
**IBM System z9
Introduction
April 2007**

**IBM System z9 Enterprise Class Update
(formerly System z9 109)
Frequently Asked Questions**

Worldwide



Table of Contents

Announcement Overview	4
Community	9
z9 EC Hardware	11
z9 EC Upgradeability	17
Availability Enhancements	19
IFL, zAAP and zIIP specialty engines	23
Clarification on a few DB2 terms for better zIIP understanding	33
z9 EC Pricing	35
Performance	40
Networking Enhancements	43
I/O Enhancements	46
On/Off Capacity on Demand (On/Off CoD) and CBU	52
Cryptographic Enhancements	56
Parallel Sysplex Enhancements	65
GDPS Enhancements	66
Operating Systems	68
z/OS Version 1 Release 9 Preview	68
z/VSE Version 4 Release 1	71
z/VM Version 5 Release 3	76
Linux for System z	80
IBM Global Financing (IGF)	85
Global Technology Services	87
Implementation Services for Parallel Sysplex	90
Statements of Direction	91

Announcement Overview

Question:

What was announced on April 27, 2006?

Answer:

IBM announced extensions to the IBM System z9™ family. First, significant new capabilities were announced for the System z9 109 (z9-109), now known as the IBM System z9 Enterprise Class (z9 EC). These capabilities include subcapacity versions, delivered through 24 new capacity settings, and the general availability of a new specialty engine for data serving and data mining, the System z9 Integrated Information Processor (zIIP). Second, following on the success and market acceptance of the z9-109 high end offering, IBM introduced a new midrange mainframe server, the IBM System z9 Business Class (z9 BC). The z9 BC is available in two models, one with a low entry point of 26 MIPS. Both models provide a wide choice of capacity settings and full support of all mainframe specialty engines: Integrated Facility for Linux® (IFLs), Internal Coupling Facilities (ICFs), IBM System z™ Application Assist Processors (zAAPs) and IBM System z9 Integrated Information Processors (zIIPs), giving clients flexibility and choice in both initial deployment and future growth.

Question:

What was announced on October 10, 2006?

Answer:

IBM announced Server Time Protocol (STP) with planned availability of January 31, 2007. The STP feature is designed to allow events occurring in different z9 EC, z9 BC, z990 and z890 servers and Coupling Facilities (CFs), to be properly sequenced in time. STP is designed to reduce or eliminate the requirement for the Sysplex Timer®. STP supports a multisite sysplex up to 100 km (62 miles) without requiring an intermediate site. Previously, an intermediate site was recommended to locate one of the Sysplex Timers, when the multisite sysplex distance exceeded 40 km (25 miles). For more information on STP please look at the Web site: www.ibm.com/systems/z/pso/stp.html. For FAQs on STP please review the Web site www.ibm.com/systems/z/faq/. IBM also announced a new feature that will enable the System z9 Business Class (z9 BC) to operate on a non raised floor.

Question:

What was announced on April 18, 2007?

Answer:

Several enhanced features with availability on May 11, 2007. With these enhancements, the System z9 offers price/performance benefits when configuring a new z9 BC, as well as offering new common features for z9 EC and z9 BC to provide flexibility and performance as you extend your system to meet business needs.

This announcement includes:

- On demand enhancements include the ability to store up to 100 On/Off Capacity on Demand (On/Off CoD) configurations (records) on the z9 EC and z9 BC Support Element.
- LPAR group capacity limit function for improved management of LPARs.

- Power consumption and temperature monitoring capability.
- zPower estimation tool.
- Hardware decimal floating point facilities.
- FCP performance improvements.
- OSA-Express2: z/VM link aggregation and failover support for the z/VM® virtual switch (VSWITCH).
- OSA-Express2: Dynamic LAN idle (performance improvement).
- OSA-Express2 and OSA-Express: Layer 3 Virtual MAC.
- OSA-Express2: Network Traffic Analyzer.
- OSA-Express2: Network Trace Analyzer and QDIO Diagnostic Synchronization to help facilitate capturing of data for tracing and trapping problems.
- z/VM integrated systems management.
- Multiple IPL path retry for ESCON® and FICON features.
- Bolt-Down Kits for additional stability where required.
- Previewing additional exploiters of the System z™ zIIP Specialty Engine.
- Hardware decimal floating point facilities.
- Lightweight Directory Access Protocol (LDAP) support for HMC user authentication.
- System-initiated CHPID reconfiguration.

Additionally, z/VM 5.3 enhancements support guest use of the Modified Indirect Data Address Word (MIDAW) facility and provides simulation and virtualization support for specialty processors, the System z9 Integrated Information Processor (zIIP) and the System z Application Assist Processor (zAAP).

Question:

What are the planned availability dates of the enhancements announced on April 27, 2006, October 10, 2006 and April 18, 2007?

Answer:

May 26, 2006:

- The 24 capacity settings of the z9 EC Models, the zIIP, and FICON Express4

January 31, 2007:

- Server Time Protocol (STP)

May 11, 2007:

- April 18, 2007 announcement features and functions

May 15, 2007:

- HMC Dual Ethernet (#0084)

June 29, 2007:

- z/VM integrated systems management for z/VM V5.3

Third Quarter 2007:

- Dynamic Lan idle for z/OS and z/OS.e V1.8 with PTF

Question:

Why was the System z9 naming changed with the April 2006 announcement?

Answer:

Fundamentally, the naming approach for the System z9 product line or family is the same. The key element of the name remains the same, System z9 (or IBM System z when referring to the mainframe class or family of products including System z9 and IBM eServer™ zSeries® etc.). To accommodate the introduction of different System z9 offerings, what has changed is the means of identifying the particular family member. Previously, the only family member available was a high-end offering, System z9 109 (z9-109), which is now known as the System z9 Enterprise Class (z9 EC) going forward. The new member of the System z9 family for the midrange will be known as the System z9 Business Class (z9 BC). This change was ultimately made so that the different products now have a more descriptive term that helps position them to our clients.

Question:

What impact did the name change have for existing System z9 109 customers?

Answer:

Our customers should have no impact as a result of this change. The process for these customers to upgrade to take advantage of the capabilities of the System z9 Enterprise Class are not affected by the change in name, except that documentation may now be listed for the “System z9 EC formerly known as z9-109”. Redirects will be employed for key Web content featuring the z9-109 so that customers can have easy access to relevant information.

Question:

How does the Mainframe Charter and the October 2004 announcements influence the System z9 announcements?

Answer:

The Mainframe Charter, announced in August 22, 2003, provided a framework for planned future investment and highlighted IBM’s intentions for delivering ongoing value to our mainframe customers.

The introduction to the market of the IBM System z9 109 (z9-109) in July 2005 and enhancements now available on the z9 EC have brought significant technology advances in the areas of performance, scalability, availability, security and virtualization. The April announcement, which included the newest member of the family, the z9 BC, extended the System z9 family to provide our latest technology to our small and medium mainframe customers.

In addition to the innovation of the z9 EC and z9 BC, we also announced additional price / performance improvements and significant advances in our drive to increase mainframe skills and broaden the mainframe community. These changes will be covered later in the Community and Pricing section of the document.

Question:

How was the z9 EC improved over the System z9 109 that was announced July 26, 2005?

Answer:

The July 26, 2005 announcement of the IBM System z9 109 (z9-109) brought significant technology advances to deliver on our mainframe strategy, while also strengthening its leadership in enterprise computing. These leadership capabilities in areas such as scalability, availability, security and resiliency, intelligent workload management, integration and virtualization made the z9-109 an ideal platform for supporting central corporate databases and mission-critical enterprise-wide applications. With the z9 EC, we have strengthened and continued that leadership as we introduce additional capacity settings with subcapacity general purpose processors (CPs), announced the hardware availability of our IBM System z9 Integrated Information Processor (IBM zIIP), and continued to help improve our FICON performance and throughput.

Listening to our customers, and realizing their requirement for more granular capacity settings within the z9-109, the z9 EC offers subcapacity general purpose processors with a larger variety of capacity choices to meet their business needs.

In January 2006 we announced the latest specialty engine – the zIIP. The zIIP, when combined with the enabling PTFs for z/OS 1.6 and DB2 Universal Database™ (UDB) for z/OS V8, is designed to help improve resource utilization and lower the cost of eligible workloads, enhancing the role of the z9 EC as the data hub of the enterprise.

The z9 EC now includes FICON Express4 for potential improved I/O capacity and performance with the next generation of FICON/FCP. This feature may offer increased channel aggregation and shorter backup windows to help reduce the cost of storage operations and infrastructure.

Question:

Who would be interested in the z9 EC with subcapacity processors?

Answer:

Offered on our IBM S/390® and zSeries midrange servers since 1999, subcapacity processors have allowed customers to choose a server sized to best meet business requirements. Smaller incremental steps between capacity settings can allow customers to manage their growth as well as their costs, in smaller increments. On the z9 EC, subcapacity processors are available on servers with eight or fewer general purpose processors. With the z9 EC four book design, additional available processors within the server can be characterized as specialty engines (ICFs, IFLs, zAAPs or zIIPs) or can be used for Capacity Backup (CBU) capacity.

This new design means that customers will now have the option of selecting a server that meets both total capacity requirements as well as being able to select the number of general purpose processors that best fits their application and processing requirements.

Question:

What software is supported on the z9 EC?

Answer:

The following table provides the minimum OS support levels for the z9 EC.

OS Software	<ul style="list-style-type: none">▪ z/OS▪ z/VM▪ Linux on System z9 (64-bit and 31-bit distribution)▪ z/VSE™▪ TPF	<ul style="list-style-type: none">▪ z/OS 1.6 and subsequent releases*▪ z/VM 5.1 and subsequent releases*▪ Red Hat RHEL 4, Red Hat RHEL 5, SUSE SLES 9, SUSE SLES 10 ▪ z/VSE 3.1 and z/VSE 4.1▪ TPF 4.1 and z/TPF 1.1
----------------	--	--

* Note: Support for z/OS 1.4 and 1.5 ended on March 31, 2007. Support for z/VM 4.4 ended September 30, 2006. Support for VSE 2.7 ended February 28, 2007.

Question:

Where can I find the FAQs on the z9 BC and STP?

Answer:

The URL: www.ibm.com/systems/z/faq/ that carried this FAQ list will also have the z9 BC FAQs.

Question:

Where can I find the data sheets for the System z9 EC and z9 BC?

Answer:

The URLs are as follows:

z9 BC Data sheet

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

z9 EC Data sheet

ibm.com/common/ssi/fcgi-bin/ssialias?infotype=pm&subtype=sp&pdf=yes&appname=STG_ZS_USEN&htmlfid=ZSD01096USEN

Community

Question:

I know the IBM System z continues to speak of progress in building a mainframe community. What is the progress in mainframe skills growth?

Answer:

Significant progress has been made toward our goal of adding 20,000 new mainframe skilled professionals in the marketplace by 2010. The IBM Academic Initiative has helped educate – 10,000 skilled students to date. Our focus is to continue to add active mainframe curriculums to schools around the globe – over 230 schools are enrolled today – and foster partnering between schools and customers to not only build, but place skills. The program is thriving as evidenced by the increase in courseware, access by schools to the university hubs, growing university ambassadors and participating professors and schools. For more information, go to ibm.com/servers/eserver/zseries/about/charter/community.html

Question:

How can I ensure my mainframe personnel have the skills they need and the right learning roadmap for my company's specific needs?

Answer:

There are a comprehensive array of offerings and materials, classes and workshops offered today that customers can leverage to learn more about the platform and to get very specific skills, whether it is systems programming, DB2[®], Linux, or Java[™], and we continue to add training based on customer needs. A new class launched late last year — Introduction to z/OS and the IBM System z Environment — was designed for those new to the mainframe environment or involved in a server consolidation, and has been extraordinarily popular with customers. We also continue to add content and courses via the skills portal, accessible via the community Web site above. The most recent addition, based on customer demand, provides customers access to 12 Academic initiative courses, and beginner through advanced curriculum.

ibm.com/servers/eserver/zseries/about/charter/skills.html

Question:

What about applications on the platform? Will we have more flexibility and choice?

Answer:

There are over 1,200 partners across the globe providing applications, tools, and solutions for the mainframe in many industries such as banking, insurance and government. Linux in particular has had tremendous growth., with close to 300 ISVs and 850 applications available today. We have invested significantly in WebSphere[®] on z/OS working with ISVs to port and enable applications. Recent surveys of existing partners indicate the majority of partners are staying current on the latest of OS within 6-9 months of release and also are maintaining or increasing their investment in the mainframe.

Partners continue to exploit the mainframe as the data server of choice, and we see more and more partners are beginning to take advantage of the architecture with both data and application on the mainframe as well. You can get more information on available applications from your IBM specialist or Business Partner or by searching the IBM Global Solutions Directory.

z9 EC Hardware

Question:

What is the machine number of the z9 EC?

Answer:

The machine number has not changed from the z9-109. It will still be 2094.

Question:

What z9 EC models are available?

Answer:

The following models were announced on July 26, 2005:

- A z9-109 S08 model can be a 1-way through 8-way - which means there are 8 processor units or PUs contained on one book.
- A z9-109 S18 model can be a 1-way through 18-way (18 PUs) contained on two books.
- A z9-109 S28 model can be a 1-way through 28-way (28 PUs) contained on three books.
- A z9-109 S38 model can be a 1-way through 38-way (38 PUs) contained on four books.
- A z9-109 S54 model can be a 1-way through 54-way (54 PUs) contained on four books.

The z9 EC will provide all these same models.

The PUs can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z9 Integrated Information Processors (zIIPs), additional System Assist Processors (SAPs), or used as additional spares.

Only eight subcapacity processors can be active on the server (and it doesn't matter which model you have). When more than eight CPs have been purchased on servers that have more than one book, a selection can be made to activate only 8 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 eight or fewer general purpose processors.

Question:

How does the z9 EC report model number and CPU ID information?

Answer:

Model information is returned via the STSI instruction. The Model Capacity Identifier field in SYSIB 1.1.1 contains model names in the form nxx, where 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 active general purpose processors (CPs). The Store CPU ID version code, which was used prior to the IBM eServer zSeries 900 (z900) to reflect the model, is always zero for the z9 EC. The Node Descriptor also reports the model, and there are console and z/OS commands which display the node descriptor.

Question

What is different about the Model S54?

Answer:

The Model S54 is an enhanced capacity model. The z9 EC is fully populated with four books and 54 PUs. You can customize the machine to be a 1 to 54-way. On the standard models, the first book has 8 configurable PUs, with 10 PUs each in books 2, 3 and 4. The S54 has 12 configurable PUs in the first per book, and 14 configurable ones in each remaining book. Like the S38, the S54 can be ordered with a minimum of 16 GB of memory up to a maximum of 512 GB.

Question:

How many spare processing units are on the z9 EC?

Answer:

There are two spare processing units on the z9 EC. These spares can be shared across the books.

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 z Application Assist Processors (zAAPs), System z9 Integrated Information Processors (zIIPs) or Internal Coupling Facility (ICFs).

Question:

How many System Assist Processors (SAPs) are on the z9 EC?

Answer:

The answer depends on the model. The minimum number of SAPs for each model is as follows:

- The z9 EC Model S08 has two.
- The z9 EC Model S18 has four.
- The z9 EC Model S28 has six.
- The z9 EC Model S38 has eight.
- The z9 EC Model S54 has eight.

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

Question:

Can I buy a z9 EC that has only IFL or ICF processors without including a general purpose processor (CP)?

Answer:

Yes. Similar to the IBM eServer zSeries 990 (z990), you can order only IFLs or ICFs in a z9 EC, using a software model of 700 with 1 to 54 IFLs or a maximum of 16 ICFs.

Question:

Will the z9 EC support Token-Ring on the Hardware Management Console (HMC)?

Answer:

While the Token-Ring is not offered as a new orderable feature on the z9 EC, current HMCs with Token-Ring may be carried forward to the z9 EC on an upgrade from a IBM eServer zSeries 900 (z900) or z990.

Question:

Will the z9 EC support carrying forward the OSA-Express Token-Ring?

Answer:

No. OSA-Express Token-Ring is not offered on the z9 EC and cannot be carried forward to the z9 EC on an upgrade from the z900 or z990.

Question:

What is the automation enhancement for LPARs?

Answer:

Customers can now use automation to change/adjust processor running time via expanded API parameters that have been expanded for task "Change LPAR Controls."

Question:

What is LPAR Group Capacity Limit?

Answer:

With the logical partition (LPAR) group capacity limit on z9 EC and z9 BC, you can now specify LPAR group capacity limits allowing you to define each LPAR with its own capacity and one or more groups of LPARs on a server. This is designed to allow z/OS® to manage the groups in such a way that the sum of the LPARs' CPU utilization within a group will not exceed the group's defined capacity. Each LPAR in a group can still optionally continue to define an individual LPAR capacity limit.

LPAR group capacity limit requires that all LPARs managed in the group are running at z/OS or z/OS.e V1.8 or later. LPAR group capacity limits may help provision a portion of a System z9 server to a group of LPARs allowing the CPU resources to float more readily between those LPARs, resulting in more productive use of "white space" and higher server utilization. Refer to the Software requirements section.

Question:

How does LPAR Group Capacity Limit work?

Answer:

It adds capability to specify capacity of the group of LPARs in MSUs per hour. PR/SM™ and WLM work together to enforce the capacity defined for the group or to enforce the capacity optionally defined for each individual LPAR. Group Capacity Limit may provide better control of CP resource consumed for WLC pricing and is exclusive to System z9. Requires at a minimum: z/OS or z/OS.e1.8

Question:

How does this work across a Sysplex?

Answer:

LPARs defined as part of a group may belong to the same sysplex or to different sysplexes. There is no requirement that all members of a sysplex be in a group, or that any members of a sysplex that are in a group be defined in the same one. If all members of sysplex reside on the same server, though, they can all be part of the same group. A group is defined for a CEC using the HMC. A group may not cross CEC boundaries.

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 or to a non-raised floor environment.

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 or non-raised floor installation.

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 z9 EC, you need a quantity of two Bolt-Down kits and for z9 BC, a quantity of one Bolt-Down kit.

(#7995) - Bolt-Down Kit, High-Raised Floor 2084 2086 2094 2096 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.

(#7996) - Bolt-Down Kit, Low-Raised Floor 2084 2086 2094 2096 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.

(#7997) - Bolt-Down Kit, Non Raised Floor 2096 2086 This feature provides frame stabilization and bolt-down hardware to help secure a frame to a concrete floor.

Question:

What is the purpose of the Power Monitor and Power Estimation capabilities?

Answer:

They are designed to help provide information about power consumption for currently installed System z9 servers and to help power and cooling planning for new System z9 servers or for upgrades to currently installed System z9 servers.

Question:

Are these capabilities a new function or a tool?

Answer:

This new tool helps monitor the power consumption and temperature of the system. The System Activity Display 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.

System z9 also provides a tool on IBM Resource Link™ which provides the user an estimate as to 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 this system.

Question:

What benefits does LDAP support for HMC user authentication provide?

Answer:

This is designed to help you manage your HMC user IDs and passwords by providing LDAP (Lightweight Directory Access Protocol) support for HMC user authentication. This allows configuring HMCs to use an LDAP server to perform userid/password authentication at logon time for additional flexibility in systems management.

Question:

How does this function work?

Answer:

This function allows the use of the current user ID and password policy for HMC user IDs and passwords, and provides one centralized user ID and password control mechanism to help meet the user's corporate security guidelines.

The user ID is defined on the HMC along with the roles to be given to the user ID. HMC settings related to the userid will continue to reside on the HMC, and the LDAP directory will be used to authenticate the user, thus reducing the need to store the user ID's password locally. SSL and non-SSL connections to the LDAP server are supported.

This function is designed to assist system administrators to easily create HMC user IDs matching existing company user names, thus reducing the need to create and distribute passwords, since this is already being managed by the corporate control mechanism.

Question:

What is z/VM integrated systems management?

Answer:

z/VM integrated systems management is designed to allow you to perform hardware and basic operating system management seamlessly using the HMC Web browser-based user interface. This may be especially useful in providing ease of use for customers new to z/VM or System z by providing the capability to manage guests without having to establish any network connections or carry out complex configuration of the system.

The HMC can automatically detect z/VM guests. The z/VM integrated systems management functions supported are profile management (disk, network adapter and memory management) and guest management (activate and deactivate guests and display guest status). HMC level 2.9.2 plus MCLs from RETAIN are required.

Question:

What new support does z/VM 5.3 provide for virtualization of the System z9 hardware resources?

Answer:

New z/VM 5.3 enhancements support guest use of the Modified Indirect Data Address Word (MIDAW) facility, which can allow more flexibility for z/VM guests and performance in certain channel programs as an alternative to data-chained channel-command words (CCWs). z/VM support for guest use of MIDAWs can allow operating systems such as z/OS to use this new aspect of z/Architecture® without regard to whether the operating systems are running in a logical partition or a virtual machine.

z/VM 5.3 also provides simulation and virtualization support for specialty processors, the System z9 Integrated Information Processor (zIIP), and the System z Application Assist Processor (zAAP), providing test and production guest support for workloads that utilize these engines. These new functions are designed to help clients extend the business value of the mainframe virtualization technology for existing and new workloads.

z9 EC Upgradeability

Question:

What are the upgrade paths from an IBM eServer zSeries 900 (z900) or a z990 to the z9 EC?

Answer:

All models of the z900 (except for the Model 100) and the z990 can upgrade directly to a z9 EC.

Question:

What are the upgrade paths from a z9 BC to the z9 EC?

Answer:

Any capacity setting on the z9 BC Model S07 can upgrade to the z9 EC Model S08.

Question:

Can I upgrade any z9 EC to a Model S54?

Answer:

Yes, but it is important to note that upgrading other z9 EC models to the S54 will be disruptive due to the difference in available PUs on the S54.

Question:

Can I upgrade from my current full capacity z9-109 to a subcapacity z9 EC?

Answer:

Yes. Each of first eight general purpose processors on the z9 EC can be divided into eight sub-units. This creates an 8 by 4 matrix of settings. As long as upgrades are positive MIPS growth, you can move around anywhere within the matrix when adding capacity. When your number of general purpose processors exceeds eight, then all of the general purpose processors must be full capacity.

Question:

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

Answer:

Yes. But remember that only eight subcapacity processors can be active on the server (and it doesn't matter which model you have). When more than eight CPs have been purchased on servers that have more than one book, a selection can be made to activate only 8 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 eight or fewer general purpose processors.

Question:

What are the migration requirements for customers with Crypto Express2 feature on z890 and z990 servers?

Answer:

The Crypto Express2 feature is supported on z990 servers and can be carried forward on an upgrade to the z9 EC.

Question:

Can I perform a miscellaneous equipment specifications (MES) upgrade involving Crypto Express2 features on the z9 EC without an outage?

Answer:

The design of the Crypto Express2 feature allows for nondisruptive upgrades. Unfortunately, there are some environments where a disruptive upgrade will occur: for example; a z9 EC with insufficient I/O slots where an additional I/O cage is required.

Availability Enhancements

Question:

What is enhanced book availability and redundant I/O interconnect that are available on the z9 EC?

Answer:

Enhanced book availability is an extension of the support for Concurrent Book Add (CBA) delivered on the z990. With proper planning, CBA is designed to allow you to concurrently upgrade a z9 EC by integrating a second, third, or fourth book into the server without affecting application processing.

The z9 EC goes an additional step by allowing a single book in a multibook server to be concurrently removed from the server and reinstalled during an upgrade or repair action. Redundant I/O interconnect provides connectivity to the server I/O resources using a second path from a different book during the action.

To help minimize the impact on current workloads and applications, you should ensure that you have sufficient inactive physical resources on the remaining books to complete a book removal. For maximum availability, it is recommended that z9 EC be purchased with one additional book. Use of enhance book availability may require acquisition of additional hardware resources.

Question:

What is the advantage of enhanced driver maintenance on the z9 EC?

Answer:

Some of the greatest contributors to downtime during planned outages are Licensed Internal Code (LIC) updates performed in support of new features and functions. When properly configured, the z9 EC is designed to support activating select new LIC levels concurrently. Concurrent activation of the select new LIC level is only supported at specific sync points. Sync points may exist throughout the life of the current LIC level. Once a sync point has passed, you will be required to wait until the next sync point supporting concurrent activation of a new LIC level. Certain LIC updates will not be supported by this function.

Enhanced driver maintenance, exclusive to the System z9, is another step IBM is taking to help reduce the duration of a planned outage.

Question:

What is the advantage of dynamic oscillator switchover?

Answer:

The z9 EC has two oscillator cards, a primary and a backup. In the event of a failure of the primary oscillator card, the backup is designed to detect the failure, switch over, and provide the clock signal to the server transparently. Previously, in the event of a failure of the active oscillator, a server outage would occur, the subsequent Power On Reset would select the backup, and the server would resume operation. Dynamic oscillator switchover is exclusive to the System z9 family.

Question:

What is the flexible memory option on z9 EC?

Answer:

The flexible memory features (FC #2802 - #2824), are offered in 16 GB increments from 32 GB to 384 GB. The features provide resources to support activation of the actual purchased memory increment in the event of a single book failure or to be available during an enhanced book availability action. Flexible memory features are only available on multi-book servers (the S18, S28, S38 and S54 models).

Question:

What is concurrent MBA fanout card hot-plug?

Answer:

A Memory Bus Adapter (MBA) fanout card is designed to provide the path for data between memory and I/O using Self-Timed Interconnect (STI) cables. With the introduction of the z9-109, and now the z9 EC, a hot-pluggable and concurrently upgradeable MBA fanout card became available. Up to 8 MBA fanout cards are available per book for a total of up to 32 MBA fanout cards on the z9 EC when four books are installed. In the event of an outage, an MBA fanout card, used for I/O, may be concurrently repaired using redundant I/O interconnect. By contrast, the MBA cards on the z990, located internal to the Multichip Module (MCM), cannot be hot-plugged and do not support concurrent repair.

Question:

What is the program directed re-IPL feature of the z9 EC?

Answer:

Program directed re-IPL is designed to allow Linux on System z to re-IPL without operator intervention. Linux on System z is designed to identify how it was IPL'd from the load device. Program directed re-IPL may request (by calling a Licensed Internal Code function) that it be reloaded from the same load device, using the same load parameters. In this way, program directed re-IPL allows Linux on System z running natively in an LPAR to trigger a re-IPL. This re-IPL is supported for both SCSI and ECKD™ devices.

z/VSE 4.1 also supports program-directed re-IPL for FCP-attached SCSI disks running natively in an LPAR. This allows an automatic installation of z/VSE on FCP-attached SCSI disks without operator intervention to re_IPI the z/VSE system.

Question:

What is Multipath Initial Program Load (IPL)?

Answer:

Multipath IPL is another enhancement in System z availability. Multipath IPL is designed to attempt the IPL on another path(s) if available, until all paths have been attempted or until the operation is successful. Without this capability, if an I/O error occurs during an IPL, the IPL will fail.

Multipath IPL is applicable to ESCON channels (CHPID type CNC) and to FICON channels (CHPID type FC). z/OS 1.6 plus PTF at a minimum on System z9 provides support for this enhancement.

Question:

What is the “Small Form Factor Pluggable (SFP)” availability feature of the FICON Express4 adapter?

Answer:

The fiber optic feature on a FICON Express4 may be swapped one at a time with a new like module, should there be a single fiber optic module failure. This swap on a single channel does not affect the traffic flowing on the other channels, and may be performed concurrently, while the other three channels are running.

Question:

What is the improved FICON error recovery feature?

Answer:

The improved FICON error recovery is designed to provide System z9 and z/OS 1.7 I/O recovery processing improvements to allow for the system to detect switch/director fabric problems that may cause FICON links to fail and recover multiple times in a short period of time.

This enhancement is designed to allow the system to detect these conditions and keep an affected path off line until an operator action is taken. This is expected to help limit the performance impacts of switch/director fabric problems.

Question:

What are the improved Channel Subsystem enhancements?

Answer:

IBM has implemented several enhancements designed to increase the availability of ESCON and FICON channels, Open Systems Adapter (OSA) ports, and InterSystem Channel-3 (ISC-3) links to facilitate recovery from errors, should they occur. This is improved hardware function and is transparent to operating systems.

Question:

What is System-Initiated CHPID reconfiguration?

Answer:

This function is designed to reduce operator effort and potentially the duration with a repair action when an ESCON or FICON channel, an OSA port, or an ISC-3 link is shared across logical partitions (LPARs) must be varied offline. This adds function to the System z9 to signal each LPAR sharing a CHPID that is to be configured physically Offline or Online from the HMC. z/OS and Linux on System z9 provide support for this enhancement.

For example, with CHPIDs being shared across multiple channel subsystems for 10 LPARs, performing a repair and verify operation on an ESCON feature could require issuing of commands in up to 10 partitions for up to 14 different channels, $14 \times 10 = 140$ different operator commands. This function may mitigate this high number of operator actions.

Support is provided by z/OS 1.6 with PTF at a minimum. For Linux on System z9, IBM is working with its distribution partners to provide this function in future distribution releases or service updates.

*Note: *All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.*

Question:

What are OSA-Express Network Trace Analyzer and QDIO Diagnostic Synchronization?

Answer:

To help facilitate capturing of data for tracing and trapping problems, the System z9 offers systems programmers and network administrators the ability to more easily solve network problems. With the introduction of the OSA-Express Network Trace Analyzer and QDIO Diagnostic Synchronization exclusively on the System z9, customers will have the ability to capture trace/trap data and forward it to z/OS 1.8 tools for easier problem determination and resolution.

The Network Traffic Analyzer function helps to provide a new diagnostic trace facility, allowing OSA packet trace records to be sent directly to the host operating system. Using Network Traffic Analyzer, along with the existing z/OS tools and service aids, z/OS system administrators will be able to monitor for ("sniff") each OSA CHPID for the network traffic.

Network Traffic Analyzer is designed to extend the serviceability features of OSA, helping provide system programmers and network administrators the capability to control, capture, and format OSA hardware packet trace information.

For OSA-Express2 features (CHPID type OSD) on z9 EC and z9 BC, requires at a minimum:

- z/OS 1.8
- z/VM 5.1 for guest exploitation

IFL, zAAP and zIIP specialty engines

Question:

What new uses of the Specialty Engines are being previewed?

Answer:

On April 18, 2007, IBM is previewing:

- an enhancement to the z/OS Communications Server that allows the IPSec processing to take advantage of zIIPs
-
-

Question:

Why is IBM extending the z/OS Communications Server to include the use of zIIP assisted IPSec?

Answer:

Network encryption is becoming more important. Business demands to protect sensitive data on the Internet have increased the requirement for users to implement end-to-end encryption on Virtual Private Networks (VPNs). At the same time as businesses are seeing an increased need for data protection, they are also increasing their use of more open networks such as the Internet. Moving secure business data and transactions from proprietary, trusted networks to more open networks creates an ever-increasing need for new technologies to protect this data.

IPSec (IP Security) can help protect data and transactions flowing over the Internet. IPSec is an open networking standard used to create security-rich connections between two points in an enterprise - this may be server-to-server, or server to network device, as long as they support the IPSec standard.

The IPSec support was integrated into z/OS Communications Server in z/OS V1.7 and provides authentication, integrity, and data privacy from z/OS to other network endpoints that support IPSec. In addition to allowing you to run host-based IPSec for secure end-to-end network flows, the z/OS V1.7 IPSec added IP filtering to protect your host. Since the IPSec support is implemented in the IP protocol layer, you can use it to help protect a variety of network traffic types to/from any application without any anticipated change to that application.

Question:

What is being previewed for the z/OS Communications Server to take advantage of the zIIP engine?

Answer:

Previewed is an enhancement to the z/OS Communications Server that allows the IPSec processing to take advantage of zIIPs, when present. The zIIP, in effect, will be a high-speed encryption engine that is designed to provide better price/performance. The new IPSec zIIP Assist function, previewed here, is intended to move most of the IPSec processing from the general purpose processors to the zIIPs. In addition to performing the encryption processing, the zIIP will also handle cryptographic validation of message integrity, and IPSec header processing. This is designed to allow you to take advantage of the cost saving benefits of the

zIIP when you implement IPsec to secure your valuable business transactions and to protect your host.

Specifically, the z/OS Communication Server (z/OS CS) is designed to interact with z/OS Workload Manager to have all of its enclave Service Request Block (SRB) work made eligible to run on the zIIP.

This capability is planned to be available August 2007 for z/OS V1.8 with PTFs and available native on z/OS V1.9 when available.*

*Note: * All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.*

Question:

Why is IPsec (network encryption) processing chosen to be eligible for execution on the zIIP engine?

Answer:

The zIIP is designed to help integrate data and data transaction processing across the enterprise and to help improve resource optimization and lower the cost of ownership for eligible data and transaction processing workloads.

Securing your data and transaction processing with end-to-end encryption with IPsec is a logical extension towards the goal of integrating the use and consolidation of data across the enterprise. Furthermore, zIIP assisted IPsec helps position IBM System z9 as a cost-effective server in environments requiring end-to-end security for IP network traffic.

Question:

What is IPsec and when would you use it ?

Answer:

IPsec (IP security) is a suite of protocols for securing Internet Protocol (IP) communications by authenticating and/or encrypting each IP packet in a data stream. IPsec also includes protocols for cryptographic key establishment. It helps allow a secure tunnel between two IP entities – helps create end-to-end network encryption.

End-to-end network encryption is becoming more pervasive due to regulatory and enterprise security policy. Furthermore, end-to-end network encryption is also becoming a requirement for companies that outsource some part of their network and want to control access to confidential data.

You can use IPsec to protect data virtually everywhere, for example: over the Internet, over an internal network, and within the data center. IPsec provides for flexible configuration that protects all connections or specific connections between a server to z/OS server, client to z/OS server, branch office gateway to z/OS server, and more. IPsec is ideal for protecting sensitive customer financial data, banking transactions, retail transactions, insurance data.

z/OS Communication Server IPsec function can secure enterprise extender (EE) traffic carrying SNA data, file transfers, CICS and DB2. It is also readily deployable when IT security policy demands encryption due to compliance regulations because there are no anticipated application changes required.

Question:

How does the z/OS Communication Server enable IPsec to be redirected to a zIIP?

Answer:

The z/OS Communication Server (z/OS CS) is planned to interact with z/OS Workload Manager to have all of its enclave Service Request Block (SRB) work directed to zIIP.

A single configuration statement (GLOBALCONFIG ZIIP IPSECURITY) within the TCP/IP profile triggers z/OS CS to request z/OS to direct this IPsec Enclave SRB processing to available zIIPs.

For more information, see the "Capacity Planning for zIIP-Assisted IPsec" whitepaper planned for August 2007.

Question:

DB2 for z/OS V8 DRDA workloads are already eligible for zIIP, are there considerations for running DRDA and IPsec together on the same zIIP?

Answer:

zIIP workload is subject to the same z/OS tuning options as other workload. Customers will be able to prioritize DRDA and IPsec workload to meet their processing needs. Often, IPsec tends to have longer running units of work. If you plan to mix IPsec and DRDA on a the same zIIP, it may be beneficial to set IPsec workloads to a lower importance and lower execution velocity than that of DRDA (by adjusting the WLM default service class from SYSOTHER).

For more information, see the "Capacity Planning for zIIP-Assisted IPsec" whitepaper planned for August 2007.

Question:

Where can I get more information on the zIIP assisted IPsec?

Answer:

For more information on zIIP assisted IPsec please go to the following Web sites:

z/OS Communications Server page

- ibm.com/software/network/commserver/zos/security/

zIIP page

- ibm.com/systems/z/ziip/

White paper - "Capacity Planning for zIIP-Assisted IPsec"

- ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP100988
- ibm.com/support/docview.wss?rs=852&uid=swg27009459

Or see your IBM representative.

Question:

Can I order subcapacity specialty engines?

Answer:

No. All specialty engines on the z9 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 z9 EC?

Answer:

The increased power of the z9 EC processor combined with other architectural advances like the increased number of Logical Partitions (LPARs) and the greatly improved internal I/O throughput gives the z9 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 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. Some environments may realize lower and other environments higher improvements.

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 z9 server that is made up entirely of IFL processors and no general purpose processors?

Answer:

Yes. A dedicated System z9 Linux server is available, meaning that all processors are IFLs on this System z9 server. Please see your IBM or Business Partner representative for more information on it.

Question:

What is a System z Application Assist Processor (zAAP) and the value to a customer?

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 (when available) 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:

How does Java running on zSeries benefit from the z9 EC?

Answer:

The increased power of the z9 EC processor combined with other architectural advances like the increased number of Logical Partitions (LPARs) and the greatly improved internal I/O throughput gives the z9 EC the ability to run increased workloads.

Since the price of the System z Application Assist Processors (zAAPs) for this solution package is similar, whether you buy it for the z9 EC or z990, this can mean customers may provide greater performance for the same number of users at a similar price.

Question:

What is the warranty period for System z9 Application Assist Processor (zAAP)? And what is the price of maintenance service for zAAPs?

Answer:

During the first 12 months following the installation of a z9 EC machine, zAAPs are under warranty. After the first 12 months the maintenance price for the zAAPs will be similar to the maintenance price for IFLs (Integrated Facilities for Linux).

Question:

What is the value proposition of the IBM zIIP?

Answer:

IBM zIIPs are intended to help customers adopt new data, transaction, and information management technology on z/OS. They are designed to help free up general purpose processor (CP) capacity that then may be utilized by other workloads.

The zIIP may help you to leverage DB2 for z/OS qualities of service for data access and information management across your enterprise by helping to make direct access to eligible workloads more cost effective and potentially reducing the need for many local copies of the data and the complexity that brings.

The zIIP may also help you to leverage IPSec network encryption to help protect your data and transactions as they flow across your network.

Question:

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

Answer:

Your new z9 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, will be provided by z/OS (and z/OS.e) V1.6 or V1.7 via a Web download. The first exploiter of zIIP will be DB2 UDB for z/OS V8 with enabling PTFs. 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 will IBM zIIPs work?

Answer:

The IBM zIIP's execution environment will accept eligible work from z/OS (and z/OS.e) 1.6 or 1.7, which will manage and direct the work between the general purpose processor and the zIIP.

The IBM zIIP is designed so that a program can work with z/OS (and z/OS.e) V1.6 or V1.7 to have all or a portion of its enclave Service Request Block (SRB) dispatched work directed to the zIIP. The z/OS operating system, acting on the direction of the program running in enclave SRB mode, controls the distribution of the work between the general purpose processor and the zIIP. Using a zIIP can help free up capacity on the general purpose processor.

Question:

How does DB2 for z/OS V8 work with zIIPs?

Answer:

The IBM zIIP is designed so that a program can work with z/OS 1.6 or 1.7 to have all or a portion of its enclave Service Request Block (SRB) dispatched work directed to the zIIP. The z/OS operating system, acting on the direction of the program running in enclave SRB mode, controls the distribution of the work between the general purpose processor and the zIIP. Using a zIIP can help free up capacity on the general purpose processor. When the zIIP is available, DB2 UDB for z/OS V8 will exploit the zIIP by sending eligible work to z/OS 1.6 or 1.7 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 z9 (or z/OS) may access a DB2 UDB for z/OS V8 database that is hosted on a z9 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 z9 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 System z9 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:

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

Answer:

Once a zIIP is installed (along with z/OS or z/OS.e 1.6 or 1.7 (with PTF) 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.
-
-

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:

If Linux on System z9 is running in an LPAR on the z9 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:

What can be the advantage of having the management of specialty engines in separate pools on z9 EC?

Answer:

On previous systems all ICFs, IFLs, and zAAPs were managed as a single type of physical resource on the server (as if they were all ICFs). This complicated capacity planning on machines that have more than one type of these processors installed and treated them as shared processors. It was difficult for customers to keep the shared logical resources within the bounds of the proper physical resources since they were all managed as this single pool of shared resources. And, they all showed up in RMF as ICFs. By separating the management of these resources on the System z9 with this new support, all logical resources (ICFs, IFLs, zAAPs and now the zIIPs) will be managed separately within their corresponding physical resource pool. RMF will report all four types of processors (including general purpose processors which have not changed here) at both the logical partition level and the physical resource level. Note that in previous machines, you were able to specify processing weights for shared logical ICFs in coupling facility partitions and/or shared IFLs in Linux mode partitions but these weights and management would all be grouped together in the physical "ICF" resource pool. By separating the "ICFs" into their separate pools and managing them independently, customers can accurately make capacity planning decisions on the correct processor type needs according to how the processors are truly being utilized.

Question:

What will be the advantage of the improved LPAR weight management of CPs, zAAPs and now the zIIP on a System z9?

Answer:

On previous systems, the introduction of zAAP created the first Logical Partition (LPAR) type that could have multiple types of shared logical processors; namely shared general purposed processors and shared zAAPs. What was not available was the ability to set a different processing weight value for the zAAPs in an LPAR. Whatever weight was set for the shared general purpose processor was effectively replicated for the shared zAAPs in the LPAR. The System z9 support allows a separate specification of a processing weight for the zAAPs, and now zIIPs in an LPAR, thus allowing flexible changes to the independent weight values within their respective processing pools.

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 monitor work priorities on zIIPs and zAAPs so that goals can be met.

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 z Application Assist Processors (zAAPs) and System z9 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 System z9 (z9 EC and z9 BC) servers. zAAPs can be simulated only on 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.

Clarification on a few DB2 terms for better zIIP understanding

Question:

What are enclave SRBs? And how are they related in the zIIP discussion?

Answer:

z/OS dispatches work in either TCB (Task Control Block) or SRB (Service Request Block) mode. DB2 parallel tasks use SRB mode and are assigned the same importance as the originating address space.

Preemptable enclaves are used to do the work on behalf of the originating TCB or SRB address space. Enclaves are grouped by common characteristics and service requests and since they are preemptible, the z/OS dispatcher (and WLM) can interrupt these tasks for more important ones (i.e. Manage a transaction end-to-end). There are two types of preemptible SRBs: client SRBs and enclave SRBs.

If a DB2 UDB for z/OS V8 request is coming in over distributed connection (i.e. DRDA over TCP/IP) then most of their work is executed in enclave SRBs.

If the request is coming over local / native connection, then that work is dispatched between TCBs, client SRBs, and enclave SRBs (star schema parallel queries and some utility index maintenance now use enclave SRBs).

Only the enclave SRB work (not the client SRB work or non-preemptible SRB or TCB work) is eligible to be redirected to the zIIP. DB2 UDB for z/OS V8 knows how its work is dispatched and directs z/OS 1.6 or later to dispatch (redirect) a portion of the eligible work to the zIIP.

Question:

What is a star schema parallel query?

Answer:

The star schema is the simplest data warehousing data model. It consists of a single “fact table” at the center and the “dimension tables” at the points, hence the star term. A fact table typically contains metrics/numeric data about the performance of a business process by a business entity over a period of time. Sales History is a good example of this. The who, what, where, and when associated with the business measurements are typically encoded data columns (dimensions) in the fact table. Each dimension table contains a set of related codes and additional descriptive data columns about each code. For example a sales office dimension table could include office code, office name, street address, city, state, zip and country. A product dimension table could include product name, price, size, weight and manufacturer.

To create a meaningful management sales report, one could present performance data from the fact table and translate coded data from the dimension tables in the result rather than just present coded data from the fact table. A single (joined) query is typically submitted to DB2 to retrieve the data. Fact tables can have millions, even billions of rows, and queries may need to process a sizable portion of the table. To reduce query elapsed time, DB2 will assess whether

the single query should be executed in parallel by multiple threads with minimal contention, and if possible, will split the single query to execute concurrently on multiple processors.

A star schema parallel query is a single query (SELECT statement) that runs against tables with a star schema relationship that DB2 chooses to execute in parallel on multiple processors. The zIIP could be utilized as one of the processors.

z9 EC Pricing

Question:

Is hardware pricing for z9 EC similar to pricing on z990?

Answer:

Yes, hardware pricing for z9 EC is very similar to z990 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:

What is the memory pricing on the z9 EC?

Answer:

System z9 EC memory will be priced as low as \$8K per GB.* (* This is the USA price, and prices will vary by country.)

Question:

How will diagonal upgrades from z990 (or z900) to z9 EC be priced?

Answer:

The platform's value proposition of net pricing diagonal upgrades (upgrades from z900 or z990) will not change. As in the past, microprocessor and memory feature conversions will be the key net priced items.

Question:

How will maintenance be priced on z9 EC?

Answer:

Monthly 24x7 ESA list prices will exist and applied against the z9 EC feature codes similar to 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:

Is software pricing on z9 EC similar to software pricing on z990?

Answer:

The approach for pricing software on the z9 EC is similar to pricing on the z990, but, in line with the pricing initiatives announced with the Mainframe Charter in August 2003. The announced MSUs for software pricing purposes on the z9 EC are on average 10% less than those for z990. With this, in many cases, software savings can be realized on the z9 EC versus the z990.

Question:

Is software pricing on z9 EC similar to software pricing on z900?

Answer:

The approach for pricing software on the z9 EC is similar to pricing on the z900, but, in line with the pricing initiatives announced with the Mainframe Charter in August 2003. The announced MSUs for software pricing purposes on the z9 EC are on average 19 less than those for z900. (The z990 offered 10% savings in software pricing MSUs when coming from a z900 and the z9 EC has an additional 10% savings for a total average of 19% over the z900). With this, in many cases, software savings can be realized on the z9 EC versus the z900.

Question:

How does a customer pay for software on the z9 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 z9 EC with the z/OS operating system. Subcapacity WLC is available on a z9 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 z9 EC in all environments, including a standalone (non-coupled) z9 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 you are running eligible 'new qualified workloads' such as WebSphere Application Server, Lotus[®] Domino[®], SAP, PeopleSoft, and Siebel. The zNALC offering continues the IBM commitment to sub-capacity 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 z9 EC server will qualify for the same announced pricing terms available on z990. 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 IBM System z9 Enterprise Class (z9 EC) and the IBM System z9 Business Class (z9 BC) servers 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 z9 EC and z9 BC 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"? Is this something new for z9 EC?

Answer:

No, software models are not new for z9 EC. Like the previous IBM z990 mainframe servers, the machine type and model of the z9 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:

Why announce zNALC?

Answer:

Currently IBM has two reduced z/OS priced offers – z/OS with NALC and z/OS.e. Both of these offers provide reduced pricing for z/OS in support of eligible new workloads running on z/OS. But, both offers have different structures and terms and conditions. Though zNALC provides similar pricing to the current NALC and z/OS.e offers, it also offers several benefits over these current new workload pricing offers.

zNALC will be available on all z/Architecture servers (z800, z900, z890, z990, z9 BC, z9 EC). Currently, NALC is available on the S/390® 9672, z900, z990, and z9 EC servers while z/OS.e is available on z800, z890, and z9 BC only. By being available on the full line of System z servers, zNALC can provide a single, consistent pricing model which can help simplify the planning of server investment and application deployment.

zNALC establishes criteria to determine which applications qualify as 'new workload' applications. In general, qualified applications are those that IBM considers 'new workload', such as Java language business applications running under WebSphere Application Server, Domino, SAP, PeopleSoft, or Siebel. NALC is applicable only to select qualifying workloads listed in its terms and conditions. With a broader description of qualifying new workloads, zNALC may be applicable to more of the applications you use.

*Note: * additional terms/conditions apply.*

Question:

What changes in pricing are available with z/VSE Version 4?

Answer:

z/VSE V4 offers a new MWLC pricing metric and a sub-capacity option for IBM System z9 EC and z9 BC servers.

Midrange Workload License Charges (MWLC) is a new pricing metric that can offer improved price/performance for z/VSE V4 customers. MWLC applies to the z/VSE V4 operating system and 12 key VSE-related middleware programs when running on IBM System z9 EC and z9 BC servers. Please note, MWLC is not available to the smallest z9 BC server, capacity setting A01 – it remains zSeries Entry License Charge™ (zELC) pricing.

Question:

When might I benefit from sub-capacity pricing with z/VSE V4?

Answer:

Customers who select the MWLC pricing metric with sub-capacity pay for software based on the highest observed rolling 4-hour average utilization of the LPAR(s) in which the product runs.

Therefore, the sub-capacity pricing may benefit customers that have white space on their server, or customers with multiple operating systems running on the same server, e.g. z/OS and z/VSE.

Sub-capacity pricing may provide a cost effective way for you to consolidate a VSE server onto a larger server running a different operating system such as z/OS.

Question:

Is MWLC available on all System z servers?

Answer:

No, MWLC is exclusive to the IBM System z9 EC and z9 BC servers only.

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 z9 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:

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 z9 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 short running batch workload (CB-S) is replaced with a Java based batch workload (CB-J). Second, all workloads were moved to more recent levels of subsystem and compiler software. Third, and most significant, the measurement environment for the LSPR now includes 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 has been enhanced to include performance ratios reflecting both “single-image” z/OS and “multi-image” z/OS environments. Traditionally, the data presented in the LSPR was based on processors configured with one z/OS image equal in size to the processor model (the exception to this was that for z990 models with greater than 16 CPs, two images of z/OS were used). Typically, zSeries processors are configured with multiple images of z/OS. Thus, the LSPR now includes 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 z9 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, and 4) overall ratio of logical engines to physical engines. 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 9 at the high end. Most systems were configured with 2 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 1.5:1 on the largest models.

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 z9 EC?

Answer:

As with the introduction of any new server, workloads with differing characteristics will see some variation in performance when moved to the z9 EC. The performance ratings for a server (MSUs) 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 z9 EC has been specifically designed to focus on new and emerging workloads while continuing to provide outstanding benefit to the predominant traditional applications. The superscalar design takes advantage of naturally occurring code sequences designed to provide all workloads with a significant boost in performance. Additionally, newer applications, such as those with compiler optimizations for the z9 EC (available with z/OS 1.7 and Java 5)

may see even higher benefits, particularly those that may be enhanced over time to exploit some of the new instructions provided with the z9 EC.

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 the goal for the z9 EC of introducing a design that can provide the capability to deliver nearly 2 times the capacity of the largest previous server and to support a much higher number of logical partitions, significant enhancements to the LPAR management algorithms have been made to help reduce the potential for increased performance variability. In the spirit of autonomic computing, the LPAR manager is designed 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 z9 EC.

Question:

What is the performance improvement a z/VSE customer might experience on the z9 EC?

Answer:

The performance ratios that a z/VSE customer workload might experience when migrating from a zSeries server to the z9 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 z990 2084-301 and the z9 EC 2094-701 is approximately 32% to 39%. z/VSE workloads could expect this same range of performance.

Question:

How do I get performance information for my TPF products running on a System z9?

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:

What capabilities are extended with the z9 EC to enhance performance for z/VM Linux guests?

Answer:

The z990 and z890 servers introduced technology designed to enhance performance by allowing adapter interrupts to be passed directly to z/VM operating system guests for HiperSockets, Fibre Channel Protocol (FCP), and OSA on z990 and z890 servers. The System z9 introduced a complementary technology, which is also available for z990 and z890.

- QDIO Enhanced Buffer-State Management (QEBSM) designed to help eliminate the overhead of hypervisor interception.
- Host Page-Management Assist (HPMA) – designed to allow the hardware to assign, lock, and unlock page frames without z/VM hypervisor assistance.

These hardware assists can allow a cooperating guest operating system to initiate QDIO operations directly to the applicable channel, without interception by z/VM, thereby helping to provide additional performance improvements. These virtualization technologies are available only to first-level guests of z/VM 5.2 and later. IBM is working with its Linux distribution partners that this function will be provided in future Linux on System z9 distribution releases or service updates.

(Minimum MCL level is required on the z990 and z890. Refer to the Preventive Planning (PSP) bucket for your z990 and z890 server for required updates.)

Question:

Can HiperSockets support Internet Protocol Version 6 (IPv6)?

Answer:

Yes. HiperSockets now supports Internet Protocol Version 6 (IPv6) designed by the Internet Engineering Task Force (IETF) to replace Internet Protocol Version 4 (IPv4) to help satisfy the demand for additional IP addresses. IPv6 provides more unique IP addresses by expanding the IP address space from 32 bits to 128 bits. IPv6 is supported by z/OS 1.7 and z/VM 5.2.

Question:

What capability is offered on the z9 EC to ease migration for customers using traditional SNA and the IBM 374X?

Answer:

The OSA for NCP feature provides LPAR to LPAR connectivity from System z9 operating systems to the Communication Controller for Linux (CCL) 1.1. Use of the OSA for NCP can help ease migration and eliminate dependencies on hardware such as the 3745/6 communication controller. It will help protect your investment in traditional SNA applications and data, while allowing you to collapse SNA within the server while exploiting and leveraging IP.

Question:

What is on the z9 EC to help to simplify networking administration and management of VLANs on the z9 EC?

Answer:

The OSA-Express2 features now support GVRP for VLAN prioritization (a component of the IEEE 802.1 standard) to help simplify networking administration and management of VLANs. Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) can be used to propagate VLAN information with switches. You may no longer be required to manually enter VLAN IDs at the switch. Support of GVRP is exclusive to System z9, is applicable to all of the OSA-Express2 features when in QDIO mode, and is supported by z/OS. z/VM 5.2 provided this support on May 26, 2006 with the PTF for APAR VM63952 and z/VM 5.1 with the PTFs for APAR VM63784 and PK08444 and continues to be supported on z/VM 5.3.

Question:

What was introduced on the z9 EC with the OSA-Express2 support for 1000BASE-T Ethernet?

Answer:

With the additional memory available on OSA-Express2, the 1000BASE-T Ethernet feature will be able to: Support large send (offloading TCP segmentation processing from host TCP/IP stack to OSA-Express2), 640 TCP/IP stacks for improved virtualization by hosting more images on System z9 and potentially reducing the number of required OSA features, and concurrent LIC update to help minimize network traffic disruption.

Question:

What is link aggregation for z/VM in Layer 2 mode?

Answer:

This support is designed to allow you to combine multiple physical OSA-Express2 ports into a single logical link for increased throughput and for near-seamless failover in the event that an OSA port in the group becomes unavailable.

Up to 8 adapters can be aggregated per virtual switch which may increase virtual switch bandwidth and provide near seamless failover in the event of a failed controller, link or switch. This new support enables increased scalability for virtual network I/O, includes support to recover from a failed external switch and is provided for Layer 2 switches only.

Link aggregation is exclusive to z9 EC and z9 BC, is applicable to the OSA-Express2 features when configured as CHPID type OSD (QDIO), and is supported by z/VM 5.3.

Question:

What is Dynamic LAN idle for z/OS and z/OS.e designed to do?

Answer:

Dynamic LAN idle is designed to reduce latency and improve performance by dynamically optimizing settings for inbound blocking of data.

System administrators can authorize the TCP/IP stack to enable a dynamic setting (instead of a static setting). The TCP/IP stack is designed to determine the best setting based on the current system and environmental conditions (inbound work load volume, CPU utilization, traffic

patterns, and so on) and can dynamically update the settings. An OSA-Express2 feature will "adapt" to the changes, avoiding thrashing and frequent updates to the OSA address table. OSA will hold packets before "presenting" the packets to the host based upon the TCP/IP settings. A dynamic setting is designed to avoid or minimize host interrupts.

Dynamic LAN idle is exclusive to System z9, is supported by the OSA-Express2 features (CHPID type OSD), and is exploited by z/OS and z/OS.e 1.8 with PTFs. z/VM 5.1 and above supports for guest exploitation.

Question:

How can OSA Layer 3 Virtual MAC benefit z/OS and z/OS.e environments?

Answer:

By providing infrastructure simplification and improved System z IP workload balancing efficiency when an LPAR is sharing the same OSA Media Access Control (MAC) address with another LPAR.

Each operating system instance can now have its own unique "logical" or "virtual" MAC (VMAC) address. All IP addresses associated with a TCP/IP stack are accessible using their own VMAC address, instead of sharing the MAC address of an OSA port and may provide improved outbound routing and simplified configuration setup.

This applies to Layer 3 mode and to an OSA port shared among Logical Channel Subsystems. OSA Layer 3 VMAC is exclusive to z9 EC and z9 BC, is applicable to the OSA-Express2 features when configured as CHPID type OSD (QDIO), and is supported by z/OS and z/OS.e 1.8 with PTFs. z/VM 5.1 and above supports for guest exploitation.

Question:

What is provided by z/VM 5.3 to support OSA-Express2 link aggregation on the z9 EC and z9 BC?

Answer:

z/VM Virtual Switch-controlled (VSWITCH-controlled) link aggregation (IEEE 802.3ad) allows you to dedicate an OSA port to the z/VM operating system when the port is participating in an aggregated group. Link aggregation (trunking) is designed to allow you to combine multiple physical OSA-Express2 ports into a single logical link for increased throughput and for nondisruptive failover in the event that a port becomes unavailable.

The OSA-Express adapters that comprise the link aggregation group must be connected to the same physical switch and LAN segment. If there is a connectivity problem with any of the OSA-Express devices within the group, the virtual switch will detect the failure and packet transmission will continue over the other devices in the group. Thus, this support provides the capability for a quicker failover for devices within the group. In the event of a failure in which the entire link aggregation group is unavailable, support for specifying a failover OSA-Express adapter will be supported. Similar to the existing failover support in z/VM 5.2, failover devices can be specified with the virtual switch definition. Only a single failover device will be deployed at a time although multiple devices can be specified for the failover capability. For optimum connectivity, the failover devices should be connected to a separate physical switch.

I/O Enhancements

Question:

What is the FICON Express4?

Answer:

FICON Express4 is a new generation of FICON/FCP features now being offered on the IBM System z9, designed to deliver increased performance compared to the FICON Express2 on the z990 and z890. FICON Express4 now offers two unrepeated distance options (4 KM and 10 KM) when using single mode fiber optic cabling, and also supports a 4 Gbps link data rate with auto-negotiation to 1 and 2 Gbps for synergy with existing switches, directors and storage devices. The FICON Express4 feature also supports pluggable optics for individual servicing of each of the four channels on a feature.

Question:

What are the key benefits of FICON Express4 support for 4 Gbps links?

Answer:

FICON Express4 can help to reduce the cost of storage operations by enabling consolidation with improved throughput and faster link data rates and help to reduce configuration complexity. FICON Express4 can also enable a more manageable migration to higher performance by supporting devices at 1, 2 or 4 Gbps. Customers may deploy 4 Gbps in the SAN backbone or to selected devices, while preserving investment in 1 or 2 Gbps connected devices.

In addition, FICON Express4 can enable a more cost effective exploitation for midrange customers with a 2 channel FICON Express4 feature available exclusively on the z9 BC. This card is available in both longwave (LX) and shortwave (SX) features.

Question:

Are FICON Express4 features supported on zSeries or S/390 processors?

Answer:

No. FICON Express4 support is exclusive to System z9. However, because it supports a 4 Gbps link data rate with auto-negotiation to 1 and 2 Gbps, zSeries or System/390™ processors operating at 1 or 2 Gbps as well as other, slower speed FICON and FCP devices are compatible with 4 Gbps capable SAN Directors, and Switches.

Question:

What IBM System Storage products will attach to FICON and FCP 4 Gbps enabled SAN?

Answer:

IBM announced an expanded range for 4 Gbps FICON and FCP enabled SAN directors and switches for enterprise and midrange environments. Announcements include:

- Cisco MDS 9513 director and 4 Gbps features for 9216A and 9216i Switches; 9506 and 9509 directors

- IBM TotalStorage® SAN32M-2 switch and SAN140M director (McDATA) 4 Gbps features and IBM SAN256M 4 Gbps feature.
- IBM TotalStorage SAN32B-2 switch and SAN 256B director (Brocade) 4 Gbps shortwave features.

IBM System Storage is ready to support System z9 and FICON Express4 with a large variety of storage options that are compatible with the new 4 Gbs FICON/FCP adapters. We have disk, tape and SAN products that either support 4 Gb/s connectivity today or will very soon. These links are all auto-negotiating so they can support intermix of 4 Gbs and 2 Gbs technology.

Question:

What has IBM announced for tape encryption and how will it benefit me?

Answer:

On August 29, 2006, IBM was first to announce tape encryption using its TS1120 tape drive. When invoked, the new feature will encrypt data as it is written to the tape media. The new feature will especially help customers protect tape cartridges that are being moved from one location to another.

Question:

I've already purchased a TS1120 tape drive; will I be able to upgrade it to support the new encryption feature?

Answer:

Yes, there is a plan to provide an upgrade for customers that previously purchased a TS1120 tape drive prior to the availability of the encryption feature.

Question:

What is the TS3400 Tape Library announced February 2007?

Answer:

The IBM System Storage TS3400 Tape Library is designed to offer the high capacity and performance advantage of the IBM System Storage TS1120 Tape Drive in a smaller automation footprint for open systems environments. The TS3400 tape library is an external 5U standalone or rack-mountable unit supporting up to two TS1120 tape drives with a data transfer rate of up to 104 MBps per drive.

The IBM Encryption Key Manager component for the Java platform can help generate and manage encryption keys for TS1120 tape drives across the enterprise. This feature uses standard key repositories and supports three different encryption key management methods: application managed, system managed or library managed.

Question:

Does the TS3400 library support attachment to System z?

Answer:

Yes, We are announcing support for open systems servers only at this time and this includes attachment to System z servers running Linux via FCP.

IBM plans to enhance the System Storage TS3400 Tape Library by adding additional support for System z (z/OS, z/VM, z/VSE and TPF). Additionally, IBM plans to support system managed encryption in a System z environment with the TS3400.

Question:

How are the Encryption keys going to be generated and managed?

Answer:

When used with z/OS, the TS1120 is designed to leverage System z's unique security and cryptographic features to provide a powerful solution for enterprise wide encryption key storage and management capabilities using ICSF.

Question:

How are the Encryption keys going to be generated and managed for z/VSE 4.1?

Answer:

A separate IBM Encryption Key Manager (EKM) component for the Java platform is available. This EKM supports the generation, communication, and storage of encryption keys for TS1120 Tape Drives across the enterprise.

z/VSE V4.1 is designed to support Systems Managed Encryption with the IBM System Storage TS1120 using an 'out of band' connection. z/VSE V4.1 support will require the EKM component running on another operating system other than z/VSE. Communication between the TS1120 and EKM is via TCP/IP.

Note: This function will not be available at z/VSE V4.1 general availability. The function is planned to be delivered later via PTF. If you plan to use this function, check the z/VSE Web site after GA for the latest status. z/VSE V3.1 support of the TS1120 Tape Drive with encryption is planned for first half 2007.

Question:

How are the Encryption keys going to be generated and managed for z/VM 5.3?

Answer:

z/VM now supports drive-based data encryption with the IBM System Storage TS1120 Tape Drive (machine type 3592, model E05). The TS1120's encryption capability and its subsystem-integration support can provide a flexible tape-data-encryption solution that provides data encryption and key management across a variety of environments with a single point of control for all encryption keys. Most importantly, this solution can help protect data on tape in a cost-effective way.

Encryption of tapes by z/VM itself requires that the IBM Encryption Key Manager be running on another operating system, using an out-of-band (such as TCP/IP) connection to the tape control unit. z/VM native support includes encryption for DDR and SPXTAPE, as well as transparent support for guests that do not provide for their own encryption (for example, CMS and Linux for System z).

z/VM also enables encryption of tapes by guests (such as z/OS) that have the ability to control the tape-encryption facilities themselves and to optionally run the Encryption Key Manager. Key management for such guests can use either an out-of-band or an in-band (such as an ESCON or FICON channel) connection between the Encryption Key Manager and the tape

control unit. With the PTF for APAR VM64063 for z/VM 5.1 and 5.2, only the Encryption Key Manager's default keys are supported for use by z/VM and by guests that do not provide for their own encryption. z/VM 5.3 expands this support to allow any key label to be used, with key labels being accessible through a key alias that is defined to z/VM.

Question:

How does FICON Express4 help business continuity and GDPS® configurations?

Answer:

FICON Express4 can help improve business continuity in two ways.

IBM already enjoys some of the lowest latency, lowest TCO, and most resilient design points for high availability configurations. With FICON Express4 configurations will have the ability to scale larger with better throughput, provide a lower TCO, and continue to offer a proven solution.

In addition, FICON Express4 continues support for FICON cascading directors. Support for FICON cascading means that a Native FICON (FC) channel or a FICON CTC can connect a server to a device or other server via two (same vendor) directors in between. FICON Cascading support is important because it can help provide high availability by reducing configuration complexity and by reducing the number of channels for interconnecting the 2 sites.

Question:

How does the z9 EC deliver increased I/O connectivity and simplify I/O configuration?

Answer:

In comparison to the z990, the z9 EC supports a 40% increase in Crypto, FICON and OSA connectivity per I/O cage by providing the capability to populate the 28 I/O slots in one I/O cage with any mix of up to 8 Crypto Express2, up to 28 FICON Express2 or FICON Express, and up to 24 OSA-Express2 or OSA-Express features. On the z990, the maximum quantity of Crypto, FICON, and OSA features in combination in one I/O cage was limited to 20 features because of power/mechanical limitations. Now up to 336 FICON channels may be installed on the z9 EC. Up to 240 FICON channels may be installed on the z990.

Question:

What new performance enhancements are available for FICON FCP?

Answer:

A FICON channel configured to support attachment to SCSI devices has been enhanced to provide up to two times the number of start I/Os per second with 4k bytes per I/O. This may provide better application performance.

FCP performance enhancements are exclusive to z9 EC and z9 BC, is applicable to the FICON Express4 and FICON Express2 features when configured as CHPID type FCP, and is supported by Linux on System z, z/VM 5.1 and later and z/VSE 3.1 and later.

Note: These results are achieved in a laboratory environment using one channel with no other processing occurring and do not represent actual field measurements.

Question:

How can Modified Indirect Data Address Word Facility (MIDAW) help improve application performance?

Answer:

MIDAW is a facility that can help improve FICON performance by reducing channel, director, and control unit overhead. Applications that use the following datasets to access information may benefit, including DB2, VSAM, Partitioned Data Set Extended (PDSE), Hierarchical File System (HFS), z/OS File System (zFS), and other extended format datasets.

Question:

What are the performance improvements found running tests on the z9 EC, MIDAW facility, FICON Express4, DB2 UDB for z/OS V8 and an IBM TotalStorage DS8000™?

Answer:

The results of internal DB2 table scan tests with Extended Format datasets on the z9 EC with the Modified Indirect Data Address Word (MIDAW) facility and the IBM TotalStorage DS8000 yielded the following results when using FICON Express4 operating at 4 Gbps on a z9 EC compared to FICON Express2 operating at 2 Gbps on a z9 109:

- A 46% improvement in throughput for all reads (270 MBps vs 185 MBps)
- A 35% reduction in response times

In addition, comparing use of the MIDAW facility with FICON Express4 operating at 4 Gbps to the use of MIDAWs with FICON Express2 operating at 2 Gbps showed a combined throughput improvement of greater than 220% for all reads (270 MBps vs 84 MBps) on DB2 table scan tests with Extended Format datasets.

These measurements are examples of what has been achieved in a laboratory environment using one FICON Express4 channel operating at 4 Gbps (CHPID type FC) on a z9 EC with z/OS 1.7 and DB2 UDB for z/OS V8. Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the numbers stated here.

Question:

What enhancements on the z9 EC help to provide increased I/O addressing capability?

Answer:

The z9 EC helps to increase the number I/O devices and amount of addressable storage with the following enhancements:

- Multiple Subchannel Sets: MSS provides a second set of subchannels for defining Parallel Access Volume (PAV) aliases. This function can help provide relief from the 64K device limit by allowing PAV aliases to be defined without making device numbers unavailable for defining additional devices. For some of our largest customers this is designed to provide greater I/O device configuration capabilities.

- 63.75K subchannels: The z9 EC addresses a maximum of 64K-1 subchannels in subchannel set 0 (zero). Previously 1024 (1K) of these subchannels were reserved for system use. IBM is making 768 of these subchannels available for customer use. This increases the storage attachment capability of the z9 EC – for example the IBM TotalStorage DS8000 Series can be defined to attach 63.75K unit addresses – so with 63.75K in the host there is symmetry between the server and the storage subsystem.
-

Question:

What is N_Port ID Virtualization announced with the System z9?

Answer:

N_Port ID Virtualization (NPIV) allows a single Fiber Channel Protocol (FCP) port on the System z9 platform to appear as multiple, distinct ports providing separate port identification and security zoning within the fabric for each operating system image as if each operating system image, or z/VM guest, had its own unique physical port. NPIV is similar to the Multiple Image Facility (MIF) available on ESCON® channels.

NPIV is an extension of the International Committee for Information Technology Standardization (INCITS) Fibre Channel standard. NPIV applies to the FICON features supported on z9 EC when a channel is configured as CHIPID type FCP, supporting connectivity to SCSI devices.

On/Off Capacity on Demand (On/Off CoD) and CBU

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). Please note that while you can not decrease the number 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.

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

Question:

What is the administrative On/Off Capacity on Demand (On/Off CoD) test?

Answer:

This feature allows the customer to order 'zero' quantity of processor features via Resource Link™, which will allow customer staff to order/test/rehearse/train/document the entire On/Off CoD process without incurring any cost. There is an unlimited number of such tests allowed and there are no time period restrictions.

Question:

What is the change available that allows adding more On/Off CoD capacity without having to restore the System z9 to the current 'purchased' capacity?

Answer:

With an On/Off CoD temporary upgrade already installed, you can now modify (add or remove) your 'temporary' capacity configuration without having to first restore the machine to its original "permanent" capacity configuration. 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 z9?

Answer:

Yes, enhancements include improvements to On/Off Capacity on Demand to assist in defining new ways to acquire temporary capacity when the business demands are changing.

The z9 provides the ability to store up to 100 On/Off CoD LICCC records on the Support Element at any given time, giving greater flexibility to quickly enable needed temporary capacity. Each record is easily identified with descriptive names, and users can select from a list of records that can be activated.

Question:

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

Answer:

There will be a 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:

On/Off CoD pricing has been enhanced to provide reduced pricing when a temporary upgrade includes activation of previously purchased but unassigned general purpose processor (CP) capacity and/or previously purchased but unassigned IFLs. Now you may activate any amount of purchased but unassigned CP capacity or purchased but unassigned IFLs on a temporary basis for a flat daily fee per processor type. There is one flat daily fee for temporary usage of purchased but unassigned CP capacity and a second, separate flat daily fee for temporary usage of purchased but unassigned IFLs. The offering is designed such that when you activate purchased but unassigned capacity you'll be charged the lesser of the flat daily fee or your standard On/Off CoD rate. Contact your IBM System z sales specialist for additional pricing information.

Question:

Can On/Off CoD utilize unassigned general purpose processor capacity as a temporary zAAP or zIIP on the z9 EC?

Answer:

Not always. If the unassigned CP capacity is a subcapacity processor, no. If the unassigned capacity is a full capacity processor, yes. Specialty engines require full capacity processors. So only an unassigned full capacity CP may be utilized as a specialty engine for On/Off CoD.

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.

Question:

Does the customer 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 the customer a priced option.

Question:

Will temporary capacity be available for specialty PUs (IFLs, ICFs, zAAPs and zIIPs) on the z9 EC?

Answer:

Yes. The z9 EC, like the z890 and z990 servers, will continue to have the ability to utilize capacity for temporary upgrades or special purpose processors through On/Off CoD: IFLs, zAAPs, zIIPs and ICFs. Please note that SAPs are not available for On/Off CoD 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 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 8 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 8 or more full capacity temporary general purpose processors).
Question:

Question:

When can I begin placing On/Off CoD orders against my z9 EC?

Answer:

On/Off CoD can be initiated as soon as the profile for the z9 EC is established. The prerequisite of establishing a profile is the signing of an On/Off CoD contract associated with ordering features 9898 and 9896.

Question:

Will z890 and z990 servers be able to utilize unassigned capacity the same as the z9 EC?

Answer:

No. Though z890 and z990 servers can utilize unassigned PUs, the z890 and z990 have no pricing distinction between unassigned (purchased but turned off) and uncharacterized (not purchased) capacity.

Question:

Will I be able to do Capacity Backup Upgrade (CBU) capability on the zIIP ?

Answer:

Yes. The z9 EC is able to activate most processor types as part of CBU upgrades: IFLs, zAAPs, ICFs, CPs and the zIIPs.

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 of installed CP processors. With the full 'matrix' upgradeability of the subcapacity general purpose CP processors you can execute CBU with processors of equal, less or greater capacity depending on your requirements. Note that you can not exceed 8 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 8, 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 a permanent engine to another engine type during CBU?

Answer:

No. All 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.

Cryptographic Enhancements

Question:

What common criteria security certifications has IBM System z9 achieved?

Answer:

On September 4, 2006, the IBM System z9 109 (z9 -109), now named the z9 EC and z9 BC all had joined the ranks of other IBM eServer zSeries servers that have achieved Evaluation Assurance Level 5 for common criteria security certification. The EAL5 ranking should give companies confidence that they can run many z/OS, z/VM and Linux-based applications containing confidential data—such as payroll, human resources, e-commerce, ERP and CRM systems—on one System z platform divided into partitions that keep each applications data secure and distinct from the others. That is, the System z architecture is designed to prevent the flow of information among logical partitions on a system, thus helping to ensure that confidential or sensitive data remains within the boundaries of a single partition.

Question:

What cryptographic capabilities were added to the z9 EC with the April 2006 announcement?

Answer:

Remote key loading is now 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 TDES (Triple DES) 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 z/OS 1.6 and 1.7 with Enhancement to Cryptographic Support for z/OS and z/OS.e 1.6/1.7 Web deliverable (z/OS.e runs on z9 BC only), z/VM 5.1, 5.2 and 5.3 for z/OS guests.

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 onsite, 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 IBM System z9 Remote Key Load function?

Answer:

IBM System z9 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.
 - ANS X9.24-2: 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 enhancement?

Answer:

International Standards Organization (ISO) 16609 supports the requirements for message authentication using T-DES symmetric encryption. Integrated Cryptographic Service Facility (ICSF) with Crypto Express2 feature, support 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 z9 EC?

Answer:

CP Assist for Cryptographic Function (CPACF) on every CP and IFL and the Crypto Express2 feature.

Question:

What cryptographic hardware is not supported on the z9 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 z9 EC.

Question:

What cryptographic hardware is offered as standard features on the z9 EC?

Answer:

The CP Assist for Cryptographic Function (CPACF) is standard on every CP and IFL, however, a no-charge enablement feature #3863 is required. The Crypto Express2 feature is an optional feature. The first order increment is two features.

Question:

What features are supported by the CP Assist for Cryptographic Function (CPACF) on the z9 EC?

Answer:

CPACF includes support of the Advanced Encryption Standard (AES) for 128-bit keys, Secure Hash Algorithm-256 (SHA-256), and Pseudo Random Number Generation (PRNG).

CPACF, supporting clear key encryption, is activated using a no charge enablement feature #3863 and offers the following on every processor (PU) identified as a general purpose processor (CP) or Integrated Facility for Linux (IFL).

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

Performance is designed to scale with the addition of PUs. SHA-1 and SHA-256 are shipped enabled on all servers and do not require the enablement feature. DES, TDES 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)
 - Designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification.
- Crypto Express2 Accelerator – to enable significant improvement in SSL acceleration on System z9, z890 and z990 servers.
 - Supports clear key RSA acceleration
 - Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol.
- Since the configuration functions are implemented in Licensed Internal Code, current Crypto Express2 features can be carried forward from z990 to the z9 EC to take advantage of the potential for increased SSL performance and configuration capability.
- 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.

Note: Crypto Express2-1P, feature # 0870, with one PCI-X adapter is not available on the z9 EC.

Question:

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

Answer:

System z9 provides the ability to configure Crypto Express2 PCI-X adapters as accelerators. When both PCI-X adapters are configured as accelerators, the Crypto Express2 feature is designed to perform up to 6000 SSL handshakes per second. This represents, approximately, a 3X performance improvement compared to the PCICA feature or the current Crypto Express2 feature on z990, on a per PCI-X adapter basis. The SSL rate was achieved with a System z9 environment with four processors and two PCI-X adapters cards (one Crypto Express2, both configured as accelerators), z/OS 1.7 with Cryptographic Support for z/OS 1.6 and 1.7 Web deliverable (will not be available after 5/26/06). It was replaced by Enhancements to Cryptographic Support for z/OS and z/OS.e 1.6/1.7 Web deliverable available May 26, 2006. Also required is ICSF FMID HCR7730.

Since the performance enhancements are implemented in Licensed Internal Code, current Crypto Express2 features that have been carried forward from z990 to System z9 platforms can take advantage of increased SSL performance and configuration capability. 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.

Previously, the Crypto Express2 feature was configured as a coprocessor. It was not optimized for SSL performance. Details are available on request.

1 Note, the previously reported SSL performance of 4995 handshakes per second was obtained on a 4-way z990 with four Crypto Express2 Coprocessors (CEX2C) features, whereas in this case the performance was measured on a 4-way System z9 platform with one Crypto Express2 feature with both PCI-X adapters configured as accelerators. It would be expected that the SSL performance on a 16-way System z9 platform with six Crypto Express2 features would be greater than that obtained on a z990, however, actual measurements have not been taken.

Question:

Will the Crypto Express2, feature code 0863, support Linux Secure Sockets Layer (SSL) and Transport Layer Security (TLS) cryptographic operations on the z9 EC?

Answer:

Yes, it supports Public Key operations with Linux on System z. IBM is working with its distribution partners to provide this function in future distribution releases, or service updates.

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 and z990 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.

CSNBSYE, CSNBSYE1, CSNBSYD, and CSNBSYD1 will provide support for clear-key AES encryption and decryption with 128-bit keys using the CPACF.

CSFNBOWH and CSNBOWH1 will provide support for SHA-1 and SHA-256 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, TDES, AES, SHA-1 and SHA-256 functions can be invoked by five (5) 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.

Question:

Is a Trusted Key Entry (TKE) workstation required for the use of the CPACF?

Answer:

No, the CP Assist for Cryptographic Function (CPACF) supports clear key functions and does not require entering of master keys.

Question:

Will there be an update to the ATS TechDocs Web site to provide additional technical information about System z9 cryptographic hardware features?

Answer:

ATS TechDocs Web site and several System z9 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 on the z9 EC?

Answer:

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

- Crypto Express2:
 - z/OS 1.6 and later with Cryptographic Support for z/OS 1.6 and 1.7 Web deliverable (will not be available after 5/26/06). It will be replaced by Enhancements to Cryptographic Support for z/OS 1.6/1.7 Web deliverable available May 26, 2006 and PTFs. z990 and z890 enhancements to cryptographic support are Web deliverable.
 - z/VM 5.1 guest exploitation, with applicable PTFs, z/VM 5.2 and later.
 - z/VSE 3.1 and later
 - Linux for System z9 – IBM is working with its distribution partners to provide this function in future distribution releases, or service updates.
 - CPACF (SHA-1, SHA-256, AES and PRNG):
 - z/OS 1.6 and later with Cryptographic Support for z/OS 1.6 and 1.7 Web deliverable (will not be available after 5/26/06). It will be replaced by Enhancements to Cryptographic Support for z/OS 1.6/1.7 Web deliverable available May 26, 2006 and PTFs. z990 and z890 enhancements to cryptographic support are Web deliverable.
 - z/VM 5.1 and later for guest exploitation.
 - z/VSE V3.1, and later, with applicable PTFs. z/VSE does not exploit SHA-256 and PRNG.
 - Linux for System z9 – IBM continues to work with its distribution partners to provide this function in future distribution releases, or service updates.
-

Question:

Will UDX's written for zSeries servers function on the z9 EC?

Answer:

Pre-existing UDXs on current systems may need to be ported to the z9 EC. Customers need to contact IBM to port existing UDXs to the Crypto Express2 environment on the z9 EC. For further information, see the answer to the next question.

Question:

Will new UDX's be supported on the z9 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) encryption performance (in accelerator mode)		X
Provides highest asymmetric (encrypted key) encryption performance (in coprocessor mode)		X
Uses CHPID numbers	N/A	N/A
Physically imbedded on each General Purpose Processor (CP)	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 performing SSL handshakes		X
RSA functions		X
High Performance SHA-1, SHA-256 Hash function	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 5.1?

Answer:

The Trusted Key Entry (TKE) 5.1 level of Licensed Internal Code (LIC) is installed in TKE workstation feature codes #0839 and #0859. The TKE 5.1 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 5.1 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 5.1 LIC is a no-charge enablement feature which is loaded prior to shipment when a TKE workstation is ordered. The TKE 5.1 LIC includes support for the Smart Card Reader.

The optional TKE features are:

- TKE 5.1 LIC (#0856) and TKE workstation (#0839)
 - TKE 5.0 LIC (#0855) and TKE workstations (#0859 and #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, a new TKE workstation will be available. The new TKE 5.1 workstation is assigned feature code 0839. Customers must use the TKE 5.1 workstation to control the z9 EC. Customers may continue to use TKE 3.x and 4.x workstations to control prior servers, but an existing workstation cannot be upgraded to TKE 5.1.

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 z9 EC TKE workstations?

Answer:

Yes. Support is available for an optional Smart Card Reader to be attached to the new TKE 5.1 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) TKE 5.1 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 z9 EC.

Parallel Sysplex Enhancements

Note: Please use the updated STP FAQs to answer your questions on the new Server Time Protocol feature of the z9 EC. These can be found at: www.ibm.com/systems/z/faq/

Question:

What improvements are provided with Coupling Facility Control Code (CFCC) Level 15?

Answer:

Coupling Facility Control Code (CFCC) Level 15 will increase the allowable tasks in the Coupling Facility (CF) from 48 to 112.

Note: When migrating CF levels, lock, list and cache structure sizes may need to be increased to support the new function. This adjustment can impact the system when it allocates structures or copies structures from one coupling facility to another at different CF levels. The coupling facility structure sizer tool is designed to size structures for you, and takes into account the amount of space needed for the current CFCC levels

GDPS Enhancements

Question:

What GDPS functions were announced in March 2007?

Answer:

GDPS/PPRC MultiPlatform Resiliency support is expanded with GDPS/PPRC V3.4. GDPS is now designed to provide management and automation for both planned and unplanned actions for Linux guests under z/VM. This new support removes the prior requirement for z/VM, supporting Linux systems running in LPARs on System z servers for planned and unplanned actions.

Incremental Resynchronization for GDPS Metro/Global Mirror in the GDPS 3.4 timeframe is designed to help improve the recovery capability of GDPS Metro/Global Mirror by allowing the Global Mirror session to be quickly reestablished between the primary and the remote recovery site if there is a failure at the intermediary site. In addition, when the failing disks are brought back into operation, they can be brought back to currency without a full copy, with only an incremental resynchronization.

Zero Suspend Flash Copy for XRC is designed to provide the ability to FlashCopy® XRC secondary devices without interrupting disaster recovery protection by not suspending and desynchronizing secondary disk.

Additional FlashCopy support for GDPS/Global Mirror is available with GDPS V3.4 to help query, manage, and monitor the status of an additional FlashCopied version of the data for GDPS/Global Mirror. This copy may be used to help performing disaster recovery testing, providing a back up copy before desynchronizing, or to back up production data without affecting the production workload.

In addition, these additional FlashCopy disks do not need to be defined in the GDPS R-Sys I/O definitions. This FlashCopy without Unit Control Block (UCB) capability can be used by systems running out of device numbers.

Metro Mirror re-establishment is designed to recover secondary devices after a FREEZE event by doing a HyperSwap™ failover. This is designed to mitigate the disaster recovery exposure by limiting the time period that devices are unsynchronized. It can also facilitate disaster recovery testing in the production environment, again by reducing the time to re-establish PPRC after a test.

GDPS/PPRC HyperSwap Manager system management provides a new single command to make the secondary disks available after a freeze event. Previously, recover required multiple panel interactions with a couple of interactions per LSS pair in the configuration. This can provide enhanced systems management and simplified operator procedures for recovery management.

GDPS Graphical user interface offers a graphical Web interface for many GDPS/PPRC standard actions and monitoring functions.

IPL Message Automation is designed to simplify recovery by further automating responses to IPL messages generated to a NIP console in a GDPS/XRC and GDPS/Global Mirror environment. Some messages that GDPS can now automate include several key messages relating to initializing a sysplex and couple data sets.

GDPS/PPRC HyperSwap coexistence provides coordination between HyperSwap and TDMF synchronize and relabel activities. With GDPS V3.4 and TDMF support, TDMF is designed to automatically issue the commands to disable and re-enable HyperSwap. This can help further automate a disk migration, speed up the process, and reduce human errors. In addition, GDPS V3.4 supports co-existence with HyperSwap function and IMS/XRF.

GDPS Health Checks extends the z/OS Health Checker infrastructure by providing GDPS specific checks for best practices. These checks help ensure that GDPS best practices are adhered to even as a customer evolves and changes their configuration. More checks are planned to be made available as they are developed.

GDPS 3.4 is available as of March 30, 2007. GDPS is designed to work in conjunction with the z9 BC, z9 EC, z990, z890, z900, and z800 servers. 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.

Operating Systems

z/OS Version 1 Release 9 Preview

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

Question:

What are the key new enhancements planned for z/OS® Version 1 Release 9 (V1.9)?

Answer:

With z/OS® V1.9, IBM plans to extend the value of the flagship mainframe operating system with improvements in all of its core competencies, including scalability, availability, and resource optimization. With increased focus on simplifying z/OS for IT professionals, plans for z/OS V1.9 include improvements to the IBM Health Checker for z/OS, the IBM Configuration Assistant for z/OS Communications Server, DFSMSrmm™, ISPF, Hardware Configuration Manager (HCM), and Coupling Facility services, as well as a new dbx GUI. Additional enhancements to z/OS are planned to make the operating system more powerful in applying centralized policy-based rules for defining and controlling how your applications behave.

Advancements in the z/OS Communications Server will help put the mainframe in a leadership position in terms of providing network solutions. With the planned ability to provide centralized policy services and policy-based routing (along with ability to apply security, intrusion detection/defense, qualities of service, as well as other features) z/OS can help you customize the network to suit the needs of your applications — doing this in a simplified, centralized, manageable, and auditable manner that is anticipated to be transparent to the application.

The z/OS Integrated Cryptographic Service Facility (ICSF) is planned to be enhanced to include the PKCS#11 standard. ICSF is the fundamental base of z/OS mainframe encryption which enables you to encrypt and decrypt data, generate and manage cryptographic keys, and perform other cryptographic functions dealing with data integrity and digital signatures. With adoption of the PKCS#11 standard, the strength of mainframe encryption and secure centralized key management can be brought to and used by Web-based application and networking environments more easily.

The z/OS Workload Manager (WLM) is planned to be enhanced with improved performance routing, priority settings, and cancel functionality, further improving on the mainframe's leadership position in workload management capabilities. With z/OS WLM, you can define business and performance goals customized for your applications. The z/OS system decides how much resource, such as CPU and storage, should be given to applications that serve the workload to meet the goal. WLM constantly monitors the system and adapts resource applications to meet application goals, taking into account not only server resources, but network traffic, router bottlenecks, application health, and transaction prioritization as well, thus providing autonomic, policy-based z/OS performance management that can be tuned to meet your applications' needs.

Previews provide insight into IBM plans and direction. Availability, prices, ordering information, and terms and conditions will be provided when the product is announced.

Question:

What is the Coexistence/Migration/Fallback policy for z/OS V1.9 now that z/OS is on an annual release cycle?

Answer:

Coexistence and migration rules should be taken into account in planning future migrations. The z/OS coexistence, migration, and fallback policy aligns with the service policy. IBM intends to continue with the practice of providing service support for each release of z/OS for three years following its general availability (GA) date. In some cases, more than three releases may be coexistence, migration, and fallback supported if IBM, at its sole discretion, chooses to provide service support for greater than three years for a release. However, any z/OS release having three or fewer months of service remaining at the time of GA of a new release will not be coexistence, migration, and fallback supported. Migration forward as well as fallback should be made within the same z/OS releases supported by the coexistence policy.

The following table shows the releases that are planned to be coexistence-supported with z/OS V1.9.

Coexistence-supported releases

Release	Coexistence-supported with release in Column
z/OS V1.9**	z/OS V1.7, z/OS V1.8, z/OS V1.9**
z/OS V1.10**	z/OS V1.8, z/OS V1.9**, z/OS V1.10**
z/OS V1.11**	z/OS V1.9**, z/OS V1.10**, z/OS V1.11**

*(**) Operating system levels beyond z/OS V1.9 represent current intentions of IBM.*

This consistent coexistence, migration, and fallback policy applies to release migrations for all configurations, whether they are: single-system configurations, individual systems within a multisystem configuration, and cases where a simultaneous IPL is used to migrate all systems in a multisystem configuration at the same time.

Question:

What are the planned end-of-service dates for the current z/OS releases?

Answer:

End-of-service for z/OS 1.4 was March 31, 2007.

End-of-service planned for z/OS 1.5 was also March 31, 2007.

Beginning with z/OS 1.6, the first of our annual releases, each release is typically planned to be in service for 3 years. End-of-service for z/OS 1.6 is planned for September 30, 2007.

End-of-service for z/OS 1.7 is planned for September 30, 2008.

Question:

What are the key dates for z/OS?

Answer:

March 30, 2007: Availability of JES3 NJE communications using TCP/IP support

March 30, 2007: Availability of new optimized LDAP directory server, called IBM Tivoli® Directory Server for z/OS

March 31, 2007: Withdrawal from service of:

- z/OS V1.4 (5694-A01)
- z/OS.e V1.4 (5655-G52)
- z/OS V1.5 (5694-A01)
- z/OS.e V1.5 (5655-G52)

June 25, 2007: Recommended last date for submitting orders for z/OS V1.7 via the CustomPac family fee-based offering, SystemPac. This date will allow for adequate order processing time.

June 30, 2007: Last date for Web download of the following Web deliverables:

- z990 Compatibility for Selected Releases (Web deliverable) (compatibility for z/OS.e V1.3, z/OS V1.3 and V1.2, and OS/390® V2.10)
- SMP/E V3R3 Web deliverable
- IBM Health Checker for V1R4/R5/R6 of z/OS and z/OS.e Web deliverable

July 23, 2007: Last date for ordering z/OS V1.7 via the CustomPac family fee-based offering, SystemPac.

September 30, 2007: Withdrawal from service of:

- z/OS V1.6 (5694-A01)
- z/OS.e V1.6 (5655-G52)

September 2007: Planned availability of z/OS V1.9

Question:

What System z servers are planned to support z/OS V1.9?

Answer:

z/OS V1.9 is planned to run on the following IBM System z™ servers:

- z9 BC
 - z9 EC
 - z990
 - z890
 - z900
 - z800
-
-

Question:

Where can I get more information on z/OS?

Answer:

Please see: <http://www-03.ibm.com/servers/eserver/zseries/zos/>

z/VSE Version 4 Release 1

Question:

What is being announced?

Answer:

z/VSE Version 4 Release 1 (z/VSE 4.1) is built on a heritage of ongoing refinement and innovation that spans four decades. It is designed to help protect and leverage customer investments in VSE information assets, and it supports:

- IBM System z9 EC, z9 BC, IBM eServer zSeries 990, 890, 900, 800 servers
- Crypto Express2 adapter and CP Assist for Cryptographic Function (CPACF)
- OSA-Express2 and FICON Express4 adapters
- Systems Managed Encryption with an IBM System Storage TS1120
- IBM System Storage DS8000 and DS6000™ (both as ECKD and FCP-attached SCSI disks)

z/VSE V4 executes in z/Architecture mode only and supports 64-bit real addressing for selected system functions. z/VSE 4.1 is designed to support up to 8 GB of real processor storage.

z/VSE V4 is not designed to support 64-bit virtual addressing or allow 64-bit addressing for user applications.

z/VSE V4 offers a new MWLC pricing metric and a sub-capacity option for IBM System z9 EC and z9 BC servers.

Midrange Workload License Charges (MWLC) is a new pricing metric that can offer improved price/performance for z/VSE V4 customers. MWLC applies to the z/VSE V4 operating system and 12 key VSE-related middleware programs when running on IBM System z9 EC and z9 BC servers. If you are migrating to MWLC from Graduated Monthly License Charges (GMLC), Growth Opportunity License Charge (GOLC), zSeries Entry License Charges (zELC), or Tiered Entry Workload License Charges (TWLC) IBM expects you to experience improved price/performance based on MWLC price points. You may gain additional price/performance improvements by implementing sub-capacity MWLC.

Question:

What is the value of z/VSE 4.1?

Answer:

z/VSE 4.1 offers value based on advanced technology and new pricing metrics.

Technical

z/VSE V4 operates in z/Architecture mode only and supports 64-bit real addressing. z/VSE 4.1 is designed to help you exploit up to 8 GB of real processor storage. For example, you may be able to use more data-in-memory techniques such as CICS Shared Data Tables, VSE Virtual Disk, or more and larger buffer pools to help improve throughput. In addition, some customers with especially large z/VSE environments may experience lower paging rates. In fact, many z/VSE 4.1 environments may be able to run without a page data set at all (the NOPDS option).

z/VSE V4 support for 64-bit real addressing applies to IBM System z9 as well as IBM eServer zSeries servers.

z/VSE 4.1 also adds new support for hardware-assisted encryption. For example, it supports configurable Crypto Express2 and CP Assist for Cryptographic Function (CPACF) enhancements. It introduces SecureFTP and supports Systems Managed Encryption using the IBM System Storage TS1120 encrypting tape.

Pricing metrics

You may be able to lower your Total Cost of Operations (TCO) using z/VSE 4.1 and the new MWLC pricing metrics (on z9 EC and z9 BC). If you are migrating to MWLC from Graduated Monthly License Charges (GMLC), Growth Opportunity License Charge (GOLC), zSeries Entry License Charges (zELC), or Tiered Entry Workload License Charges (TWLC), IBM expects you to experience improved price/performance based on MWLC price points. You may gain additional price/performance improvements by implementing sub-capacity MWLC.

The new MWLC pricing metrics apply only when z/VSE 4.1 is running on IBM System z9 EC (formerly z9 109) or z9 BC servers.

Question:

What are some of the differences between z/VSE 4.1 and z/VSE 3.1?

Answer:

Here are just a few of the differences:

- z/VSE V4 operates in z/Architecture mode only and supports 64-bit real addressing. z/VSE 4.1 is designed to support up to 8 GB of processor storage.
- z/VSE V3 operates in ESA/390 mode only and real addressing is limited to 31-bit real addressing. z/VSE V3.1 supports up to 2 GB of real processor storage.
- z/VSE 4.1 is available with full-capacity and sub-capacity MWLC pricing on System z9 EC and z9 BC servers.
- z/VSE 3.1 is available with TWLC pricing on System z9 EC and z9 BC servers.
- In addition, z/VSE 4.1 is designed with new security logging and reporting capabilities, support for added hardware-assisted encryption, and several user requirements.

Question:

What does 64-bit mean for z/VSE?

Answer:

z/VSE V4 supports 64-bit real addressing for selected system functions. z/VSE 4.1 is designed to support up to 8 GB of real processor storage.

Support for 64-bit real addressing may help you exploit more of the large real processor storage available on even the smallest IBM System z servers. For example, you may choose to utilize data-in-memory techniques such as CICS Shared Data Tables, VSE Virtual Disk, or more and larger buffer pools. Customers with especially large z/VSE environments may experience lower paging rates. In fact, many z/VSE 4.1 environments may be able to run without a page data set (the NOPDS option).

z/VSE V4 exploitation of z/Architecture and 64-bit real addressing is designed to be transparent to user applications.

z/VSE V4 is not designed to support 64-bit virtual addressing. It is not designed to allow 64-bit addressing for user applications. z/VSE V4 supports 31-bit addressing for virtual addresses or data spaces. The maximum size of individual address spaces and data spaces remains limited to 2 GB.

Question:

Are there any prerequisites for z/VSE 4.1?

Answer:

z/VSE V4 operates in z/Architecture mode only. That means it requires an IBM System z9 EC (formerly the z9 109), z9 BC, IBM eServer zSeries 990, 890, 900, or 800 server.

Many VSE customers traditionally run multiple VSE guests under VM. When running under VM, z/VSE 4.1 requires z/VM 5.2 or later.

Question:

When is z/VSE 4.1 available for ordering? When is it generally available for shipment? What are the dates for z/VSE 3.1 end-of-marketing (EoM) and end-of-service (EoS)?

Answer:

z/VSE 4.1 can be ordered beginning the day of announcement (January 9, 2007).

General availability (GA) of z/VSE 4.1 (English version) is planned for March 16, 2007. GA of the Japanese (Kanji) version is planned one month later on April 20, 2007.

z/VSE 3.1 allows customers with Multiprise® 3000 and S/390 G5/G6 (not supported by z/VSE V.1 because they aren't z/Architecture servers) to get relatively current with a supported version and release. As of January 9, 2007 no date has been announced for z/VSE V3.1 EoS.

You can find the status of z/VSE 4.1, z/VSE 3.1, or any VSE version or release at:

www-03.ibm.com/servers/eserver/zseries/zvse/about/status.html

Question:

Where do I get more information on VSE V4?

Answer:

The z/VSE Web site is a comprehensive source of information on z/VSE 4.1, News, Product Announcements, z/VSE solutions, Events, VSE in general, product documentation, selected downloads, plus links to related IBM products, partners, and Independent Software Vendors (ISVs).

ibm.com/servers/eserver/zseries/zvse/

Question:

What changes in pricing are available with z/VSE Version 4?

Answer:

z/VSE V4 offers a new MWLC pricing metric and a sub-capacity option for IBM System z9 EC and z9 BC servers.

Question:

Is MWLC available on all System z servers?

Answer:

The MWLC pricing metric for z/VSE and its sub-capacity option are for IBM System z9 EC and z9 BC servers only.

Question:

If I acquire a new z9 BC server, what pricing metric is available for z/VSE V3?

Answer:

z/VSE V3 on a z9 BC continues to be eligible for the Tiered EWLC (TWLC) pricing metric. When you upgrade to z/VSE V4, the new MWLC pricing metric may be an option.

Question:

If I'm running z/VSE V4 on other than a z9 EC or z9 BC, what pricing metric is available?

Answer:

For supported servers, z/VSE V4 has the same pricing metrics as z/VSE V3. For example, on a z890, Tiered EWLC is available, and for z990 the Group pricing and ELC are available.

Question:

What is z/VSE sub-capacity pricing?

Answer:

IBM z9 servers with sub-capacity MWLC pricing may license and pay for MWLC-priced programs based on the utilization of the Logical Partitions (LPARs) and/or z/VM guest machines where the program executes.

Sub-capacity MWLC pricing is only available on servers where z/VSE V4 is the only VSE operating system and prior versions of VSE use and license have been discontinued.

Please note, if z/VSE V4 under MWLC is to run as a guest under VM, then it must be z/VM 5.2 or later.

Question:

How will sub-capacity pricing be measured?

Answer:

System z9 servers with sub-capacity MWLC pricing may license and pay for MWLC-priced programs based on the utilization of the Logical Partitions (LPARs) and/or z/VM guest machines where the program executes.

In a sub-capacity MWLC environment, LPAR and/or z/VM guest utilization will be measured based on the highest observed rolling 4-hour average utilization. LPAR utilization will be reported each month using the Sub-Capacity Reporting Tool (SCRT) and pricing will be adjusted monthly based upon the Sub-Capacity Report. If you elect sub-capacity MWLC, then you must generate sub-capacity reports and submit them to IBM each month, among other requirements.

Question:

Where can I find more information?

Answer:

Here are a list of urls:

MWLC announcement: **ibm.com**/servers/eserver/zseries/swprice/announce.html

MWLC Web site: **ibm.com**/zseries/swprice/mwlc.html

z/VSE: **ibm.com**/servers/eserver/zseries/zvse/

z/VM Version 5 Release 3

Question:

What did IBM announce on February 6, 2007?

Answer:

IBM continued the evolution of its premier and world-class System z virtualization by announcing a new release of z/VM, Version 5 Release 3 (5.3) with enhancements for scalability, security, and virtualization technology to help support increased workloads on System z servers and enhance its security characteristics.

For more information on z/VM 5.3, see software announcement 207-019 (US) dated February 6, 2007 or ibm.com/eserver/zseries/zvm/.

Question:

When is 5.3 planned to be available?

Answer:

June 29, 2007

Question:

Whom should the new z/VM 5.3 be marketed to?

Answer:

z/VM 5.3 and the IBM Integrated Facility for Linux (IFL) should be marketed to clients who want to run multiple Linux server images on System z servers. This includes clients who want to:

- Run multiple Linux server images hosted by z/VM 5.3 on IFL processors without necessarily increasing the IBM software charges for z/OS, z/OS.e, TPF, z/TPF, VSE/ESA, z/VSE, or other IBM applications running on System z standard processors (CPs).
- Run Linux applications on current or new System z servers within an enterprise. The optional RACF® feature can help provide additional security in a Linux environment.
- Access other System z data and applications from Linux applications running as guests of z/VM on the same server.
- Run more Linux logical servers than can be currently provided by LPARs (up to 60 on z9 EC, up to 30 on z9 BC, z990, and z890, and up to 15 on z900 and z800 servers) or take advantage of the resource-sharing capabilities that z/VM provides for multiple Linux server images.
- Use additional System z hardware that is not supported by Linux on System z, but is supported by z/VM, such as:
 - Exploitation of large real multiprocessor configurations (for example, 32-way) for Linux application workloads that "scale out" on multiple, relatively small n-way virtual machines
 - Exploitation of large real storage for 31-bit Linux systems
 - Exploitation of expanded storage across a set of Linux server images
 - Dynamically configure I/O

Question:

Do z/VM 5.2 and 5.3 support the increase of up to 60 LPARs?

Answer:

Yes. The use of up to 60 LPARs is exclusive to the z9 EC environment and is supported by z/VM 5.1, and later, satisfying the statement of direction in Software Announcement 203-128, dated May 13, 2003.

Question:

z/VM 5.3 includes enhancements in scalability, security, and virtualization technology. Please explain these enhancements.

Answer:

Scalability and constraint relief enhancements include:

- Support for larger Logical Partitions (LPARs) with more than 128 GB of memory
- Up to real 32 processors in a single z/VM image
- Support for the Collaborative Memory Management Assist (CMMA) by which host and guest exchange information to optimize their use and management of memory
- Enhanced memory utilization using Virtual Machine Resource Manager (VMRM) between z/VM and Linux guests
- Support for the Hyper Parallel Access Volume (HyperPAV) function of IBM System Storage devices
- Improved FlashCopy support allows:
 - Specification of up to 12 target minidisks
 - Determine the status of FlashCopy requests
 - Exploit hardware asynchronous cache destage and discard
- Support for the IBM System Storage SAN Volume Controller (SVC) Storage Engine 2145 allows Linux on System z guests of z/VM V5 (all releases) to access the IBM System Storage DS4000 and other OEM SCSI disk devices.

Virtualization technology and Linux enablement

- New guest support for IBM System z Application Assist Processors (zAAPs) and System z9 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 System z9 (z9 EC and z9 BC) servers. zAAPs can be simulated only on 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.
- Usability enhancements for the virtual switch (VSWITCH) and guest LAN environments including enhanced ease-of-use for Virtual LAN (VLAN) and promiscuous mode configuration changes

- Support for IEEE 802.3ad link aggregation and failover support is designed to allow all the OSA-Express2 features that are associated with a virtual switch to be grouped and used as a single “fat pipe”
- Guest use of Modified Indirect Data Address Words (MIDAWs) to allow more flexibility and performance in certain channel programs, as a higher-performing alternative to data-chained channel-command words (CCWs)
- Guest access to the system ASCII console to facilitate recovery of the guest during an emergency
- Additional enhancements for SCSI disk support for Linux users
- Secure Sockets Layer (SSL) server support for additional Linux distributions
- Link aggregation

Network virtualization

- Enhanced virtual network management
- Enhanced failover support for IPv4 and IPv6 devices
- Virtual IP Address (VIPA) support for IPv6

Security enhancements

- New LDAP server and associated client utilities
 - Enhanced system security with support for longer passwords (password phrases)
 - z/VM SSL server enhancements designed to improve security
 - Tape data protection with support for encryption
-

Question:

What new functions are provided with Systems Management in z/VM 5.3?

Answer:

Systems management improvements include:

- Systems Management API has been enhanced with a sockets-based server and new functions to manage virtual machines
 - z/VM systems management functions to be performed from the Hardware Management Console (HMC)
 - New function level for Directory Maintenance Facility (DirMaint™), FL530
 - Enhancements to the Performance Toolkit for VM™
 - Improved guest configuration with a new COMMAND statement
-

Question:

What does improved memory management of guests provide for the z9 EC and new z9 BC customer?

Answer:

z/VM 5.3 adds support for the Collaborative Memory Management Assist (CMMA) on the z9 EC and z9 BC. This VM support, in conjunction with CMMA exploitation in guest operating systems such as Linux for System z, allows the z/VM 5.3 Control Program (CP) host and its guests to communicate attributes for specific 4K-byte blocks of guest memory. This exchange of information allows both host and guest to optimize their use and management of memory, in the following ways:

- CP knows when a Linux application releases storage, allowing CP to select those pages for removal at a higher priority, and to reclaim the page frames without the overhead of paging-out their data content to expanded storage or disk.
- CP recognizes clean disk cache pages, the contents of which Linux is able to reconstruct, allowing CP to bypass paging-out the data contents when reclaiming the backing frames for these pages. If Linux or its application subsequently tries to refer to the discarded page, Linux is notified that the page has been discarded and can reread the contents from disk or otherwise reconstruct them.
- The guest further benefits from the Host Page-Management Assist (HPMA) announced in Hardware Announcement 105-241, dated July 27, 2005. In conjunction with CMMA, HPMA allows the machine to supply fresh backing page frames for guest memory when the guest reuses a previously discarded page, eliminating the need for z/VM hypervisor interception and resolution of these host page faults.

This satisfies the Statement of Direction made in Software Announcement 205-168, (US) dated July 27, 2005.

Refer to the Preventive Service Planning (PSP) bucket for your z9 EC or z9 BC server for required updates and the minimum MCL level. The PSP buckets can be found on Resource Link at: <https://techsupport.services.ibm.com/server/390.psp390>

IBM is working with its Linux distribution partners to provide CMMA exploitation in future Linux on System z distribution releases or service updates.

Question:

Where can I go for more detailed information about z/VM V5 and Linux and System z technology?

Answer:

See the following links for additional information:

General z/VM information:

- z/VM 5.1 resources: [ibm.com/eserver/zseries/zvm/zvm510](https://www.ibm.com/eserver/zseries/zvm/zvm510)
- z/VM 5.2 resources: [ibm.com/eserver/zseries/zvm/zvm520](https://www.ibm.com/eserver/zseries/zvm/zvm520)
- z/VM V.3 resources: [ibm.com/eserver/zseries/zvm/zvm530](https://www.ibm.com/eserver/zseries/zvm/zvm530)

FAQs:

- z/VM 5.1 [ibm.com/servers/eserver/zseries/faq/pdf/zvm5_1_faq.pdf](https://www.ibm.com/servers/eserver/zseries/faq/pdf/zvm5_1_faq.pdf)
- z/VM 5.2 [ibm.com/servers/eserver/zseries/faq/pdf/zvm5_2_faq.pdf](https://www.ibm.com/servers/eserver/zseries/faq/pdf/zvm5_2_faq.pdf)
- z/VM 5.3 [ibm.com/servers/eserver/zseries/faq/pdf/zvm5_3_faq.pdf](https://www.ibm.com/servers/eserver/zseries/faq/pdf/zvm5_3_faq.pdf)

Linux for System z

Question:

What is the difference between the Linux 64-bit and 31-bit distributions?

Answer:

Linux 64-bit distributions support the IBM System z9 and eServer zSeries architecture, which is 64-bit enabled. They are supported on the recently introduced IBM System z9 Business Class (z9 BC) and IBM System z9 Enterprise Class (z9 EC), the IBM eServer zSeries 990 (z990), IBM eServer zSeries 890 (z890), IBM eServer zSeries 900 (z900) and IBM eServer zSeries 800 (z800) models. These systems are also capable of running 31-bit Linux distributions, and can run 31 and 64-bit distributions concurrently.

Linux 31-bit distributions are supported on the IBM System z9 Business Class (z9 BC), IBM System z9 Enterprise (z9 EC), the IBM eServer zSeries servers (z990, z890, z900, z800), and the 9672 S/390 Parallel Enterprise Server™ Generation 5 and 6, and the S/390 Multiprise 3000 servers. System z hardware can also run code built for the 31-bit mainframe systems.

Question:

How can Linux on System z help to take back control of your IT infrastructure?

Answer:

Linux running on System z servers can be the focal point for IT infrastructure simplification by providing virtual server growth on a single System z9 instead of physical expansion on distributed servers. It is also possible to simplify the distributed server environment through centralization to one operating environment instead of several diverse hardware/software platforms.

Linux on System z can be the foundation for an on demand IT infrastructure, because it combines the legendary scalability, security and reliability of IBM mainframes with the flexibility and open standards of the Linux operating system, and can help to optimize response to business challenges while reducing maintenance effort.

For more details read the Linux on System z brochure:

ibm.com/systems/z/library/brochures/pdf/gm130240.pdf

Question:

What are advantages that are unique to Linux on System z as opposed to Linux running on other hardware platforms?

Answer:

It is a synergistic combination of Linux, IBM System z servers, z/VM® virtualization and systems management capabilities, unavailable on other platforms, that provides customers unique value propositions to better control their IT infrastructure.

- Linux benefits from the IBM mainframe security, robustness and processes
- Linux benefits from the IBM mainframe virtualization technology – over 35 years of constant innovation
- Linux benefits from the IBM mainframe workload and systems management capabilities

Question:

What are the Linux benefits from the IBM mainframe virtualization technology?

Answer:

IBM mainframe virtualization technology and Linux enable a substantial IT infrastructure optimization. Customers can run from tens to hundreds of virtual Linux servers on a single IBM mainframe – easily deployed and maintained without adding costs associated with new server footprints in distributed environments.

The benefits are:

- Virtual growth instead of physical expansion on Intel or RISC servers
- Efficient resource utilization because test, development, education, backup and other servers can share resources, applications or utility code, which can reduce software licensing costs
- Network simplification through a highly virtualized internal network. Security is also enhanced since the network is internal to the System z, and not exposed and vulnerable to eavesdropping devices commonly called “sniffers”.
- Easier systems management through Single Point of Control for administration and operation

The advanced mainframe virtualization technology is a major reason that sets Linux on the mainframe apart from other operating environments.

Question:

Where can I get commercial Linux for System z distributions?

Answer:

Commercial Distributions are available from IBM Linux Distribution Partners Red Hat (www.redhat.com), and Novell SUSE LINUX (www.suse.com). Commercially available Linux distributions can include the Linux operating system enabled for specific hardware platforms, an assortment of device drivers, routines for installation, and value add programs like Web servers and shells.

Typically these distributions are delivered over the Internet or packaged on a physical deliverable. IBM suggests that these distributions be your first choice for production environments because of their availability of service and support from the Linux Distributor, or IBM Support Line.

Question:

On which machine types is the Integrated Facility for Linux (IFL) available?

Answer:

The IFL is available for all System z servers currently in production. They are the recently introduced IBM System z9 BC and IBM System z9 EC and the eServer zSeries models z900, z890, z900 and z800.

Question:

Where can I get the most current information on IBM middleware for Linux on System z?

Answer:

For the most current information refer to the Web site at:

ibm.com/linux/matrix/linuxmatrixhwz.html

Question:

What software products /solutions are available to be used with Linux on System z?

Answer:

Software products are available to support major themes such as IT infrastructure simplification or business integration.

The software products covers the areas of:

- **Application Development**
Mainframe servers are capable of isolating multiple execution environments with LPAR and z/VM functionality. One or more virtual Linux servers can be easily created for each developer, isolated from other developer or production servers.
- **Application Deployment**
Some products provide an application development and deployment environment to develop on non-Linux and deploy the application on Linux. Some offer migration and bridging components to move modern enterprise applications to Linux.
- **Application Serving for z/OS and z/VSE**
Linux opens the world for mainframe customers to a broader application portfolio – in the form of ISV and Open Source applications. For most current information on available software products, refer to: ibm.com/systems/z/linux/apps/all.html
- **Business Connectors**
Business connectors are the “code” for the integration of the various systems, including: DB2 Connect™, CICS Transaction Gateway, IMS Connect, WebSphere MQ, and VSE Connectors.
- **Database Serving**
DB2, Oracle, Informix, and MySQL are available databases for the mainframe Linux environment.
- **E-mail and Collaboration**
Linux on IBM System z provides an attractive platform which can support a variety of customer requirements from simple messaging to sophisticated collaboration / groupware implementation. IBM Lotus Domino is a security-rich collaboration platform that helps users work more efficiently, supporting industry standards.
- **Infrastructure Servers**
Linux on IBM System z excels in doing the work of many smaller servers on a single mainframe server. Many of the smaller servers are the so called ‘Infrastructure’ servers such as file, print, DNS, FTP, NFS, security, etc.

- **Security**
IBM System z servers running Linux offer an attractive foundation on which to build a comprehensive IT security infrastructure. Numerous security applications and tools are available from IBM and other ISVs that provide a complete and cost-effective security infrastructure for the Linux on System z environment.
- **SNA Network Communication**
IBM Communication Controller for Linux (CCL) allows customers to keep their business data and applications on the mainframe while moving Network Control Program (NCP) functions, and the associated CPU cycles, to Linux on System z.
 - CCL offers a smooth transition and attractively priced alternative for aging 3745 and 3746 Communication Controller networking hardware. CCL provides a migration path for enterprises using traditional Systems Network Architecture (SNA) including SNA Network Interconnect (SNI) to communicate with their customers and business partners.
 - IBM Communications Server for Linux (CS) enables applications to run over multiple network protocols, acting as a multiprotocol gateway. CS also works as a Telnet server, providing systems network architecture (SNA) network access to client applications running anywhere in your TCP/IP network. Also included is Enterprise Extender, a leading-edge solution for accessing SNA applications over TCP/IP networks.
- **Systems Management**
Linux on IBM System z can help to optimize IT resources, simplify systems management and reduce the complexity, costs, and efforts associated with today's heterogeneous IT environment.
 - IBM Tivoli products for Linux on IBM System z include a comprehensive suite of solutions that include products to monitor and manage workloads.
 - IBM Virtualization Engine™, and Infrastructure Services can help to position System z to play a central role as the hub for enterprise-wide infrastructure management by making it easier to manage the heterogeneous IT assets.
 - In addition, providing the functions necessary to manage and maintain the virtual Linux servers on System z is a key element in running a highly efficient and successful optimized infrastructure.
 - Workload Management for Linux on System z with z/VM
 - Performance Toolkit for VM
- **Web Application Serving**
Linux on IBM System z is an excellent choice for customers who wish to re-host Web applications that have been running on competitive hardware. WebSphere on mainframe Linux is no more difficult to deploy than WebSphere on distributed systems. Customers still use the logical distributed architecture model, but gain the advantage of a single physical model that is closer to the existing data since it is on the same mainframe.
Mainframe technology and business connectors helps customers get rapid data access to help integrate Linux, z/OS, z/VM, and z/VSE systems and the vast amount of data and applications they support.

Question:

Does IBM maintain a list of software vendor products that are being offered for Linux on System z?

Answer:

Yes. You can access a list of software developer products available for Linux on System z at:

- ibm.com/systems/z/linux/apps/all.html
- ibm.com/servers/eserver/zseries/solutions/s390da/linuxproduct.html

FYI, the information on these Web pages has been provided to IBM by the Software Developers and is subject to change. Any dates presented as future deliverable dates for General Availability are also subject to change as well at any time. Any questions on these products or specific delivery dates should be addressed to the supplier of those products.

IBM Global Financing (IGF)

Question:

How does financing fit with the on demand capabilities of the System z9 products?

Answer:

IGF can finance the two capacity on demand offerings — Capacity Upgrade on Demand and On/Off Capacity on Demand. IGF financing can assist customers to more easily acquire incremental server resources in response to demand changes.

Capacity Upgrade on Demand financing

When customers know they're going to need additional capacity in the future, but don't know when, Capacity Upgrade on Demand financing can provide the perfect balance of flexibility and affordability. Customers get the ability to activate pre-installed additional processors at any time during their original lease term. An extension to the original lease term keeps payments low and reflects the longer life expectancy of the system. The cost of the upgrade is rolled in as an affordable addendum lease, so the customer gets all the performance and capacity benefits of a higher-capacity system, for just a modest increase to their monthly payment.

On/Off Capacity on Demand financing

With On/Off Capacity on demand financing, your customer's system is installed with both activated and inactivated processor capacity. The customer pays a monthly fee over a fixed term for only the base capacity activated at the time of installation.

When additional capacity is turned on and off as needed, they pay an additional usage fee for the incremental increase in capacity, billed in processor-days (with a minimum increment of one processor-day). These additional fees can be assessed as one-time charges, or financed on a short-term revolving type of credit arrangement for a period of twelve months, so the customer can "smooth out" the billing for peak capacity to a more manageable payment structure. At the end of the short-term financing period the remaining balance, if any, can be paid off in one lump sum, or included in a recalculated lease extension for the remaining term of the base agreement, or rolled into a new term.

Question:

What capabilities exist for financing non-HW items that are part of a business transformation project utilizing a System z9?

Answer:

IBM Project Financing is a single-source financing solution for large-scale, multivendor business transformations that can include consulting services, infrastructure investment and business process implementation. Highly customizable structures can be tailored to even the most complex project plans, and can cover the entire project's life cycle. IBM Project Financing facilitates project approval, preserves cash flow and credit lines, and lets you better match costs to projected benefits to enable self-funding.

Question:

What about financing non-hardware items like this outside the scope of a large project?

Answer:

IGF can finance all hardware, software and services in a single contract.

Question:

How do I engage IBM Global Financing?

Answer:

The first step is to contact your IBM Global Financing representative. If you don't already have a representative, please click here: "[ibm.com/financing/servers](https://www.ibm.com/financing/servers)" to find one.

Question:

Where can I find current IBM Global Financing pricing information?

Answer:

Please contact your IBM Global Financing representative for complete pricing information. If you don't already have a representative, please click here "[ibm.com/financing](https://www.ibm.com/financing)" to find one.

Global Technology Services

Fiber Cabling for your IBM System z9

Question:

There is a broad range of connectivity on the System z platforms. How can I help to be sure I have the right cables to attach to my infrastructure?

Answer:

We agree — there is a wide variety of connectivity. Also, the fiber requirements are changing and getting more complex. In order to help manage the complexity and plan your environment, cables attaching your mainframe to your IT infrastructure zSeries fiber cabling service and enterprise fiber cabling offered are available from IBM Global Services. These services offerings allow you to plan for your environment as a whole instead of product by product. Without a cabling strategy for your data center you may spend unnecessary time and money initiating changes. The zSeries Fiber Cabling Service helps you effectively manage your connectivity needs on your zSeries and System z9 platforms.

Question:

Will there be a specific System z9 Fiber Cabling Service offering?

Answer:

No. The zSeries Fiber Cabling Service will support both the zSeries and the System z9 platform.

Question:

Why is IBM offering the zSeries fiber cabling service offerings?

Answer:

The complexity of specifying fiber cabling connectivity has increased dramatically over the past ten years because of new I/O function, higher bandwidth, significant growth in fiber cabling technology, a migration to small form factor connectors, and the emergence of new cabling standards. The number of cabling options now requires detailed planning. A customer must make cabling decisions based on goals that best fit their business and environment. The combination of the right IBM fiber cabling service offering and trained cabling specialists can help resolve the cabling issues and find the right solution for the customer's server installation.

Question:

What are the zSeries fiber cabling services options?

Answer:

The zSeries fiber cabling service offering has three options to choose from. Each option has been created to help handle different customer situations for planning and managing individual jumper cable situations.

- Option 1 is a comprehensive fiber optic jumper cable package offering fiber cable planning, new cables, installation and documentation. This option applies to new server

installations as well as server upgrades where either an entire new set of fiber cables is desired or an add of I/O ports is part of the upgrade.

- Option 2 is a fiber cable migration offering which can be used when reusing existing fiber cables in a server upgrade scenario. This option provides for the planning, relabeling, moving and documentation of existing fiber cables to the new system.

Option 1 (new cables) and Option 2 (reused cables) will be frequently used together for a typical server upgrade scenario. Option 1 and 2 are included in a single services contract for ease-of-use and fast deployment.

- Option 3 provides new cables and installation in situations where the fiber cable planning has been completed and a cable list is already available. Option 3 is always a separate IBM Networking Services contract.

Fiber cables are available for ESCON, FICON, OSA, ISC-3, Sysplex Timer and ETR.

Question:

Can I use my currently installed IBM jumper cables with the System z9 platform?

Answer:

The answer, today, is “maybe” depending on the application, the customer upgrade strategy, the connectors supported on the system and on the devices, as well as the customer cabling strategy. There are also other alternatives that should be considered even if reuse of existing cables is possible. For example, new jumper cables might be more desirable than reusing existing cables with two meter conversion kits. New jumper cables with the correct fiber connectors on both ends eliminates one connection in the optical path (link reliability), avoids the 2 meter cable/connector “mess” immediately under the server floor (manageability), can help eliminate the stress on the connector when moving the cable (cable failure), and helps provide for easier connectivity changes (faster upgrade and change times). Similarly, using MCP cables when reusing multimode fiber optic cables with long wavelength (LX) transceivers can facilitate the use of existing fiber optic cables but may impact less loss budgets and may be affected by link data rates. (Note that MCP cables are not supported for data rates over 1 Gbps link data rates only). Another alternate solution is an enterprise fiber trunking solution to help make the fiber infrastructure even more useable, changeable, upgradeable and manageable. Today, the customer needs a fiber cabling strategy to guide decisions and provide a workable infrastructure. IBM Networking Services can help.

Question:

I have heard about a structured fiber optic cabling infrastructure. What is this?

Answer:

A structured fiber cabling environment means fiber optic trunks where, for example, up to 72 FICON channel fiber cables can be contained in approximately a one-inch trunk. Other hardware would include patch panels, patch cables (between patch panels), under floor boxes, fiber cable harnesses (six channels per harness), mounting brackets and associated hardware are involved. A structured infrastructure may eliminate the hundreds (thousands) of jumper cables that may currently be in the data center. It is a very manageable infrastructure which moves cabling problems away from the server or switch and into the patch panel which, by its very nature, allows flexible change. FICON is the key protocol being trunked today. Key

words for a structured environment are useable, changeable, upgradeable, and manageable. Any data center with two or more servers should consider a structured infrastructure.

Question:

I have heard about a structured fiber optic cabling infrastructure and the standard TIA-942. What is this?

Answer:

The purpose of the TIA-942 Telecommunications Infrastructure Standard for Data Centers is to provide guidelines for the design of a data center or computer room. It is intended for use by designers who need a comprehensive understanding of data center design, including facility planning, the cabling system, and network design. Structured cabling is one of the design elements of this standard.

Question:

What are the enterprise fiber cabling service options?

Answer:

The enterprise fiber cabling service offering has two options to choose from. Each option has been created to handle a different customer situation for planning and managing individual trunked cable situations.

- Option 1 is the zSeries fiber optic trunk cabling package. This option offers a trunking solution for a single System z platform being installed into an already structured (fiber trunk) environment. This option includes System z planning, fiber optic trunk commodities, installation, and documentation. It should be noted that the planning, procurement and installation of individual jumper cables, if required for the server installation, can also be included with this option which allows all fiber connectivity requirements to be fulfilled in a single IBM Networking Services contract.
 - Option 2 is the enterprise fiber cabling option which was formally called Fiber Transport Services. This option is designed to provide a custom structured (trunked) solution for small, medium or large enterprises. Whereas the zSeries fiber optic trunk cabling package was limited to a single zSeries server, this option will provide a fiber optic infrastructure for the entire data center allowing the interconnection of servers, switches, and devices. This option includes fiber optic trunk planning, fiber optic trunk commodities, installation, and documentation for the data center. The planning, procurement and installation of individual jumper cables, if required, can also be included with this option which allows all fiber connectivity requirements to be packaged into a single IBM Networking Services contract.
-

Question:

Will there be a separate System z9 Fiber Optic Trunk Cabling Package for the z9 EC?

Answer:

No. Customers should use the zSeries Fiber Optic Trunk Cabling Package when they need a trunking solution for a single z9-109.

Implementation Services for Parallel Sysplex

Question:

What did IBM Global Technology Services (GTS) announce on April 18, 2007?

Answer:

IBM GTS announced the "Implementation Services for Parallel Sysplex (6948-74Z)" on April 18, 2007.

Question:

What does the "Implementation Services for Parallel Sysplex" provide?

Answer:

The "Implementation Services for Parallel Sysplex" Service Product consists of eight modules. Each module is separately priced according to a predefined scope of work and can be acquired individually. The eight modules are:

- z/OS and Sysplex infrastructure review
 - z/OS and Sysplex infrastructure planning, design and implementation
 - z/OS and Sysplex availability and performance assessment
 - z/OS Parallel Sysplex configuration assistance
 - z/OS Sysplex operator training
 - z/OS Sysplex maintenance strategy
 - z/OS Sysplex maintenance environment design and implementation
 - z/OS Sysplex production cutover and deployment planning
-
-

Question:

Why did IBM come up with this Service Product?

Answer:

IBM came up with this Service Product to help customers to exploit the capabilities and benefits which a Sysplex environment can offer, which includes data sharing, high availability and incremental growth.

Question:

Why did IBM offer this Service Product in a modular fashion?

Answer:

IBM offered this Service Product in a modular fashion to ease customers in budgeting and buying what's required. Customers at different stages of Sysplex implementation will require assistance in different disciplines.

Question:

Where can I get more information about this Service Product?

Answer:

You can either visit our Web site at ibm.com/services/server or contact an IBM representative to get more information on the offering.

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 Dynamic ICF expansion?

Answer:

IBM intends to remove the Dynamic ICF expansion function from System z and zSeries processors.

Question:

What are the Statements of Direction (SOD) concerning future enhancements to the z/OS XML System Services component of z/OS?

Answer:

On April 18, 2007, there are SODs as follows:

IBM is intending to enable the z/OS XML to take advantage of zAAPs.

IBM is intending to enable the z/OS XML to fully take advantage of zIIPs.

IBM intends to extend and expand on the use of z/OS XML System Services enabled for zAAP specialty processors as the basis for future enhancements:

- IBM intends to enhance the XML Toolkit for z/OS so that eligible workloads may exploit the z/OS XML component
 - IBM intends to add validating parsing to the z/OS XML component
-
-

Question:

What additional detail is available for the SODs concerning z/OS XML System Services?

Answer:

In z/OS V1.8, IBM introduced a new system component of z/OS, z/OS XML System Services (z/OS XML), a system-level XML parser integrated with the base z/OS operating system and designed to deliver an optimized set of services for parsing XML documents (z/OS XML has also been made available on z/OS V1.7). The initial beneficiaries of this system component were middleware and applications requiring high performance non-validating XML parsing. z/OS XML may currently be accessed by an Assembler programming interface and one of the first exploiters, DB2 9 for z/OS, uses this Assembler interface for XML native support. IBM plans to add C/C++ support for z/OS XML with z/OS V1.9 (satisfying the Statement of Direction in Software Announcement 206-039, dated February 28, 2006).

IBM is announcing its intent to enable the z/OS XML component to take advantage of zAAPs. This future enhancement means that middleware and applications requesting z/OS XML System Services (for example DB2 processing via local connection) will have z/OS XML System Services processing execute on the zAAP. Specifically, all z/OS XML System Services parsing executing in TCB mode will be eligible to run on the zAAP.

In addition, IBM is announcing its intent to enable the z/OS XML component to fully take advantage of zIIPs, when present. With respect to DB2, z/OS XML processing may be partially directed to zIIPs when utilized as part of a distributed request (like DB2 DRDA). The future enhancement will further benefit DB2 by directing the full amount of the z/OS XML System Services processing to zIIPs when it is utilized as part of any workload running in enclave SRB mode (like DRDA). Specifically, all z/OS XML System Services parsing that is executed in enclave SRB mode will be eligible to run on the zIIP.

zAAPs and zIIPs are designed to help free-up general computing capacity and lower total cost of operation for select new workloads such as Java, business intelligence (BI), ERP, CRM and IPSec encryption on the mainframe. IBM does not impose software charges on zAAP and zIIP capacity. Collectively, z/OS XML System Services support of zAAP and zIIP means that you can have the advantages of XML processing on z/OS with Total Cost of Ownership (TCO) benefits of either the zIIP or the zAAP processor regardless of the invocation environment.

As part of a comprehensive plan, IBM intends to extend and expand on the use of z/OS XML System Services enabled for zAAP specialty processors as the basis for future enhancements:

- IBM intends to enhance the XML Toolkit for z/OS so that eligible workloads may exploit the z/OS XML component - this extends zAAP exploitation to the XML Toolkit for z/OS.
 - IBM intends to add validating parsing to the z/OS XML component - this extends zAAP exploitation for XML validating parsing as well.
-

Question:

What statements of direction has IBM made for IBM Tivoli?

Answer:

To support customers who are moving to secure data transmission across open networks using IPSec, IBM Tivoli intends to deliver the ability to monitor IPSec in OMEGAMON XE for Mainframe Networks 4.1 this year. IBM Tivoli OMEGAMON XE for Mainframe Networks will provide users the ability to monitor the use of IP filters and the performance of IPSec tunnels for the TCP/IP stacks on a z/OS system. Users will be able to identify potential network attacks or configuration problems with IP filters and IPSec security associations (SAs). The information provided may be used to perform problem determination and to identify possible corrective actions.

Question:

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

Answer:

The z/OS 1.9 preview announcement letter 207-018, dated February 6, 2007, contained the following Statement of General Direction:

z/OS.e 1.8 (5655-G52) is planned to be the last release of z/OS.e. Marketing, ordering, support, and service for z/OS (5694-A01) remain unaffected.

z/OS.e 1.8 remains orderable until its planned withdrawal from marketing in October 2007. In accordance with the z/OS (5694-A01) and z/OS.e service policy (to provide service support for each release for three years following its general availability date), IBM intends to withdraw service for z/OS.e 1.8 in September 2009. For information about other z/OS.e releases, refer to the Related information section for withdrawal from service dates.

IBM plans to provide a new pricing alternative for z/OS, System z New Application License Charges (zNALC). zNALC will replace New Application License Charges (NALC) and z/OS.e, and is intended to be IBM's strategic z/OS offering for new workloads. zNALC is planned to be available March 16, 2007. Refer to Software Announcement 207-006 for additional information.

Question:

IBM made the following statements of direction for Server Time Protocol in the October 10, 2006 announcement have they been fulfilled by the April 18, 2007?

Answer:

No, the following two statements of direction are not fulfilled by the April 18, 2007 announcement:

1. Enhanced Coordinated Server Time Accuracy: IBM intends to enhance the accuracy of initializing and maintaining Coordinated Server Time to an international time standard such as Coordinated Universal Time (UTC). The then current server is planned to have the capability of attaching to an external time source, such as a Global Positioning System (GPS) receiver.
 2. Network Time Protocol (NTP) client support: IBM intends to enhance the STP design to provide Network Time Protocol (NTP) client capability, so that Coordinated Server Time may be initialized and maintained to time provided by an NTP server. The purpose of this function is to allow customers to have the same time across an enterprise comprised of heterogeneous platforms.
-

Question:

What statements of direction has IBM made for ICB-3?

Answer:

IBM intends to phase out the Integrated Cluster Bus-3 links (ICB-3 links) over time. IBM plans to support ICB-3 links through the lifecycle of the System z9 platform.

Question:

What statements of direction has IBM made for IBM TotalStorage?

Answer:

IBM plans to enhance the System Storage TS3400 Tape Library by adding support for System z (z/OS, z/VM, z/VSE and TPF). Additionally, IBM plans to support system-managed encryption in a System z environment with the TS3400.



Copyright IBM Corporation 2007
IBM Corporation
New Orchard Rd.
Armonk, NY 10504

Printed in the United States of America, 04/2007

All Rights Reserved

IBM, the IBM logo, IBM eServer, e-business logo, CICS, DB2, DB2 Connect, DB2 Universal Database, Domino, DRDA, DS4000, DS6000, DS8000, ECKD, ESCON, FICON, FlashCopy, GDPS, HiperSockets, HyperSwap, IMS, Lotus, MQSeries, Multiprise, OMEGAMON, On Demand Business logo, Parallel Sysplex, Performance Toolkit for VM, QMF, Resource Link, RMF, S/390, S/390 Parallel Enterprise Server, Sysplex Timer, System Storage, System z, System z9, Tivoli, TotalStorage, Virtualization Engine, VM/ESA, VSE/ESA, VTAM, WebSphere, z/Architecture, z/OS, z/VM, z/VSE, zSeries, and zSeries Entry License Charge are trademarks or registered trademarks of the International Business Machines Corporation.

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.

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.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product and service names may be trademarks or service marks of others

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

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

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