
**IBM System z9
Introduction
October 2006**

IBM System z9 Business Class (z9 BC)
Frequently Asked Questions

Worldwide



z9 BC FAQ PACKAGE

Table of Contents

Announcement Overview	3
Community	7
z9 BC Hardware	9
z9 BC Upgradeability	12
Availability Enhancements	13
IFL, zAAP and zIIP specialty engines	15
Clarification on a few DB2 terms for better zIIP understanding	21
z9 BC Pricing	23
Performance	27
Networking Enhancements	30
I/O Enhancements	32
On/Off Capacity on Demand (On/Off CoD) and CBU	36
Cryptographic Enhancements	40
Parallel Sysplex Enhancements	48
GDPS Enhancements	49
IBM Global Financing (IGF)	50
Fiber Cabling for your IBM System z9	52
z/VSE V4	56
Announcement Overview - z/VM	58
Statements of Direction	60

Announcement Overview

Question:

What was announced April 27, 2006?

Answer:

First, IBM announced the IBM System z9™ Business Class (BC) which brings many of the System z9 platform values to a wider audience, both midrange and small enterprise businesses. The z9 BC, with two models, offers a low entry point (26 MIPS), granular growth, more specialty engines, subcapacity pricing, On/Off Capacity on Demand enhancements and improved I/O designed to meet existing and changing demands for traditional and new workloads.

IBM also introduced the hardware availability of our new System z9 Integrated Information Processor (IBM zIIP) and improvements to our FICON® technology along with both 2-port LX and 4-port cards.

New capabilities were announced for the IBM System z9 109 (z9-109), now known as the System z9 Enterprise Class (z9 EC). These capabilities include subcapacity versions, delivered through 24 new capacity settings, and the general availability of the zIIP.

Question:

What was announced 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 Sysplex Timers. 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/psa/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:

Why was the System z9 naming been 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 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 will now be known as the IBM 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 new 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 August 22, 2003, provided a framework for planned future investments 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. This announcement, which includes the newest member of the family - the z9 BC, will extend the System z9 family to provide our latest technology to our small and medium mainframe customers.

In addition to the innovation we have announced with the introduction of the z9 EC and z9 BC, we have 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:

What is important about the System z9 Business Class announcement on April 27, 2006?

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 make the z9-109 an ideal platform choice for supporting central corporate databases and mission-critical enterprise-wide applications.

The z9 BC offers extensive growth options and excellent price/performance for those customers requiring a lower-capacity entry point, more granular growth options and fewer engines than offered with the z9 EC.

System z9 BC leadership in its technical capabilities being delivered, as compared to the IBM eServer zSeries 890 (z890) include new levels of granularity, additional engines, availability of the new IBM zIIP and continued improvements to FICON, security and availability.

In accordance with our customers' requirements, the z9 BC offers:

- 73 capacity settings for choosing the right size mainframe and helping to optimize software costs.
- Additional Processor Units (PUs), up to 7, to grow new workloads to help meet business needs.
- The zIIP, when combined with the enabling PTFs for z/OS® (and z/OS.e) 1.6 or 1.7 and DB2® UDB z/OS V8 (planned to be available via web deliverable in June 2006), is designed to help improve resource utilization and lower the cost of eligible workloads, and position the z9 BC as the data hub of the enterprise.
- Enhancements to FICON Express4 for potential improvements to I/O capacity and performance with the next generation of FICON/FCP. This new feature may offer increased channel aggregation and shorter backup windows to help reduce the cost of storage operations and infrastructure. A new 2-port LX card, available exclusively on the z9 BC, may help reduce costs associated with migration from ESCON® to FICON. IBM System Storage™ has announced plans to offer 4 GB devices in 2Q06.

Question:

What is important about the System z9 Business Class announcement on October 10, 2006?

Answer:

The October 10, 2006 announcement of the z9 BC brought many of the System z9 platform values to a wider audience, both midrange and small enterprises. The z9 BC offers a low entry point, very granular growth, more specialty engines, subcapacity pricing, On/Off Capacity on Demand enhancements and improved I/O designed to meet existing and changing demands for traditional and new workloads. While a raised floor environment is well-suited for the z9 BC, IBM is now offering a non raised floor option that may help to ease the cost of entry into the z9 BC by avoiding the need to build a raised floor for some infrastructures.

Question:

What are the planned availability dates for the features and functions announced on April 27th that are associated with IBM System z9 BC?

Answer:

All features and functions for the z9 BC, including all MES upgrades and upgrades to the z9 EC Model S08, that were announced on April 27, 2006 are currently available. The October 10th announcements of the STP feature and the non raised floor support are available November 17th, 2006.

Question:

What software is supported on the z9 BC?

Answer:

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

OS Software	<ul style="list-style-type: none">▪ z/OS (and z/OS.e)▪ z/VM▪ Linux[®] on System z9 (64-bit and 31-bit distribution)▪ z/VSE™▪ VSE/ESA™▪ TPF	<ul style="list-style-type: none">▪ z/OS 1.4 and subsequent releases*▪ z/VM 4.4 and subsequent releases *▪ Red Hat RHEL 3, Red Hat RHEL 4, SUSE SLES 8, SUSE SLES 9▪ z/VSE 3.1▪ VSE/ESA 2.7 *▪ TPF 4.1 and z/TPF 1.1
-------------	--	---

* Note: Support for z/OS V1.4 and 1.5 will end on March 31, 2007. Support for z/VM V4.4 will end September 30, 2006. Support for VSE 2.7 will end February 28, 2007

Question:

Where can I find the FAQs on the new z9 EC?

Answer:

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

Question:

Where can I find the data sheets associated with the z9 BC, z9 EC and the z/VM 5.2 (also announced on April 27, 2006)?

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

z/VM 5.2 Data sheet

ibm.com/common/ssi/fcgi-bin/ssialias?infotype=pm&subtype=sp&pdf=yes&appname=SS_ZS_USEN&htmlfid=ZSD01145USEN.

Community

Question:

I know 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 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 initiative, based on customer demand, is intended to provide 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 1200 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 System z operating systems 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 Global Solutions directory [IBM Global Solutions Directory](#).

z9 BC Hardware

Question:

What is the machine type number for the z9 BC?

Answer:

The machine type number is 2096.

Question:

What models of the z9 BC are available?

Answer:

z9 BC offers two standard models, the R07 and S07 with up to 8 Processor Units (PUs) versus 5 PUs on z890.

- Model R07 maximum PUs by type are: up to 3 General Purpose Processors, 1 standard System Assit Processor (SAP), 6 Specialty engines (Integrated Facility for Linux [IFL], Internal Coupling Facility [ICF], zIIP, and System z Application Assist Processor [zAAP]), 3 additional optional SAPs
 - Model S07 maximum PUs by type are: up to 4 General Purpose Processors, 1 standard SAP, 7 Specialty engines (IFL, ICF, zIIP, and zAAP), 3 additional optional SAPs
-
-

Question:

How does the z9 BC granularity differ from z890?

Answer

The z890 offered 28 capacity settings, and the z9 BC offers 73 capacity settings.

- Model R07 offers 20 options
 - Model S07 offers 53 options, starting with more capacity than offered on the R07.
-
-

Question:

What is different about the z9 BC R07 capacity setting A01?

Answer:

The smallest R07, capacity setting A01, allows us to offer a processor at 26 MIPS, comparable to the z890 capacity setting 110. This entry processor is sized and priced aggressively (from both a hardware and software perspective) to meet our smallest mainframe customer requirements for a lower entry point. The A01 has most of the same features as the other 19 capacity settings; however it has smaller I/O configuration/expansion capabilities. Customers interested in the A01 may not need horizontal growth of their I/O capabilities and so its I/O capacity more closely matches that of the z890 capacity setting 110. The A01 has one I/O cage and can support up to 240 ESCON channels, 64 FICON Express4, up to 24 ports using OSA-Express2, up to 8 ICB-4 links and up to 8 Crypto Express2 adapters.

Question:

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

Answer:

Model information is returned via the STSI instruction. The Model Capacity Indicator field in SYSIB 1.1.1 contains model names in the form nxx, where n indicates the subcapacity size (A-Z) and where xx indicates the number of active Central Processors.

Question:

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

Answer:

One standard SAP and an additional three optional ones can be configured.

Question:

Can I buy a z9 BC that has only IFL or ICF processors without including a general purpose processor?

Answer:

Yes. Similar to the z890, you can order only IFL or ICFs in a z9 BC Model S07. The Model R07 must have at least one general purpose processor.

Question:

Will the z9 BC 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 BC, current HMCs with Token-Ring may be carried forward to the z9 BC on an upgrade from a z800 or z890.

Question:

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

Answer:

No. OSA-Express Token-Ring is not offered on the z9 BC and cannot be carried forward to the z9 BC on an upgrade from the z800 or z890.

Question:

What is the new separate Processor Unit (PU) management enhancement that provides increased flexibility for managing PUs on z9 BC?

Answer:

Processor Units (PUs) defined as Internal Coupling Facility (ICF) processors, Integrated Facility for Linux (IFL) processors, System z9 Integrated Information Processors (zIIPs), or System z Application Assist Processors (zAAPs) are now managed separately. In the past, ICF processors, IFL processors, and zAAPs were grouped together for allocation within and across the LPARs.

The separate management of PU types is designed to enhance and simplify capacity planning and management of the configured LPARs and their associated processor resources.

Separate PU management is exclusive to the System z9 and is supported by z/OS.

Question:

What does the improved LPAR weight management of general purpose processors (CPs) and zAAPs provide?

Answer:

The ability to specify a separate LPAR weight for shared zIIPs and zAAPs helps to enhance and simplify capacity planning and management of the configured LPARs and their associated processor resources.

For LPARs that have CPs, zIIPs and/or zAAPs configured, a new zAAP/zIIP weight specification is provided to allow a new unique LPAR weight specification for shared zAAPs to be defined. The existing LPAR shared processor weight specification is now applied only to the CPs configured to the LPAR.

In the past, the existing shared processor weight specification was applied to both the shared CPs and to shared zAAPs and zIIPs configured to the LPAR.

Improved weight management of CPs, zIIPs, and zAAPs is exclusive to the System z9 and is transparent to the operating systems.

z9 BC Upgradeability

Question:

What are the upgrade paths from a z800 or a z890 to the new z9 BC?

Answer:

The z800 model 004 and all models of the z890 can upgrade directly to either model of the z9 BC.

Question:

What is the upgrade path from the z9 BC to a z9 EC?

Answer:

The z9 BC Model S07 can upgrade directly to a z9 EC Model S08.

Question:

What is the upgrade path for the z9 BC Model R07?

Answer:

The z9 BC model R07 can upgrade within the Model R07 and directly to the Model S07.

Question:

What are the migration requirements for customers with Crypto Express2 features on z890 servers?

Answer:

The Crypto Express2 feature is supported on z890 servers and can be carried forward on an upgrade to the z9 BC.

Question:

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

Answer:

The design of the Crypto Express2 feature allows for nondisruptive upgrades.

Question:

What are the migration requirements for customers with FICON Express2 features on z890 servers?

Answer:

The FICON Express2 feature is supported on z890 servers and can be carried forward on an upgrade to the z9 BC upgrades.

Availability Enhancements

Question:

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

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 BC is designed to support activating a selected new LIC level concurrently. Concurrent activation of the selected new LIC level is only supported at specific sync points (points in the maintenance process when LIC may be applied concurrently - MCL service level). 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.

Contact your IBM representative to help you determine the appropriate configuration. With proper planning, you may be able to avoid planned outages when installing driver maintenance.

Question:

What is the redundant I/O interconnect that is available on the z9 BC?

Answer:

Redundant I/O interconnect is designed to help maintain critical connections to devices. The z9 BC is designed such that in the event of a failure of an MBA fanout card or STI cable and subsequent repair, you do not lose access to any I/O. Redundant I/O interconnect is exclusive to the System z9.

Question:

What is the advantage of dynamic oscillator switchover?

Answer:

The z9 BC 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.

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. On z9 BC, an MBA fanout card is available that supports both concurrent maintenance and concurrent install. Up to 8 MBA fanout cards may be installed per system. In the event of an outage, an MBA fanout card, used for I/O, may be concurrently repaired using redundant I/O interconnect. This helps to assure that you don't lose access to any of the I/O cards.

Question:

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

Answer:

Program-directed re-IPL is designed to allow Linux on System z9 to re-IPL without operator intervention. Linux on System z9 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 z9 running natively in an LPAR to trigger a re-IPL. This re-IPL is supported for both SCSI and ECKD™ devices.

Question:

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

Answer:

The fiber optic features on a FICON Express4 may be swapped one at a time with a new module, should there be a fiber optic module failure. This swap on a single channel does not effect 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 V1.7 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 offline until an operator action is taken. This is expected to help limit the performance impacts of switch/director fabric problems. The improved FICON error recovery function is expected to be available in the first half of 2006, and available on z/OS V1.4 and higher with APAR OA13644.

IFL, zAAP and zIIP specialty engines

Question:

Can I order subcapacity specialty engines?

Answer:

No. All specialty engines on the z9 BC 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 running on System z benefit from the new z9 BC?

Answer:

The 37% increased power of the z9 BC engine compared to z890, combined with other architectural advances like the greatly improved internal I/O throughput, gives the z9 BC the ability to run increased workloads with better potential for price performance.

With the new pricing of the Integrated Facility for Linux (IFL) and since the price for Linux software is typically processor-based, 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 compared to a z890.

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.

Question:

Can I buy a z9 BC server that is made up entirely of IFL processors and no general purpose processors?

Answer:

Yes. A dedicated z9 BC Linux server is available with the z9 BC Model S07. 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 enable customers to strategically integrate their Java technology-based Web applications with their core business database environment by providing a more cost-effective, specialized z/OS Java execution environment.

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 new z9 BC?

Answer:

The 37% increased power of the z9 BC engine compared to z890, combined with other architectural advances like improved internal I/O throughput, gives the z9 BC the ability to run increased workloads.

With the new pricing of the System z Application Assist Processors (zAAPs), customers may receive greater performance for the same number of users as they would on a z890.

Question:

What is the warranty period for System z Application Assist Processor (zAAP)?

Answer:

zAAPs are under warranty during the first 12 months following the installation of a z9 BC.

Question:

What is the value proposition of the IBM zIIP?

Answer:

IBM zIIPs are designed to help free up general purpose processor (CP) capacity that then may be utilized by other workloads. The zIIP may also help you to leverage the z9 BC, z/OS and 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.

Question:

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

Answer:

Your new z9 BC 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) 1.6 or 1.7 via a Web download. The zIIP will be supported by DB2 UDB for z/OS V8 with enabling PTFs.

Question:

How will IBM zIIPs work?

Answer:

The IBM zIIP's execution environment will accept eligible work from z/OS (and z/OS.e) V1.6 or 1.7, which will manage and direct the work between the general purpose processor and the zIIP. DB2 UDB for z/OS V8 will exploit the zIIP capability for eligible workloads.

Question:

Tell me more about how the IBM zIIPs will work.

Answer:

The IBM zIIP is designed so that a program can work with z/OS (and z/OS.e) V1.6 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 V1.6 or 1.7 that can be directed to the zIIP.

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

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 BC. 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 BC 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 the prerequisite software), monitoring the zIIP activity will be similar to monitoring zAAP activity. Specifically it is intended to allow you to:

- 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 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 result 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:

Can my application on z/OS take advantage of both zAAPs and zIIPs?

Answer:

The use of zAAPs and zIIPs by a single transaction flow is not mutually exclusive. The two specialty engines are designed to run two different types of work. The zAAP processor provides a specialized processor for running Java applications and the zIIP is designed to redirect workload that is typically more closely associated with operating system processing. When available, it will be possible for transactions and/or batch to use a zIIP to handle eligible portions of the DRDA over TCP/IP processing for DB2 UDB z/OS V8, while also running Java stored procedures on a zAAP on the same z9 BC.

Question:

If Linux on System z9 is running in an LPAR on the z9 BC, 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, it can take advantage of redirection of processing to the zIIP.

Question:

What will be the advantage of have the management of specialty engines in separate pools on z9 BC?

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 new 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 "ICF"s 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 general purpose processors, zAAPs and now the new 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 a 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 new z/OS support for WLM goal mode management for zAAP processors be extended to the zIIP?

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

Question:

Where can I get more information on the zIIP?

Answer:

For more information on zIIP please look at the Website: www.ibm.com/systems/z/ziip. For the zIIP FAQs please review the Website: www.ibm.com/systems/z/faq/.

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.

Preemptible 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 V1.6 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 BC Pricing

Question:

Is hardware pricing for z9 BC similar to pricing on z890?

Answer:

Yes, hardware pricing for z9 BC is similar to z890 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 new concerning memory pricing on the z9 BC?

Answer:

System z9 BC 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 z890 or z800 Model 004 to z9 BC be priced?

Answer:

The platform's value proposition of net pricing diagonal upgrades (upgrades from z800 model 004 or z890) 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 BC?

Answer:

Monthly 24x7 ESA list prices will exist and be applied against the z9 BC feature codes similar to z890.

Question:

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

Answer:

The price of the IFL, the zAAP and the zIIP will be US \$95,000 in the United States. Pricing may vary in other countries.

Question:

Is software pricing on z9 BC similar to software pricing on z890?

Answer:

The approach for pricing software on the z9 BC is similar to pricing on the z890 and in line with the pricing initiatives announced with the Mainframe Charter in August 2003. The announced MSUs for software pricing purposes on the z9 BC are on average 10% less than those for z890. With this, in many cases, software savings can be realized on the z9 BC versus the z890.

Question:

Is system software pricing on z9 BC similar to pricing on z800?

Answer:

The z800 offered zSeries Entry License Charges™ (zELC) software pricing or EWLC and EWLC Tiered pricing. The z9 BC continues with the EWLC and EWLC Tiered approach for pricing software. The pricing initiatives announced with the Mainframe Charter in August 2003 result in MSUs for software pricing purposes on the z9 BC that are on average 19% less than those for z800. With this, in many cases, software savings can be realized on the z9 BC versus the z800.

Question:

How is a customer charged for IBM monthly license charge software on the new z9 BC server?

Answer:

IBM extends Entry Workload License Charge (EWLC) and Tiered Entry Workload License Charge (TWLC) pricing metrics to the IBM System z9 BC servers. In addition, IBM announces zSeries Entry License Charges (zELC) pricing only for the smallest z9 BC server, z9 BC capacity setting A01.

With the exception of the z9 BC capacity setting A01 server, EWLC and TWLC are the only pricing metrics available on a standalone z9 BC server.

For a z9 BC server participating in a qualified Parallel Sysplex® environment, you may choose either EWLC and EWLC Tiered pricing, aggregated Parallel Sysplex License Charges (PSLC), or aggregated Workload License Charges (WLC). For more information on PSLC or WLC, refer to: ibm.com/zseries/swprice

Question:

What are all the different software price types available for the z9 BC?

Answer:

IBM monthly license charge software pricing for the z9 BC supports customer requirements for the various ways that the server may be used.

- **Entry Workload License Charges (EWLC):** for z800, z890 and z9 BC customers. EWLC enables qualifying customers to pay for subcapacity eligible IBM software based on the utilization of the LPAR or LPARs where that product executes. This subcapacity pricing provides the potential to lower software charges on a standalone z800, z890 or z9 BC.
- **EWLC Tiered price structure:** offers flat pricing based on z9 BC server capacity using a tiered structure. Non-subcapacity based products, VSE/ESA™, z/VSE™ and z/VM middleware products will have an EWLC Tiered price.
- **zSeries Entry License Charges (zELC):** used only for the smallest z9 BC capacity setting - the Model R07 capacity setting A01.
- **Parallel Sysplex License Charges (PSLC):** may be used when a z9 BC is coupled to a Parallel Sysplex environment with PSLC eligible systems and pricing aggregation is applicable.

- **Workload License Charges (WLC):** may be used when z9 BC is coupled to a Parallel Sysplex environment with WLC for full capacity or subcapacity eligible systems and pricing aggregation is applicable. EWLC compared to WLC: EWLC and Workload License Charges (WLC) are the two subcapacity capable monthly license charge pricing metrics from IBM. EWLC is similar to subcapacity WLC, in terms of implementation and mechanics. Both pricing metrics offer LPAR-based pricing for subcapacity eligible software products, based on the highest rolling 4-hour average utilization of the LPAR or LPARs where the eligible product executes.

Note: A z9 BC coupled to a Parallel Sysplex may also elect standalone pricing with EWLC and EWLC Tiered price structure. Other price types (e.g. flat MLC, Value Unit, Engine-Based Value Units) may be applicable for individual SW product pricing.

Question:

How will my operating systems be priced on the z9 BC server?

Answer:

- The z/OS operating systems and many key IBM middleware products have EWLC pricing.
 - z/VSE, VSE/ESA and other IBM middleware products will be priced with EWLC Tiered Price structure
 - z/VM V5 is a one time charge priced product with engine based value unit pricing.
-

Question:

Will I be able to run z/OS.e on a z9 BC server?

Answer:

Yes, z/OS.e is a specially priced offering of z/OS that is only available for z9 BC, z890 and z800 customers. This product offering is intended to help customers exploit the fast growing world of on demand business by making the deployment of new application workloads, e.g. WebSphere® and Enterprise Applications, on the z9 BC more attractively priced.

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 on a given z9 BC model. This requirement is at a server level, so a customer could have an LPAR with more zIIPs than general purpose processors, as long as there are enough general purpose processors in the entire server to meet the one for one requirement.

Question:

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

Answer:

Yes. One zIIP and one zAAP can share one general purpose processor 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 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.

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 System z9 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 the latest 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 general purpose processors, 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 System z9.

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 exactly 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” as configurations deviate from 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 BC?

Answer:

As with the introduction of any new server, workloads with differing characteristics will see some variation in performance when moved to the z9 BC. 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 BC 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 to provide all workloads with a

significant boost in performance. Additionally, newer applications, such as those with compiler optimizations for the z9 BC (available with z/OS 1.7 and Java 5), are likely to see even higher benefits, particularly those that may be enhanced over time to exploit some of the new instructions provided with the z9 BC.

Question:

What is the performance improvement a VSE customer might experience on the z9 BC?

Answer:

The performance ratios that a VSE customer workload might experience when migrating from a zSeries server to the z9 BC is represented by the range of ratios seen by a comparable z/OS migration.

For example, if the published ratio in the LSPR between a zSeries and System z9 is approximately 32% to 39%. 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 V1.1 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 BC 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 planned to be made 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 V5.2. IBM is working with its Linux distribution partners to include this function in future Linux on System z9 distribution releases or service updates.

(Minimum MCL level is required on the z990 and z890. Refer to the Preventative Planning (PSP) bucket for your z990 and z890 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 BC 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) V1.1. Use of the OSA for NCP can help ease migration and eliminate dependencies on hardware such as the 3745/6 communication controller. It can 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 BC to help to simplify networking administration and management of VLANs on the z9 BC?

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 and z/VM. z/VM 5.2 will provide this support on May 26, 2006 with the PTF for APAR VM63952 and z/VM V5.1 with the PTFs for APAR VM63784 and PK08444.

Question:

What was introduced on the System z9 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 minimize network traffic disruption.

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 (4KM and 10KM) 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.

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 helping 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 port FICON Express4 feature available exclusively on the new z9 BC.

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

Answer:

On April 25, 2006 IBM also announced an expanded range for 4 Gbps FICON and FCP enabled SAN directors and switches for enterprise and midrange environments. Announcements included:

- 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. (IBM SAN256M 4 Gbps feature statement of direction, January 2006)
- IBM TotalStorage SAN32B-2 switch and SAN 256B director (Brocade) 4 Gbps shortwave features planned to be expanded with 4 Gbps long wave features in 2Q/2006.

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 Gbps FICON/FCP adapters. We have disk, tape and SAN products that either support 4 Gbps connectivity today or will very soon. These links are all auto-negotiating so they can support intermix of 4 Gbps and 2 Gbps 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 on 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 plan to provide an upgrade for customers that previously purchased a TS1120 tape drive prior to the availability of the encryption feature.

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 does FICON Express 4 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, an most resilient design points for high availability configurations. With this announcement IBM 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:

Does the z9 BC deliver increased I/O connectivity and simplify I/O configuration?

Answer:

In comparison to the z890, the z9 BC supports a 40% increase in Crypto, FICON and OSA connectivity per I/O cage. Previously, on the z890, 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 112 FICON channels may be installed in the 28 I/O slots available on the z9 BC. Up to 80 FICON channels may be installed on the z990.

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 BC help to provide increased I/O addressing capability?

Answer:

The z9 BC 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 BC 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 BC - 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 BC 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:

Yes. 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.

For the z9 BC model R07, you can not order On/Off COD capacity that exceeds the highest LSPR rating of the model R07.

Note that On/Off CoD for specialty engines will always be full capacity processors.

Question:

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

Answer:

This new feature allows the customer to order ‘zero’ quantity 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 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:

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

Answer:

There will be a new API provided within the existing HMC SNMP command which is designed to enable customers to use other automation code (which conforms to the new 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 processor 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 (CP) capacity as a temporary zAAP or zIIP on the z9 BC?

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 general purpose processor (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 engines (IFLs, ICFs, zAAPs and zIIPs) on the z9 BC?

Answer:

Yes. The z9 BC, like the z890 and z990 servers, will continue to have the ability to utilize capacity for temporary upgrades or specialty purposes 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 processors?

Answer:

Yes. Remember that you no exceed twice the 'purchased' capacity nor can you decrease the number of installed general purpose processors. Also, for the z9 BC model R07, you can not order On/Off COD capacity that exceeds the highest LSPR rating of the model R07.

Question:

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

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 new zIIP ?

Answer:

Yes. The z9 BC is able to activate most PU types as part of CBU upgrades: IFLs, zAAPs, ICFs, general purpose processors and the new zIIPs.

Question:

Can I order subcapacity CBU 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 processors you can execute CBU with processors of equal, less or greater capacity depending on your requirements.

Question:

With the z890 I needed to order an additional engine to my configuration when I did CBU. Will I need to do that with the z9 BC?

Answer:

No, there is no longer a requirement to add an additional engine when you execute CBU. You may add CBU capacity by increasing the capacity setting of your active CP(s) if that satisfies your CBU capacity requirement.

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.

Question:

Can I CBU a z9 BC R07 to a z9 BC S07 configuration?

Answer:

No, there is no ability to CBU a Model R07 to a Model S07.

Cryptographic Enhancements

Question:

What common criteria security certifications has IBM System z9 achieved?

Answer:

On March 24, 2006, the IBM z9 109 (z9-109) joined the ranks of other 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. The z9 BC was designed for EAL5 certification.

Question:

What new cryptographic capabilities are part of this 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 V1.6 and V1.7 with Enhancement to Cryptographic Support for z/OS and z/OS.e 1.6/7 Web deliverable, and z/VM 5.1 and 5.2 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 BC?

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 BC?

Answer:

PCI Cryptographic Coprocessor (PCICC), PCI X Cryptographic Coprocessor (PCIXCC), CMOS Cryptographic Coprocessor Facility (CCF) and the PCI Cryptographic Accelerator (PCICA) features are not supported on the z9 BC.

Question:

Is the cryptographic hardware offered as standard features on the z9 BC?

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 BC?

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.
 - 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 z890 to the z9 BC to take advantage of increased SSL performance and the new configuration capability.
- Up to eight Crypto Express2 features per server.

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 a performance improvement compared to the PCICA feature or the current Crypto Express2 feature on z890, 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 will be replaced by Enhancements to Cryptographic Support for z/OS and z/OS.e 1.6/1.7 Web deliverable (planned to be 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 z890 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.

Question:

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

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 PCI-X Cryptographic Coprocessor (PCIXCC) feature are planned to be supported by the Crypto Express2 feature.

The following ICSF callable services will be available with CPACF (and continue to support DES and TDES):

- 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.
-
-

Question:

What other means are available for customers to utilize CP Assist for Cryptographic Function (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 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
Provides highest SSL handshake performance (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:

What releases of operating systems are required to support the Crypto Express2 hardware feature on the z9 BC?

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 and z/OS.e 1.6/1.7 Web deliverable (planned to be available May 26, 2006) and PTFs. Compatibility support is provided on z/OS 1.4 and 1.5. z990 and z890 enhancements to cryptographic support are Web deliverable.
 - z/VM 5.1 for z/OS and Linux guests, with applicable PTFs.
 - VSE/ESA V2.7 and later (z/VSE V3.1 for configurable Crypto Express2)
 - 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 and z/OS.e 1.6/1.7 Web deliverable (planned to be available May 26, 2006) and PTFs. Compatibility support is provided on z/OS 1.4 and 1.5. z990 and z890 enhancements to cryptographic support are Web deliverable.
 - z/VM 4.4 and later for guests.
 - z/VSE 3.1, and later, with applicable PTFs.
 - 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 BC?

Answer:

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

Question:

Will new UDX's be supported on the z9 BC?

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:

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:

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

Answer:

Yes, a new TKE workstation will be available. The new TKE 5.0 workstation is assigned feature code 0839. Customers must use the TKE 5.0 workstation to control the System z9. 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 5.0.

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 BC TKE workstations?

Answer:

Yes. Support is available for an optional Smart Card Reader to be attached to the new TKE 5.0 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 5.0 workstations per system will be supported on the System z9. 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 BC.

Parallel Sysplex Enhancements

Note: Please use the updated STP FAQs to answer your questions on the new Server Time Protocol feature of the z9 BC. These can be found at:

www.ibm.com/systems/z/faq/

GDPS Enhancements

Question:

What GDPS functions were announced in January 2006?

Answer:

IBM announced extensions to the HyperSwap™ functionality. HyperSwap is the technology within GDPS that designed to allow z/OS and z/VM systems to “switch” to secondary disks in the event of a disk failure or disk maintenance with only a momentary pause in application availability. Virtual servers, (Linux servers for example), running as z/VM guests, can also benefit from GDPS and the extended HyperSwap function. GDPS added a new HyperSwap trigger designed to allow tailoring for unplanned HyperSwap conditions when a disk subsystem is experiencing soft failures, such as a non-responsive disk subsystem.

GDPS/HyperSwap Manager: within a single Data Center allows GDPS/PPRC HyperSwap Manager configured within a single data center to invoke HyperSwap for individual disk subsystems that need to be switched, helping to simplify systems management and improve availability.

GDPS Enhanced Recovery Support: GDPS/PPRC enhancement helps to reduce potentially long or variable recovery times when there is a failover by helping to ensure that the secondary copy of the data at the recovery site and the data in the coupling facility are time consistent. This helps to enable consistent application restart times and reduce the need for special recovery procedures.

z/OS 1.7 XRC+ Support: System Logger is planned to provide new support for XRC+ by allowing asynchronous writes to staging data sets for logstreams improving the throughput for high-volume logging applications, such as WebSphere, CICS, and IMS. The use of asynchronous writes can allow the use of XRC for some application for which it was not previously practical, which may significantly help to improve availability and recovery.

GDPS/Global Mirror Support: The latest member of the GDPS suite of offerings, GDPS/Global Mirror offers a multi-site, end-to-end disaster recovery solution for your IBM z/OS systems and open systems data using IBM TotalStorage Global Mirror technology.

These functions are available with GDPS V3.3. Contact gdps@us.ibm.com to get more information on GDPS solutions.

IBM Global Financing (IGF)

Question:

How does financing fit with the on demand capabilities of the z9 BC product?

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-hardware items that are part of a business transformation project utilizing a z9 BC?

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://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://ibm.com/financing) to find one.

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 System z platforms.

Question:

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

Answer:

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

Question:

Why is IBM offering the System z 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 System z fiber cabling services options?

Answer:

The System z 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 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?

Answer:

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

z/VSE V4

Question:

Did the April 27, 2006 announcement contain anything about VSE?

Answer:

The April 27 announcement included important VSE content. IBM previewed a new version of z/VSE. The z/VSE V4 preview introduced two very significant new innovations for VSE.

First, z/VSE V4 will be the first version of VSE to exploit z/Architecture. z/VSE V4 is designed to exploit 64-bit real addressing and support more than 2 GB of processor storage. z/VSE V4 will execute in z/Architecture mode only and will run exclusively on IBM System z servers and such other servers as IBM may specify.

Second, IBM plans to offer attractive new z/VSE V4 pricing for IBM System z. The z/VSE V4 preview contains this Statement of Direction (SOD): “It is IBM's intent to provide new software pricing for z/VSE V4 when running on select processors, subject to applicable terms and conditions. IBM expects this new software pricing metric to provide more granularity and a subcapacity pricing option.”

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

New pricing metrics are expected to provide worthwhile new options for eligible VSE customers. Subcapacity pricing effectively increases the software price granularity available. For example, z/VSE V4 customers on IBM System z9 who select the subcapacity option will pay only for the capacity actually used, not necessarily the total capacity of the server installed. That means z/VSE V4 customers may be able to install more capable IBM System z9 servers than they might have selected otherwise.

Question:

When will z/VSE V4 become available? What will the new price points be?

Answer:

That information is not available at this time. Previews do not contain specific dates or pricing information. Previews are provided to help customer in their IT planning. Of course, any statements regarding IBM’s future direction or plans are subject to withdrawal or change without notice.

Question:

Will I be able to run z/VSE V4 on my Multiprise 3000?

Answer:

No. z/VSE V4 is designed to run in z/Architecture mode. It will run exclusively on IBM System z servers. IBM System z includes IBM System z9, z990, z890, z900, and z800 servers. Unlike today’s z/VSE V3, z/VSE V4 will not run on Multiprise® 3000 or S/390 G5/G6 processors.

Question:

What about migration? Will it be difficult to migrate to z/VSE V4?

Answer:

It is reasonable to believe that migration to z/VSE V4 will be relatively straightforward. Support for z/Architecture is designed to be transparent to systems and application programs. In addition, the Fast Service Upgrade (FSU) process will be supported from z/VSE V3 and VSE/ESA V2.7.

Question:

Can I get new pricing on z/VSE V3? If I install z/VSE V4 on my z800, can I get the new pricing?

Answer:

No. The new pricing metrics apply exclusively to z/VSE V4 on IBM System z9. Other versions of VSE (for example, z/VSE V3 or VSE/ESA V2.7) will continue to be priced as today. z/VSE V4 running on other servers (i.e z990, z890, z900, and z800) is expected to be priced much like z/VSE V3 is priced today.

Question:

Is the z/VSE V4 Preview all that was announced for VSE?

Answer:

No. z/VSE V3 supports some important new IBM System z technology. z/VSE V3 (and later) supports CPACF enhancements, Crypto Express2 in both accelerator and coprocessor mode, and N_Port ID Virtualization. In addition, VSE/ESA V2.7 (and later) supports FICON Express4, OSA Express2 1000 BASE-T ethernet, and up to 60 LPARs.

Note: z9 BC supports up to 30 LPARs, z9 EC supports up to 60 LPARs.

Announcement Overview - z/VM

Question:

What additional features are being announced for z/VM on April 27, 2006?

Answer:

The new additional features being added to z/VM V5.2 with the PTF for APAR VM63952 to exploit the z9 EC and the new z9 BC plus other enhancements include:

- Improved memory management for Linux guests with the PTF for VM63856
 - Exploitation of selected features of the z9 EC and the new z9 BC, including:
 - Support for new instructions
 - Support for FICON Express4 (4 Gbps FICON)
 - Support of HiperSockets using IPv6 protocol by TCP/IP and guest LAN
 - Simplified VLAN management with support for GVRP
 - Enhancements for FCP N_Port ID virtualization (NPIV)
 - Additional exploitation of the IBM DS6000 and DS8000 series
 - Support for Parallel Access Volumes (PAVs) as minidisks
 - Ease-of-use enhancement for the Performance Toolkit for VM
-
-

Question:

How will these new enhancements be made available and when?

Answer:

These enhancements are planned to be delivered as a CORrective PTF in the z/VM service stream with the PTFs for APARs VM63952 and VM63856. The PTF for APAR VM63952 is planned to be available on May 26, 2006 and the PTF for APAR VM63856 is planned to be available on June 30, 2006.

Question:

Was there any announcement regarding z/VM support for the IBM TotalStorage DS4000 Series Midrange Disk Systems?

Answer:

Yes. The z/VM V5.2 support for the IBM DS4000 series midrange disk systems as announced in Software Announcement 205-168, dated July 27, 2005, is being withdrawn, effective immediately. z/VM V5.2 continues to support both the IBM TotalStorage DS6000 series and the IBM TotalStorage DS8000 series on System z servers configured for ECKD and SCSI use to meet the data requirements for z/VM, including the Linux on System z environment or with z/OS, z/OS.e, z/TPF, and z/VSE. The DS6000 is designed and priced to support lower total cost of ownership and provide a highly-available, robust storage solution for System z environments that do not require the full scalability of the DS8000.

Question:

Has z/VM been evaluated for common criteria security certification?

Answer:

Yes. On October 26, 2005, the German Federal Office of Information Security (Bundesamt für Sicherheit in der Informationstechnik, BSI) issued its certification that z/VM V5.1 conforms to the requirements of the Controlled Access Protection Profile (CAPP) and the Labeled Security Protection Profile (LSPP), both at Evaluation Assurance Level 3+.

Question:

Will V5.1 be available for ordering after the availability of V5.2?

Answer:

Yes. z/VM V5.1 and its priced, optional features of DirMaint™, RACF®, and the Performance Toolkit for VM™ will continue to be orderable after the availability of z/VM V5.2. However, z/VM V5.1 and its priced, optional features, DirMaint, RACF, and the Performance Toolkit, should only be ordered if you require a z/VM release with Common Criteria (ISO/IEC 15408) certification against the Labeled Security Protection Profile (LSPP) and the Controlled Access Protection Profile (CAPP), both at the EAL3+ assurance level. The Recommended Service Update (RSU) Level 5103RSU (PTF UM97510) plus the RACF for z/VM PTF for APAR VM63613 is required to maintain certification. If you apply any additional service to z/VM V5.1, other than PTF UM97510 and the PTF for RACF APAR VM63613, you no longer will have a certified system. z/VM V5.1 and its priced, optional features of DirMaint, RACF, and the Performance Toolkit for VM are planned to be withdrawn from marketing on September 30, 2007 and to be supported until September 30, 2007.

Question:

Where can I get more information on these new enhancements for z/VM V5.2?

Answer:

There is a complete set of Frequently Asked Questions for z/VM planned to be available on the z/VM Web site at: www.ibm.com/zvm520/

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 Server Time Protocol in the October 10th announcement?

Answer:

There were two statements of direction in the announcement for STP:

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 did IBM make on July 27, 2005 for future z/VM support on the System z9 platform?

Answer:

IBM intends to provide future enhancements to z/VM supporting the following System z9 functions:

- System and guest exploitation of HiperSockets IPv6
- Improved memory management between z/VM and Linux on System z
- Simplified networking administration and management of VLANs with support for GARP VLAN Registration Protocol (GVRP) using OSA-Express2
- Capability to allow guest operating systems and z/VM users to query virtual port names when using N_Port ID Virtualization

Question:

What statements of direction has IBM made for IBM TotalStorage?

Answer:

IBM intends to provide the following enhancements:

IBM TotalStorage DS8000 Series of disk products: These enhancements for use in the System z9 environment include:

- IBM TotalStorage z/OS Global Mirror (XRC) and IBM TotalStorage Global Mirror (asynchronous PPRC) capability on the IBM TotalStorage DS8000 Series to enable the use of consistency groups that can span both remote mirroring products for enhanced business continuity. This capability is intended to help customers maintain and leverage their investment in z/OS Global Mirror (XRC) while implementing Global Mirror (asynchronous PPRC) for heterogeneous server environments.
- A continuous data protection function on the IBM TotalStorage DS8000 Series for System z9 environments. This capability is intended to provide support for enhanced data resiliency and business continuity.

IBM TotalStorage SAN Volume Controller (SVC): To enhance the interoperability of the IBM TotalStorage SAN Volume Controller, IBM is planning to provide SVC support for Linux on System z9 environments. This is designed to allow System z9 hosts and supported open system hosts to share storage pools created by the SAN Volume Controller.

IBM TotalStorage Virtual Tape Server (VTS) - **business continuity enhancements:**

- **Provision** of an enhanced Import/Export function designed to allow sets of cartridges to be interchanged between VTSs and/or Peer-to-Peer (PtP) VTSs or exported for storage in a remote vault. This enhancement is designed to allow virtual volumes to be directed to a VTS volume pool, subsequently exported on physical cartridges under host control, and imported into another VTS under host control.

Extension of the PtP VTS - support for full-duplex virtual volume replication between three sites. DFSMS parameters are designed to provide policy-based copy management to control the source and target VTSs, the preferred routing, and the mode of operation.

- **IBM TotalStorage SAN File System (SFS):** SFS was designed to help you manage the growing complexities and rising costs associated with dramatic growth in disk storage requirements. IBM continues to expand on the concept of system-managed storage by working to extend some of the storage management capabilities of DFSMS to open system environments to support improvements in storage asset utilization and administrator productivity.

IBM plans to increase the range of IBM TotalStorage SAN File System supported hosts to include the 31-bit SUSE Linux Enterprise Server 8 (SLES 8) client distribution running under z/VM on System z. This support is designed to apply exclusively to hosts acting as the SFS client in a System z server running SUSE SLES 8. This extension of host support for SFS is intended to support increased data sharing between UNIX®, Linux, Microsoft® Windows®, and mainframe 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:

Did the April 27, 2006 announcement contain anything about VSE?

Answer:

The April 27 announcement included important VSE content. IBM previewed a new version of z/VSE which contains this Statement of Direction (SOD): “It is IBM's intent to provide new software pricing for z/VSE V4 when running on select processors, subject to applicable terms and conditions. IBM expects this new software pricing metric to provide more granularity and a subcapacity pricing option.”

New pricing metrics are expected to provide worthwhile new options for eligible VSE customers. Subcapacity pricing effectively increases the software price granularity available. For example, z/VSE V4 customers on IBM System z9 who select the subcapacity option will pay only for the capacity actually used, not necessarily the total capacity of the server installed. That means z/VSE V4 customers may be able to install more capable IBM System z9 servers than they might have selected otherwise.



(C) (c) Copyright IBM Corporation 2006
IBM Corporation
New Orchard Rd.
Armonk, NY 10504

Printed in the United States of America, 10/06
All Rights Reserved

IBM, the IBM logo, IBM @server, IBM eServer, e-business logo, CICS, DB2, Domino, DRDA, ECKD, ESCON, FICON, GDPS, HiperSockets, HyperSwap, IMS, Lotus, MQSeries, On Demand Business logo, Parallel Sysplex, Resource Link, RMF, S/390, Sysplex Timer, System Storage, System z, System z9, TotalStorage, VM/ESA, VSE/ESA, 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 is a trademark of Intel Corporation in the United States, other countries, or both.

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

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries. All other registered trademarks and trademarks are the properties of their respective companies.

This publication was produced in the United States.

IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

You can find additional information via IBM's World Wide Web server at ibm.com.

IBM Hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply. Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

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

z/VSE can execute in 31-bit mode only. It does not implement z/Architecture, and specifically does not implement 64-bit mode capabilities. z/VSE is designed to exploit selected features of IBM z9 and zSeries hardware.