single-image” z/OS and “multi-image” z/OS environments when System z9 was introduced. Typically, zSeries processors are configured with multiple images of z/OS. Thus, the LSPR continues to include a table of performance ratios based on average multi-image z/OS configurations for each processor model as determined from the profiling data. Since the multi-image z/OS table is much more representative of the vast majority of customer configurations, it is used as the basis for setting MIPS and MSUs for the z10 EC.

Question:

What multi-image configurations are used to produce the LSPR multi-image table?

Answer:

A wide variety of multi-image configurations exist. The main variables in a configuration typically are: 1) number of images, 2) size of each image (number of logical engines), 3) relative weight of each image, 4) overall ratio of logical engines to physical engines, 5) the number of books and 6) the number of ICFs/IFLs. The configurations used for the LSPR multi-image table are based on the average values for these variables as observed across a processor family. It was found that the average number of images ranged from 5 at low-end models to 8 at the high end. Most systems were configured with two major images (those defined with >10% relative weight). On low- to midrange models, at least one of the major images tended to be configured with a number of logical engines close to the number of physical engines. On high-end boxes, the major images were generally configured with a number of logical engines well below the count of physical engines reflecting the more common use of these processors for consolidation. The overall ratio of logical to physical engines (often referred to as “the level of over-commitment” in a virtualized environment) averaged as high as 5:1 on the smallest models, hovered around 2:1 across the majority of models, and dropped to 1.3:1 on the largest models. The majority of models were configured with one book more than necessary to hold the enabled processing engines, and an average of 2 ICFs/IFLs were installed.

Question:

Which LSPR table should I use for capacity sizing?

Answer:

For high-level sizing, most users will find the multi-image table to reflect configurations closest to their own. This is simply due to the fact that most systems are run with multiple z/OS images. However, the most accurate sizing requires the zPCR tool, which can be customized to match a specific multi-image configuration rather than the average configurations reflected in the multi-image LSPR table.

Question:

If I compare the two tables, why are the capacity ratios for some models higher in the single-image table while other models have higher ratios in the multi-image table?

Answer:

Just as capacity ratios are sensitive to workload characteristics (note the varying capacity ratios within a table associated with different workloads), capacity ratios will also be sensitive to the configuration of z/OS images on a processor. If one compares a processor configured only with a single, large z/OS image to the same processor configured with multiple z/OS images, there are both pluses and minuses that come into play. There is a cost incurred to manage multiple z/OS images and their associated logical processors. There is also a cost incurred as the size of a z/OS image increases. Thus, if one compares a configuration of a single large z/OS image to a configuration of multiple but smaller z/OS images, the net result can vary as the magnitude of the pluses and minuses will vary. The sensitivity of the multi-image configurations to the number of images, size of each image, relative weights and overall logical: physical ratio will cause a fair amount of variability in the capacity ratios of these configurations. The multi-image table provides a representative view of these ratios based on average configurations. However, "your mileage will vary" applies here as configurations deviate from the average. The zPCR tool can provide capacity ratios customized to specific configurations.

Question:

How much variability in performance should I expect when moving a workload to a z10 EC?

Answer:

As with the introduction of any new server, workloads with differing characteristics will see some variation in performance when moved to the z10 EC. The performance ratings for a server are determined by averaging the performance of a variety of workloads that represent what we understand to be the major components of our customers' production environments. While the ratings provide good "middle-of-the-road" values, they do represent an average, and by definition some workloads fall higher than the average and some workloads fall below. The z10 EC has been specifically designed to focus on new and emerging workloads where the speed of the processor is a dominant factor in performance. The result is a quantum jump in clock speed - the z10 EC runs at 4.4 GHz compared to the z9 EC which ran at 1.7 GHz. The storage hierarchy design of the z10 EC is also improved over z9 EC, however, the improvement is somewhat limited by the laws of physics so the latencies have increased relative to the clock speed. Thus, workloads that are CPU-intensive will tend to run above average while workloads that are storage-intensive will tend to run below average, and the spread around the average will likely be larger than seen in recent processors. Additionally, newer applications, such as those with compiler optimizations for the z10 EC may see even higher benefits, particularly those that may be enhanced over time to exploit some of the new instructions provided with the z10 EC.

The LSPR measurements can provide an indication of the potential variability when moving z/OS workloads to a z10 EC. For example, using the single-image z/OS measurements on a 2097-716 versus a 2094-716, we saw performance ratios of: a) 1.51x for the average workload mix, b) 1.62x for the highest workload ODE-B (CPU-intensive), and c) 1.42x for the lowest workload OLTP-W (storage-intensive). The variation of individual jobs or transactions can be even larger, for example, the average job in our CB-L workload improved 1.58x but the range in individual job improvement was from 1.2x to 2.1x.

Question:

Once my workload is up and running on a z9 EC, how much variability in performance will I see?

Answer:

Minute-to-minute, hour-to-hour and day-to-day performance variability generally grows with the size (capacity) of the server and the complexity of the LPAR configuration. With its improved processor speed and the capability to be configured with up to 64 engines, the z10 EC has the capability to deliver nearly 1.7 times the capacity of the largest previous server. Significant enhancements to the z/OS dispatcher and the PR/SM™ management algorithms (see HiperDispatch discussion below) have been made to help reduce the potential for increased performance variability. In the spirit of autonomic computing, PR/SM and the z/OS dispatcher cooperate to automatically place and dispatch logical partitions to help optimize the performance of the hardware, and minimize the interference of one partition to another. However, while the average performance of workloads is expected to remain reasonably consistent when viewed at small increments of time or by individual jobs or transactions, performance could potentially see more variation than in the past simply due to the expected larger and more complex LPAR configurations that can be supported by the z10 EC.

Question:

What is HiperDispatch and how does it impact performance?

Answer:

HiperDispatch is the z/OS exploitation of PR/SM's new Vertical CPU Management (VCM) capabilities and is exclusive to the z10 EC. Rather than dispatch tasks randomly across all logical processors in a partition, z/OS will assign tasks to groups of about four logical processors, and dispatch work to a "high priority" subset of the logicals. PR/SM provides processor topology information and updates to z/OS, and ties the high priority logical processors to physical processors. HiperDispatch can lead to improved efficiencies in both the hardware and software in the following two manners: 1) work may be dispatched across fewer logical processors therefore reducing the "multi-processor (MP) effects" and lowering the interference among multiple partitions; 2) specific z/OS tasks may be dispatched to a small subset of logical processors which 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.

Question:

What kind of performance improvement can I expect to see from HiperDispatch?

Answer:

The magnitude of the potential improvement from HiperDispatch is related to: a) the number of physical processors, b) the size of the z/OS images in the configuration, c) the logical:physical overcommit ratio and, d) the memory reference pattern or storage hierarchy characteristics of the workload. Generally, a configuration where the largest z/OS image fits within a book will

see minimal improvement. Workloads that are fairly CPU-intensive (like batch applications) will see only small improvements even for configurations with larger z/OS images since they typically have long-running tasks that tend to stick on a logical engine anyway. Workloads that tend to have common tasks and high dispatch rates as often seen in transactional applications may see larger improvements, again depending on the size of the z/OS images involved. LPAR configurations that are over committed, i.e. have higher logical to physical ratios, may see some improvement although the benefit of dispatching to a reduced number of logicals overlaps with benefits already available with IRD and various automation techniques that tend to reduce the number of online logical processors to match capacity needs. The range in benefit is expected to be from 0% to 10% following the sensitivities described above; specifically, configurations with z/OS images small enough to fit in a book or running batch-like workloads will tend to fall at the low-end of the range, multi-book configurations with z/OS images in the 16-way to 32-way range and running transactional workloads will tend to fall toward the middle of the range, and very large multi-book configurations with very large z/OS images and running workloads with intense memory reference patterns will tend to fall toward the high end of the range.

Question:

What is the performance improvement a z/VSE customer might experience on the z10 EC?

Answer:

The performance ratios that a z/VSE customer workload might experience when migrating from a zSeries server to the z10 EC is represented by the range of ratios seen by a comparable z/OS migration. For example, the published ratio in the LSPR between the z9 EC 2094-704 and the z10 EC 2097-704 is approximately 41% to 61%. z/VSE workloads could expect this same range of performance.

Question:

What is the performance improvement a z/VM customer might experience on the z10 EC?

Answer: (Version for the case where we only publish a single z/VM number in LSPR)

The performance ratios that a z/VM customer workload might experience when migrating to z10 EC from older processors will vary. For z/VM, the published z10 EC to z9 EC ratio in the LSPR for a 16-way configuration is approximately 1.30. However, customers may experience a range of performance. It is suggested that you consult the full range of LSPR workloads. The 30% improvement is in the lower end of the improvement range. The lower the memory or processor over-commitment, the greater the potential for the improvement to be in the upper end of the range. A z/VM partition with a lower number of logical processors will also have more potential for being in the upper end of the range.

Question:

Where can I read more about the performance of z/VM in relationship to the z10 EC processor family?

Answer:

The z/VM Performance Resources Page, located at <http://www.vm.ibm.com/perf/>, contains information on z/VM Performance. It will be updated by March 21, 2008 to include additional information on z/VM workloads running on the z10 EC processor.

Question:

How do I get performance information for my TPF products running on a System z10 EC?

Answer:

For the TPF V4 and z/TPF 1.1 products, the TPF Workload Specifics ITRRs should be used. For more information please contact your TPF Support Representative, or send a request to tpfqa@us.ibm.com.

Networking Enhancements

Question:

Tell me about the new OSA Express3 10 GbE feature and when will it be available?

Answer:

The new OSA-Express3 10 GbE LR is designed to increase the throughput for standard frames (1492 byte) and jumbo frames (8992 byte) to help satisfy the bandwidth requirements of your applications. This increase in performance (compared to OSA-Express2 10 GbE) has been achieved through an enhancement to the architecture that supports direct host memory access by using a data router, eliminating 'store and forward' delays. OSA-Express3 10 GbE is planned to be available second quarter 2008.

Question:

What advantage will I get with the HiperSockets Multi Write Facility?

Answer:

HiperSockets eliminates the need to utilize I/O subsystem operations and the need to transverse an external network connection to communicate between logical partitions in the same z10 EC server. HiperSockets is designed to offer significant value in server consolidation connecting many virtual servers, and can be used instead of certain coupling link configurations in a Parallel Sysplex environment.

HiperSockets has been enhanced on the z10 EC to support multiple output buffers on a single SIGA write instruction. This operation is beneficial for the streaming of bulk data over a HiperSockets link between two logical partitions and is designed to offer performance improvements. This function is planned to be made available on z/OS V1.9 in 2Q2008 via a PTF.

Question:

What will HiperSockets Layer2 support provide for me?

Answer:

HiperSockets is a virtual network and was originally designed to only read IP addresses. Some non-IP based applications such as SNA or NetBios could not use HiperSockets because a Layer 2 network and Layer 2 addresses are required for these applications. This new capability on the z10 EC means that it can host both IP and non-IP workloads.

Question:

What networking functions are planned for z/OS?

Answer:

The z/OS Communications Server designs for z/OS V1.10 include:

- New Defense Filtering capabilities designed to provide a mechanism for users to block detected attacks by dynamically installing defensive filters in a TCP/IP stack.
- The z/OS Load Balancing Advisor and Agent are planned to be enhanced to allow users to exploit the AT-TLS feature to secure connections that carry SASP flows.

- Network TCP/IP stack performance improvements in multiple areas, including CPU consumption, cache line contention, and common storage utilization.
- New support to help you coordinate LU name assignments among TN3270 servers in Sysplex
- Enhancements to SNA networking functions, and the z/OS FTP Server and Client

I/O Enhancements

Question:

What is the performance enhancement introduced for Fibre Channel Protocol (FCP) on the z10 EC?

Answer:

The 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. A significant increase in I/O operations per second for small block sizes can also be expected with FICON Express2.

These results were achieved in a laboratory environment using one channel configured as CHIPID type FCP with no other processing occurring and do not represent actual field measurements.

Question:

What is InfiniBand?

Answer:

InfiniBand is an industry-standard specification that defines a first order interconnection technology which is used to interconnect servers, communications infrastructure equipment, storage and embedded systems. InfiniBand is a fabric architecture that leverages switched, point-to-point channels with data transfers of up to 120 gigabits per second, both in chassis backplane applications as well as through copper and optical fiber connections.

Question:

Why was InfiniBand implemented on the z10 EC and what on the z10 EC takes advantage of it?

Answer:

The goal of System z was to introduce an industry-standard high speed host bus physical interface to replace the Self-Timed Interconnect (STI) proprietary host bus interface. At the same time, System z was looking for an industry-standard protocol to provide enhanced performance for Parallel Sysplex coupling links. The InfiniBand host bus physical interface supports 12x Double Data Rate (12x IB-DDR) with a link speed of 6 gigabytes per second (GBps) when attached to a z10 EC and supports 12x Single Data Rate (12x IB-SDR) with a link speed of 3 Gbps when a z10 EC is attached to a z9 EC or z9 BC. The InfiniBand technology supports a distance of up to 150 meters (492 feet) for 12x IB-SDR or 12x IB-DDR links, delivering greater flexibility compared to the Integrated Cluster Bus-4 (ICB-4) link speed of 2 Gbps and distance of up to 10 meters (33 feet).

Server Time Protocol (STP) will use InfiniBand Coupling links to exchange timekeeping messages required for time synchronization.

Question:

You now have three different types of connectivity from the book. Why?

Answer:

When customers require connectivity to ICB-4 coupling links on z9 EC, z9 BC, z990, and z890 from z10 EC, the Memory Bus Adapter (MBA) with its 2 GBps connection is required. When the speed of ICB-4 (2 GBps) and the distance (10 meters) is insufficient for customers' Parallel Sysplex configurations, then the Host Channel Adapter (HCA) 6 GBps interface and its distance of 150 meters will provide benefit. When connectivity to the I/O cages is desired, there is a unique HCA to help satisfy the bandwidth requirements for security (Crypto Express2), coupling at extended distance (ISC-3), SANs (FICON), and LANs (the Open Systems Adapters).

Question:

Remind me again about MIDAW – is it available on the z10 EC and what is the benefit of having it?

Answer:

The Modified Indirect Data Address Word (MIDAW) Facility was introduced with the System z9 servers and is also available on the z10 EC. It was designed to provide more capacity over native FICON channels for programs that process data sets exploiting striping and compression (such as DB2, VSAM, PDSE, HFS and zFS) by reducing channel, director and control unit overhead.

Question:

Does the z10 EC offer an alternative solution for SNA configurations that require NCP functions?

Answer:

Yes. The OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features have the capability to provide channel connectivity from System z operating systems to IBM Communications Controller for Linux on System z (CLL) using the Open Systems Adapter for Network Control Program (OSA for NCP) supporting the Channel Data Link Control (CDLC) protocol. For SNA configurations that require NCP functions, CLL can be considered as a migration strategy to replace IBM Communications Controllers (374x). The CDLC connectivity option enables TPF environments to exploit CLL.

Question:

Will FICON Express2 and FICON Express be available on the z10 EC?

Answer:

FICON Express2 and FICON Express will be available as an MES but not on a new build.

They are not orderable on the z10 EC models, but if they are moved from another system and installed at the time of an upgrade to the z10 EC, they may be retained.

Capacity on Demand Offerings

Question:

How does the new Capacity on Demand architecture benefit me?

Answer:

The new Capacity on Demand architecture offers you increased flexibility and capabilities over previous systems. Unlike previous systems, where only one temporary entitlement record (TER) could be active at a given time, with the new Capacity on Demand architecture, you can have up to four different TERs active at the same time. In addition, the new architecture allows for concurrent permanent upgrades while temporary capacity is active.

Question:

What are the major differences from the System z9 on demand offerings?

Answer:

The z10 EC Capacity on Demand (CoD) implementation allows more flexibility and control over the way you can add temporary upgrades as well as the capability to permanently increase system capacity.

The z10 EC allows you to install up to four temporary records on the CEC and to have any or all these records active at any given time. This means that On/Off CoD can be active with up to three other offerings simultaneously. Each record provides variability in the amount of resources that can be activated and can be controlled and updated independently of each other.

Also, you have the ability to store up to 200 temporary records on the z10 EC Support Element (SE) to provide as much flexibility in temporary record definition as you need. These new capabilities provide the flexibility to monitor the state of each record as well as the ability to add capabilities to individual records concurrently, eliminating the need to order new temporary records as new customer upgrade scenarios develop.

The Customer Initiated Upgrade (CIU) application on Resource Link is designed to allow you to respond to sudden increased capacity requirements by requesting a permanent z10 EC processor unit (PU) or memory upgrade through the Web. This permanent upgrade can occur while the requested resources are activated through a temporary On/Off CoD. Orders of all PU types and memory upgrades that can be delivered by Licensed Internal Code-Control Code (LIC-CC) can be delivered through the CIU application. Permanent upgrades may be performed up to the maximum available resources on the z10 EC.

With the proper contracts in place, all temporary capacity offerings, and the ability to replenish these offerings are available through Resource Link.

Question:

What is the new Capacity for Planned Event (CPE)?

Answer:

CPE is temporary access to dormant PUs, intended to replace capacity lost throughout 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 z10 EC. CPE is similar to CBU in that it is intended to replace lost capacity; however, it differs in its scope and intent. Where CBU addresses disaster recovery scenarios that can take up to three months to remedy, CPE is intended for short-duration events lasting up to three 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).

Question:

How will CPE be priced?

Answer:

There is a fixed price for each CPE record (which is for a three day event). For the price you can enable any amount of dormant capacity on the server and dynamically move between various configurations within a 72 hour time line starting from the initial activation. There are no additional software charges for capacity activated by CPE.

Question:

How does On/Off CoD work?

Answer:

On/Off CoD is available for up to twice the 'purchased' capacity for a given machine based on the LSPR mixed workload multi-image ITRR (Internal Throughput Rate Ratio). On/Off CoD upgrades are allowed for any processor configuration, up to the limit, as long as the number and / or capacity level of the processors is increased. Upgrades that decrease the number of processors, or decrease the capacity level of processors are not allowed. When ordering CP capacity, you will order your upgrades based on a percentage increase over the currently purchased capacity. Other engines are ordered in full engine increments.

Note that On/Off Capacity on Demand for specialty engines will always be full capacity processors.

Question:

Is there still an administrative On/Off Capacity on Demand (On/Off CoD) test?

Answer:

No. With z10 EC, the need for administrative On/Off CoD tests has been eliminated. Unlike early versions of System z9, On/Off CoD upgrades are not immediately activated on download, so you can perform and test the entire On/Off CoD process, short of engine activations, with standard On/Off CoD records and not incur any charges.

Question:

What is the change available that allows me to change my On/Off CoD capacity without having to order and download a new On/Off CoD record?

Answer:

With the new Capacity on Demand record structure, instead of ordering a separate On/Off CoD record for each possible configuration, you can now order a single reusable record that identifies the maximum possible configuration you want to activate. You order CP capacity based on a percentage increase over the currently purchased capacity. Other engines are ordered in full engine increments. When you need to activate temporary capacity, you identify the desired target configuration. Need more processing capacity? Select a new larger target configuration. You can also decrease the amount of temporary capacity active by selecting a lower target. When you are done, simply return to your base configuration. Your On/Off CoD record is still available and you can use it for additional capacity at any time. You are still restricted to an upper bound for temporary upgrade size of no more than double the machine's purchased capacity configuration. You will be charged for the additional capacity on a 24 hour basis, and if you increase capacity multiple times during a 24 hour period, the charge will apply to the greatest amount of capacity activated.

Question:

Can I have more than one record describing my On/Off CoD configuration loaded on my z10 EC?

Answer:

Yes, you can download and stage up to 200 separate Capacity on Demand records on the Support Element. Multiple records may also be loaded at the same time, however, only one On/Off CoD record may have active capacity at a given time.

Question:

What is the API that is available for On/Off CoD?

Answer:

There is an API provided within the existing HMC SNMP command which is designed to enable customers to use other automation code (which conforms to the API) to enable activation of On/Off CoD on the HMC without human intervention. This will allow for flexibility of operation of the On/Off CoD function.

Question:

What is unassigned capacity?

Answer:

Unassigned capacity can be either general purpose or IFL capacity that has been purchased by the end user, but for business reasons has been unassigned (turned off) by the end user so no business processing can be executed on the unassigned capacity.

Question:

What is the hardware price of On/Off CoD when using unassigned capacity compared to capacity that has not been previously purchased?

Answer:

At this time there are no additional hardware or maintenance charges for activating unassigned CP or IFL capacity with an On/Off CoD record. You are still responsible for any additional software charges that may result from activating unassigned capacity. IBM, at its discretion, reserves the right to add maintenance charges at a later time.

Question:

Can On/Off CoD use unassigned IFL or CP capacity for any purpose other than its original intended purpose?

Answer:

Yes. The processor reserved as the unassigned IFL capacity can be activated as any other On/Off CoD processor type. Note: if you activate unassigned capacity for any purpose other than its original intended purpose, you will be billed as if the capacity were unowned.

Question:

Do I need to follow a different process to utilize unassigned capacity with On/Off CoD?

Answer:

No. The order process is the same. Resource Link will give you a priced option.

Question:

Will temporary capacity be available for specialty PUs (IFLs, ICFs, zAAPs, zIIPs, SAPs) on the z10 EC?

Answer:

Yes. All engine types are available for temporary upgrades.

Question:

Can I order On/Off CoD if I have subcapacity CP processors?

Answer:

Yes. Again, On/Off CoD is available for up to twice the 'purchased' capacity for a given machine and you can not decrease the number or the capacity level of installed general purpose processors. With the full 'matrix' upgradeability of the subcapacity processors you can temporarily change capacity with processors of equal, less or greater capacity depending on your requirements. When you use subcapacity settings for your temporary On/Off CoD capacity, you can not exceed 12 general purpose processors. (NOTE – you could have subcapacity general purpose processors and, when you execute On/Off CoD, as long as you don't exceed more than twice your purchased capacity, have 12 or more full capacity temporary general purpose processors).

Question:

When can I begin placing On/Off CoD orders against my z10 EC?

Answer:

On/Off CoD can be initiated as soon as the profile for the z10 EC is established. The prerequisite of establishing a profile is the signing of the necessary contract supplements associated with ordering features 9900 and 9896.

Question:

Will my System z9 processors be able to use the new Capacity on Demand features?

Answer:

No These new Capacity on Demand features are only available on the z10 EC.

Question:

Will I be able to do Capacity Backup Upgrade (CBU) capability on any engine type?

Answer:

Yes. The z10 EC is able to activate all processor types as part of CBU upgrades: IFLs, zAAPs, zIIPs, ICFs, CPs and SAPs.

Question:

Can I order CBU processors if I have subcapacity processors?

Answer:

Yes. CBU is available when you are running with subcapacity processors but you can not decrease the number or the capacity level of installed CP processors. Note that you can not exceed 12 subcapacity processors (CP or CBU). A customer no longer has to increase the quantity of CPs for CBU, but can just increase the capacity of the existing CP count by ordering CBU CP features with greater capacity. When the quantity of CBU processors exceeds 12, all CBU processors will be full capacity.

Question:

Can I add CBU capacity by selecting CBU engines that have more capacity than my permanent configurations and have less actual engines than my original permanent capacity?

Answer:

No, you can not reduce the number of engines in your permanent configuration when you add CBU capacity.

Question:

Can I convert an active permanent engine to another engine type during CBU?

Answer:

No. All active permanent engines must remain as part of the CBU environment (although they may change in capacity) and you can not convert them to another type during the CBU event. Unassigned engines may be used as another engine type during a CBU event.

Question:

What Capacity on Demand features can I order on IBM Resource Link?

Answer:

From Resource Link you can order permanent processor and memory upgrades, On/Off CoD, CPE, and CBU records.

Question:

Is FC 9898, the permanent upgrade feature, still a prerequisite to On/Off CoD?

Answer:

No. You can now enable your machine for temporary On/Off CoD upgrades without having to enable the machine for permanent upgrades.

Question:

If I have Capacity on Demand features on my system now will they be lost if I upgrade to a z10 EC?

Answer:

A: The CBU, CIU Enablement and On/Off Capacity on Demand Enablement features can be brought forward to the z10 EC. Any previously staged upgrades ordered through CIU for permanent or On/Off CoD upgrades will be lost.

Question:

Has the way I order Capacity on Demand features changed?

Answer:

Permanent processor and memory upgrades, CBU and the new Capacity for Planned Event may be ordered either through your IBM or BP sales representative, or Resource Link. On/Off Capacity on Demand orders must still be placed through Resource Link only.

Question:

Tell me about the Capacity Provisioning capabilities of the z10 EC?

Answer:

A new Capacity Provisioning Manager planned for z/OS V1.10, and available on z/OS V1.9 with PTF, can monitor z/OS systems on System z10 servers. Activation and deactivation of temporary capacity can be suggested or performed automatically based on user-defined schedules and workload criteria.

In addition, with z/OS V1.10, IBM plans to introduce the Capacity Provisioning Control Center, a tool for managing capacity provisioning for System z10 servers. It is designed to manage provisioning policies and domain configurations. Provisioning policies specify the criteria for capacity increases and decreases, while domain configurations specify systems to be observed and servers to be managed. Initial support is planned for a policy definition application which requires a workstation running Windows® XP.

Specifically, the Capacity Provisioning Control Center will provide the following functions:

- Create and edit Capacity Provisioning policies
- Create and edit Capacity Provisioning domain configurations
- Connect to the Provisioning Manager; Display the status of the Provisioning Managers
- Install Capacity Provisioning policies and domain configurations into the Provisioning Manager

Cryptographic Enhancements

Question:

Remote key loading capability was added to the z9 EC with the April 2006 announcement. Is Remote Key loading capability available on z10 EC?

Answer:

Yes, Remote key loading is available for capable Automated Teller Machines (ATM) and Point of Sale (POS) systems. Remote key loading refers to the process of loading DES (Data Encryption Standard) keys and T-DES (Triple DES) keys to ATMs and POS systems from a central administrative site without the need for personnel to visit each machine to manually load keys. In the past, key loading has been done by manually loading each of the two clear text key parts individually and separately into ATMs and POSs. Remote key loading provides a more cost-effective way of managing the terminals than by having several people travel to the ATM and POS system with key parts. This function is supported by Enhancement to Cryptographic Support for z/OS and z/OS.e 1.6/1.7 Web deliverable, Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e V1.7 and V1.8 Web deliverable, z/VM 5.2 and 5.3 for guest exploitation. Linux can also take advantage of the robust set of functionality provided by the Common Cryptographic Architecture (CCA).

Question:

What advantages does remote key loading provide for businesses that manage ATMs and POS systems?

Answer:

By providing a way to load ATM and POS keys without sending staff on-site, remote key loading can reduce downtime due to key errors, reduce service call and key management costs and improve the ability to manage ATM/POS conversions and upgrades.

Question:

What ATMs and POS systems will support the z10 EC Remote Key Load function?

Answer:

IBM System z10 Remote Key Load capability will support ATMs and POS systems that adhere to the following standards:

- ISO/IEC 11770-3: Information Technology, Security Techniques, Key Management, Part 3: Mechanisms Using Asymmetric Techniques.
- ANSI X 9.24-2: Retail Financial Services, Symmetric Key Management, Part 2: Using Asymmetric Techniques for the Distribution of Symmetric Keys.

Question:

What is ISO 16609 CBC Mode T-DES MAC?

Answer:

International Standards Organization (ISO) 16609 supports the requirements for message authentication using T-DES symmetric encryption. Integrated Cryptographic Service Facility (ICSF) with the Crypto Express2 feature, supports basic mechanisms in Remote Key loading. The implementation offers a secure bridge between the highly secure Common Cryptographic Architecture (CCA) environment and the various formats and encryption schemes offered by ATM vendors. Refer to Application Programmers Guide, SA22-7522, for additional details.

Question:

What cryptographic hardware is supported on the z10 EC?

Answer:

CP Assist for Cryptographic Function (CPACF), which is shared between two Processor Units (PUs), and the Crypto Express2 feature.

Question:

What cryptographic hardware is not supported on the z10 EC?

Answer:

PCI Cryptographic Coprocessor (PCICC), PCI X Cryptographic Coprocessor (PCIXCC), CMOS Cryptographic Coprocessor Facility (CCF), PCI Cryptographic Accelerator (PCICA) and the Crypto Express2-1P features are not supported on the z10 EC.

Question:

What cryptographic hardware is offered as standard features on the z10 EC?

Answer:

The CP Assist for Cryptographic Function (CPACF) is a standard feature on the z10 EC. A no-charge enablement feature #3863 is required to utilize the CP Assist for Cryptographic Function and for export control. The Crypto Express2 feature is an optional feature. The first order increment is two features.

Question:

What features are currently supported by the CP Assist for Cryptographic Function (CPACF) on the z10 EC?

Answer:

CPACF includes support of the DES, T-DES, Advanced Encryption Standard (AES), Secure Hash Algorithms and Pseudo Random Number Generation (PRNG).

CPACF, supporting clear key encryption, is activated using a no charge enablement feature #3863 and offers the following:

- Data Encryption Standard (DES)
- Triple Data Encryption Standard (T-DES)
- Advanced Encryption Standard (AES) - 128 bit
- SHA-1
- SHA-256
- Pseudo Random Number Generation (PRNG)

CPACF has been enhanced to include the following on CPs and IFLs:

- Advanced Encryption Standard (AES) for 192 bit and 256 keys
- SHA-224, SHA-384 and SHA-512 for message digest.

CPACF throughput scales linearly as CPACFs are added. SHA-1, SHA-256, and SHA-512 are shipped enabled on all servers and do not require the enablement feature. DES, T-DES and AES functions require enablement of the CPACF function (Feature Code 3863) for export control. CPACF cryptographic functions are aimed at encryption, decryption and hashing of data transferred over open networks and data sent to storage.

Question:

What features are available with Crypto Express2?

Answer:

Crypto Express2 is configurable. The Crypto Express2 feature has two PCI-X adapters, and each can be defined as either a Coprocessor or as an Accelerator

- The Crypto Express2 Coprocessor (default):
 - Supports Secure key encrypted transactions
 - Supports highly secure cryptographic functions, use of secure encrypted key values, and User Defined Extensions (UDX)
 - Is designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification.
 - Crypto Express2 Accelerator – to enable SSL acceleration on z10 EC.
 - Supports clear key RSA acceleration
 - Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol.
 - Up to eight Crypto Express2 features per server.
 - All logical partitions (LPARs) in all Logical Channel Subsystems (LCSSs) have access to the Crypto Express2 feature, up to 32 LPARs per feature.
-

Question:

What are the functional differences between the Coprocessor and Accelerator configurations on Crypto Express2 features?

Answer:

The z10 EC provides the ability to configure Crypto Express2 PCI-X adapters as accelerators or as coprocessors. When both PCI-X adapters are configured as accelerators, the Crypto Express2 feature is designed to perform up to 6000 SSL handshakes per second. The SSL rate was achieved with a z10 EC environment with four processors and two PCI-X adapters cards (one Crypto Express2, both configured as accelerators), Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e V1.7 and V1.8 Web deliverable.

These measurements are examples of the maximum transactions per second achieved in a laboratory environment with no other processing occurring and do not represent actual field measurements. Details are available upon request.

Since the performance enhancements are implemented in Licensed Internal Code, Crypto Express2 features have been carried forward from z9 EC to z10 EC platforms.

Question:

Will the Crypto Express2, feature code 0863, support Linux running on System z on the z10 EC?

Answer:

Yes, Linux running on the z10 EC support both secure and clear key operations. The Secure Sockets Layer (SSL) and Transport Layer Security (TLS) cryptographic operations are supported in both modes along with extensive secure key cryptographic support provided by the Common Cryptographic Architecture (CCA) library now available for Linux running on z10 EC. IBM is working with its distribution partners to continue providing cryptographic functions in future distribution releases, or service updates.

Question:

Will IBM offer, at some point in the future, support for RSA Keys greater than 2048 bits in length?

Answer:

RSA Keys, up to 4096 bits, are supported on z10 EC. The RSA services in the CCA API are extended to support RSA keys with modules lengths up to 4096 bits. The services affected include key generation, RSA-based key management, digital signatures, and other functions related to these services.

Crypto Express2, feature code 0863, along with Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e V1.7 and V1.8 Web deliverable are required. RSA keys greater than 2048 bits, in length, are supported when running z/VM 5.2 and above for guest exploitation.

Question:

What Integrated Cryptographic Service Facility (ICSF) services are available with CP Assist for Cryptographic Function (CPACF)?

Answer:

All critical Integrated Cryptographic Service Facility (ICSF) services that currently execute on z890, z990, z9 EC and z9 BC PCIX Cryptographic Coprocessor (PCIXCC) feature are planned to be supported by the Crypto Express2 feature.

The following ICSF callable services will be available with CPACF and Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e V1.7 and V1.8 Web deliverable.

CSNBSYE, CSNBSYE1, CSNBSYD, and CSNBSYD1 will provide support for clear-key AES encryption and decryption with 256-bit keys using the CPACF.

CSFNBOWH and CSNBOWH1 will provide support for SHA-1 and SHA-512 using CPACF.

These services continue to support DES and TDES.

Question:

What other means are available for customers to utilize CPACF?

Answer:

For IBM and customer written programs, CPACF function for DES, T-DES, AES-128, AES-192 and AES-256, SHA-1, SHA-224, SHA-256, SHA-384 and SHA-512 functions can be invoked by instructions as described in the z/Architecture® Principles of Operation, SA22-7832-02. As a group, these instructions are known as the Message Security Assist (MSA). These are all problem state instructions and are all in RRE format.

Linux running on z10 EC can also utilize CPACF Function. IBM is working with its distribution partners to provide support for the new cryptographic functions in future distributions releases, or service updates.

Question:

Is a Trusted Key Entry (TKE) workstation mandatory for the use of the CP Assist for Cryptographic Function (CPACF) and Crypto Express2 feature?

Answer:

The CP Assist for Cryptographic Function (CPACF) supports clear key functions and does not require entering of master keys. The TKE workstation is an optional feature. It is not mandatory for use with the Crypto Express2 feature, however, the TKE workstation offers security-rich local and remote key management, and also provides authorized persons a method of operational and master key entry, identification, exchange, separation, and update.

Question:

Will there be an update to the ATS TechDocs Web site to provide additional technical information about z10 EC cryptographic hardware features?

Answer:

ATS TechDocs Web site and several z10 EC cryptographic technical papers will be updated as appropriate. The ATS TechDocs Web site URL is ibm.com/support/techdocs/atmastr.nsf

Question:

What releases of operating systems are required to support the Crypto Express2 hardware feature and CP Assist for Cryptographic Function (CPACF) on the z10 EC?

Answer:

The minimum software support requirements for Crypto Express2 and CP Assist for Cryptographic Function (CPACF) features are as follows:

- Crypto Express2:
 - Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e V1.7 and V1.8 Web deliverable.
 - z/VM 5.2 guest exploitation.
 - z/VSE 3.1
 - z/TPF V1.1 (acceleration mode only) with APAR PJ30717
 - Linux on System z. The latest levels of Novell SUSE Linux Enterprise Server and Red Hat Enterprise Linux distributions include the support necessary to utilize the Crypto Express2. Note: z/VSE supports clear key RSA operations only. Linux on System z and z/VM V5.2, and later, support clear and secure key operations.

- CPACF (DES, TDES, SHA-1, SHA-256, SHA-384, SHA-512, AES-128, AES-192, AES-256 and PRNG):
 - z/OS V1.7 with any of the following:
 - Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 Web download (no longer available),
 - Enhancements to Cryptographic Support for z/OS and z/OS.e V1R6/R7 Web deliverable (no longer available),
 - Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e 1 V1.7 and V1.8 Web deliverable
 - z/OS V1.8 or higher
 - z/VM V5.2 for guest exploitation.
 - z/VSE V3.1 and IBM TCP/IP for VSE/ESA™ V1.5e with PTFs.
 - z/TPF V1.1
 - TPF V4.1
 - Linux for System z - The latest levels of currently available distributions from Novell SUSE and Red Hat include the support necessary to utilize CP Assist for Cryptographic Function.
 - Enhancements to CP Assist for Cryptographic Function (CPACF) on the z10 EC requires at a minimum:
 - Cryptographic Support for z/OS V1.7, V1.8, V1.9 and z/OS.e 1 V1.7 and V1.8 Web deliverable.
 - z/VM V5.2 for guest exploitation.
 - z/VSE V4.1 and IBM TCP/IP for VSE/ESA V1.5e with PTFs.
 - Linux on System z - IBM is working with its Linux distribution partners to support CPACF enhancements in future Linux on System z releases.
-

Question:

What z/OS cryptographic support is planned?

Answer:

With z/OS V1.10, IBM plans to update the z/OS Integrated Cryptographic Services Facility (ICSF) with the following functionality:

- 4096-bit RSA key support. IBM plans to provide 4096-bit RSA support on System z servers. The servers must have the 4096-bit RSA signature generation and verification support available with feature 0863 installed, and the Crypto Express Coprocessor with microcode level MCL006-MCL009 on these servers: z10 EC, z9 EC and z9 BC.
- ISO Format-3 PIN Block support that meets the ISO 9564-1 Banking standard is planned by IBM. Feature 0863 must be installed, and the Crypto Express2 Coprocessor with microcode level MCL006-MCL009 on these servers: z10 EC, z9 EC and z9 BC.

Additionally, z/OS V1.10 support for these functional enhancements to System SSL is planned:

- Utilize hardware support for RSA digital signature generate and verification and RSA encrypt and decrypt available on z10 EC, z9 EC and z9 BC servers with feature 0863 installed with the latest Crypto Express2 Coprocessor with microcode level 3.30, which is provided by MCL006-MCL009.

Question:

Will UDX's written for zSeries servers function on the z10 EC?

Answer:

Customers should always contact IBM to understand the effects of a new environment on their UDX's. For further information, see the answer to the next question.

Question:

Will new UDX's be supported on the z10 EC?

Answer:

Yes. If you wish to inquire further about UDX support for the Crypto Express2 feature, please contact one of the following IBM representatives: Leo Moesgaard (Leo.Moesgaard@dk.IBM.com) or David Evans (Davee@us.IBM.com)

Question:

What are the functions and attributes of the CP Assist for Cryptographic Function (CPACF) and Crypto Express2 features?

Answer:

The following table highlights the features or attributes:

Functions or attributes	CPACF	Crypto Express 2
Supports z/OS applications using ICSF	X	X
SSL handshake capability - performance best when both PCI-X adapters are configured in accelerator mode		X
Provides highest symmetric encryption performance (clear key)	X	
Provides highest symmetric encryption performance (secure key)		X
Provides highest asymmetric (clear key) performance (in accelerator mode)		X
Provides highest asymmetric (encrypted key) performance (in coprocessor mode)		X
Uses CHPID numbers	N/A	N/A
Shared physically between two (2) Processor Units	X	
Requires CP Assist for Cryptographic Function (CPACF) enablement	X	X
Requires ICSF to be active, for z/OS users		X
Requires system master keys to be loaded (in coprocessor mode)		X
Offers user programming function support (UDX) (in coprocessor mode)		X
Usable for data privacy - encryption and decryption processing	X	X

Usable for data integrity - hashing and message authentication	X	X
Usable for financial processes and key management operations (in coprocessor mode)		X
Crypto performance RMF monitoring		X
System (master) key storage (in coprocessor mode)		X
Retained key storage (in coprocessor mode)		X
Designed for tamper-resistant hardware packaging (in coprocessor mode)		X
Designed for FIPS 140-2 Level 4 certification (in coprocessor mode)		X
Supports SSL functions	X	X
Supports Linux applications	X	X
RSA functions		X
High Performance SHA-1, SHA-128, SHA-224,SHA-256, 384, 512 Hash functions	X	
Clear key DES/T-DES	X	
Clear key RSA		X

Question:

Does the use of RSA Retained private keys limit availability?

Answer:

Yes. The use of retained private keys creates an application single point of failure. Since RSA Retained private keys cannot be copied, backed up or scaled from a performance perspective, these keys should only be used if mandated by the customers' security policy. For those customers that require a private key that is intended to be shared across logical partitions, they should use RSA keys encrypted under a host master key instead of a retained key. The use of the RSA keys encrypted under a host master key is designed to prevent the loss of the key associated with the RSA Retained private key specific to the Crypto Express2 feature.

Question:

What functions are supported by TKE 5.2?

Answer:

The Trusted Key Entry (TKE) 5.2 level of Licensed Internal Code (LIC) is installed in TKE workstation feature code #0839. The TKE 5.2 LIC continues to support the ability to store key parts on diskettes or paper, or optionally on smart cards, or to use a TKE authority key stored on a diskette, or optionally on a smart card, and to log on to the Cryptographic Coprocessors using a passphrase, or optionally a logon key pair. The benefits of TKE 5.2 LIC include service mode support and usability enhancements including a service user to improve access to operations and a new task layout display to maintain TKE console consistency. TKE 5.2 LIC is a no-charge enablement feature which is loaded prior to shipment when a TKE workstation is ordered. The TKE 5.2 LIC includes support for the Smart Card Reader.

The optional TKE features are:

- TKE 5.2 LIC (#0857) and TKE workstation (#0839)

- TKE Smart Card Reader (#0887)
 - TKE additional smart cards (#0888)
-
-

Question:

Will there be an upgrade to the Trusted Key Entry (TKE) workstation?

Answer:

Yes, the current TKE workstation for z10 EC will be upgradeable. The TKE workstation is assigned feature code #0839 with LIC level 5.2. Customers must use the TKE workstation to control the z10 EC. Customers may continue to use the previous TKE , feature code #0859, with an upgrade of LIC level 5.2.

Question:

Will customers have the option to order Trusted Key Entry (TKE) workstation with Token-Ring or Ethernet?

Answer:

No, a TKE workstation with Token Ring is not offered. Ethernet is the only option offered. This satisfies the Statement of General Direction in Hardware Announcement 104-115, dated April 7, 2004, Hardware Announcement 104-117, dated April 7, 2004, and Hardware Announcement 104-118, dated April 7, 2004.

Question:

Will Smart card reader support be available with z10 EC TKE workstations?

Answer:

Yes, support is available for an optional Smart Card Reader to be attached to the new TKE LIC level 5.2 workstation. Customers may also carry forward optional Smart Card Reader features.

Question:

Will the number of TKE workstations be limited to one per system?

Answer:

Up to three (3) LIC level TKE 5.2 workstations per system will be supported on the z9 EC. This means that up to three (3) smart card reader features, each feature consisting of two smart card readers, can be attached to the z10 EC.

Parallel Sysplex Enhancements

Note: Please use the updated STP FAQs to answer your questions on the Server Time Protocol feature of the z10 EC. These can be found at: www.ibm.com/systems/z/faq/

Question:

What advantages will I have with InfiniBand Coupling Links (IFBs)?

Answer:

IFBs provide a high speed coupling connection up to 150 meters. While ICB-4 still provide the fastest coupling service times for distances up to 10 meters, IFBs can provide significantly improved service times compared to ISC-3s for distances up to 150 meters.

Question:

How far back can a z10 EC connect to in a Parallel Sysplex environment?

Answer:

System z10 servers are designed to coexist in the same Parallel Sysplex environment with (n-2) server families. This allows System z10 servers to coexist with the IBM System z9, as well as the IBM z890 and z990 server families.

Question:

Do the IBM System z9 servers support IFB links?

Answer:

Yes, the IBM System z9 servers can be configured with IFB links. Although with this announcement, the System z9 servers can only be configured as stand-alone Coupling Facilities, IBM intends in the future to support any System z9 configuration with IFBs links.

Question

What is new with Server Time Protocol (STP)?

Answer:

IBM has announced NTP Client function for Server Time Protocol (STP).

This support will allow customers to initialize the time of an STP-only Coordinated Timing Network to the time provided by a Network Time Protocol (NTP) server, and maintain time accuracy. This allows an enterprise comprised of heterogeneous platforms to track to the same time source. NTP Client support is also available on the IBM System z9 servers.

NTP Client support resolves a Statement of Direction made in Hardware Announcement Letter 106-715, October 10, 2006.

Question:

What Parallel Sysplex and Parallel Sysplex-related improvements are planned for z/OS?

Answer:

With z/OS V1.10, IBM plans to introduce the following enhancements:

- In z/OS V1.10, XCF and XES plan to extend and enhance their existing health checks to provide new and improved checks to detect single points of failure. Improved CF structure and duplexing checks are planned to help you avoid sysplex related problems. New checks

are also planned for the Sysplex Failure Manager action specifications to help you improve sysplex availability.

- HMC plans support for configuration packages similar to those supported by HCD, to allow a subset of configuration to be created. For example, this support is designed to allow you to create and send a configuration package for a sysplex or a single site to another location.
- z/OS Communications Server designs for z/OS V1.10 include new support to help you coordinate LU name assignments among TM3270 servers in sysplex. With this enhancement, one TN3270E Telnet Server in the group acts as an LU Name Server and allocates shared LU names to other TN3270E Telnet Servers within the group. This allows load balancing across multiple TN3270E Telnet Servers with consistent configurations. High availability of the LU Name Server service is provided with automated takeover and recovery.
- Statement of Direction - an IBM z/OS Management Facility plans to provide the infrastructure, services, and user interfaces to support a modern, Web-browser based management console for z/OS. The z/OS Management Facility is intended to enable system programmers, administrators and operators to manage and administer a mainframe system by simplifying day to day operations and administration of a z/OS system. This first release of the z/OS Management Facility will provide job and process management and Parallel Sysplex management support.
- In z/OS V1.10, IBM intends to deliver the final phase of Consoles Enhancements. In this final phase, consoles processing will be designed to reduce serialization contention by reducing the scope of serialization for many operations from a console class to an individual console. Additionally, support will be provided to increase the maximum number of MCS, SMCS, and subsystem consoles in a sysplex from 99 per sysplex to 99 active consoles per system.
- In z/OS V1.10, RACF® design is planned to be changed to help you preserve RACF database data integrity in a sysplex. This is designed to help improve availability by eliminating potential causes of database corruption.
- A z/OS UNIX System Services function will be designed to allow you to change sysplex root data sets dynamically, without a sysplex-wide IPL. New commands are expected to eliminate a cause for planned outages and to facilitate migration of sysplex roots from HFS to zFS.
- Additional improvements are planned for the Load Balancing Advisor and Load Balancing Agent, DFSMSHsm™ CDS backup, GRS and RRS.

IBM Systems Director Active Energy Manager (AEM) on Linux for System z V3.1

Question:

What is the difference between Active Energy Manager for Linux on System z V3.1 and the System z environmental monitor, nicknamed the "mainframe gas gauge"?

Answer:

There are a number of differences between the Active Energy Manager and the System z environmental monitor.

- AEM V3.1 has functions that extend the System z on board environmental/energy monitor:
 - a. The System z environmental monitor, nicknamed the "mainframe gas gauge," of the System z System Activity Display (SAD) has a single screen for statistics with a graph for processor utilization. AEM V3.1 provides added value in that statistics are captured from System z10 and the data is readily available for trending analysis.
 - b. The System z on board energy monitor is a single point in time snapshot of energy in KW, cooling in BTU/hr, inlet temperature and utilization in percent by Processor Unit (PU). AEM V3.1 provides the ability to view trends associated with System z10 over a period of time.
- AEM for Linux on System z V3.1 is a cross platform energy management solution that can be used to not only monitor System z10, but also monitor and manage energy use for IBM BladeCenter[®], IBM POWER[™] Systems and IBM System x[™]. The System z environmental monitor is specifically for the mainframe.
 - AEM V3.1 monitors energy and thermal data from System z10. Energy and thermal monitoring support with the HMC SAD was introduced on System z9.
 - AEM V3.1 provides a source of energy management data that can be exploited by Tivoli enterprise solutions such as IBM Tivoli Monitoring and IBM Tivoli Usage and Accounting Manager. The System z environmental monitor offers a Power Estimator Tool for future planning of system configurations including the energy needed to both run and cool System z.
 - AEM V3.1 is a software offering that runs on Linux on System z. The environmental monitor runs on all System z9 servers with driver 67 and all System z10 servers and is independent of operating systems.

Question:

Where can I get more information on AEM for Linux on System z ?

Answer:

More information on AEM can be found at:

<http://www-03.ibm.com/systems/management/director/extensions/actengmrg.html>

System z and Cell Broadband Engine Architecture

Question:

What is Cell Broadband Engine™ (Cell/B.E.) Architecture?

Answer:

Cell Broadband Engine (Cell/B.E.) is a microprocessor architecture that was jointly developed by Sony Computer Entertainment, Toshiba and IBM. Cell/B.E. technology combines the general-purpose Power Architecture® core of modest performance with streamlined co-processing elements which can greatly accelerate multimedia or vector processing applications, as well as many other forms of dedicated or numerically intensive computation. The first commercial application of Cell/B.E. was in Sony's PLAYSTATION®3 game console.

The Cell/B.E. microprocessor often delivers more performance than conventional microprocessors while being highly scalable and easier to program than exotic technologies like graphics processing units (GPU), digital signal processors (DSP) or field-programmable gate arrays (FPGA).

Today IBM sells the Cell/B.E. technology in the form of the IBM BladeCenter. On 29 August 2007, IBM announced the IBM BladeCenter QS21. With two Cell/B.E. processors each blade can deliver a peak performance of approximately 460 GFLOPS. At 1.05 Giga Floating Point Operations per Second (GigaFLOPS) per watt, it is one of the most power efficient computing platforms to date.

Question:

What types of applications are appropriate for Cell/B.E. technology?

Answer:

Applications that can take advantage of SIMD processing are ideal for Cell/B.E. technology. SIMD (Single Instruction, Multiple Data) is a technique employed to achieve data level parallelism, as in a vector processor. An application that may take advantage of SIMD is one where the same value is being added (or subtracted) to a large number of data points, a common operation in many multimedia applications. One example would be changing the brightness of an image. Each pixel of an image consists of three values for the brightness of the red, green and blue portions of the color. To change the brightness, the R G and B values are read from memory, a value is added (or subtracted) from them, and the resulting values are written back out to memory. With an SIMD processor, such as the Cell /B.E. processor, there are two improvements to this process. For one, the data is understood to be in blocks, and a number of values can be loaded all at once. Instead of a series of instructions saying "get this pixel, now get the next pixel", an SIMD processor will have a single instruction that effectively says "get lots of pixels" ("lots" is a number that varies from design to design). For a variety of reasons, this can take much less time than "getting" each pixel individually, as with traditional CPU design.

Applications most commonly associated with Cell/B.E. technology and SIMD processing, are the Graphics and Video workloads. However, many applications and workloads across a wide range of commercial and scientific uses are realizing tremendous benefit such as algorithmic trading, portfolio risk analysis, medical and bioinformatics, scientific computing, physics data manipulation applications, audio and signal processing and compression/decompression.

Question:

What does "integration of System z and Cell/B.E. technology mean?"

Answer:

IBM intends to integrate the Cell/B.E. technology into the System z ecosystem. This will allow the implementation of applications which combine processing on both System z and Cell/B.E. while extending the mainframe Qualities of Services (QoSs) to facilitate the management of a pool of Cell/B.E. processors from a System z management interface. This will be done in a way that incorporates workload management, virtualization, security, availability, monitoring and chargeback.

Question:

What value will this solution provide to System z customers?

Answer:

Enabling System z and Cell/B.E. processors to work tightly together will provide System z customers the ability to enhance the capabilities of existing applications or to create new solutions which leverage System z mission critical transactional and database processing with high performance computations. Previously customers have had to move computationally intense applications to distributed platforms with a different set of system management tools and organizational governance. This increased complexity and costs. The IBM BladeCenter QS21 (Cell/B.E. Processor), provides highly competitive price/performance and performance/watt which can reduce the number of High Performance Cluster (HPC) servers required. The integration with System z QoS can reduce the operational complexity with a uniform management structure. The combined benefits of a System z and Cell/B.E. processor solution will ultimately deliver to our clients, faster processing of critical business functions with enhanced decision making capability, achieved with an overall improved total cost of ownership.

Question:

What value will this solution provide to Cell/B.E. BladeCenter customers?

Answer:

Enabling Cell/B.E. processors to work within the System z ecosystem brings tremendous benefits to the overall function of a typical HPC cluster. Immediate benefit will be realized from gaining real-time access to the mission critical data stored on the mainframe. In addition, BladeCenter QS21 customers will be able to take advantage of the strengths offered by today's mainframe based workload management, systems management and virtualization functionality.

Question:

When will the System z and Cell/B.E. processor solution be available?

Answer:

Customers can begin to use System z and Cell/B.E. processing today. The IBM BladeCenter development kit functions are not yet in place to realize the ultimate "integrated" solution, proofs of concept can be built using existing programming functionality. As a case in point, IBM has successfully executed a proof of concept that has the System z (DB2 / z/OS and Linux on System z) and Cell/B.E. BladeCenter QS21 technologies working together. We are able to demonstrate a multi-threaded Java z/OS application running on z/OS extracting a portfolio of stock option data from a DB2 / z/OS table and passing it to the Cell/B.E. technology to run a European Options Monte Carlo simulation with 20,000 simulations. The results are then returned and stored in DB2 / z/OS and visualized by System z.

Over time, enhancements will be made to System z and Cell/B.E. operating systems, development toolkits, connectivity capability and processor technology, each bringing more functionality and improved performance to already exciting solution..... A 'marriage' of two technologies that perfectly compliment each other!

GDPS Enhancements

Question:

What GDPS® functions were announced with GDPS v3.5 on February 26, 2008?

Answer:

GDPS has announced

- Distributed Cluster Management (DCM). This is designed to provide coordinated disaster recovery across System z and non-System z servers by integrating with distributed cluster managers. GDPS V3.5 support includes integrated with Veritas Cluster Server™ (VCS). Initial support will be provided by GDPS/PPRC and GDPS/XRC.
- GDPS/PPRC Multiplatform Resiliency for System z expanded to include Red Hat Enterprise Linux 4
- Expanding graphical user interface (GUI). The GUI has been available for GDPS/PPRC since March 2007 with GDPS V3.4. GDPS V3.5 expands this support by providing a GUI interface to manage Peer-to-Peer Virtual Tape System (PtP VTS), as well as providing the GUI interface for GDPS/PPRC HyperSwap™ Manager and GDPS/Global Mirror.
- Greater disk efficiency with support for FlashCopy® Space Efficient.
- Improved performance and system management with support for z/OS Global Mirror Multi-Reader.
- New service support policy extending GDPS service for up to three years.

GDPS is previewing the following functions that will be available with GDPS V3.5:

- Increased availability with GDPS/MzGM support for z/OS Metro/Global Mirror Incremental Resync
- Improved recovery time with CBU and On/Off CoD enhancements
- Improved price performance with System Data Mover offload to zIIP

IBM is announcing the following Statements of Direction:

- Distributed Cluster Management integration with Tivoli System Automation Application Manager. Initial support will be provided by GDPS/PPRC in 2008.
- GDPS cooperation with Tivoli Business Continuity Process Manager. This provides an integration of GDPS configuration data in CMDB with an integration of GDPS as operational manager in a workflow support for Business Continuity.

For a complete list of other supported hardware platforms and software prerequisites, refer to the GDPS Web site <http://www.ibm.com/systems/z/gdps/>

Contact gdps@us.ibm.com to get more information on GDPS solutions.

IBM Global Financing (IGF)

Question:

What can IGF offer customers interested in acquiring new z10 EC technology?

Answer:

IBM Global Financing (IGF), the leasing and lending business segment of IBM can offer financing terms and conditions that in many cases can make it possible for qualified clients to acquire the new z10 EC with leases at lower monthly payments than they currently make today.

IGF can help customers:

- Upgrade to the new IBM System z10 EC technology: Swap out older equipment and provide attractive lease financing on new equipment with non-disruptive migrations
- With financing for a field upgrade: Grow within new z10 EC technology with easy upgrade lease financing and extensions, often at the same payment level or less.
- A Sale Leaseback: IGF can purchase existing technology from the customer, and lease the base machine back to the client or upgrade it in place with a lease package for the new z10 EC technology.


Question:

How can leasing with IGF lower a customer's monthly payments?

Answer:


- I. The customer only pays for what he or she uses. In fair market value leases, customers are effectively paying for the purchase price of the machine less the residual value (RV). The lessee benefits from a lower monthly payment. Full risk with respect to the residual value of the machine rests with the financing company, who in turn is looking to utilize the remaining RV at the end of lease. On the other hand, customers taking a loan to buy the equipment outright must pay based on the total 100% value of the equipment. And they also assume the technology obsolescence risk relating to asset ownership. Leasing is an excellent way to acquire usage of the equipment without actually owning the equipment.
- II. A z10 EC upgrade can even lower a customer's existing monthly payments. When a customer decides to upgrade to a new z10 EC, the new lease on his older machine is extended to be co-terminus with the lease structured for the upgrade components. The benefit to the customer is that he now receives new technology easily and often can keep lease payments the same or actually lower his monthly payment.

Original Lease



(Monthly payment reduced)

z10 EC upgrade at Month 18



(Figure 1) 36 month original term, vs. 36 month extension with upgraded z10 EC technology.

In the above figure, the customer achieves a lower monthly payment through:

- Recalculation of the old base machine to a new 36th month term
- Inclusion of the cost of the new technology upgrade into the lease
- Resulting in the customer taking advantage of the economics of the enhanced residual value of the old base machine.

With IGF, the benefits of financing an upgrade need not be a one-time phenomenon. Customers have the ability to repeat the financing of an upgrade every time there is a IBM System z upgrade and reap the benefits.

Question:

What are the customer advantages to financing new z10 EC technology with IGF?

Answer:

Leasing with IGF:

- Reduces the total cost of ownership and improves price performance
 - Offers payment and term flexibility tailored to match either project or revenue generation milestones
 - Preserves customer cash and current credit lines for more strategic investments such as facility expansion, increased R&D, or receivable financing
 - Allows upgrades to future technology enhancements
 - Provides end of lease options that can offer seamless planning for replacement technology and disposal services for old equipment in an environmentally safe fashion
-
-

Question:

How can leasing with IGF assist a customer in meeting their growth objectives?

Answer:

Leasing helps a customer to grow by:

- Preserving cash and current credit lines for other strategic business needs
- Permitting customers to buy what they need for successful implementation, not just what they can afford in their current budget
- Providing access to new lines of credit to help accelerate project implementation
- Promoting upgrades as needed with timely lease extensions and upgrade financing

Question:

Can IGF help lower total cost of ownership?

Answer:

By leasing with IGF customers transfer technology risk and its implied cost to IGF, while preserving end of lease purchase, return or continuation options for themselves. By leasing rather than purchasing, the savvy customer can acquire the use of the equipment without actually owning the asset, leveraging residual value to drive total cost of 'ownership' downwards. This savings can be extended to follow on replacement equipment leases, providing a virtuous circle of TCO savings over multiple technology refresh cycles.

Question:

How does IGF support the New Enterprise Data Center?

Answer:

The New Enterprise Data Center is all about helping IT create business value for the enterprise. IBM has the products and services to help clients improve alignment of IT infrastructure with business and financial priorities. IGF can play a very valuable practical role assisting clients with funding solutions for New Enterprise Data Center engagements. IGF not only can offer leasing on new hardware, helping clients avoid technical obsolescence and improve their financial metrics, but it can also build these capabilities into a total financing line for a client, including financing of services and software delivered over time as the project evolves, helping the client manage overall cash flow and better align their investments to projected benefits. For example, IGF OIO financing and Project Financing offerings provide single-sourced, highly-customized financing across a project's entire lifecycle making it easier for clients to manage both up-front investment and ongoing operating costs. These financial offerings integrate all the elements of a total business transformation solution (hardware, software and services) under a single contractual umbrella, making it highly attractive to clients and a powerful competitive differentiator for IBM. Such custom sculpted cash flow structures from IGF can help accelerate project implementation, preserve cash and provide an additional layer of flexibility to ensure linkage of IT to business and financial priorities over time.

Question:

Can IGF finance IT services?

Answer:

Yes. IGF can finance credit qualified customers' acquisition of hardware, software and IT services either in standalone transactions or in conjunction with an equipment hardware lease or as part of a large infrastructure project. Regardless of scope or size customers can choose to include IT services in their monthly payment schedule.

Question:

Can IGF finance other products with z10 EC? (hardware, software and services)

Answer:

Yes, IGF can finance all the components of an IT solution for the customer. Components may be IBM products and services, or non IBM IT related products and services. Total solution financing from IGF offers one-stop convenient financing for credit qualified customers, regardless if they buy direct from IBM or from IBM Business Partners.

Question:

I have a major customer in the Financial Services Sector. Can't they access cheaper capital than IGF? Why would they want to lease?

Answer:

IGF typically borrows under the excellent IBM Corporation credit rating, but it is true that some banks, may have a cost of capital equal or less than IGF. However, banks also prefer to invest their capital in higher performing assets rather than in IT infrastructure which, by definition, is non performing. By using IGF not only do banks preserve capital for more lucrative investment opportunities such as business loans and supporting their own customer's needs, but they also remove any risk of technology obsolescence associated with rapidly evolving IT equipment.

Question:

How is IGF able to offer such competitive rates?

Answer:

For credit qualified customers IBM Global Financing (IGF) offers attractive lease and loan pricing, oftentimes at below market rates in selected markets, along with equitable fair contract terms and conditions. IGF is able to be a manufacturer's source of reliable credit for IBM customers looking to access technology at attractive rates primarily because of our steady, conservative and vigorous approach to risk, the strong credit rating of IBM, and our understanding of technology and proven ability to extract the most value out of that technology.

Question:

What offerings does IGF have for major infrastructure projects?

Answer:

For enterprise clients who are partnering with IBM on large infrastructure projects, IBM can assemble an Open Infrastructure Offering (OIO) that combines hardware, software, professional services, maintenance, financing and disaster recovery into a single, customized agreement. Every OIO agreement is tailored to a specific customer situation and provides the flexibility to make changes in response to changing needs, including the ability to substitute new technologies as needed. It combines a simplified acquisition process with a consolidated monthly bill. An OIO can reduce overall IT expenses and optimize asset management for in-place and future assets. A single contract, predefined interest rates and monthly billings reduce financial risks and make budget planning easier through improved IT expense forecasting.

Question:

How does IGF support business partners?

Answer:

IGF works with IBM business partners who wish to offer our financing programs to their customers. In most of the developed IT markets worldwide, IGF enables IBM Business Partners to offer IGF financing with a Web based tool called Rapid Online Financing (RoF). Partners can use RoF to credit qualify their customer, price out financing proposals, and deliver contracts, all within minutes. In addition, IGF supports IBM Business Partners with working capital inventory financing programs to assist them to do more business in their markets for the IBM Corporation.

Question:

How can customers engage IGF financing?

Answer:

Customers may contact IGF by engaging with their local IBM Client Rep or Business Partner; their supporting IGF financial sales executive, or by visiting www.ibm.com/financing.

Statements of Direction

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

Question:

What statements of direction has IBM made for 1x IB-DDR to be introduced to support extended distance for Coupling using InfiniBand?

Answer:

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

The common fiber optic cabling environment is designed to facilitate a migration from ISC-3 (2 Gbps link data rate) to 1x IB-DDR (5 Gbps link data rate) in a Parallel Sysplex environment on System z10.

Question:

What statements of direction has IBM made for phasing out ICB-4 links?

Answer:

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

Question:

What statements of direction has IBM made for support of CHPID type FCV to be removed?

Answer:

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

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

Question:

What statements of direction has IBM made for System z FICON features to no longer support buffer credits for performance at extended distance?

Answer:

The FICON Express4 features are intended to be the last features to support extended distance without performance degradation. IBM intends to not offer FICON features with buffer credits for performance at extended distances. Future FICON features are intended to support distances up to 10 km without performance degradation. Extended distance solutions may include FICON directors or switches (for buffer credit provision) or Dense Wave Division Multiplexers (for buffer credit simulation).

Question:

What statements of direction has IBM made about System z9 to support 12x IB-SDR on z9 EC and z9 BC general purpose servers?

Answer:

Support of 12x InfiniBand at single data rate (12x IB-SDR) - a link data rate of 3 GigaBytes per second (3 Gbps) - is planned to be made available on z9 EC and z9 BC servers whether the Processor Units (PUs) are defined as Central Processors (CPs), IFLs, zAAPs or zIIPs.

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

Question:

What statements of direction has IBM made for z/VM LPAR enhancements?

Answer:

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

Question:

What statements of direction has IBM made for additional support for managing z/VM systems?

Answer:

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

Question:

What statements of direction has IBM made for z/OS environments?

Answer:

IBM intends to expand z/OS and IBM System Storage support for Extended Address Volumes with larger volume sizes and to allow additional data set types to reside in the cylinders after the first 65,520 cylinders.

IBM intends to remove the following z/OS Communications Server function: Network Database (NDB) function, BIND DNS 4.9.3, Boot Information Negotiation Layer (BINL), and the DHCP server function.

IBM intends to provide support within z/OS that will allow authorized applications to query, change and perform basic operational procedures against the installed System z hardware base.

As part of its long-term commitment to simplifying z/OS and System z operations and management and increasing the efficiency of IT professionals, IBM intends to introduce an IBM z/OS Management Facility, which will be designed to provide the infrastructure, services and user interfaces to support a browser based graphical user interface needed to support a management console for z/OS. The z/OS Management Facility is intended to enable technology administrators and operators to more easily manage and administer a z/OS system. The IBM OMEGAMON z/OS Management Console provides important z/OS display and monitoring capabilities; it is planned to coexist and integrate with the z/OS Management Facility.



Copyright IBM Corporation 2008
IBM Corporation
New Orchard Rd.
Armonk, NY 10504

Printed in the United States of America, 02/2008

All Rights Reserved

IBM, the IBM logo, IBM eServer, BladeCenter, CICS, Cool Blue, DB2, DB2 Universal Database, DFSMS, DFSMSHsm, Domino, DRDA, DS8000, ESCON, FICON, FlashCopy, GDPS, HiperSockets, HyperSwap, IMS, Lotus, OMEGAMON, OS/390, Parallel Sysplex, PartnerWorld, Power Architecture, PR/SM, RACF, Resource Link, RMF, System Storage, System x, System z, System z9, System z10, Tivoli, VSE/ESA, VTAM, WebSphere, z9, z10, z/Architecture, z/OS, z/VM, z/VSE, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.