


IBM System z10 Enterprise Class Technical Introduction

The server's role in a dynamic infrastructure

Key functional elements and enhancements

Hardware and software capabilities



Bill White
Per Fremstad
Parwez Hamid
Fernando Nogal
Karl-Erik Stenfors

Redbooks



International Technical Support Organization

**IBM System z10 Enterprise Class Technical
Introduction**

November 2009

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Archived

Third Edition (November 2009)

This edition applies to the IBM System z10 Enterprise Class server.

© Copyright International Business Machines Corporation 2008, 2009. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team who wrote this book	ix
Become a published author	x
Comments welcome	xi
Chapter 1. Introducing the System z10 Enterprise Class	1
1.1 Wanted: an infrastructure (r)evolution	2
1.2 z10 EC comparison	7
1.3 z10 EC server enclosure	8
1.3.1 Central Electronic Complex	9
1.3.2 I/O subsystem	11
1.3.3 I/O connectivity	12
1.4 Performance	14
1.5 Capacity On Demand	16
1.6 Software	16
Chapter 2. Hardware overview	17
2.1 z10 EC highlights	18
2.2 Models and model upgrades	18
2.3 The frames	21
2.4 CEC cage and books	22
2.5 Multi-chip module (MCM)	24
2.6 z10 EC processor chip	24
2.7 Processor unit (PU)	25
2.8 Memory	26
2.9 I/O system structure	28
2.10 I/O cages and features	29
2.10.1 ESCON channels	31
2.10.2 FICON Express8	32
2.10.3 FICON Express4	32
2.10.4 FICON Express2	33
2.10.5 FICON Express	34
2.10.6 OSA-Express3	34
2.10.7 OSA-Express2	35
2.11 Cryptographic functions	36
2.11.1 CP Assist for Cryptographic Function	37
2.11.2 Crypto Express2 feature	38
2.11.3 Crypto Express3 feature	38
2.11.4 TKE workstation	39
2.12 Coupling and clustering	39
2.12.1 ISC-3	39
2.12.2 ICB-4	40
2.12.3 Internal Coupling (IC)	40
2.12.4 Parallel Sysplex InfiniBand (PSIFB) coupling	40
2.12.5 System-Managed CF Structure Duplexing	40
2.12.6 Coupling Facility Control Code (CFCC) level 16	41

2.13	Time functions	41
2.13.1	External time reference (ETR)	41
2.13.2	Server Time Protocol (STP)	41
2.13.3	Network Time Protocol (NTP) support	42
2.14	HMC and SE	42
2.15	Power and cooling	43
2.15.1	Hybrid cooling system	43
2.15.2	Internal Battery Feature	44
2.15.3	IBM Systems Director Active Energy Manager	44
Chapter 3. Key functions and capabilities		45
3.1	Virtualization	46
3.1.1	Hardware virtualization	46
3.1.2	Software virtualization	49
3.2	Technology improvements	49
3.2.1	Microprocessor enhancements	50
3.2.2	Granular capacity and capacity settings	51
3.2.3	Memory enhancements	52
3.2.4	Connectivity enhancements	54
3.2.5	Cryptography enhancements	61
3.2.6	Hardware Management Console enhancements	61
3.3	Common time functions	62
3.3.1	Sysplex Timer	62
3.3.2	Server Time Protocol (STP)	62
3.4	Capacity on Demand (CoD) enhancements	64
3.4.1	Permanent and temporary upgrades	65
3.4.2	z/OS capacity provisioning	67
3.5	Throughput optimization enhancements	68
3.6	Reliability, availability, and serviceability improvements	69
3.7	Parallel Sysplex technology	70
3.8	Summary	72
Chapter 4. Software support		75
4.1	Software support summary	76
4.2	Support by operating system	78
4.2.1	z/OS	78
4.2.2	z/VM	80
4.2.3	z/VSE	82
4.2.4	Linux on System z	83
4.2.5	TPF and z/TPF	84
4.3	Support for selected functions	85
4.3.1	Single system image	85
4.3.2	z/VM-mode LPAR	86
4.3.3	Dynamic PU exploitation	86
4.3.4	zAAP on zIIP capability	86
4.3.5	Large memory	87
4.3.6	Dynamic LPAR memory upgrade	87
4.3.7	Hardware decimal floating point	88
4.3.8	High Performance FICON for System z10	88
4.3.9	Cryptographic support	91
4.3.10	z/OS ICSF	92

4.4 z/OS considerations	93
4.5 Coupling Facility and CFCC considerations	95
4.6 IOCP	95
4.7 Worldwide portname (WWPN) prediction tool.	96
4.8 ICKDSF	96
4.9 Software licensing considerations.	97
4.9.1 Workload License Charges (WLC)	97
4.9.2 System z New Application License Charges (zNALC)	98
4.9.3 Select Application License Charges (SALC).	98
4.9.4 Midrange Workload Licence Charges.	99
4.9.5 System z International Program License Agreement (IPLA).	99
4.10 References	100
Appendix A. Frequently asked questions	101
Appendix B. Channel options	109
Related publications	113
IBM Redbooks publications	113
Online resources	113
Other publications	113
How to get Redbooks publications.	114
Help from IBM	114
Index	115

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	OMEGAMON®	System i®
CICS®	OS/390®	System p®
DB2 Connect™	Parallel Sysplex®	System Storage™
DB2®	Passport Advantage®	System x®
Domino®	Power Systems™	System z10™
DRDA®	POWER®	System z9®
DS8000®	PR/SM™	System z®
Dynamic Infrastructure®	Processor Resource/Systems Manager™	Tivoli®
ESCON®	RACF®	TotalStorage®
FICON®	Rational Rose®	WebSphere®
HiperSockets™	Rational®	z/Architecture®
IBM Systems Director Active Energy Manager™	Redbooks®	z/OS®
IBM®	Redbooks (logo)  ®	z/VM®
IMS™	Resource Link™	z/VSE™
Language Environment®	S/390®	z9®
Lotus®	Service Request Manager®	zSeries®
MQSeries®	Sysplex Timer®	

The following terms are trademarks of other companies:

AMD, the AMD Arrow logo, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

InfiniBand, and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Ambassador, and the LSI logo are trademarks or registered trademarks of LSI Corporation.

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

Red Hat, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

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

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

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

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

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

Preface

This IBM® Redbooks® publication introduces the IBM System z10™ Enterprise Class server, which is based on z/Architecture®. It builds on the inherent strengths of the System z® platform, delivering new technologies and virtualization that are designed to offer improvements in price and performance for key workloads, as well as enabling a new range of solutions. The z10 EC further extends System z's leadership in key capabilities with the delivery of expanded scalability for growth and large-scale consolidation, availability to help reduce risk and improve flexibility to respond to changing business requirements, and improved security. The z10 EC is at the core of the enhanced System z platform that is designed to deliver technologies that businesses need today along with a foundation to drive future business growth.

This book provides basic information about z10 EC capabilities, hardware functions and features, and associated software support. It is intended for IT managers, architects, consultants, and anyone else who wants to understand the new elements of the z10 EC.

The changes to this edition are based on the System z hardware announcement dated October 20, 2009.

This book is an introduction to the z10 EC mainframe. Readers are not expected to be generally familiar with current IBM System z technology and terminology.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Poughkeepsie Center.

Bill White is a Project Leader and Senior System z Networking and Connectivity Specialist at the International Technical Support Organization, Poughkeepsie Center.

Per Fremstad is an IBM Certified Senior IT Specialist from the IBM Systems and Technology Group in IBM Norway. He has worked for IBM since 1982 and has extensive experience with mainframes and z/OS®. Per also works extensively with Linux® on System z and z/VM®. During the past 25 years, he has worked in various roles within IBM and with a large number of customers. He frequently teaches about z/OS and z/Architecture subjects and has been actively teaching at Oslo University College for the last four years. Per holds a BSc from the University of Oslo, Norway.

Parwez Hamid is an Executive IT Consultant working for the IBM Server and Technology Group. During the past 36 years he has worked in various IT roles within IBM. Since 1988 he has worked with a large number of IBM mainframe customers and spent much of his time introducing new technology. Currently, he provides pre-sales technical support for the IBM System z product portfolio and is the Lead System z Technical Specialist for the UK and Ireland. Parwez has co-authored a number of ITSO Redbooks publications and prepares technical material for the world-wide announcement of System z servers. Parwez works closely with System z product development in Poughkeepsie, New York, and provides input and feedback for future product plans. Additionally, Parwez is a member of the IBM IT Specialist Professional Certification Board in the UK and is also a technical staff member of the IBM UK Technical Council, which is made of senior technical specialists representing all

IBM client, consulting, services, and product groups. Parwez teaches and presents at numerous IBM user group and IBM internal conferences.

Fernando Nogal is an IBM Certified Consulting IT Specialist working as an STG Technical Consultant for the Spain, Portugal, Greece, and Israel IMT. He specializes in on demand infrastructures and architectures. In his 27 years with IBM, he has held a variety of technical positions, mainly providing support for mainframe customers. Previously, he was on assignment to the Europe Middle East and Africa (EMEA) zSeries® Technical Support group, working full time on complex solutions for e-business on zSeries. His job included, and still does, presenting and consulting for architectures and infrastructures, and providing strategic guidance to System z customers regarding the establishment and enablement of e-business technologies on System z, including the z/OS, z/VM, and Linux environments. He is a zChampion and a core member of the System z Business Leaders Council. An accomplished writer, he has authored and co-authored 19 IBM Redbooks publications and several technical papers. Other activities include chairing a Virtual Team for IBM interested in e-business on System z and serving as a University Ambassador. He travels extensively on direct customer engagements and as a speaker at IBM and customer events and trade shows.

Karl-Erik Stenfors is a Senior IT Specialist in the Product and Solutions Support Centre (PSSC) in Montpellier, France. He has more than 40 years of experience in the large systems field as a Systems Programmer and as a consultant with IBM customers and, since 1986, with IBM. His areas of expertise include IBM System z hardware and operating systems, including z/VM, z/OS and Linux. He teaches at numerous IBM user group and internal conferences. He is currently working with the System z lab in Poughkeepsie, providing customer requirement input to create an IBM System vision for the future, through the zChampions workgroup.

A special thanks to Ivan Dobos, Wolfgang Fries, Marian Gasparovic, Brian Hatfield, and Dick Jorna for their efforts in creating the ground work for this edition of this publication.

Thanks to the following people for their contributions to this project:

Connie Beuselink, Ellen Carbarnes, William Clark, Darelle Gent, Michael Gerhart, Gary King, Jeff Kubala, Scott Langenthal, Kenneth Oakes, Patrick Rausch, Charles Webb, and Frank Wisnewski
IBM Poughkeepsie

Alan Altmark, Les Geer, Reed Mullen, Damian Osisek, Brian Valentine, and Steve Wilkins
IBM Endicott

Harv Emery and Greg Hutchison
IBM Gaithersburg

Monika Zimmermann and Uwe Zumpe
IBM Boeblingen, Germany

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at: ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Archived



Introducing the System z10 Enterprise Class

The announcement of the System z10 Enterprise Class (z10 EC) is arguably the most exciting piece of news on the mainframe scene in recent years. At its heart is the z10 Enterprise chip, which at 4.4 GHz is the fastest quad-core processor in the industry. The z10 EC server can be configured up to a 64-way, with up to 1.5 TB of memory. It also has new connectivity options that enhance the z10 EC server's open characteristics. It delivers, in a single footprint, unprecedented performance and capacity growth while drawing upon the rich heritage of previous z/Architecture servers. The z10 EC is a well-balanced, general-purpose server that is equally at ease on compute-intensive workloads as it is with I/O-intensive workloads.

The System z server design continues to follow the fundamental principle of being able to simultaneously support a large number of heterogeneous workloads and provide the highest quality of service. The workloads in themselves have changed a lot, and the design must adapt to this change.

The last couple of decades have witnessed an explosion in applications, architectures, and platforms. A lot of experimentation occurred in the marketplace. With the generalized availability of the internet and the appearance of commodity hardware and software, several patterns have emerged that have gained center stage.

Multi-tier application architectures and their deployment on heterogeneous infrastructures are common today. When these applications are mission critical, however, it takes a great amount of effort to ensure that the infrastructure provides the required qualities of service, and careful engineering of the application's several tiers is required to provide the robustness, scaling, consistent response, and other characteristics demanded by the users and lines of business.

Providing the required service level in a distributed environment implies acquiring and installing extra equipment and software to ensure availability and security, and additional manpower to configure, administer, troubleshoot, and tune such a complex set of separate and diverse environments. Often, by the end of the distributed equipment's life cycle its residual value is null, requiring new acquisitions and software licences, re-certification, and so on. Back to square one. In today's resource-constrained environments there must be a better way.

The z10 EC offers an extensive software portfolio that spans from IBM WebSphere®, full support for service-oriented architecture (SOA), Web services, J2EE, Linux, and open standards, to the more traditional batch and transactional environments such as Customer Information Control System (CICS®) and Information Management System (IMS™). For instance, considering just the Linux on System z environment, more than 3,000 applications are offered by over 400 independent software vendors (ISVs).

IBM has a holistic approach to System z design that includes hardware, software, and procedures and takes into account a wide range of factors, including compatibility and investment protection, thus ensuring a tighter fit with the IT requirements of the entire enterprise.

1.1 Wanted: an infrastructure (r)evolution

Exploitation of information technology (IT) by enterprises continues to grow and the demands placed upon it are increasingly complex. The world is not stopping—in fact, business pace is accelerating. The pervasiveness of the Internet fuels ever-increasing utilization modes and users. And the most rapidly growing type of user is not people, but devices. All sorts of services are being offered and new business models are being implemented. The demands placed on the network and computing resources will reach a breaking point unless something changes.

Awareness that the very foundation of IT infrastructures is not up to the job is growing. Most existing infrastructures are too complex, too inefficient, and too inflexible. How, then, can those infrastructures evolve and what must they become in order to avoid the breaking point? And, while they are evolving, the need to improve service delivery, manage the escalating complexity, and maintain a secure enterprise continues to be felt. To compound it, there is a daily pressure to cost-effectively run a business, while supporting growth and innovation. Aligning IT with the goals of the business is an absolute top priority.

In the IBM vision of the future, transformation of the IT delivery model is strongly based on new levels of efficiency and service excellence for businesses, driven by and from the data center.

To achieve success in the transformation of their IT model, and truly maximize the benefits of this new approach, organizations must develop and follow a plan for their transformation, or journey, towards that goal. IBM has developed a roadmap to help enterprises build such a plan. The roadmap lets IT free itself from operational complexity and reallocate scarce resources to drive business innovation. The roadmap follows a model based on an infrastructure supporting a highly dynamic, efficient, and shared environment. This is, indeed, a new view of the data center. It allows IT to better manage costs, improve operational performance and resiliency, and more quickly respond to business needs.

By implementing this evolved infrastructure, organizations can better position themselves to adopt and integrate new technologies, such as Web 2.0 and cloud computing, and deliver dynamic and seamless access to IT services and resources.

Clouds, as seen from their users' side, offer services through the network. User requirements are in the functionality but also in the availability, ease of access, and security areas, so much so that organizations may decide to adopt private clouds, while also exploiting public or hybrid clouds. From the service-provider viewpoint, guaranteeing availability and security, along with repeatable and predictable response times, requires a very flexible IT infrastructure and advanced resource management.

IBM calls this evolved environment a Dynamic Infrastructure® and the IBM System z10 is at its core. Due to its advanced characteristics, the mainframe already provides many of the qualities of service and functions required, as we discuss next.

Through its own transformation and engagements with thousands of enterprise clients, IBM has identified three stages of adoption along the way, which are described in this section:

- ▶ Simplified
- ▶ Shared
- ▶ Dynamic

Simplified

In this stage, to drive new levels of economics in the data center, operational issues are addressed through consolidation, virtualization, energy offerings, and service management. Most enterprises start their journey here.

The z10 EC supports advanced server consolidation and offers the best virtualization in the industry. Up to 60 logical partitions (LPARs) can be used. Each one can run any of the supported operating systems:

- ▶ z/OS
- ▶ z/VM
- ▶ z/VSE™
- ▶ z/TPF
- ▶ Linux on System z

As experience demonstrates, these can be run at up to 100% sustained utilization levels, although most clients prefer to leave a bit of white space and run at 90% or slightly under.

The Processor Resource/Systems Management (PR/SM™) function, responsible for hardware virtualization of the server, is always active and has been enhanced to provide additional performance benefits. PR/SM technology on the z10 EC has received Common Criteria EAL5¹ security certification. Each logical partition is as secure as an isolated server.

The z10 EC also offers software virtualization through z/VM. z/VM's virtualized z/Architecture servers, known as virtual machines, support all operating systems and other software supported on a logical partition. In fact, a z/VM virtual machine is the functional equivalent of a real server.

z/VM's extreme virtualization capabilities, which have been perfected since its introduction in 1967, enable virtualization of thousands of distributed servers on a single z10 EC server. IBM is conducting a very large internal consolidation project, which aims to consolidate approximately 3,900 distributed servers into approximately 30 mainframes, using z/VM and Linux on System z. The project expects to achieve reductions of over 80% in the use of space and energy. So far, expectations are being fulfilled. Similar results have been publicly presented by various clients, and these reductions directly translate into significant monetary savings.

Consider also the potential gains in software licensing. The pricing model for many distributed software products is linked to the number of processors or processor cores. Consolidating under z/VM and exploiting the specialized Integrated Facility for Linux (IFL) processors can achieve a large reduction in the number of used cores. For a discussion of available processor types see "PU characterization" on page 25.

¹ Evaluation Assurance Level with specific Target of Evaluation, Certificate for System z10 EC published October 29, 2008.

In addition to server consolidation and image reduction by vertical growth under z/VM, z/OS provides a highly sophisticated environment for application integration and co-residence with data, especially for the mission-critical applications.

The z10 EC expands the subcapacity settings offer to up to 12 central processors (CPs), delivering the scalability and granularity to meet the needs of medium-sized enterprises, while also satisfying the requirements of large enterprises having large-scale, mission-critical transaction and data-processing requirements. The first 12 processors can be configured at three different sub-capacity levels, giving a total of one hundred distinct capacity settings in the system, and providing for a range of 1:140 in processing power.

In the same footprint, the z10 EC 64-way server can deliver up to 70% more capacity than the largest z9® EC (the largest z9 EC is a 54-way). The z10 EC continues to offer all the specialty engines available with System z9®.

Most upgrades are concurrent to the hardware. As we describe later, the z10 EC reaches new availability levels by eliminating various pre-planning needs and other disruptive operations.

Summing up these characteristics leads to an interesting result:

- Capacity range and flexibility
 - + A processor equally able to handle compute-intensive and I/O-intensive workloads
 - + Specialty engines for improved price/performance
 - + Extreme virtualization
-
- = A very wide scope of applications that can be considered for server consolidation and application integration.

This presents a significant opportunity for most enterprises to simplify their IT infrastructures. The mainframe's inherent reliability, security, and availability, as well as its operational model, can now be of benefit to other, up to now distributed, applications.

Further simplification is possible by exploiting the z10 EC HiperSockets™² and z/VM's virtual switch functions. These may be used, at no additional cost, to replace physical routers, switches, and their cables, while eliminating security exposures and simplifying configuration and administration tasks. In some real simplification cases cables have been reduced by 97%.

IT operational simplification benefits also from the intrinsic autonomic characteristics of the z10 EC, the consolidation and reduction of the number of system images, and the management best practices and products developed and available for the mainframe, in particular for the z/OS environment.

Shared

By shifting the focus from operational management to service management, this stage creates a shared IT infrastructure that can be provisioned and scaled rapidly and efficiently. Organizations can create virtualized resource pools for server platforms, storage systems, networks, and applications, delivering IT capabilities to users in a more flexible way.

Many clients use their mainframe and application investments to support future business growth and to provide an important competitive advantage. Having chosen the mainframe as the platform to support their environment, these clients are making *on demand* business a reality. Yet other clients consider the mainframe based on the combined price-performance improvements of software and hardware.

² For a description of HiperSockets see "HiperSockets" on page 14. The z/VM virtual switch is a z/VM system function that uses memory to emulate switching hardware.

An important point is that the z10 *stack* consists of much more than just a server. This is because of the total-systems view that guides System z development. The *z-stack* is built around services, systems management, software, and storage. It delivers a complete range of policy-driven functions, pioneered and most advanced in the z/OS environment, including:

- ▶ Access management to authenticate and authorize who can access specific business services and associated IT resources.
- ▶ Utilization management to drive maximum use of the system. Unlike other classes of servers, z10 is designed to run at 100% of utilization 100% of the time, based on the varied demands of its users.
- ▶ Just-in-time capacity to deliver additional processing power and capacity when needed.
- ▶ Virtualization security to enable clients to allocate resources on demand without fear of security risks.
- ▶ Enterprise-wide operational management and automation, leading to a more autonomic environment.

In addition to the hardware-enabled resource sharing, other uses of virtualization include:

- ▶ Isolating production, test, training, and development environments
- ▶ Supporting back-level applications
- ▶ Enabling parallel migration to new system or application levels, and providing easy back-out capabilities

The resource-sharing abilities of the z/VM operating system can drive additional savings by:

- ▶ Allowing dormant servers that do not use resources to be activated when required. This can help reduce hardware, software, and maintenance costs.
- ▶ Pooling resources such as processor, I/O facilities, and disk space. Virtual servers can be dynamically provisioned out of these pools and, when their useful life ends, the resources are returned to the pools and recycled, with the utmost security.
- ▶ Offering very fast virtual server provisioning. A complete server can be deployed and ready for use in just a few minutes, using resources from the pool and image cloning.
- ▶ Eliminating the need to re-certify servers for specific purposes. Environments are certified to the virtual server. This must be done only once, even if the server requires scaling up, because the underlying hardware and architecture does not change. Significant reductions in time and manpower can be achieved.
- ▶ Use virtualized resources to test hardware configurations without incurring the cost of buying the actual hardware, and providing the flexibility to easily optimize these configurations.

Dynamic

At this stage, organizations achieve alignment with business goals and can respond dynamically as business needs arise. Opposite from the *break/fix* mentality gripping many data centers, this new environment creates an infrastructure that is economical, integrated, agile and responsive, having harvested new technologies to support the new types of business enterprises. Social networks, highly integrated Web 2.0 applications, and cloud computing deliver a rich environment and real-time information, as needed.

System z is the premier server offering from IBM, and the result of sustained and continuous investment and development policies. Commitment to IBM Systems design means that z10 EC brings all this innovation while helping customers leverage their current investment in the mainframe, as well as helping to improve the economics of IT.

The System z10 EC continues the evolution of the mainframe, building upon the z/Architecture definitions. The System z10 EC extends and integrates key platform characteristics such as dynamic and flexible partitioning, resource management for mixed and unpredictable workload environments, availability, scalability, clustering, and security and systems management with emerging e-business on demand application technologies, such as WebSphere, Java™, and Linux.

There is no conflict between application modular design and integrated application deployment. In fact, the availability, as a deployment choice, of an environment with shared resources, rich operational characteristics, and time-tested capabilities is highly desirable.

All of these technologies and improvements come into play when the z10 EC is at the heart of the SOA solutions for an enterprise. In particular, the high-availability, security, and scalability requirements of an Enterprise Service Bus (ESB) make its deployment on a mainframe environment highly advisable.

IBM mainframes traditionally provide an advanced combination of reliability, availability, security, scalability, and virtualization. The z10 EC has been designed to extend these capabilities and is optimized for today's business needs. The z10 EC is a platform of choice for the integration of the new generations of applications with existing applications and data.

z10 at the core of a dynamic infrastructure

A dynamic infrastructure is able to rapidly respond to sudden requirements, even unforeseen ones. It is resilient, highly automated, optimized, and efficient and offers a catalog of services while allowing granular metering and billing of those services.

The z10 EC enhances the availability and flexibility of just-in-time deployment of additional server resources, known as Capacity on Demand (CoD). CoD provides flexibility, granularity, and responsiveness by allowing the user to dynamically change capacity when business requirements change. With the proper contracts, up to eight temporary capacity offerings can be installed on the server. Additional capacity resources can be dynamically activated, either fully or in part, by using granular activation controls directly from the management console, without the having to interact with IBM Support.

IBM has further enhanced and extended the z10 EC leadership with improved access to data and the network. The following list indicates several of many enhancements:

- ▶ Tighter security with a CP Assist for Cryptographic Function (CPACF) protected key and longer personal account numbers for stronger protection of data
- ▶ Enhancements for improved performance connecting to the network
- ▶ Increased flexibility in defining your options to handle backup requirements
- ▶ Enhanced time accuracy to an external time source

A fast-growing number of enterprises are reaching the limits of available physical space and electrical power at their data centers. The extreme virtualization capabilities of the System z10 Enterprise Class enable the creation of dense and simplified infrastructures that are highly secure and can lower operational costs.

In summary, System z10 characteristics and qualities of service offer an excellent match to the requirements of a dynamic infrastructure, and this is why it is claimed to be at the core of such an infrastructure.

System z10 can improve the integration of people, processes, and technology to help run the business more cost effectively while also supporting business growth and innovation. It is,

thus, the most powerful tool available to reduce cost, energy, and complexity in enterprise data centers.

Tivoli adds value to System z10 and the data center

Clients will find that the linchpin of the data center is Tivoli's Service Management portfolio, which includes IBM Tivoli Service Management Center (ITSMCz) software.

ITSMCz automates the management of complex IT disciplines, including storage, databases, and new software deployments. It also addresses business mandates such as reducing IT costs and managing IT operations and regulatory and compliance requirements.

Visibility, control, and automation are built into the ITSMCz:

- ▶ **Visibility:** Clients have a single, integrated view of their critical applications on their mainframe, showing the linkages between IT assets and business applications.
- ▶ **Control:** Clients can customize their view to support functions such as business services, services requests, finance, security, IT production, support, and operational control.
- ▶ **Automation:** Based on a client's requirements, ITSMCz combines process automation software such as IBM Tivoli's Change and Configuration Management Database, Application Dependency Discovery Manager, Business Service Manager, and Service Request Manager®.

Storage is part of the System z10 stack

The ongoing synergy between System z and System Storage™ design teams has resulted in several compelling enhancements in the last few years, including the Modified Indirect Data Address Word (MIDAW) facility, which helps improve channel efficiency and throughput for Extended Format data sets including DB2® and VSAM), and High Performance FICON® for System z (zHPF), which improves performance for small data transfers of online transaction processing (OLTP) workloads such as DB2, VSAM, PDSE, and zFS.

Recent advances in IBM System Storage disk technology give clients the opportunity to take advantage of the IBM disk offerings' increased function and value, especially in the area of secure data encryption. Those offerings include updated business continuity features that make the most of the new mainframe's power.

Also for the System z10, the IBM System Storage Virtual Tape solution delivers improved tape processing while supporting business continuity and security through innovative enhancements.

1.2 z10 EC comparison

The System z10 Enterprise Class is a follow-on to the System z9 Enterprise Class (z9 EC), which was announced in July 2005. The z10 EC employs leading-edge silicon on insulator (CMOS 11S-SOI) and other technologies, such as InfiniBand, to provide advantages such as very high frequency chips, additional granularity options, improved availability, and enhanced on demand options. In addition, it supports the latest offerings for data encryption.

Five models of the z10 EC are offered. These are named E12, E26, E40, E56, and E64. The names represent the maximum number of configurable processors in the model.

The z10 EC system architecture ensures continuity and upgradability from the z9 EC design. Upgrading from IBM zSeries Model 990 servers is also possible.

Figure 1-1 provides a comparison of System z10 Enterprise Class with previous System z servers along four major attributes:

- ▶ Single engine processing capacity
- ▶ Number of engines
- ▶ Memory
- ▶ I/O bandwidth




Figure 1-1 System z design comparison

1.3 z10 EC server enclosure

In this section we briefly review the most significant characteristics of the System z10 Enterprise Class. Chapter 2, “Hardware overview” on page 17, provides further details.

The z10 EC server is a twin-frame system. It has a machine type designation of 2097. Figure 1-2 shows an external view of the server.



Figure 1-2 System z10 Enterprise Class external view

The frames in the z10 EC are known as the A frame and the Z frame.

The A frame contains:

- ▶ The central electronic complex (CEC)
- ▶ Modular cooling units
- ▶ One I/O cage
- ▶ Power supplies
- ▶ An optional internal battery feature (IBF)

The Z frame contains:

- ▶ Two system support elements (SEs)
- ▶ Zero, one, or two additional I/O cages
- ▶ Power supplies
- ▶ An optional IBF

The two redundant SEs are used to configure and manage the z10 EC server (for example, defining the I/O configuration and configuring the logical partitions).

1.3.1 Central Electronic Complex

The CEC is housed in its own cage. The cage houses from one to four processor books that are fully interconnected. Each book contains a multi-chip module (MCM), memory and I/O cage connectors, and (optionally) coupling link connectors.

The z10 EC is built on the proven superscalar microprocessor architecture already deployed on the z9 EC. However, the new processor unit (PU) chip has several distinctive innovations, notably in error checking and correcting, and new specialized circuitry, for instance, to support

decimal floating point operations. And, of course, it has a 4.4 GHz high-speed quad-core design.

Each book has one MCM that houses five PU chips and two storage control (SC) chips. Each PU chip has either three or four enabled cores. The MCM continues to be cooled by modular refrigeration units (MRUs) with air-cooling backup.

In any model of the server, two cores are designated as spares, and each individual core can be transparently spared, as contrasted with previous systems where the chip was the sparing unit.

Memory has been increased, as compared with the z9 EC. In each book, up to 384 GB can be installed, but because 16 GB are part of the base and reserved for the hardware system area (HSA), the maximum amount of purchasable memory is 1520 GB, just shy of 1.5 TB. Plan-ahead memory, a new capability whereby memory can be installed but not enabled for use until needed, further enhances system availability for continuous operations.

PU characterization

At server initialization time, each purchased PU is *characterized* as one of a variety of types. It is also possible to dynamically characterize PUs. A PU that is not characterized cannot be used. A PU may be characterized as follows:

- CP** Central processor: the standard z10 EC processors. For use with any supported operating system and user applications.
- ICF** Internal coupling facility: used for z/OS clustering. ICFs are dedicated for this purpose and exclusively run the Coupling Facility Control Code (CFCC).
- IFL** Integrated Facility for Linux: exploited by Linux and for z/VM processing in support of Linux. z/VM is often used to host multiple Linux virtual machines (called guests). It is not possible to IPL operating systems other than z/VM or Linux on an IFL.
- SAP** System assist processor: offloads and manages I/O operations. Several are standard with the z10 EC. More may be configured if additional I/O processing capacity is needed.
- zAAP³** System z10 Application Assist Processor: exploited under z/OS for designated workloads, which include the IBM JVM and some XML System Services functions.
- zIIP³** System z10 Integrated Information Processor: exploited under z/OS for designated workloads, which include various XML System Services, IPsec off-load, certain parts of DB2 DRDA®, star schema, HiperSockets for large messages, and the IBM GBS Scalable Architecture for Financial Reporting.

Note: Work dispatched on zAAP and zIIP does not incur any IBM software charges.

It is also possible to run a zAAP-eligible workload on zIIPs^a if no zAAPs are installed on the server. This capability is offered to enable optimization and maximization of investment on zIIPs.

a. This capability is available with z/OS V1.11 (and z/OS V1.9 and V1.10 with PTF for APAR OA27495) on all z9 and z10 servers. Some additional restrictions apply.

CP Assist for Cryptographic Function (CPACF)

The z10 EC continues to use the Cryptographic Assist Architecture first implemented on z990. Further enhancements have been made to the z10 EC CPACF.

³ z/VM V5 R3 and later support zIIP and zAAP processors for z/OS guest workloads.

CPACF is physically implemented in the quad-core chip by the Compression and Cryptography Accelerator (CCA). Each of the two CCAs is shared by two cores. CPACF supported protocols include:

- ▶ Data Encryption Standard (DES).
- ▶ Triple Data Encryption Standard (TDES).
- ▶ Secure Hash Algorithm (SHA):
 - SHA-1: 160 bit
 - SHA-2:
 - 224 bit
 - 256 bit
 - 384 bit
 - 512 bit
- ▶ Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys.
- ▶ Pseudo Random Number Generation (PRNG).

Note that PRNG is also a standard function supported on the Crypto Express features.
- ▶ Random number generation long (RNGL): 8 bytes to 8096 bytes.
- ▶ Random number generation (RNG) with up to 4096-bit key RSA support.
- ▶ The CPACF functions are supported by z/OS, z/VM, z/VSE, and Linux on System z.

1.3.2 I/O subsystem

As with its predecessors, the z10 EC server has a dedicated subsystem to manage all input/output operations. Known as the channel subsystem, it is composed of:

SAP	System assist processor: a specialized processor that uses the installed PU cores ⁴ . Its role is to offload I/O operations and manage channels and the I/O operations queues. It relieves the other PUs of all I/O tasks, allowing them to be dedicated to application logic. An adequate number of SAP processors is automatically defined, depending on the number of installed books. These are part of the base configuration of the server.
HSA	Hardware system area: A reserved part of the system memory, it contains the I/O configuration and is used by SAPs. On the z10 EC a fixed amount of 16 GB is reserved, which is not part of the customer-purchased memory. This provides for greater configuration flexibility and higher availability by eliminating some planned and pre-planned disruptive situations.
Channels	Small processors that communicate with the I/O control units (CUs). They manage the data transfer between memory and the external device. Channels are contained in the I/O card features.
Channel path	The means that the channel subsystem uses this to communicate with the I/O devices. Due to I/O virtualization, multiple independent channel paths can be established on a single channel, allowing the channel to be shared ⁵ between multiple logical partitions and each partition to have its unique channel path.
Subchannels	Appears to a program as a logical device and contains the information required to perform an I/O operation. One subchannel exists for each I/O device addressable by the channel subsystem.

⁴ Each z10 EC PU can be characterized as one of six different configurations. For more information see "PU characterization" on page 25.

The z/Architecture specifies an I/O subsystem to which all I/O processing is offloaded. This is a big contributor to the performance and availability of the system and strongly contrasts with architectures of the other servers.

The z10 EC I/O subsystem direction is evolutionary, drawing on development from the z990 and z9 EC. It continues to be based on I/O cages, I/O cards, and I/O buses. The I/O subsystem is supported by a new I/O bus, and includes the InfiniBand infrastructure (replacing the self-timed interconnect features found in the prior System z servers). This new infrastructure is designed to reduce overhead and latency and provide increased data throughput.

InfiniBand

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

A single connection is capable of carrying several types of traffic, such as communications, management, clustering, and storage. Additional characteristics include low processing overhead, low latency, and high bandwidth. Thus, it can become quite pervasive.

InfiniBand is very scalable, as experience proves, from two-node interconnects to clusters of thousands of nodes, including high-performance computing clusters. It is a mature and field-proven technology, used in thousands of data centers.

InfiniBand is being exploited by the z10 EC server. Internally, in the server, the new cables from the CEC cage to the I/O cages now carry the InfiniBand protocol. For external usage, Parallel Sysplex® InfiniBand (PSIFB) links are introduced. They are used to interconnect System z servers in a Parallel Sysplex.

1.3.3 I/O connectivity

The z10 EC generation of the I/O platform, particularly through the exploitation of InfiniBand, OSA-Express3, FICON Express8, and High Performance FICON for System z (zHPF), is intended to provide significant performance improvements over the previous I/O platform used for FICON Express4 and OSA-Express2.

I/O cage

The z10 EC has a CEC cage and, as a minimum, one I/O cage in the A frame. The Z frame can accommodate two additional I/O cages, bringing the total for the system to three. The I/O cages can accommodate the following feature types:

- ▶ ESCON®
- ▶ FICON Express8, FICON Express4, FICON Express2, and FICON Express
- ▶ OSA-Express3 and OSA-Express2
- ▶ Crypto Express3 and Crypto Express2
- ▶ Coupling links (ISC-3)

It is possible to populate the 28 I/O slots in one I/O cage with any mix of the above-mentioned cards.

⁵ The function that allows sharing I/O paths across logical partitions is known as the multiple image facility (MIF).

ESCON channels

The Enterprise Systems Connection (ESCON) channels support connectivity to ESCON disks, tapes, and printer devices. Historically, they represent the first use of optical I/O technology on the mainframe. They are much slower than FICON channels, which are the preferred technology. There are no changes in the ESCON support as compared with z9 EC.

FICON channels

Fibre Connection (FICON) channels follow the Fibre Channel (FC) standard and support data storage and access requirements as well as the latest FC technology in storage and access devices. FICON channels support the following protocols:

- ▶ Native FICON, Channel-to-Channel (CTC) connectivity, and zHPF traffic to FICON devices such as disks, tapes, and printers in z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments.
- ▶ Fibre Channel Protocol (FCP) in z/VM and Linux on System z environments support connectivity to disks and tapes through Fibre Channel switches and directors. z/VSE supports FCP for SCSI disks only. The FCP channel can connect to FCP SAN fabrics and access FCP/SCSI devices.

It is possible to choose any combination of the FICON Express8, FICON Express4, FICON Express2, and FICON Express features. Depending on the feature, auto-negotiated link data rates of 1, 2, 4, or 8 Gbps are supported. FICON Express8 provides significant improvements in start I/Os and data throughput.

Open Systems Adapter (OSA)

The Open Systems Adapter features provide local networking (LAN) connectivity and comply with IEEE standards. In addition, OSA features assume several functions of the TCP/IP stack that would normally be performed by the processor. This can provide significant performance benefits.

The z10 EC can have up to 24 OSA features (96 ports). It is possible to choose any combination of OSA-Express2 and OSA-Express3 features. For example, 1000BASE-T Ethernet supporting 10, 100, and 1000 Mbps over copper cabling or Gigabit Ethernet and 10 Gigabit Ethernet for the multimode and single mode fiber optic cabling environments.

Crypto

The Crypto Express features (Crypto Express3 and Crypto Express2) provide for tamperproof, high-performance cryptographic operations. Each feature has two PCI-X or PCI Express adapters. Each of the adapters can be configured as either a coprocessor or an accelerator:

- ▶ Crypto Express Coprocessor: for secure key-encrypted transactions (default)
 - Designed to support security-rich 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 Express Accelerator: for Secure Sockets Layer (SSL) acceleration
 - Designed to support high-performance clear key RSA operations
 - Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol

Support for 13-digit through 19-digit personal account numbers is provided for stronger protection of data.

The tamper-resistant hardware security module, which is contained in the Crypto Express features, is designed to meet the FIPS 140-2 Level 4 security requirements for hardware security models.

The configurable Crypto Express features are supported by z/OS, z/VM, and Linux on System z. z/VSE supports clear-key RSA operations only.

Coupling links

Coupling links are used in Parallel Sysplex cluster configurations of System z servers. The links provide high-speed bidirectional communication between members of the Sysplex. The z10 EC supports internal coupling links for memory-to-memory transfers, Integrated Cluster Bus-4 (ICB-4) for data sharing at distances up to 10 meters, 12x InfiniBand for distances up to 150 meters (492 feet), and InterSystem Channel-3 (ISC-3) and 1x InfiniBand for unrepeated distances up to 10 km (6.2 miles).

HiperSockets

The HiperSockets function is an integrated function of the z10 that provides users with attachments to up to 16 high-speed *virtual* local area networks with minimal system and network overhead.

HiperSockets is a function of the virtualization Licensed Internal Code (LIC) and performs memory-to-memory data transfers in a totally secure way. HiperSockets eliminates having to utilize I/O subsystem operations and having to traverse an external network connection to communicate between logical partitions in the same z10 EC server. Therefore, HiperSockets offers significant value in server consolidation by connecting virtual servers.

1.4 Performance

The z10 EC Model E64 is designed to offer approximately 1.7 times more capacity than the z9 EC Model S54 system. Uniprocessor performance has also increased significantly. A z10 EC Model 701 offers performance improvements of up to 1.62 times the z9 EC Model 701.

On average, the z10 EC can deliver up to 50% more performance in an n-way configuration than an IBM System z9 EC n-way. However, variations on the observed performance increase are dependent upon the workload type.

IBM continues to measure performance of the systems by using a variety of workloads and publishes the results in the Large Systems Performance Reference (LSPR) report. The LSPR is available at:

<http://www.ibm.com/servers/eserver/zseries/lspr/>

The MSU ratings are available at:

<http://www.ibm.com/servers/eserver/zseries/library/swpriceinfo>

LSPR workload suite

The LSPR workload suite comprises the following workloads:

- ▶ Traditional online transaction processing workload OLTP-T (formerly known as IMS)
- ▶ Web-enabled online transaction processing workload OLTP-W (also known as Web/CICS/DB2)
- ▶ A heavy Java-based online stock trading application WASDB (previously referred to as Trade2-EJB).

- ▶ Batch processing, represented by the CB-L (commercial batch with long-running jobs or CBW2)
- ▶ A new ODE-B Java batch workload, replacing the CB-J workload

The traditional Commercial Batch Short Job Steps (CB-S) workload (formerly CB84) has been dropped.

The LSPR provides performance ratios for individual workloads and for the default mixed workload, which is composed of equal amounts of four of the workloads described above (OLTP-T, OLTP-W, WASDB, and CB-L). The z10 EC LSPR tables continue to rate all z/Architecture processors running in LPAR mode and 64-bit mode. The single-number metrics are based on a combination of the default mixed workload ratios, typical multi-LPAR configurations, and expected early-program migration scenarios. In addition to z/OS workloads used to set the single-number metrics, the z10 EC LSPR tables contain information pertaining to Linux and z/VM environments.

Capacity ratio estimates

Figure 1-3 shows the estimated capacity ratios for z10 EC and z9 EC. The capacity estimate is based on the LSPR workload suite described previously.




Figure 1-3 z10 EC to z9 EC performance comparison

The LSPR contains the internal throughput rate ratios (ITRRs) for the z10 EC and the previous generations of processors based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user might experience varies depending on factors such as the amount of multiprogramming in the user's job stream, the I/O configuration, and the workload processed.

Workload performance variation

Because of the nature of the z10 EC multi-book system and resource management across those books, high performance variability similar to that seen on the z990 and z9 EC is expected. This variability can be observed in several ways. The range of performance ratings across the individual workloads is likely to have a large spread. The customer impact of this increased variability is seen as increased deviations of workloads from single-number metric-based factors such as MIPS, MSUs, and CPU time chargeback algorithms.

HiperDispatch

HiperDispatch is a function exclusive to the z10 servers. This function greatly enhances the processor affinity support to further improve the usable capacity of the server. In addition to the PR/SM enhancements, it requires support by z/OS. Now, PR/SM and z/OS cooperate on affinity assignments, which are done at the task level rather than just at the partition level.

1.5 Capacity On Demand

New with the System z10 Enterprise Class is the possibility of doing just-in-time deployment of capacity resources. This function is designed to provide more flexibility to dynamically change capacity when business requirements change. No interaction is required with IBM at the time of activation. It is possible to:

- ▶ Define one or more flexible configurations that can be used to solve multiple temporary situations. Previously, only one configuration was possible.
- ▶ Have multiple configurations active at once, and the configurations themselves have flexible selective activation of only the needed resources.
- ▶ Purchase capacity either before or after execution for On/Off Capacity on Demand. This capacity is represented by tokens that are consumed at execution time.
- ▶ Add permanent capacity to the server when temporary changes are active.

1.6 Software

The z10 EC is supported by a large set of software, including ISV applications. This section only lists the supported operating systems. Exploitation of some features may require the latest releases. Chapter 4, “Software support” on page 75, provides further information.

Operating system support for the System z10 EC includes:

- ▶ z/OS Version 1 Release 7 with IBM Lifecycle Extension and z/OS Version 1 Release 8 with IBM Lifecycle Extension
Note that z/OS.e is not supported.
- ▶ z/OS Version 1 Release 9 or later
- ▶ z/VM Version 5 Release 3 or later
- ▶ z/VSE Version 4 Release 1 or later
- ▶ TPF Version 4 Release 1 and z/TPF Version 1 Release 1
- ▶ Linux on System z distributions:
 - Novell SUSE: SLES⁶ 9, SLES 10, and SLES 11
 - Red Hat: RHEL⁷ 4 or RHEL 5

⁶ SLES is the abbreviation for Novell SUSE Linux Enterprise Server.

⁷ RHEL is the abbreviation for Red Hat Enterprise Linux

Hardware overview

The System z10 Enterprise Class server is the next step in the evolution of the mainframe family. It continues this evolution by introducing several innovations and expanding existing functions, building upon the z/Architecture.

This chapter expands upon the overview of key hardware elements of the System z10 Enterprise Class server provided in Chapter 1, “Introducing the System z10 Enterprise Class” on page 1, and compares them with the System z9 EC server, where relevant.

This chapter discusses the following topics:

- ▶ 2.1, “z10 EC highlights” on page 18
- ▶ 2.2, “Models and model upgrades” on page 18
- ▶ 2.3, “The frames” on page 21
- ▶ 2.4, “CEC cage and books” on page 22
- ▶ 2.5, “Multi-chip module (MCM)” on page 24
- ▶ 2.6, “z10 EC processor chip” on page 24
- ▶ 2.7, “Processor unit (PU)” on page 25
- ▶ 2.8, “Memory” on page 26
- ▶ 2.9, “I/O system structure” on page 28
- ▶ 2.10, “I/O cages and features” on page 29
- ▶ 2.11, “Cryptographic functions” on page 36
- ▶ 2.12, “Coupling and clustering” on page 39
- ▶ 2.13, “Time functions” on page 41
- ▶ 2.14, “HMC and SE” on page 42
- ▶ 2.15, “Power and cooling” on page 43

2.1 z10 EC highlights

The major System z10 Enterprise Class improvements over its predecessors include:

- ▶ Increased total system capacity in a 64-way server and additional subcapacity settings, offering increased levels of performance and scalability to help enable new business growth
- ▶ Quad-core 4.4 GHz processor chips that can help improve the execution of processor-intensive workloads
- ▶ Up to 1.5 TB of available real memory per server for growing application needs (with up to 1 TB real memory per logical partition)
- ▶ Just-in-time deployment of capacity resources, which can improve flexibility when making temporary or permanent changes, and plan-ahead memory for nondisruptive memory upgrades
- ▶ A new 16 GB fixed hardware system area (HSA) that is managed separately from customer-purchased memory
- ▶ Exploitation of InfiniBand technology
- ▶ Improvements to the I/O subsystem and new I/O features
- ▶ Additional security options for the CP Assist for Cryptographic Function (CPACF)
- ▶ A new HiperDispatch function for improved efficiencies in hardware and z/OS software
- ▶ Hardware decimal floating point on each core on the processor unit (PU)
- ▶ Server Timer Protocol (STP) enhancements for time accuracy, availability, and systems management with message exchanges using ISC-3 or 1x InfiniBand connections

In all, these enhancements provide customers with options for continued growth, continuity, and ability to upgrade.

For an in-depth discussion of the IBM System z10 Enterprise Class functions and features see the *IBM System z10 Enterprise Class Technical Guide*, SG24-7516.

2.2 Models and model upgrades

The System z10 Enterprise Class has been assigned a machine type (M/T) of 2097, which uniquely identifies the server. The server is offered in five different models. These models are named E12, E26, E40, E56, and E64. The model determines the maximum number of processor units (PUs) available for characterization. PUs are delivered in single-engine increments. The first four models utilize a 17-PU multi-chip module (MCM), of which 12 to 14 PUs are available for characterization. The fifth model, E64, utilizes one 17-PU MCM and three 20-PU MCMs to provide up to 64 configurable PUs.

As with the System z9 EC, spare PUs and system assist processors (SAPs) are integral to the server. Refer to Table 2-1 for a model summary including SAPs and spare PUs for the different models. For an explanation of PU characterization see “PU characterization” on page 25.

The five z10 EC server orderable models are:

- ▶ The z10 EC Model E12 has one book with 17 PUs, of which 12 can be characterized. The five remaining PUs are three SAPs and two spares.
- ▶ The z10 EC Model E26 has two books with 17 PUs in each book for a total of 34 PUs, of which 26 can be characterized. The eight remaining PUs are six SAPs, three in each book, and two spares, one in each book.
- ▶ The z10 EC Model E40 has three books with 17 PUs in each book for a total of 51 PUs, of which 40 can be characterized. The eleven remaining PUs are nine SAPs, three in each book, and two spares, one in book 0 and one in book 1.
- ▶ The z10 EC Model E56 has four books with 17 PUs in each book for a total of 68 PUs, of which 56 can be characterized. The 12 remaining PUs are 10 SAPs, three each in books 0, 1, and 2, and one in book 3, and two spares in book 3.
- ▶ The z10 EC Model E64 has four books with 17 PUs in book 0 and 20 PUs in books 1, 2, and 3 for a total of 77 PUs, of which 64 can be characterized. The 13 remaining PUs are eleven SAPs, one in book 0, three each in books 1 and 2, and four in book 3, and two spares, one in book 0 and one in book 1.

Table 2-1 Model summary

Model	Books/PUs	CPs	Standard SAPs	Spares
E12	1/17	0–12	3	2
E26	2/34	0–26	6	2
E40	3/51	0–40	9	2
E56	4/68	0–56	10	2
E64	4/77	0–64	11	2

The z10 EC offers 100 different capacity levels, which span a range of approximately 1 to 140. This is discussed in 3.2.2, “Granular capacity and capacity settings” on page 51.

Figure 2-1 summarizes the upgrade paths to z10 EC.

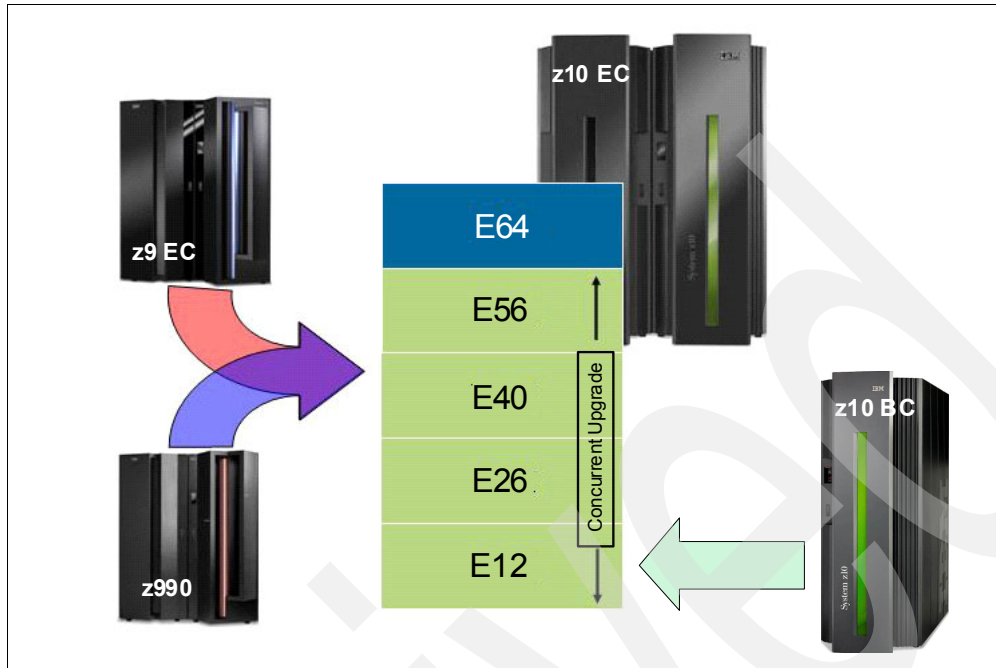


Figure 2-1 System z10 Enterprise Class upgrade paths

z10 EC upgrades

Model upgrades within the server family are accomplished by installing additional books. Books, being on separate power boundaries, are physically isolated from each other, thereby allowing them to be plugged and unplugged independently. Refer to Table 2-2 for upgrades available within the family.

Table 2-2 z10 EC to z10 EC upgrade paths

Model	E12	E26	E40	E56	E64
E12	—	Yes	Yes	Yes	Yes
E26	—	—	Yes	Yes	Yes
E40	—	—	—	Yes	Yes
E56	—	—	—	—	Yes

All z10 EC to z10 EC model upgrades are concurrent except when the target is the model E64. This is a non-concurrent upgrade because model E64 uses a different set of MCMs.

Upgrades to z10 EC from z10 BC

A z10 BC server can be upgraded to a z10 EC Model E12. This upgrade is disruptive.

Upgrades to z10 EC from z9 EC

Upgrades are also available from the currently installed z9 EC servers. The five model numbers for the z9 EC servers are S08, S18, S28, S38, and S54. These upgrades are disruptive.

Upgrades to z10 EC from z990

Similarly, the upgrades from currently installed z990 servers to z10 EC servers are also offered for all the models of the z990. The four model numbers for the z990 servers are A08, B16, C24, and D32. These upgrades are disruptive.

2.3 The frames

The System z10 Enterprise Class server is always a two-frame system. The frames are called the A frame and the Z frame. Refer to Figure 2-2 for an internal front view of the two frames. Several hardware elements pointed out are described later in this chapter.




Figure 2-2 z10 EC internal front view

The z10 EC is slightly bigger than the z9 EC. Refer to Table 2-3 for the physical dimensions of the system and its frames.

Table 2-3 z10 EC physical dimensions

Frame cover	Width mm (in)	Depth mm (in)	Height mm (in)
System with covers	1565 (61.60)	1854 (71.0)	2013.2 (79.26)
System with covers and reduction	1565 (61.60)	1854 (71.0)	1785.0 (70.30)
Each frame with one side cover and without packaging	780 (30.75)	1270 (50.0)	2013.2 (79.26)

2.4 CEC cage and books

The z10 EC server has a multi-book system structure similar to the z9 EC server. A book looks like a box and plugs into one of the four slots in the central electronic complex (CEC) cage of the z10 EC server. The CEC cage is located in the in A frame of the z10 EC server. Refer to Figure 2-2 on page 21 for a pictorial view of the CEC cage and the location of the four books. Each book contains:

- ▶ A multi-chip module (MCM). Each MCM includes five quad-core processor unit (PU) chips and two storage control (SC) chips. MCMs are further described in 2.5, “Multi-chip module (MCM)” on page 24. Refer to Table 2-1 on page 19 for the model summary and the relation between the number of books and number of available PUs.
- ▶ A minimum of 32 and a maximum of 384 GB of physical memory.
- ▶ A combination of up to eight InfiniBand Host Channel Adapter (HCA2-Optical or HCA2-Copper) fanout cards and memory bus adapter (MBA) fanout cards. Each of the cards has two ports, thereby supporting up to 16 connections. HCA2-Copper connections are for links to the I/O cages in the server, and the HCA2-Optical and MBA connections are to external servers (coupling links). MBA cards are used for ICB-4 links.
- ▶ Three distributed converter assemblies (DCAs) that provide power to the book. Loss of a DCA leaves enough book power to satisfy the book’s power requirements. The DCAs can be concurrently maintained.

Note: IBM has issued the following Statement of General Direction:

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

Figure 2-3 shows a view of a z10 EC book without the containing box.




Figure 2-3 z10 EC book structure and components

The z10 EC offers a significant increase in system scalability and opportunity for server consolidation by providing a multi-book system structure. As shown in Figure 2-4, all books are interconnected in a star configuration with high-speed communications links through the L2 caches, which allows the system to be operated and controlled by the PR/SM facility as a symmetrical, memory-coherent multiprocessor.

The ring topology employed in the z9 EC server is not used in the z10 EC. The point-to-point connection topology allows direct communication among all books. It is designed to get the maximum benefit of the improved processor clock speed.

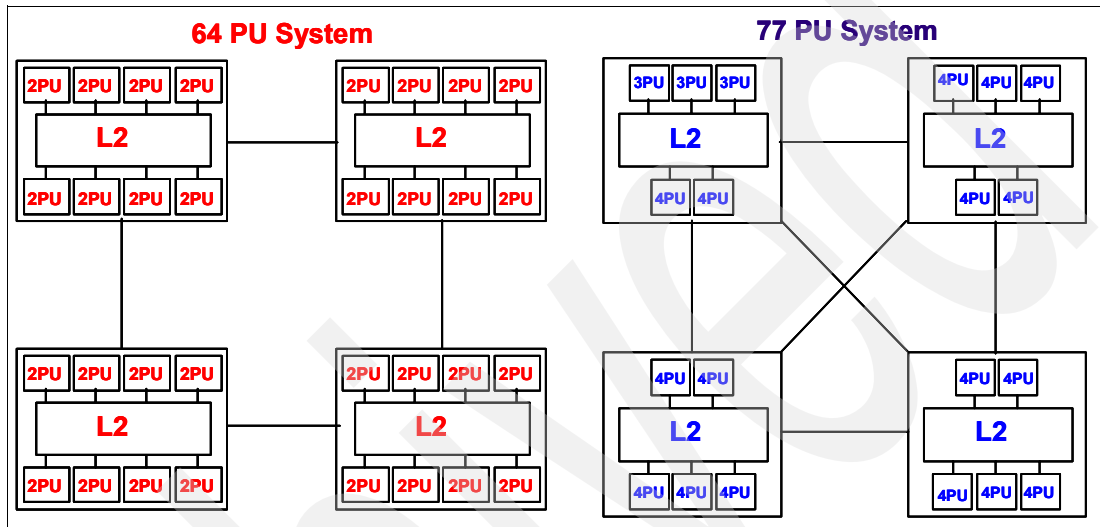


Figure 2-4 z9 EC versus z10 EC inter-book communication structure

Table 2-4 compares the book characteristics of the z9 EC to the z10 EC.

Table 2-4 z9 EC and z10 EC books

Characteristic	z9 EC	z10 EC
SMP configuration	4 books, 64 PUs	4 books, 77 PUs
Topology	Dual ring	Fully connected
Jumper books	Yes	No
Max memory	512 GB	1.5 TB
Cache levels	L1, L2	L1, L1.5, L2

2.5 Multi-chip module (MCM)

The multi-chip module is a high-performance, glass-ceramic module, providing the highest level of processing integration in the industry. It is the heart of the server (Figure 2-5).




Figure 2-5 z10 EC multi-chip module

The z10 EC MCM has seven chip sites. All chip types on the MCM use Complementary Metal Oxide of Silicon (CMOS) 11s chip technology. CMOS 11s is a state-of-the-art microprocessor technology based on 10-layer copper interconnections and silicon-on insulator technologies. The chip lithography line width is 0.065 μm (65 nm). The chip contains close to 1 billion transistors in a 450 mm² die.

There is one MCM per book and the MCM contains all of the processor chips and L2 cache of the book. The z10 EC server has five PU chips per MCM and each PU chip has up to four PUs (cores). Two MCM options are offered: with 17 or 20 processor units. All the models employ an MCM size of 17 PUs except for the model E64, which has one book with a 17 PU MCM and three books with 20 PU MCMs, for a total of 77 PUs.

The MCM also has two storage control (SC) chips. Each SC chip packs 24 MB of SRAM cache, interface logic for 20 cores, and SMP fabric logic into 450 mm². The two SC chips are configured to provide a single 48 MB cache shared by all 17 or 20 cores on the module, yielding outstanding SMP scalability on real-world transaction processing workloads.

There are four SEEPROM (S) chips, of which two are active and two are redundant, that contain product data for the MCM, chips, and other engineering information. The clock functions are distributed across PU and SC chips.

2.6 z10 EC processor chip

The z10 EC features an all-new high-frequency (4.4 GHz) four-core processor chip, a new microprocessor design, a robust cache hierarchy, and an SMP design optimized for enterprise database and transaction processing workloads, as well as for emerging workloads such as Java and Linux.

It leverages leading-edge technology and circuit design techniques while building on the rich heritage of mainframe system design, including industry-leading reliability, availability, and

serviceability. New functional features enable increased software efficiency and scalability while maintaining full compatibility with existing software. Further detail is given in 3.2.1, “Microprocessor enhancements” on page 50.

2.7 Processor unit (PU)

A PU is the generic term for the z/Architecture processor on the multi-chip module (MCM). A PU is imbedded in a System z chip core. Each PU is a superscalar processor with the following attributes:

- ▶ The basic cycle time is approximately 230 picoseconds.
- ▶ Up to two instructions can be decoded per cycle.
- ▶ Up to three instructions can be executed (finished) per cycle.
- ▶ Instructions are completed in the order in which they appear in the instructions stream. A high-frequency, low-latency, mostly in-order pipeline, providing robust performance across a wide range of workloads, is used.
- ▶ Memory accesses might not be in the same instruction order (out-of-order operand fetching).
- ▶ Most instructions flow through a pipeline with different numbers of steps for various types of instructions. Several instructions may be in progress at any moment, subject to the maximum number of decodes and completions per cycle.

Each PU has an L1 cache divided into a 64 KB cache for instructions and a 128 KB cache for data. Each PU also has an L1.5 cache. This cache is 3 MB in size. This implementation optimizes performance of the system for high-frequency, very fast processors.

Each L1 cache has a translation look-aside buffer (TLB) of 512 entries associated with it. In addition, a secondary TLB is used to further enhance performance. This structure supports large working sets, multiple address spaces, and a two-level virtualization architecture.

Hardware fault detection is imbedded throughout the design and combined with comprehensive instruction-level retry and dynamic CPU sparing. Those provide the reliability and availability required for true mainframe quality.

The z10 EC processor provides full compatibility with existing software for ESA/390 and z/Architecture, while extending the Instruction Set Architecture (ISA) to enable enhanced function and performance. Over 50 new hardware instructions support more efficient code generation, particularly for Java and C++ programs.

Decimal floating-point hardware fully implements the new IEEE 754r standard, helping provide better performance and higher precision for decimal calculations, an enhancement of particular interest to financial institutions.

On-chip cryptographic hardware includes extended key and hash sizes for the Advanced Encryption Standard (AES) and Secure Hash Algorithm (SHA) algorithms.

PU characterization

Processor units are ordered in single increments. The internal server functions, based on the configuration ordered, *characterize* processors into various types during initialization of the processor—often called a power-on reset (POR) operation. Characterizing PUs dynamically without a POR is possible. A processor that is not characterized cannot be used.

At least one CP must be purchased with, or before, a zAAP or zIIP can be purchased. Customers can purchase one zAAP, one zIIP, or both, for each CP (assigned or unassigned) on the server. However, a logical partition definition can contain more zAAPs or zIIPs than CPs. For example, in a server with two CPs a maximum of two zAAPs and two zIIPs can be installed. A logical partition definition for that server could contain up to two logical CPs, two logical zAAPs, and two logical zIIPs.

Converting a processor from one type to any other type is possible. These conversions happen concurrently with the operation of the system.

Notes: The addition of ICFs, IFLs, zAAP, zIIPs, and SAPs to a server does not change the server capacity setting or its MSU rating (only CPs do).

IBM does not impose any software charges on work dispatched on zAAP and zIIP processors.

2.8 Memory

Maximum physical memory sizes are directly related to the number of books in the system. Each book may contain a maximum of 384 GB of physical memory. Up to 1520 GB (~1.5 TB) of physical memory can be purchased. This is equal to four books x 384 GB minus 16 GB reserved for the hardware system area.

Memory can be purchased in increments of 16 GB up to a total size of 256 GB. From 256 GB, the increment size doubles to 32 GB until 512 GB. From 512 GB to 944 GB, the increment is 48 GB, and beyond that, up to 1520 GB, a 64 GB increment is used. Memory for books can be ordered as follows:

- ▶ A 1-book system (z10 EC Model E12) may contain from 64 GB to 384 GB of physical memory. Memory is orderable in 16 GB and 32 GB increments to up to 352 GB.
- ▶ A 2-book system (z10 EC Model E26) may contain from 128 GB to 768 GB of physical memory. Memory is orderable in 16 GB, 32 GB, and 48 GB increments to up to 752 GB.
- ▶ A 3-book system (z10 EC Model E40) may contain from 192 GB up to a maximum of 1152 GB of physical memory. Memory is orderable in 16 GB, 32 GB, 48 GB, and 64 GB increments to up to 1136 GB.
- ▶ A 4-book system (z10 EC Model E56 or z10 EC S64) may contain from 288 GB up to a maximum of 1536 GB of physical memory. Memory is orderable in 16 GB, 32 GB, 48 GB, and 64 GB increments to up to 1520 GB.

The fact that the maximum amount of orderable memory in each of the models is not equal to the maximum supported amount of physical memory is explained by 16 GB of physical memory being set aside for the HSA and, when certain amounts of memory are exceeded, the change in memory increment sizes.

Physically memory is organized as follows:

- ▶ A book always contains a minimum of 64 GB.
- ▶ A book may have more memory installed than enabled. The excess amount of memory can be installed by a Licensed Internal Code load when required by the installation.
- ▶ Memory upgrades are satisfied from already-installed unused memory capacity until exhausted. When no more unused memory is available from the installed memory cards, either the cards must be upgraded to a higher capacity or the addition of a book with additional memory is necessary.

When activated, a logical partition can use memory resources located in any book. No matter in which book the memory resides, a logical partition has access to that memory if so allocated. Despite the book structure, the z10 EC is still a Symmetric Multi-Processor (SMP).

A memory upgrade is concurrent when it requires no change of the physical memory cards. A memory card change is disruptive when no use is made of Enhanced Book Availability. Refer to *IBM System z10 Enterprise Class Technical Guide, SG24-7516*, for a description of Enhanced Book Availability.

For a model upgrade that results in the addition of a book, the minimum memory increment is added to the system. Remember that the minimum physical memory size in a book is 64 GB. During a model upgrade, the addition of a book is a concurrent operation. The addition of the physical memory that resides in the added book is also concurrent.

Concurrent memory upgrade

Memory can be upgraded concurrently using Licensed Internal Code - Configuration Control (LIC-CC) if physical memory is available as described previously. The plan-ahead memory function available with the z10 server provides the ability to plan for nondisruptive memory upgrades by having the system pre-plugged, based on a target configuration. Pre-plugged memory is enabled through a LIC-CC order placed by the customer.

Hardware system area

The hardware system area (HSA) is a reserved memory area that is used for several internal functions, but the bulk is used by channel subsystem functions. The HSA has grown with each successive mainframe generation. On previous servers, model upgrades and also new logical partition definitions or changes required pre-planning and were sometimes disruptive because of changes in HSA size. For further information and benefits see 3.2.3, "Memory enhancements" on page 52.

2.9 I/O system structure

Refer to Figure 2-6 for the I/O system structure overview for the z9 EC server and for the z10 EC server.

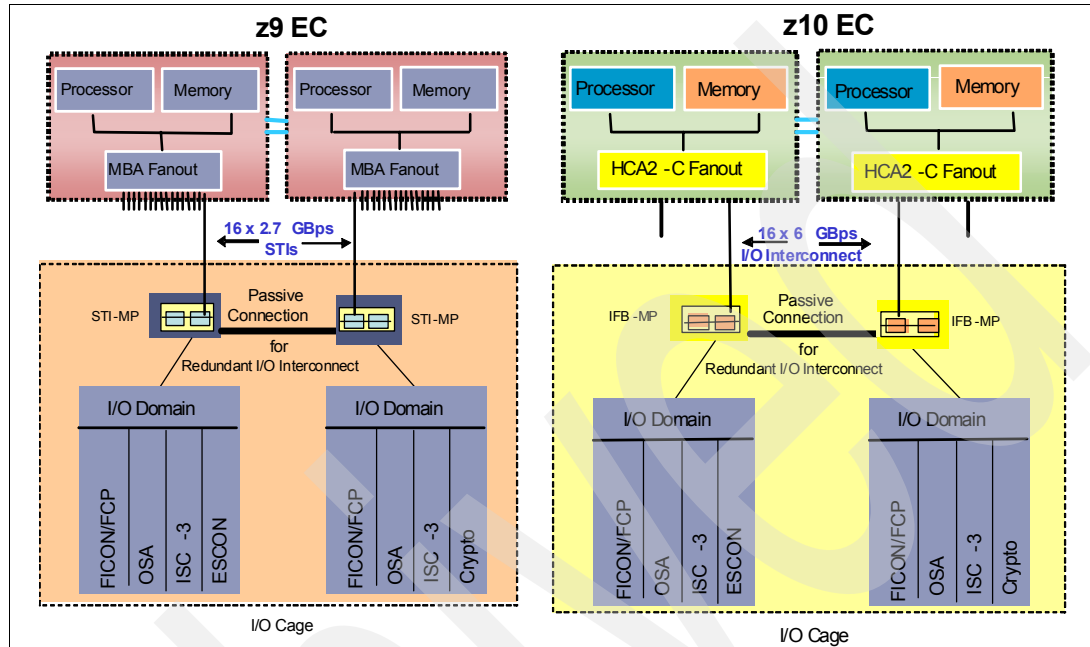


Figure 2-6 z9 EC and z10 EC system structure for I/O

The z10 EC has several types of fanout cards residing on the front of the book package:

- ▶ An InfiniBand HCA2-C (copper) fanout that supports ESCON, FICON, OSA, ISC-3, Crypto Express2, and Crypto Express3 features in the I/O cages
- ▶ An memory bus adapter (MBA) fanout that is used for ICB-4 connections at a distance of up to 10 meters (33 feet)

The z10 EC supports up to eight fanouts (HCA-C, HCA-O, HCA2-O LR, or the MBA) for each book, with a maximum of 24 for a 4-book system. Each fanout comes with two ports, giving a maximum of 48 ports for I/O connectivity.

The z10 EC exploits InfiniBand (IFB) connections to I/O cages, driven from the Host Channel Adapter (HCA2-C) fanout cards that are located on the front of the book. The HCA2-C fanout card is designated to connect to an I/O cage by a copper cable. The two ports on the fanout card are dedicated to I/O. This is different from the z9 EC, which uses self-timed interconnect (STI) connections driven from the MBA fanouts to connect to the I/O cages.

The z10 EC server has up to eight fanout cards (numbered D1, D2, and D5 to DA) per book, each driving two IFB cables, resulting in up to 16 IFB connections per book (16 STI connections with the z9 EC server).

In a system configured for maximum availability, alternate paths maintain access to critical I/O devices, such as disks, networks, and so on.

Refer to *System z Connectivity Handbook*, SC24-5444, for a more detailed description of the I/O interfaces.

Coupling connectivity

In addition to the HCA2-C fanout card, the z10 EC has three additional fanout cards, the HCA2-O, HCA2-O LR, and the MBA fanout. These cards are exclusively used for coupling link connectivity in a Parallel Sysplex configuration.

The HCA2-O provides optical connections for InfiniBand I/O interconnect (Parallel Sysplex using InfiniBand (PSIFB)) between:

- ▶ A z10 EC and another z10 EC or z10 BC. This connection has a maximum link data rate of up to 6 GB per second.
- ▶ A z10 EC and a System z9 server. This connection has a maximum data rate of up to 3 GB per second.

Note: The InfiniBand link data rate of 6 GBps or 3 GBps does not represent the performance of the link. The actual performance depends on many factors, such as latency through the adapters, cable lengths, and the type of workload. Although the link data rate can be higher with InfiniBand coupling links than with ICB links, the service times of coupling operations are greater.

The HCA-O LR card supports a maximum data link rate of 5 Gbps on a coupling link connection between two z10 servers. HCA2-O LR is exclusive to System z10 servers.

The MBA fanout card provides coupling links (ICB-4 only) to either z10 servers or to a z9, z990, or z890 server. This allows connecting to existing servers in existing Parallel Sysplex environments.

As indicated previously, the HCA2-C provides copper connections for InfiniBand I/O interconnect from book to I/O cards in I/O cages.

2.10 I/O cages and features

Each book has up to eight dual-port fanout cards to transfer data. Each port has a bi-directional bandwidth of 6 GBps. Up to 16 IFB I/O interconnect connections provide an aggregated bandwidth of up to 96 GBps per book.

The HCA2-C IFB I/O interconnect connects to an I/O cage that may contain a variety of channels, coupling link, OSA-Express, and cryptographic features.

The z10 EC server holds a minimum of one I/O cage at the bottom of the A frame and two optional I/O cages in the Z frame. Refer to Figure 2-2 on page 21, where all three I/O cages are shown.

Each I/O cage supports up to seven I/O domains and a total of 28 I/O card slots. Each I/O domain supports four I/O features (ESCON, FICON, OSA, or ISC). See Figure 2-7 for a pictorial view of an I/O cage.

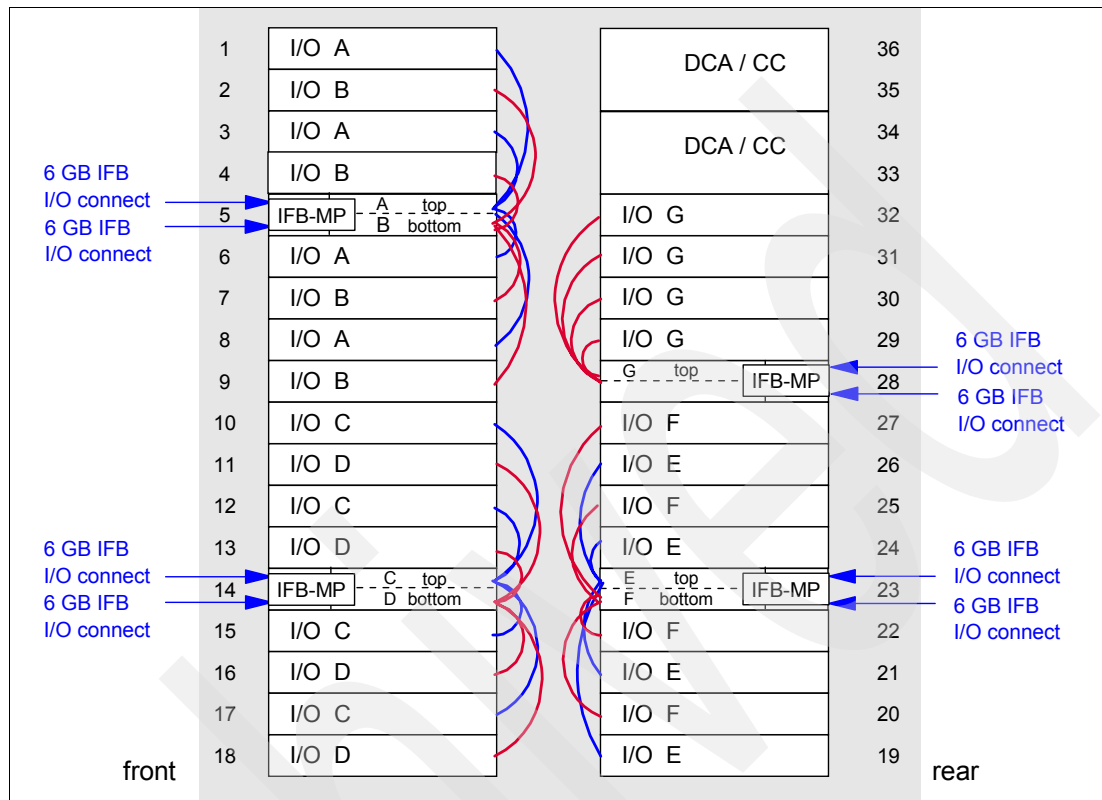


Figure 2-7 z10 EC I/O cage

The different I/O domains (A, B, C, D, E, F, and G) and the InfiniBand MultiPlexer (IFB-MP), which connects to the CEC cage as well as to the I/O feature itself, are shown. Up to four of the 32 slots in the I/O cage are occupied by the IFB-MB.

The following I/O features can be ordered for a new z10 EC server:

- ▶ ESCON
- ▶ FICON Express8 LX (long wavelength - 10 km)
- ▶ FICON Express8 SX (short wavelength)
- ▶ OSA-Express3 10 GbE LR (long reach)
- ▶ OSA-Express3 10 GbE SR (short reach)
- ▶ OSA-Express3 GbE LX (long wavelength)
- ▶ OSA-Express3 GbE SX (short wavelength)
- ▶ OSA-Express3 1000BASE-T Ethernet
- ▶ OSA-Express2 1000BASE-T Ethernet
- ▶ Crypto Express3
- ▶ ISC-3 (peer mode only)
- ▶ ICB-4 (not available on model E64)

The following features are not orderable for a z10 EC, but if present in a z9 EC or z990 server may be *carried forward* when upgrading to a z10 EC:

- ▶ FICON Express4 LX (4 km and 10 km)
- ▶ FICON Express4 SX
- ▶ FICON Express2 (LX and SX)
- ▶ FICON Express (LX and SX)
- ▶ OSA-Express2 10 GbE Long Reach
- ▶ OSA-Express2 LX (long wavelength)
- ▶ OSA-Express2 SX (short wavelength)
- ▶ Crypto Express2

The following z9 EC and z990 features are *not* supported on a z10 EC:

- ▶ FICON (pre-FICON Express)
- ▶ OSA-Express
- ▶ ICB-2
- ▶ ICB-3
- ▶ ISC-3 Links in Compatibility Mode
- ▶ PCIXCC and PCICA
- ▶ Parallel channels (use an ESCON converter)

For a list of the z10 EC supported I/O features and their characteristics refer to Appendix B, “Channel options” on page 109.

2.10.1 ESCON channels

ESCON channels support the ESCON architecture and directly attach to ESCON-supported I/O devices.

16-port ESCON feature

The 16-port ESCON feature occupies one I/O slot in an I/O cage. Each port on the feature uses a 1300 nanometer (nm) light-emitting diode (LED) transceiver, designed to be connected to 62.5 μ m multimode fiber optic cables only.

Up to a maximum of 15 ESCON channels per feature are active. There is a minimum of one spare port per feature to allow for channel sparing in the event of a failure of one of the other ports.

ESCON channel port enablement feature

The 15 active ports on each 16-port ESCON feature are activated in groups of four ports through LIC-CC. Each port operates at a data rate of 200 Mbps.

The first group of four ESCON ports requires two 16-port ESCON features. This is for redundancy reasons. After the first pair of ESCON cards is fully allocated (by seven ESCON port groups, using 28 ports), single cards are used for additional ESCON ports groups.

Ports are activated equally across all installed 16-port ESCON features for high availability.

Note: It is the intent of IBM for ESCON channels to be phased out. System z10 EC and System z10 BC will be the last servers to support greater than 240 ESCON channels. We recommend that you review the usage of your installed ESCON channels and wherever possible migrate to FICON channels.

The PRIZM Protocol Converter Appliance from Optica Technologies Incorporated provides a FICON-to-ESCON conversion function that has been System z qualified. For more information see:

<http://www.opticatech.com/>

Note: IBM cannot confirm the accuracy of compatibility, performance, or any other claims by vendors for products that have not been System z qualified. Questions regarding these capabilities and device support should be addressed to the suppliers of those products.

2.10.2 FICON Express8

Two types of FICON Express8 transceivers are supported on new build z10 EC servers—one long wavelength (LX) laser version and one short wavelength (SX) laser version:

- ▶ FICON Express8 10KM LX feature
- ▶ FICON Express8 SX feature

Each port supports attachment to the following:

- ▶ FICON/FCP switches and directors that support 2 Gbps, 4 Gbps, or 8 Gbps
- ▶ Control units that support 2 Gbps, 4 Gbps, or 8 Gbps

Note: FICON Express4, FICON Express2, and FICON Express features are withdrawn from marketing.

When upgrading to a System z10, replace your FICON Express, FICON Express2, and FICON Express4 features with FICON Express8 features. The FICON Express8 features offer better performance and increased bandwidth.

FICON Express8 10KM LX feature

The FICON Express8 10KM LX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector and auto-negotiated link speeds of 2 Gbps, 4 Gbps, and 8 Gbps up to an unrepeated maximum distance of 10 km (6.2 miles).

FICON Express8 SX feature

The FICON Express8 SX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector and auto-negotiated link speeds of 2 Gbps, 4 Gbps, and 8 Gbps up to an unrepeated maximum distance of up to 500 meters at 2 Gbps, 380 meters at 4 Gbps, or 150 meters at 8 Gbps.

2.10.3 FICON Express4

Three types of FICON Express4 transceivers are supported on z10 EC servers only if carried over during an upgrade—two long wavelength (LX) laser versions and one short wavelength (SX) laser version:

- ▶ FICON Express4 10KM LX feature
- ▶ FICON Express4 4KM LX feature
- ▶ FICON Express4 SX feature

Note: FICON Express4 features will be the last features to negotiate down to 1 Gbps. FICON Express4 features for z10 EC have been withdrawn from marketing.

Each port supports attachment to the following items:

- ▶ FICON/FCP switches and directors that support 1 Gbps, 2 Gbps, or 4 Gbps
- ▶ Control units that support 1 Gbps, 2 Gbps, or 4 Gbps

FICON Express4 10KM LX feature

The FICON Express4 10KM LX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector and link speeds of 1 Gbps, 2 Gbps, or 4 Gbps up to an unrepeated maximum distance of 10 km (6.2 miles).

FICON Express4 4KM LX feature

The FICON Express4 4KM LX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector and link speeds of 1 Gbps, 2 Gbps, or 4 Gbps up to an unrepeated maximum distance of 4 km (2.5 miles).

Interoperability of 10 km transceivers with 4 km transceivers is supported, provided that the unrepeated distance between the two transceivers does not exceed 4 km.

FICON Express4 SX feature

The FICON Express4 SX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector, and supports auto-negotiated link speeds of 1 Gbps, 2 Gbps, and 4 Gbps up to an unrepeated maximum distance of up to 860 meters operating at 1 Gbps, 500 meters operating at 2 Gbps, or 380 meters operating at 4 Gbps.

2.10.4 FICON Express2

The FICON Express2 feature is supported on a z10 EC only if carried over on an upgrade. Two types of FICON Express2 channel transceivers are supported on z10 EC servers when carried forward on an upgrade—a long wavelength (LX) laser version and a short wavelength (SX) laser version.

- ▶ FICON Express2 LX feature
- ▶ FICON Express2 SX feature

Each port supports attachment to the following items:

- ▶ FICON/FCP switches and directors that support 1 Gbps or 2 Gbps
- ▶ Control units that support 1 Gbps or 2 Gbps

FICON Express2 LX feature

The FICON Express2 LX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector and link speeds of 1 Gbps or 2 Gbps up to an unrepeated maximum distance of 10 km (6.2 miles).

FICON Express2 SX feature

The FICON Express2 SX feature occupies one I/O slot in the I/O cage. It has four ports, each supporting an LC duplex connector and auto-negotiated link speeds of 1 Gbps or 2 Gbps up to an unrepeated maximum distance of up to 860 at 1 Gbps or up to 500 at 2 Gbps.

2.10.5 FICON Express

FICON Express features may be carried forward to the z10 EC when upgrading from a z9 or z990 server. Two types of FICON Express channel transceivers are supported—a long wavelength (LX) laser version and a short wavelength (SX) laser version:

- ▶ FICON Express LX feature
- ▶ FICON Express SX feature

FICON Express LX feature

The FICON Express LX feature occupies one I/O slot in the I/O cage. It has two ports, each supporting an LC duplex connector and a link speed of 1 Gbps.

Each port supports attachment to the following items:

- ▶ FICON LX Bridge one-port feature of IBM 9032 ESCON Director at 1 Gbps *only*
- ▶ FICON/FCP switches and directors that support 1 Gbps
- ▶ Control units that support 1 Gbps

Note: FICON Express2 and FICON Express4 features do not support FCV mode. FCV mode is available on z10 EC only if the FICON Express LX feature is carried over on upgrades.

The z10 is intended to be the last server to support FICON Express LX feature and CHPID type FCV.

FICON Express SX feature

The FICON Express SX feature occupies one I/O slot in the I/O cage. It has two ports, each supporting an LC duplex connector and a link speed of 1 Gbps. Each port supports attachment to the following items:

- ▶ FICON/FCP switches and directors that support 1 Gbps
- ▶ Control units that support 1 Gbps

Refer to *System z Connectivity Handbook*, SC24-5444, or *FICON Planning and Implementation Guide*, SG24-6497, for more details.

2.10.6 OSA-Express3

This section describes the connectivity options offered by the OSA-Express3 features. The following OSA-Express3 features can be installed in z10 EC servers:

- ▶ OSA-Express3 10 Gigabit Ethernet (GbE) Long Reach (LR)
- ▶ OSA-Express3 10 Gigabit Ethernet Short Reach (SR)
- ▶ OSA-Express3 Gigabit Ethernet long wavelength (GbE LX)
- ▶ OSA-Express3 Gigabit Ethernet short wavelength (GbE SX)
- ▶ OSA-Express3 1000BASE-T Ethernet

OSA-Express3 10 GbE LR feature

The OSA-Express3 10 GbE LR feature occupies one slot in an I/O cage and has two ports that connect to a 10 Gbps Ethernet LAN through a 9 μ m single-mode fiber optic cable terminated with an LC Duplex connector. The feature supports an unrepeated maximum distance of 10 km.

The OSA-Express3 10 GbE LR feature replaces the OSA-Express2 10 GbE LR feature, which is no longer orderable.

OSA-Express3 10 GbE SR feature

The OSA-Express3 10 GbE SR feature occupies one slot in the I/O cage. It has two ports that connect to a 10 Gbps Ethernet LAN through a 62.5 μm or 50 μm multi-mode fiber optic cable terminated with an LC Duplex connector. The maximum supported unrepeated distance is 33 m on a 62.5 μm multi-mode fiber optic cable, and 300 m on a 50 μm multi-mode fiber optic cable.

OSA-Express3 GbE LX feature

The OSA-Express3 GbE LX occupies one slot in the I/O cage. It has four ports that connect to a 1 Gbps Ethernet LAN through a 9 μm single-mode fiber optic cable terminated with an LC Duplex connector, supporting an unrepeated maximum distance of 5 km (3.1 miles). A multimode (62.5 or 50 μm) fiber optic cable can be used with this feature. The use of these multimode cable types requires a mode conditioning patch (MCP) cable at each end of the fiber optic link. Use of the single-mode to multi-mode MCP cables reduces the supported distance of the link to a maximum of 550 meters (1084 feet).

OSA-Express3 GbE SX feature

OSA-Express3 GbE SX occupies one slot in the I/O cage. It has four ports that connect to a 1 Gbps Ethernet LAN through 50 or 62.5 μm multi-mode fiber optic cable terminated with an LC Duplex connector over an unrepeated distance of 550 meters (for 50 μm fiber) or 220 meters (for 62.5 μm fiber).

OSA-Express3 1000BASE-T Ethernet feature

OSA-Express3 1000BASE-T occupies one slot in the I/O cage. It has four ports that connect to a 1000 Mbps (1 Gbps), 100 Mbps, or 10 Mbps Ethernet LAN. Each port has an RJ-45 receptacle for UTP Cat5 cabling, which supports a maximum distance of 100 meters.

2.10.7 OSA-Express2

This section describes the connectivity options offered by the OSA-Express2 features. The following OSA-Express2 features can be installed on z10 EC servers:

- ▶ OSA-Express2 Gigabit Ethernet (GbE) long wavelength (LX)
- ▶ OSA-Express2 Gigabit Ethernet short wavelength (SX)
- ▶ OSA-Express2 1000BASE-T Ethernet

The OSA-Express2 Gigabit Ethernet 10 GbE LR feature is available only if carried forward by an upgrade.

OSA-Express features installed on previous servers are *not* supported on a z10 EC and *cannot* be carried forward on an upgrade.

OSA-Express2 10 GbE LR feature

The OSA-Express2 10 GbE LR feature occupies one slot in an I/O cage and has one port that connects to a 10 Gbps Ethernet LAN through a 9 μm single mode fiber optic cable terminated with an LC Duplex connector. The feature supports an un-repeated maximum distance of 10 km.

OSA-Express2 GbE LX feature

The OSA-Express2 GbE LX feature occupies one slot in an I/O cage and has two independent ports. Each port supports a connection to a 1 Gbps Ethernet LAN through a 9 μm single-mode fiber optic cable terminated with an LC Duplex connector. This feature utilizes a long wavelength laser as the optical transceiver.

A multimode (62.5 or 50 μm) fiber cable may be used with the OSA-Express2 GbE LX feature. The use of these multimode cable types requires a mode conditioning patch (MCP) cable to be used at each end of the fiber link. Use of the single-mode to multimode MCP cables reduces the supported optical distance of the link to a maximum end-to-end distance of 550 meters.

OSA-Express2 GbE SX feature

The OSA-Express2 GbE SX feature occupies one slot in an I/O cage and has two independent ports. Each port supports a connection to a 1 Gbps Ethernet LAN through a 62.5 μm or 50 μm multi-mode fiber optic cable terminated with an LC Duplex connector. The feature utilizes a short wavelength laser as the optical transceiver.

OSA-Express2 1000BASE-T Ethernet feature

The OSA-Express2 1000BASE-T Ethernet occupies one slot in the I/O cage. It has two ports connecting to either a 1000BASE-T (1000 Mbps), 100BASE-TX (100 Mbps), or 10BASE-T (10 Mbps) Ethernet LAN. Each port has an RJ-45 receptacle for UTP Cat5 cabling, which supports a maximum distance of 100 meters.

For details about all OSA-Express features see the *IBM System z Connectivity Handbook*, SG24-5444 or *OSA-Express Implementation Guide*, SG24-5948.

2.11 Cryptographic functions

The z10 EC server includes both standard cryptographic hardware and optional cryptographic features to provide flexibility and growth capability. IBM has a long history of providing hardware cryptographic solutions. Use of the cryptographic hardware function requires support by the operating system. For the z/OS operating system, the Integrated Cryptographic Service Facility (ICSF) is a base component that provides the administrative interface and a large set of application interfaces to the hardware.

Cryptographic support on the z10 EC includes:

- ▶ CP Assist for Cryptographic Function
- ▶ Crypto Express2 and Crypto Express3 cryptographic adapter features
- ▶ Trusted key entry workstation feature

2.11.1 CP Assist for Cryptographic Function

Figure 2-8 shows the layout of the z10 EC Compression and Cryptographic Accelerator (CCA). The chip contains two CCAs and each two cores on the chip share the encrypting and hashing engines of a CCA (each core has a dedicated compression engine). Every processor in the z10 EC server characterized as a CP or an IFL has access to the CP Assist for Cryptographic Function (CPACF).

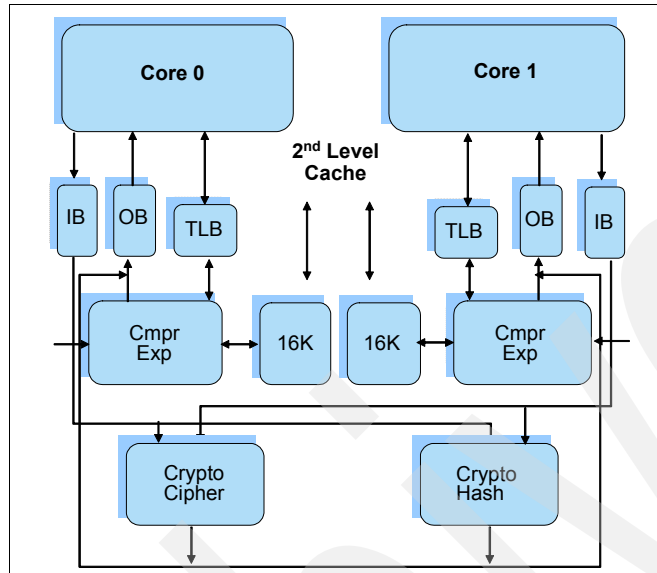


Figure 2-8 z10 EC Compression and Cryptographic Accelerator

The assist provides high-performance hardware encrypting and decrypting support for clear key operations and is designed to scale with PU performance enhancements. Special instructions are used with the cryptographic assist function.

CPACF offers a set of symmetric cryptographic functions for high encrypting and decrypting performance of clear key operations for SSL, VPN, and data storing applications that do not require FIPS 140-2 level 4 security. The cryptographic architecture includes support for:

- ▶ Data Encryption Standard (DES) data encrypting and decrypting
- ▶ Triple Data Encryption Standard (TDES) data encrypting and decrypting
- ▶ Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys
- ▶ Pseudo random number generation (PRNG)

Note that PRNG is also a standard function supported on the Crypto Express2 and Crypto Express3 features.

- ▶ Random number generation long (RNGL) of 8 bytes to 8192 bytes
- ▶ Random number generation (RNG) with up to 4096-bit key RSA support
- ▶ MAC message authorization both single key and double key
- ▶ Personal identification number (PIN) generation, verification, and translation functions
- ▶ Hashing algorithms: SHA-1 and SHA-2 support for SHA-224, SHA-256, SHA-384, and SHA-512

SHA-1 and SHA-2 support for SHA-224, SHA-256, SHA-384, and SHA-512 are shipped enabled on all servers and do not require the CPACF enablement feature. The CPACF functions are supported by z/OS, z/VM, and Linux on System z.

An enhancement to CPACF is designed to facilitate the continued privacy of cryptographic key material when used for data encryption. CPACF ensures that key material is not visible to applications or operating systems during encryption operations.

Protected key CPACF is designed to provide substantial throughput improvements for large-volume data encryption as well as low latency for encryption of small blocks of data. Furthermore, the information management tool, IBM Encryption Tool for IMS and DB2 Databases, improves performance for protected key applications.

2.11.2 Crypto Express2 feature

The Crypto Express2 feature provides cryptographic functions on the z10 EC server. The Crypto Express2 feature has two PCI-X adapters. Each of the PCI-X adapters can be configured as either a coprocessor or an accelerator:

- ▶ **Crypto Express2 Coprocessor:** for secure key encrypted transactions (default)
 - Designed to support security-rich cryptographic functions, use of secure encrypted key values, and user-defined extensions (UDX). Offers functions similar to the previous PCICC or PCIXCC cryptographic feature, including the secure key functions. This mode is intended for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification.
- ▶ **Crypto Express2 Accelerator:** for Secure Sockets Layer (SSL) acceleration
 - Designed to support clear key RSA operations
 - Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol

To support reliability, availability, and serviceability (RAS) requirements, the initial purchase must contain two Crypto Express2 features. Additional features may be added in increments of one.

Support for 14-digit, 15-digit, 16-digit, 17-digit, 18-digit, and 19-digit personal account numbers is provided for stronger protection of data. The Integrated Cryptographic Service Facility (ICSF), a z/OS component, has been enhanced to exploit this capability.

The configurable Crypto Express2 feature is supported by z/OS, z/VM, z/VSE, and Linux on System z.

2.11.3 Crypto Express3 feature

The Crypto Express3 feature (FC 0864) has two PCI Express cryptographic adapters. The Crypto Express3 feature (FC 0871) has one PCI Express cryptographic adapter. Each of the PCI Express cryptographic adapters can be configured as a cryptographic coprocessor or a cryptographic accelerator.

The Crypto Express3 feature is the newest state-of-the-art generation cryptographic feature. Like its predecessors, it is designed to complement the functions of CPACF. This new feature is tamper-sensing and tamper-responding. It provides dual processors operating in parallel supporting cryptographic operations with high reliability.

The Crypto Express3 feature contains all the functions of the Crypto Express2 and introduces a number of new functions, including:

- ▶ Dynamic power management designed to keep within the temperature limits of the feature and at the same time maximize RSA performance.
- ▶ Up to 32 LPARs in all logical channel subsystems have access to the feature.

- ▶ Improved RAS over previous crypto features due to dual processors and the service processor.
- ▶ Function update while installed using secure code load.
- ▶ When a PCI-E adapter is defined as a coprocessor lock-step checking by the dual processors enhances error detection and fault isolation.
- ▶ Dynamic addition and configuration of the Crypto Features to LPARs without an outage.
- ▶ Updated cryptographic algorithms used to load the LIC from the TKE workstation.
- ▶ Support for smart card applications using Europay, MasterCard, and VISA specifications.

The Crypto Express3 feature is designed to deliver throughput improvements for both symmetric and asymmetric operations.

2.11.4 TKE workstation

The trusted key entry (TKE) workstation offers security-rich local and remote key management, providing authorized persons with a method of operational and master key entry, identification, exchange, separation, and update.

The TKE workstation supports connectivity to an Ethernet local area network operating at 10, 100, or 1000 Mbps.

An optional smart card reader can be attached to the TKE 5.2 or later workstation to allow for the use of smart cards that contain an embedded microprocessor and associated memory for data storage. Access to and the use of confidential data on the smart cards is protected by a user-defined personal identification number. The latest version of the TKE LIC is 6.0, introducing enhancements and usability features.

2.12 Coupling and clustering

In the past, Parallel Sysplex support has been provided over several types of connection, ISC, ICB, and IC, each of which involves unique development effort for the support code and for the hardware (except IC).

Coupling connectivity on z10 EC in support of Parallel Sysplex environments can now use the new InfiniBand (PSIFB) connections for Parallel Sysplex. PSIFB supports longer distances between servers compared with ICB. Customers who use ISC-3 coupling links might be able to benefit from migrating to PSIFB by performing link consolidation.

InfiniBand technology will allow, over time, moving all of the Parallel Sysplex support to a single type of interface that provides high-speed interconnection at short distances (replacing ICB) and longer distance fiber optic interconnection (replacing ISC).

2.12.1 ISC-3

InterSystem Channel-3 (ISC-3) links provide the connectivity required for data sharing between the Coupling Facility and the System z servers directly attached to it. The ISC-3 feature is available in peer mode only and can be used to connect to other System z servers. ISC-3 supports a link data rate of 2 Gbps. STP message exchanges can flow over ISC-3.

2.12.2 ICB-4

An Integrated Cluster Bus-4 (ICB-4) connection consists of one link that attaches directly to an MBA's STI port in the system. The ICB-4 does not connect to a card in the I/O cage. ICB-4 supports a link data rate of 2 GBps.

2.12.3 Internal Coupling (IC)

The internal coupling channel emulates the coupling facility connection in Licensed Internal Code (LIC) between images within a single system. It operates at memory speed and no hardware is required.

2.12.4 Parallel Sysplex InfiniBand (PSIFB) coupling

PSIFB coupling links are high-speed links that are available on System z10 and System z9 servers. There are two types of Host Channel Adapter (HCA) fanouts used for PSIFB coupling links on the z10 EC:

- ▶ HCA2-O fanout supporting InfiniBand 12x Double-Data Rate (12x IB-DDR) and 12x InfiniBand Single-Data Rate (12x IB-SDR)
- ▶ HCA2-O Long Reach (LR) fanout supporting 1x IB-DDR and 1x IB-SDR

Also see "Coupling links" on page 55.

PSIFB coupling link using a HCA2-O fanout

A PSIFB coupling link using a HCA2-O fanout operates at 6 GBps if used between two z10 servers and at 3 GBps when connecting a z10 to a z9 server. The link speed is auto-negotiated to the highest common rate. The HCA2-O fanout uses a fiber optical cable that is connected to a System z10 or z9 server. The maximum supported distance is 150 m.

PSIFB coupling link using a HCA2-O LR fanout

A PSIFB LR coupling link using HCA2-O LR operates at 2.5 Gbps or 5 Gbps between two z10 servers. HCA2-O LR uses a fiber optical cable to connect two System z10 servers. The maximum unrepeated distance supported is 10 km. When using repeaters (DWDM) the maximum distance is up to 100 km.

Time source for STP traffic

PSIFB can be used to carry Server Time Protocol (STP) timekeeping information.

2.12.5 System-Managed CF Structure Duplexing

System-Managed Coupling Facility (CF) Structure Duplexing provides a general-purpose, hardware-assisted, easy-to-exploit mechanism for duplexing CF structure data. This provides a robust recovery mechanism for failures (such as loss of a single structure or CF or loss of connectivity to a single CF) through rapid failover to the other structure instance of the duplex pair.

Customers interested in deploying System-Managed CF Structure Duplexing should read the technical paper *System-Managed CF Structure Duplexing*, ZSW01975USEN, which you can access by selecting **Learn More** on the Parallel Sysplex Web site:

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

2.12.6 Coupling Facility Control Code (CFCC) level 16

CFCC level 16 is available for the System z10 server and contains the following improvements:

- ▶ System-Managed CF Structure Duplexing enhancements

Prior to CFCC level 16, Systems-Managed CF Structure Duplexing required two protocol exchanges to occur synchronously to CF processing of the duplexed structure request. CFCC level 16 allows one of these requests to be asynchronous to CF processing. This implies that the CF-to-CF exchange can occur without z/OS waiting for acknowledgement. This allows faster service time, with more benefits as the coupling facilities are further apart, such as in a multi-site Parallel Sysplex. Both coupling facilities must be at CFCC level 16 for these enhancements to occur.

- ▶ List Notification improvements

Today, when a list changes its state from empty to non-empty all its connectors are notified. The first connector notified reads the new message but subsequent readers will find nothing. CFCC level 16 approaches this differently to improve CPU utilization. It only notifies one connector in a round-robin fashion and if the shared queue (as in IMS Shared Queue and WebSphere MQ Shared Queue) is read within a fixed period of time, the other connectors do not have to be notified. If the list is not read again within the time limit, the other connectors are informed.

No significant CF structure sizing changes are expected when going from CFCC level 15 to CFCC level 16. However, we strongly recommend to use the CFSizer tool available at:

<http://www.ibm.com/systems/z/cfsizer/>

2.13 Time functions

Time functions are used to provide an accurate time-of-day value and to ensure that the time-of-day value is properly coordinated among all of the systems in a complex. This is critical for Parallel Sysplex operation.

2.13.1 External time reference (ETR)

Two external time reference cards are a standard feature of the z10 EC server. The ETR cards contain the ETR ports for Sysplex Timer® connection and provide a dual-path interface to the IBM Sysplex Timers, which may be used for timing synchronization between systems.

The two ETR cards are located in the processor cage of the z10 EC server.

2.13.2 Server Time Protocol (STP)

Server Time Protocol is a server-wide facility that is implemented in the Licensed Internal Code of System z. The STP presents a single view of time to PR/SM and provides the capability for multiple servers and CFs to maintain time synchronization with each other. A System z or CF may be enabled for STP by installing the STP feature.

The STP feature is intended to be the supported method for maintaining time synchronization between System z servers and CFs.

For additional information about STP, refer to *Server Time Protocol Planning Guide*, SG24-7280, and to *Server Time Protocol Implementation Guide*, SG24-7281.

2.13.3 Network Time Protocol (NTP) support

NTP support is available on the z10 EC server and has been added to the STP code on System z9. This implementation answers the need for a single time source across the heterogeneous platforms in the enterprise. With this implementation the System z10 and the System z9 servers support the use of NTP as time sources.

2.14 HMC and SE

The Hardware Management Console (HMC) is used to manage, monitor, and operate one or more IBM System z servers and their associated logical partitions. The HMC is attached to a LAN, as is the server's support element (SE). The HMC communicates with each Central Processor Complex (CPC) through the CPC's SE. When tasks are performed on the Hardware Management Console, the commands are sent to one or more support elements, which then issue commands to their CPCs. The HMC and SE Version 2.10.2 support the System z10 servers.

In the last several years, the HMC has been enhanced to support many new functions and tasks to extend the management capabilities of the platform. This is true with the z10 servers and will continue in the future. Many of the more recent capabilities benefit from the various network connectivity options available, such as:

- ▶ HMC/SE LAN only
- ▶ HMC to a corporate intranet
- ▶ HMC to intranet and Internet

The HMC consists of:

- ▶ Processor or system unit, including two Ethernet LAN adapters, capable of operating at 10, 100, or 1000 Mbps, and a DVD RAM to install LIC
- ▶ Flat panel display
- ▶ Keyboard
- ▶ Mouse

The System z10 is supplied with a pair of integrated ThinkPad SEs. One is always active while the other is strictly an alternate. Power for the SEs is supplied by the server power supply, and there are no additional power requirements. Unlike previous servers, the internal LAN for the SEs on the z10 server connect to the bulk power hub in the Z frame. There is an additional connection from the hub to the HMC utilizing the VLAN capability of the server.

The HMC and SE Version 2.10.2 for System z10 servers introduces these enhancements:

- ▶ Digitally signed firmware

A critical issue with firmware upgrades is security and data integrity. Procedures are in place to use a process to digitally sign the firmware update files on the HMC, the SE, and the TKE. Using a hash-algorithm, a message digest is generated that is then encrypted with a private key to produce a digital signature. This operation ensures that any changes made to the data will be detected during the upgrade process. It helps ensure that no malware can be installed on System z products during firmware updates. It enables, with other existing security functions, System z10 CPACF functions to comply with Federal

Information Processing Standard (FIPS) 140-2 Level 1 for Cryptographic LIC changes. The enhancement follows the System z focus of security for the HMC and the SE.

- ▶ Optional user password on disruptive confirmation

The requirement to supply a user password on disruptive confirmation is optional. The general recommendation remains to require a password.

- ▶ Improved consistency of confirmation panels on the HMC and the SE

Attention indicators are on the top of panels and there will be a list of objects affected by the action, target as well as secondary objects, for example, LPARs if the target is CPC.

- ▶ Serviceability enhancements for FICON channels

Simplified problem determination to more quickly detect fiber optic cabling problems in a Storage Area Network.

All FICON channel error information is forwarded to the HMC, thus facilitating detection and reporting trends and threshold for the channels including aggregate views including data from multiple systems.

2.15 Power and cooling

The z10 EC server footprint is slightly bigger than the z9 EC footprint. The width of the frames is identical but the depth of the z10 EC server frames is 71.0 in (1803 MM), as compared with the 62.1 in (1577 mm) of the z9 EC frame.

The power service specifications are the same, but the power consumed by the z10 EC server can be greater. A fully loaded z10 EC server maximum consumption is 31.7 KW. Refer to Table 2-5 for the electrical service requirements for the different configurations.

Table 2-5 Electrical service requirements

Model	1 I/O cage	2 I/O cages	3 I/O cages
E12	2x60A	2x60A	2x60A
E26	2x60A	4x60A	4x60A
E40	4x60A	4x60A	4x60A
E56	4x60A	4x60A	4x60A
E64	4x60A	4x60A	4x60A

2.15.1 Hybrid cooling system

The z10 EC server has a hybrid cooling system that is designed to lower power consumption. It is an air-cooled system, assisted by refrigeration. Refrigeration is provided by a closed-loop liquid cooling subsystem. The entire cooling subsystem has a modular construction. Its components and functions are found throughout the cages.

Refrigeration cooling is the primary cooling source and is backed up by an air-cooling system. If one of the refrigeration units fails, backup blowers are switched on to compensate for the lost refrigeration capacity with additional air cooling. At the same time, the oscillator card is set to a slower cycle time, slowing the system down by up to 10% of its maximum capacity to allow the degraded cooling capacity to maintain the proper temperature range. Running at a slower cycle time, the MCMs produce less heat. The slowdown process is done in steps, based on the temperature in the books.

2.15.2 Internal Battery Feature

The Internal Battery Feature (IBF) is an optional feature on the z10 EC server. Refer to Figure 2-2 on page 21 for a pictorial view of the location of this feature. This optional IBF provides the function of a local uninterrupted power source.

The IBF further enhances the robustness of the power design, increasing power line disturbance immunity. It provides battery power to preserve processor data in case of a loss of power on all four AC feeds from the utility company. The IBF can hold power briefly during a *brownout*, or for orderly shutdown in case of a longer outage. The IBF provides up to 10 minutes of full power, depending on the I/O configuration.

2.15.3 IBM Systems Director Active Energy Manager

IBM Systems Director Active Energy Manager™ (AEM) is an energy management solution building block that returns true control of energy costs to the customer. It enables you to manage actual power consumption and resulting thermal loads that IBM servers place on the data center. It is an industry-leading cornerstone of the IBM energy management framework. In tandem with chip vendors Intel® and AMD and consortiums such as Green Grid, AEM advances the IBM initiative to deliver price performance per square foot.

AEM runs on Windows®, Linux on System x®, Linux on System p®, and Linux on System z. Refer to its documentation for more specific information.

How AEM works

The following list is a brief overview of how AEM works:

- ▶ Hardware, firmware, and systems management software in servers and blades can take inventory of components.
- ▶ AEM adds power draw up for each server or blade and tracks that usage over time.
- ▶ When power is constrained, AEM allows power to be allocated on a server-by-server basis. Note the following information:
 - Care should be taken that limiting power consumption does not affect performance.
 - Sensors and alerts can warn the user if limiting power to this server could affect performance.
 - The System z10 EC server does not support power capping.
- ▶ Certain data can be gathered from the System z10 HMC:
 - System name, machine type, model, serial number, firmware level
 - Ambient and exhaust temperature
 - Average and peak power (over a 1-minute period)
 - Other limited status and configuration information

Key functions and capabilities

The System z10 Enterprise Class is the follow-on to the System z9 EC (z9 EC) server. As its predecessor, it also offers five hardware models, but has a more powerful uniprocessor, more processor units, and new functions and features.

The z10 EC represents a new level of microprocessor technology, made possible through advances in the design and manufacturing processes.

Virtualization has been improved, additional security options are provided, the I/O subsystem and several I/O features have been enhanced, and new I/O features are offered.

Based on customer requirements and changes in the market demands, the z10 EC offers some important enhancements and additional flexibility to the Capacity on Demand functions.

The z10 EC improves the reliability, availability, and serviceability of the server compared with previous System z servers by introducing solutions to minimize not only unplanned outages, but also to decrease the need for planned outages.

This chapter builds upon the information presented in Chapter 1, “Introducing the System z10 Enterprise Class” on page 1, and Chapter 2, “Hardware overview” on page 17. It discusses the following topics:

- ▶ 3.1, “Virtualization” on page 46
- ▶ 3.2, “Technology improvements” on page 49
- ▶ 3.3, “Common time functions” on page 62
- ▶ 3.4, “Capacity on Demand (CoD) enhancements” on page 64
- ▶ 3.5, “Throughput optimization enhancements” on page 68
- ▶ 3.6, “Reliability, availability, and serviceability improvements” on page 69
- ▶ 3.7, “Parallel Sysplex technology” on page 70
- ▶ 3.8, “Summary” on page 72

3.1 Virtualization

Virtualization is a key strength of the z10 EC server. Virtualization is embedded in the z/Architecture, which includes a precise and model-independent definition of the hardware-to-software interface. It is deeply built into the server implementation, which supports virtualization through both hardware and software.

Virtualization creates the appearance of multiple concurrent servers by sharing the existing hardware. Its major goal is to fully utilize the server resources, thus lowering the total amount of resources needed and their cost. Virtualization can be seen as an application with very demanding performance and security requirements. The z10 EC is able to handle tens, hundreds, even thousands, of virtual servers, so a very high context switching rate is to be expected, and accesses to the memory, caches, and virtual I/O devices must be kept completely isolated. The z/Architecture, the z10 EC, the z10 BC, and their predecessors have been *designed* to meet those requirements with very low overhead and the highest security certification in the industry: common criteria EAL5 with specific target of evaluation (logical partitions). This design has been proved in many customer installations in recent decades.

Virtualization requires a hypervisor. A hypervisor is control code that manages multiple independent operating system images. Hypervisors can be implemented in software or hardware, and System z has both. In System z the hardware hypervisor is implemented in firmware and is called Processor Resource/Systems Manager™ (PR/SM). PR/SM is part of the base server and does not require any software to run. The z/VM operating system implements the software hypervisor. z/VM requires some PR/SM functions.

Virtualization provides totally secured environments for the virtualized servers and IBM publishes a System Integrity Statement for both z/OS and z/VM:

- ▶ For z/OS:

http://www-03.ibm.com/servers/eserver/zseries/zos/racf/zos_integrity_statement.html

- ▶ For z/VM:

<http://www.vm.ibm.com/security/zvminteg.html>

3.1.1 Hardware virtualization

PR/SM was first implemented in the mainframe in the late 1980s. It manages subsets of the server resources known as logical partitions (LPARs) and virtualizes processors, memory, and I/O features. Some features are purely virtual implementations. HiperSockets works like a LAN but does not use any I/O hardware.

Up to 60 LPARs can be defined. In each, a supported operating system can be run. The LPAR definition includes a number of *logical* PUs, memory, and I/O devices.

z/VM-mode partitions

System z10 has a partition mode called z/VM-mode. The processor types that can be configured to a z/VM-mode partition are:

- ▶ CPs
- ▶ IFLs
- ▶ zIIPs
- ▶ zAAPs
- ▶ ICFs

z/VM V5R4 and later versions support this mode that provides increased flexibility and simplifies systems management by allowing z/VM to manage guests to perform the following tasks all in the same VM LPAR:

- ▶ Operate Linux on System z on IFLs.
- ▶ Operate z/VSE and z/OS on CPs.
- ▶ Offload z/OS system software overhead, such as DB2 workloads, on zIIPs.
- ▶ Provide an economical Java execution environment under z/OS on zAAPs.

This support fulfills an IBM Statement of General Direction issued at the time of the original z10 EC announcement.

Logical processors

Logical processors are perceived by the operating systems as real ones. They assume the following types:

- ▶ CPs
- ▶ zAAPs
- ▶ zIIPs
- ▶ IFLs
- ▶ ICFs

SAPs are never part of an LPAR configuration.

PR/SM is responsible for honoring requests for logical processor work by dispatching logical processors on physical processors of the same type. Under certain circumstances logical zAAPs and zIIPs can be dispatched on physical CPs. Physical processors can be shared across LPARs, but can also be dedicated to an LPAR. However, an LPAR must have its logical processors either all shared or all dedicated.

PR/SM ensures that, when switching a physical processor from one logical processor to another, processor state is properly saved and restored, including all the registers. Data isolation, integrity, and coherence are strictly enforced at all times.

Logical processors can be dynamically added to and removed from LPARs. Operating system support is required in order to take advantage of this capability. z/OS requires that reserved processors be defined to the LPAR for the addition to take place. z/VM is able to dynamically recognize and add additional processors.

Memory

To ensure security and data integrity, memory cannot be shared by active LPARs. In fact, a strict isolation is maintained. When an LPAR is activated, its defined memory is allocated in blocks, which must be a multiple of a given value. This value depends on the total allocation and varies between 256 MB and 2 GB. Thus, memory can be serially reused.

Using the plan-ahead capability, memory can be physically installed but not enabled until it is necessary. z/OS and z/VM support dynamically increasing the size of the LPAR.

LPAR memory is said to be virtualized in the sense that in all LPARs memory addresses start at zero. This should not be confused with the operating system virtualizing its LPAR memory. The z/Architecture has a robust virtual storage architecture that allows, per LPAR, the definition of an unlimited number of address spaces and the simultaneous use *by each program* of up to 1,023 of those address spaces. Each address space can be up to 16 EB (1 exabyte = 2^{60} bytes). Thus, the architecture has no real limits. Practical limits are determined by the available hardware resources, including disk storage for paging.

Isolation of the address spaces is strictly enforced by the Dynamic Address Translation hardware mechanism, which also validates the right to read or write in each page frame by comparing the page key with the key of the program requesting access. Three addressing modes, 24-bit, 31-bit, and 64-bit, are simultaneously supported. Definition and management of the address spaces is under operating system control. This mechanism has been in use since the System 370, and memory keys were part of the original System 360 design.

Operating systems may allow sharing of address spaces, or parts thereof, across multiple processes. For instance, under z/VM, a single copy of the read-only part of a kernel can be shared by all virtual machines using that operating system, resulting in large savings of real memory and improvements of performance.

I/O virtualization

The z10 EC supports four channel subsystems with 256 channels each, for a total of 1024 channels. In addition to dedicated use of channels and I/O devices to an LPAR, I/O virtualization allows concurrent sharing of channels, and the I/O devices accessed through these channels, by several active LPARs. The function is known as Multiple Image Facility (MIF). The sharing channels may belong to different channel subsystems, in which case they are known as spanned channels.

Data streams for the sharing LPARs are carried on the same physical path with total isolation and integrity. For each active LPAR that has the channel configured online, PR/SM establishes one logical channel path. For availability reasons, multiple logical channel paths should exist for critical devices (for instance, disks containing vital data sets).

When isolation is required, configuration rules allow restricting the access of each logical partition to particular channel paths and specific I/O devices on those channel paths.

Many installations use the Parallel Access Volume (PAV) function, which allows accessing a device by several different addresses (normally one base address and three aliases), thus increasing the throughput of the device but using more device addresses. HyperPAV takes the technology a step further by allowing the I/O Supervisor (IOS) in z/OS to dynamically create PAV structures depending on the current I/O demand in the system, thus lowering the need for manually tuning the system for PAV use.

For large installations, which usually have a large number of devices, the total number of device addresses can be very high. Thus, the concept of *channel sets* was introduced with System z9. Each channel can address two sets of 64 K device addresses, allowing the base addresses to be defined on *set 0* (IBM reserves 256 subchannels on set 0) and the aliases on *set 1*. In total, 130,816 subchannel addresses are available *per channel*.

Channel sets are exploited by the Peer-to-Peer Remote Copy (PPRC) function by the ability to have the PPRC primary devices defined in channel set 0, while secondary devices can be defined in channel set 1, thus providing more connectivity through channel set 0.

To further reduce the complexity of managing large I/O configurations System z introduces Extended Address Volumes (EAV). EAV is designed to build very large disk volumes using virtualization technology. By being able to extend the disk volume size a customer may potentially need less volumes to hold his data, therefore making systems and data management less complex.

The health checker function in z/OS V1.10 introduces a health check in the I/O Supervisor that can help system administrators identify single points of failure in the I/O configuration.

The dynamic I/O configuration function is supported by z/OS and z/VM. It provides the capability of concurrently changing the currently active I/O configuration. Changes can be

made to channel paths, control units, and devices. The existence of a fixed HSA area in the z10 EC greatly eases the planning requirements and enhances the flexibility and availability of these reconfigurations.

3.1.2 Software virtualization

Software virtualization is provided by the z/VM product. Strictly, it is a function of the CP component of z/VM. Starting in 1967, IBM has continuously provided software virtualization in its mainframe servers.

z/VM uses the resources of the LPAR in which it is running to create functional equivalents of real System z servers, which are known as virtual machines (VMs) or *guests*. In addition, z/VM is able to emulate I/O peripherals including, for instance, printers by using spooling techniques and LAN switches and disks by exploiting memory.

z/VM allows very fine-grained allocation of resources, for example, in the case of processor sharing, the minimum is approximately 1/10,000 of a processor. Another example: Disks can be subdivided into independent areas, known as *minidisks*, each of which is exploited by its users as a real disk, only smaller. Minidisks are shareable, and can be used for all types of data and also for temporary space in a pool of on demand storage.

Under z/VM, virtual processors, virtual central and expanded storages, and all the virtual I/O devices of the VMs are dynamically definable (provisionable). z/VM supports the concurrent addition (but not deletion) of memory to its LPAR and immediately makes it available to guests. Guests themselves may support the dynamic addition of memory. All other changes are concurrent. To render these concurrent definitions also nondisruptive requires support by the operating system running in the VM, which is also the case when running in an LPAR.

Although z/VM imposes no limits on the number of defined virtual machines, the number of active virtual machines is limited by the available resources. On a large server, such as the z10 EC, thousands of virtual machines can be activated.

It is beyond the scope of this book to provide a more detailed description of z/VM or other highlights of its capabilities. For a deeper discussion of z/VM see *Introduction to the New Mainframe: z/VM Basics*, SG24-7316, downloadable from:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247316.pdf>

3.2 Technology improvements

The technology improvements for the z10 EC fall into five categories:

- ▶ Microprocessor
- ▶ Capacity
- ▶ Memory
- ▶ Connectivity
- ▶ Cryptography

These are intended to provide a more scalable, flexible, manageable, and secure consolidation and integration platform contributing to a lower total cost of ownership.

3.2.1 Microprocessor enhancements

The System z10 Enterprise Class has a newly developed microprocessor chip and a newly developed infrastructure chip. Both of those chips use CMOS 11 technology and represent a major step forward in technology utilization for the System z products, resulting in increased packaging density.

Like for the z9, the microprocessor chip and the infrastructure chip for the z10 EC are packaged together on a new multi-chip module (MCM). The MCM contains five microprocessor chips and two infrastructure chips, while the z9 MCMs included 16 chips in total. Each microprocessor chip contains four cores. The MCM is installed inside a book, and the z10 EC can contain from one to four books. The book also contains the memory arrays, I/O connectivity infrastructure, and various other mechanical and power controls.

The book is connected to the I/O cages through one or more cables. As new standards are making their way on to the z10 EC, these cables are now using the standard InfiniBand protocol to transfer large volumes of data between the memory and the I/O cards located in I/O cages.

z10 EC processor chip

The z10 EC chip provides more functions per chip—four cores on a single chip—thanks to technology improvements that allow designing and manufacturing more transistors per square inch. This translates into using fewer chips to implement the needed functions, which helps enhance system availability.

Both chips were developed in close cooperation with the System p development organization that has designed the new Power6 chip. It could be said that the new z10 EC microprocessor chip and the Power6 chip share a lot of DNA—they are siblings, but not identical twins. They have common characteristics but differ in many ways and will continue to do so.

The z10 microprocessor chip has a significant new design when compared with the z9. The System z microprocessor development has been following the same basic design set since the 9672-G4 (announced in 1997) until the z9. That basic design has now been stretched to its maximum, so a fundamental change was necessary.

The processor chip, shown in Figure 3-1, also includes two co-processors for hardware acceleration of data compression and cryptography, I/O bus and memory controllers, and an interface to a separate storage controller/cache chip.

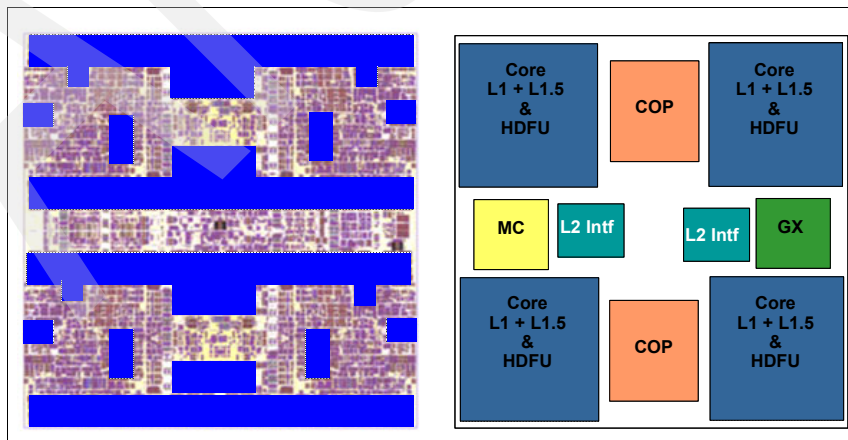


Figure 3-1 z10 EC Enterprise Quad-Core microprocessor chip

On-chip cryptographic hardware includes extended key and hash sizes for the AES and SHA algorithms.

Hardware decimal floating point function

The z10 EC microprocessor implements a hardware decimal floating point function, designed to speed up such calculations and provide the necessary precision demanded mainly by the financial institutions sector. The decimal floating point hardware fully implements the new IEEE 754r standard.

New machine instructions

The z/Architecture offers a rich CISC Instruction Set Architecture (ISA). The z10 EC offers 894 instructions, of which 668 are implemented entirely in hardware. Multiple arithmetic formats are supported.

The z10 EC architectural extensions include over 50 new instructions, the bulk of which was designed in collaboration with software developers to improve compiled code efficiency. These should particularly be of benefit to Java-based, WebSphere-based, and Linux-based workloads. New instructions are grouped under the following categories:

- ▶ Storage immediate operations
- ▶ Combined comparison and conditional branch, based on the comparison result
- ▶ Combined comparison and trap to exception handlers, based on the comparison result
- ▶ Operations on storage operands defined relative to the current instruction address
- ▶ Combined rotate and logical bit operations
- ▶ Extensions to existing instructions
- ▶ Instructions for enablement of software/hardware cache optimization
- ▶ Support of large page frames

3.2.2 Granular capacity and capacity settings

The z10 EC expands the offer on sub-capacity settings. Finer granularity in capacity levels allows the growth of installed capacity to more closely follow the enterprise growth, for a smoother, pay-as-you-go investment profile. The many performance and monitoring tools available on System z environments, coupled with the flexibility of the capacity on demand options (see 3.4, “Capacity on Demand (CoD) enhancements” on page 64) provide for managed growth with capacity being available when needed.

The z10 EC offers four distinct capacity levels for CPs (full capacity and three sub-capacities). A processor characterized as anything other than a CP is always set at full capacity. There is, correspondingly, a different pricing model for non-CP processors regarding purchase and maintenance prices, as well as various offerings for software licensing costs.

A capacity level is a setting of each CP to a sub-capacity of the full CP capacity. Full capacity CPs are identified as CP7. On the z10 EC server, 64 CPs can be configured as CP7. Besides full capacity CPs, three sub-capacity levels (CP6, CP5, and CP4), each for up to twelve CPs, are offered. The four capacity levels appear in hardware descriptions as feature codes on the CPs. These feature codes (FC) are:

- ▶ CP7 is FC 6810
- ▶ CP6 is FC 6809
- ▶ CP5 is FC 6808
- ▶ CP4 is FC 6807

Granular capacity adds 36 sub-capacity settings to the 64 capacity settings that are available with full capacity CPs (CP7). Each of the 36 sub-capacity settings only apply to up to twelve CPs, independent of the z10 EC model installed. Figure 3-2 shows the relative capacity of subcapacity models.




Figure 3-2 z10 EC granular capacity for up to 12 CPs

If more than twelve CPs are configured for the server they will all be full capacity, because all CPs must be on the same capacity level. The capacity indicator numbers are:

- ▶ 701 to 764 for capacity level CP7
- ▶ 601 to 612 for capacity level CP6
- ▶ 501 to 512 for capacity level CP5
- ▶ 401 to 412 for capacity level CP4

Information about CPs in the remainder of this chapter applies to all CP capacity levels, CP7, CP6, CP5, and CP4, unless otherwise indicated.

Note: The actual throughput that a user will experience may vary, depending on considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, and the workload being processed.

To help size a System z server to fit your requirements, IBM provides a free tool that reflects the latest IBM LSPR measurements, called the IBM Processor Capacity Reference (zPCR). The tool can be downloaded from:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381>

3.2.3 Memory enhancements

The z10 EC has greatly increased the available memory capacity over previous servers. The system can now have up to 1,520 GB of usable memory installed. The logical partitions can now be configured with 1 TB of memory, and z/OS will support the new memory size starting with Version 1 Release 8. In fact, z/OS V1R8 and later support up to 4 TB of main storage. For the first time, the hardware systems area (HSA) is fixed in size (16 GB) and not included in the memory for which the customer orders and pays.

Note that the z/Architecture simultaneously supports 24-bit, 31-bit, and 64-bit addressing modes. This provides backwards compatibility and investment protection.

Hardware system area

The z10 EC has a fixed-size hardware system area. This is intended to improve the server availability. Because the HSA is big enough to accommodate all possible configurations for all the logical partitions, several operations that were disruptive on previous servers due to HSA size are now concurrent. In addition, some planning needs are eliminated.

The HSA has a fixed size of 16 GB and resides in a separately reserved area of memory separate from customer-purchased memory.

A fixed large HSA enables dynamic addition and removal of the following features without planning:

- ▶ New logical partition to new or existing channel subsystem (CSS)
- ▶ New CSS (up to four can be defined)
- ▶ New subchannel set (up to two can be defined)
- ▶ Maximum number of devices in each subchannel set
- ▶ Dynamic I/O enabled as a default
- ▶ Logical processors by type
- ▶ Cryptographic processors

Plan-ahead memory

Planning for future memory requirements and installing dormant memory in the server allows future upgrades to be done concurrently and, with appropriate operating system support, nondisruptively.

If a customer can anticipate an increase of the required memory, a target memory size can be configured along with a starting memory size. The starting memory size will be activated and the remainder will be inactive. When additional physical memory is required, it is fulfilled by activating the appropriate number of planned memory features. This activation is concurrent and can be nondisruptive to the applications depending on the operating system support. Both z/OS and z/VM support this function.

Plan-ahead memory should not be confused with flexible memory support. Plan-ahead memory is for a permanent increase of installed memory, whereas flexible memory provides a temporary replacement of a part of memory that becomes unavailable.

Flexible memory

Flexible memory was first introduced on the z9 EC as part of the design changes and offerings to support enhanced book availability. Flexible memory is used to temporarily replace the memory that becomes unavailable when performing maintenance on a book. On z10 EC, the additional resources required for the flexible memory configurations are provided through the purchase of planned memory features along with the purchase of memory entitlement. Flexible memory configurations are available on multi-book models E26, E40, E56, and E64 and range from 32 GB to 1136 GB, depending on the model.

Contact an IBM representative to help determine the appropriate configuration.

Large page support

The size of pages and page frames has been 4 KB for a long time. The z10 EC server introduces the capability of, in addition to pages of 4 KB, having large pages with the size of 1 MB. This is a performance item addressing particular workloads and relates to large main storage usage. Large page support is exclusive to System z10 servers. Both page frame sizes can be simultaneously used.

Large pages cause the translation lookaside buffer (TLB) to better represent the working set and suffer fewer misses by allowing a single TLB entry to cover more address translations. Exploiters of large pages are better represented in the TLB and are expected to perform better.

This support is primarily of benefit for long-running applications that are memory access intensive. Large pages are not recommended for general use. Short-lived processes with small working sets are normally not good candidates for large pages and would see little to no improvement. The use of large pages must be decided based on knowledge obtained from measurement of memory usage and page translation overhead for a specific workload.

Large pages are treated as fixed pages and are never paged out. They are only available for 64-bit virtual private storage such as virtual memory located above 2 GB.

3.2.4 Connectivity enhancements

The I/O cages on the z10 EC are similar to the ones used in the z9, and the I/O cards are the same as those used in the z9, with some exceptions. See Table B-1 on page 109 for more details. Five new OSA-Express3 cards and other connectivity changes are also introduced on the z10 EC. The new cards are listed in 2.10.6, “OSA-Express3” on page 34.

The physical connection between the processor and memory and the I/O cages uses a new cable to support the InfiniBand technology. Prior to and including the z9, STI cables were specifically developed for this function. With new InfiniBand cables the bandwidth per cable increases from 2.7 GB per second to 6 GB per second.

Standard InfiniBand cables and protocol can be used for Parallel Sysplex coupling links for servers that are up to 150 meter apart. InfiniBand offers a longer distance than the existing ICB cable and increased bandwidth, similar to the bandwidth obtained with the cable used internally in the server. The System z9 servers can be upgraded to use the new coupling link and can participate in a Parallel Sysplex using this new technology.

Advantages of InfiniBand

InfiniBand addresses the challenges that IT infrastructures face as more demands are placed on the interconnect with ever-increasing requirements for computing and storage resources. InfiniBand has a number of advantages, such as:

- ▶ **Superior performance:** InfiniBand has a defined road map to 120 Gbps—the fastest support specification of any industry-standard interconnect.
- ▶ **Reduced complexity:** InfiniBand allows for the consolidation of multiple I/Os on a single cable or backplane interconnect. InfiniBand also consolidates the transmission of clustering, communications, storage, and management data types over a single connection.
- ▶ **Highest interconnect efficiency:** InfiniBand was developed to provide efficient scalability of multiple systems. InfiniBand provides communication-processing functions in hardware—relieving the CPU of this task—and enables the full resource utilization of each node added to the cluster. In addition, InfiniBand incorporates Remote Direct Memory Access, which is an optimized data transfer protocol that further enables the server processor to focus on application processing.
- ▶ **Reliable and stable connections:** InfiniBand provides reliable end-to-end data connections and defines this capability to be implemented in hardware. In addition, InfiniBand facilitates the deployment of virtualization solutions, which allow multiple applications to run on the same interconnect with dedicated application partitions.

The Server Time Protocol (STP) can also benefit from this coupling technology. STP timing signals can be transported over PSIFB coupling links.

Coupling links

The five coupling link options for communication in a Parallel Sysplex environment are:

- ▶ Internal Coupling links (ICs), which are used for internal communication between Coupling Facilities (CFs) defined in LPARs and z/OS images on the same server.
- ▶ Integrated Cluster Bus-4 (ICB-4), which supports a link data rate of 2 gigabytes per second (GBps) and is used for z/OS-to-CF communication over short distances, using 10-meter (33-foot) copper cables, of which 3 meters (10 feet) are used for internal routing and strain relief. The ICB-4 is used to connect a z10 EC to other System z servers.
- ▶ InterSystem Channel-3 (ISC-3), which supports a link data rate of 2 Gbps and is used for z/OS-to-CF communication at unrepeated distances up to 10 km (6.2 miles) using 9 μ m single mode fiber optic cables and repeated distances up to 100 km (62 miles) using System z-qualified DWDM equipment. ISC-3s are supported exclusively in peer mode.
- ▶ InfiniBand (HCA2-O) coupling links (12x IB-SDR or 12x IB-DDR) are used for z/OS-to-CF communication at distances up to 150 meters (492 feet) using industry standard OM3 50 μ m fiber optic cables.
 - 2x InfiniBand coupling links support single data rate (SDR) at 3 GBps when System z10 is connected to System z9.
 - 2x InfiniBand coupling links support double data rate (DDR) at 6 GBps for a System z10-to-System z10 connection.
- ▶ InfiniBand (HCA2-O LR) coupling links (1x IB-SDR or 1x IB-DDR) for z/OS-to-CF communication at unrepeated distances up to 10 km (6.2 miles) using 9 μ m single mode fiber optic cables and repeated distances up to 100 km (62 miles) using System z-qualified DWDM equipment.

Note: The InfiniBand coupling link data rate (6 GBps, 3 GBps, 5 Gbps, or 2.5 Gbps) does *not* represent the performance of the link. The actual performance depends on many factors, including latency through the adapters, cable lengths, and the type of workload.

When comparing coupling links data rates, InfiniBand (12x IB-SDR or 12x IB-DDR) might be higher than ICB-4 and InfiniBand (1x IB-SDR or 1x IB-DDR) might be higher than that of ISC-3, but with InfiniBand the service times of coupling operations are greater and the actual throughput might be less than with ICB-4 links or ISC-3 links.

Refer to the *Coupling Facility Configuration Options* white paper for a more specific explanation regarding the use of the current ICB-4 or ISC-3 technology versus migrating to InfiniBand coupling links. The white paper is available at:

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

FICON enhancements

The z10 EC server offers several exclusive enhancements to FICON Express8, FICON Express4, and FICON Express2 features.

High performance FICON for System z (zHPF)

High performance FICON for System z10 brings improvement in performance and reliability, availability, and serviceability (RAS). Several enhancements have been made to the z/Architecture and the FICON interface architecture in order to provide optimizations for

online transaction processing (OLTP) workloads. When exploited by the FICON channel, the z/OS operating system, and the control unit, FICON channel and control unit overhead can be reduced and performance improved. Additionally, the changes to the architectures provide end-to-end system enhancements to improve RAS.

zHPF channel programs can be exploited by z/OS OLTP I/O workloads, such as DB2, VSAM, PDSE, and zFS, which transfer small blocks of fixed size data (4 K blocks). zHPF exploitation requires implementation by the DS8000® Disk Storage systems.

The FICON Express8, FICON Express4, and FICON Express2 features (CHPID type FC) will support both the existing FICON protocol and the zHPF protocol concurrently in the server Licensed Internal Code. Support is exclusive to System z10.

For more information about FICON channel performance, see the technical papers on the System z I/O connectivity Web site at:

http://www-03.ibm.com/systems/z/hardware/connectivity/ficon_performance.html

Extended distance FICON

Exploitation of an enhancement to the industry standard FICON architecture (FC-SB-3) can help avoid degradation of performance at extended distances by implementing a new protocol for *persistent* information unit (IU) pacing. Control units that exploit the enhancement to the architecture can increase the pacing count (the number of IUs allowed to be in flight from channel to control unit). Extended distance FICON also allows the channel to remember the last pacing update for use on subsequent operations to help avoid degradation of performance at the start of each new operation.

Improved IU pacing can help optimize the utilization of the link (for example, help keep a 4 Gbps link fully utilized at 50 km) and allows channel extenders to work at any distance, with performance results similar to those experienced when using emulation.

The requirements for channel extension equipment are simplified with the increased number of commands in flight. This may benefit z/OS Global Mirror (Extended Remote Copy, XRC) applications, as the channel extension kit is no longer required to simulate specific channel commands. Simplifying the channel extension requirements may help reduce the total cost of ownership of end-to-end solutions.

Extended Distance FICON is transparent to operating systems and applies to all the FICON Express2, FICON Express4, and FICON Express8 features carrying native FICON traffic (CHPID type FC). For exploitation, the control unit must support the new IU pacing protocol.

Exploitation of extended distance FICON is supported by the IBM System Storage DS8000 series with an appropriate level of Licensed Machine Code (LMC).

FICON name server registration

The FICON channel now provides the same information to the fabric as is commonly provided by open systems, registering with the name server in the attached FICON directors. This enables the quick and efficient management of storage area network (SAN) and performing problem determination and analysis.

Platform registration is a standard service defined in the Fibre Channel - Generic Services 3 (FC-GS-3) standard (INCITS (ANSI) T11.3 group). It allows a platform (storage subsystem, host, and so on) to register information about itself with the fabric (directors).

This z10 exclusive function is transparent to operating systems and applicable to all FICON Express8, FICON Express4, FICON Express2, and FICON Express features (CHPID type FC).

FCP enhancements for small block sizes

The Fibre Channel Protocol (FCP) Licensed Internal Code has been modified to help provide increased I/O operations per second for small block sizes.

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

This FCP performance improvement is transparent to operating systems and applies to all the FICON Express8, FICON Express4, and FICON Express2 features when configured as CHPID type FCP, communicating with SCSI devices.

For more information about FCP channel performance, see the performance technical papers on the System z I/O connectivity Web site at:

http://www-03.ibm.com/systems/z/hardware/connectivity/fcp_performance.html

SCSI IPL base function

The SCSI Initial Program Load (IPL) enablement feature, first introduced on z990 in October of 2003, is no longer required. The function is now delivered as a part of the server Licensed Internal Code. SCSI IPL allows an IPL of an operating system from an FCP-attached SCSI disk.

N_Port ID Virtualization (NPIV)

NPIV is designed to allow the sharing of a single physical FCP channel among operating system images, whether in logical partitions or as z/VM guests in virtual machines. This is achieved by assigning a unique World Wide Port Name (WWPN) for each operating system connected to the FCP channel. In turn, each operating system appears to have its own distinct WWPN in the SAN environment, hence enabling separation of the associated FCP traffic on the channel.

Access controls based on the assigned WWPN can be applied in the SAN environment, using standard mechanisms such as zoning in SAN switches and logical unit number (LUN) masking in the storage controllers.

Worldwide portname prediction tool

A part of the installation of your IBM System z10 server is the planning of the SAN environment. IBM has made available a standalone tool to assist with this planning prior to the installation.

The tool, known as the worldwide port name (WWPN) prediction tool, assigns WWPNs to each virtual Fibre Channel Protocol (FCP) channel/port using the same WWPN assignment algorithms that a system uses when assigning WWPNs for channels utilizing N_Port Identifier Virtualization (NPIV). Thus, the SAN can be set up in advance, allowing operations to proceed much faster once the server is installed.

The WWPN prediction tool takes a .csv file containing the FCP-specific I/O device definitions and creates the WWPN assignments that are required to set up the SAN. A binary configuration file that can be imported later by the system is also created. The .csv file can either be created manually or exported from the Hardware Configuration Definition/Hardware Configuration Manager (HCD/HCM).

Fiber Quick Connect for FICON LX

Fiber Quick Connect (FQC), an optional feature on z10 EC, is now being offered for all FICON LX (single-mode fiber) channels, in addition to the current support for ESCON (62.5 μ m multimode fiber) channels. FQC is designed to significantly reduce the amount of time required for on-site installation and setup of fiber optic cabling.

FQC facilitates adds, moves, and changes of ESCON and FICON LX fiber optic cables in the data center, and may reduce fiber connection time by up to 80%. FQC is for factory installation of IBM Facilities Cabling Services - Fiber Transport System (FTS) fiber harnesses for connection to channels in the I/O cage. FTS fiber harnesses enable connection to FTS direct-attach fiber trunk cables from IBM Global Technology Services.

FQC supports all of the ESCON channels and all of the FICON LX channels in all of the I/O cages of the server.

Open Systems Adapter enhancements

The z10 EC offers five new OSA-Express3 features. Other capabilities are mentioned that can help consolidate or simplify the data center environment.

OSA-Express3 features highlights

Exclusive to z10 are five new OSA-Express3 features. See 2.10.6, “OSA-Express3” on page 34. When compared with similar OSA-Express2 features, which they replace, the new features provide important benefits, such as:

- ▶ Doubling the density of ports: This reduces the number of CHPIDs to manage and the number of required I/O slots, which may reduce the number of I/O cages or I/O drawers. Up to 96 LAN ports versus 48 are now available.
- ▶ Designed to reduce the minimum round-trip networking time between systems (reduced latency): Improvements of up to 45% on round-trip time at the TCP/IP application layer may be realized with the OSA-Express3 10 GbE and OSA-Express3 GbE features.
- ▶ Designed to improve throughput (mixed inbound/outbound) for standard and jumbo frames.

These enhancements are because of a new function present in all OSA-Express3 features: the data router. With OSA-Express3, what was previously done in firmware is now performed in hardware. Additional logic in the IBM ASIC handles packet construction, inspection, and routing, thereby allowing packets to flow between host memory and the LAN at line speed without firmware intervention.

With the data router, the *store and forward* technique in DMA is no longer used. The data router enables a direct host memory-to-LAN flow. This avoids a *hop* and is designed to reduce latency and to increase throughput for standard frames (1492 byte) and jumbo frames (8992 byte).

Open Systems Adapter for NCP

The Open Systems Adapter for NCP (OSN) support available with OSA-Express3 Gigabit Ethernet, OSA-Express3 1000BASE-T Ethernet, OSA-Express2 Gigabit Ethernet, and OSA-Express2 1000BASE-T Ethernet features has the capability to provide channel connectivity from System z operating systems to IBM Communication Controller for Linux on System z (CCL) using the Open Systems Adapter for the Network Control Program (OSA for NCP) supporting the Channel Data Link Control (CDLC) protocol.

If SNA solutions that require NCP functions are required, CCL can be considered as a migration strategy to replace IBM Communications Controllers (374x). The CDLC connectivity option enables TPF and z/TPF environments to exploit CCL.

VLAN support

Virtual local area network (VLAN) is a function of OSA features that takes advantage of the IEEE 802.q standard for virtual bridged LANs. VLANs allow easier administration of logical groups of stations that communicate as though they were on the same LAN. In the virtualized environment of System z many TCP/IP stacks can exist, potentially sharing OSA features.

VLAN provides a greater degree of isolation by allowing contact with a server from only the set of stations comprising the VLAN.

VLAN is supported by z/OS, z/VM, and Linux on System z.

VMAC support

When sharing OSA port addresses across LPARs, VMAC support enables each operating system instance to have a unique virtual MAC (VMAC) address. All IP addresses associated with a TCP/IP stack are accessible using their own VMAC address, instead of sharing the MAC address of the OSA port. Advantages include a simplified configuration setup and improvements to IP workload load balancing and outbound routing.

This support is available for Layer 3 mode and is exploited by z/OS and by z/VM for guest exploitation.

QDIO data connection isolation for the z/VM environment

New workloads increasingly require multi-tier security zones. In a virtualized environment, an essential aspect is to protect workloads from intrusion or exposure of data and processes from other workloads.

The Queued Direct Input/Output (QDIO) data connection isolation enables:

- ▶ Adherence to security and HIPPA-security guidelines and regulations for network isolation between the instances sharing physical network connectivity
- ▶ Establishing security zone boundaries that have been defined by the network administrators
- ▶ A mechanism to isolate a QDIO data connection (on an OSA port) by forcing traffic to flow to the external network, ensuring that all communication flows only between an operating system and the external network

Internal *routing* can be disabled on a per-QDIO connection basis. This support does not affect the ability to share an OSA-Express port. Sharing occurs as it does today, but the ability to communicate between sharing QDIO data connections can be restricted through the use of this support.

QDIO data connection isolation applies to the z/VM environment, when using the Virtual Switch (VSWITCH) function, and to all of the OSA-Express3 and OSA-Express2 features (CHPID type OSD) on System z10 and to the OSA-Express2 features on System z9. z/OS supports a similar capability. See “QDIO interface isolation for z/OS”.

QDIO interface isolation for z/OS

Some environments require strict controls for routing data traffic between servers or nodes. In certain cases, the LPAR-to-LPAR capability of a shared OSA port can prevent such controls from being enforced. With interface isolation, internal routing can be controlled on an LPAR basis. When interface isolation is enabled, the OSA will discard any packets destined for a z/OS LPAR that is registered in the OAT as isolated.

QDIO interface isolation is supported by Communications Server for z/OS V1R11 and all OSA-Express3 and OSA-Express2 features on System z10.

QDIO optimized latency mode

QDIO optimized latency mode (OLM) can help improve performance for applications that have a critical requirement to minimize response times for inbound and outbound data. OLM optimizes the interrupt processing as follows:

- ▶ For inbound processing, the TCP/IP stack looks more frequently for available data to process, ensuring that any new data is read from the OSA-Express3 without requiring additional program controlled interrupts (PCIs).
- ▶ For outbound processing, the OSA-Express3 also looks more frequently for available data to process from the TCP/IP stack, thus not requiring a Signal Adapter (SIGA) instruction to determine whether more data is available.

HiperSockets enhancements

HiperSockets has been called the *network in a box*. z10 EC provides HiperSockets enhancements in two areas. HiperSockets Multiple Write Facility allows for zIIP-Assisted HiperSockets for large messages.

HiperSockets Layer 2 support

With this support the HiperSockets internal networks on System z10 EC can support two transport modes: Layer 2 (link layer) as well as the current Layer 3 (network or IP layer). Traffic can be Internet Protocol (IP) Version 4 or Version 6 (IPv4, IPv6) or non-IP (such as AppleTalk, DECnet, IPX, NetBIOS, SNA, or others). HiperSockets devices are now protocol-independent and Layer 3 independent. Each HiperSockets device has its own Layer 2 Media Access Control (MAC) address, which is designed to allow the use of applications that depend on the existence of Layer 2 addresses such as Dynamic Host Configuration Protocol (DHCP) servers and firewalls.

Layer 2 support can help facilitate server consolidation. Complexity can be reduced, network configuration is simplified and intuitive, and LAN administrators can configure and maintain the mainframe environment the same as they do a non-mainframe environment.

HiperSockets Layer 2 support is exclusive to System z10, and is supported by Linux on System z, and by z/VM for guest exploitation.

HiperSockets Multiple Write Facility

HiperSockets performance has been enhanced to allow for the streaming of bulk data over a HiperSockets link between logical partitions (LPARs). The receiving LPAR can now process a much larger amount of data per I/O interrupt. This enhancement is transparent to the operating system in the receiving LPAR. HiperSockets Multiple Write Facility, with fewer I/O interrupts, is designed to reduce CPU utilization of the sending and receiving LPAR.

The HiperSockets Multiple Write Facility is supported in the z/OS environment.

zIIP-Assisted HiperSockets for large messages

In z/OS, HiperSockets has been enhanced for zIIP exploitation. Specifically, the z/OS Communications Server allows the HiperSockets Multiple Write Facility processing for outbound large messages originating from z/OS to be performed on a zIIP.

zIIP-Assisted HiperSockets can help make highly secure and available HiperSockets networking an even more attractive option. z/OS application workloads based on XML, HTTP, SOAP, Java, and traditional file transfer can benefit from zIIP enablement by lowering general-purpose processor utilization for such TCP/IP traffic.

When the workload is eligible, the TCP/IP HiperSockets device driver layer (write) processing is redirected to a zIIP, which will unblock the sending application.

zIIP Assisted HiperSockets for large messages is available with z/OS V1.10 (plus PTF UK37306) on System z10 servers.

3.2.5 Cryptography enhancements

The z10 EC delivers cryptographic facilities similar to the z9. The implementation of the processor part of the cryptographic technology (CPACF) packaging is changed, but the change does not affect the functions provided and the way to use them from the applications. The same cryptographic features can be used on the z10 EC as those used on the z9.

CPACF enhancements

The cryptographic functions include improvements designed to facilitate continued privacy of cryptographic keys. CPACF helps to ensure that keys are not visible to applications and operating systems when used for encryption.

CPACF is designed to provide significant throughput improvements for encryption of large volumes of data as well as low latency for encryption of small blocks of data. Furthermore, enhancements to the information management tool, IBM Encryption Tool for IMS and DB2 Databases, is designed to improve performance for protected key encryption applications.

Support for 13-digit through 19-digit personal account numbers

Credit card companies sometimes perform card security code computations based on personal account number (PAN) data. The Integrated Cryptographic Service Facility (ICSF) callable services have been enhanced to support 13-digit through 19-digit PANs. Support for 13-digit through 19-digit PANs is exclusive to System z10 and is offered by z/OS and z/VM for guest exploitation.

3.2.6 Hardware Management Console enhancements

HMC/SE Version 2.10.2 is the current version available for the System z10 servers. The HMC application has several enhancements, such as:

- ▶ **Digitally signed firmware:** One critical issue with firmware upgrades is security and data integrity. Procedures are in place to use a process to digitally sign the firmware update files on the HMC, the SE, and the TKE.

Using a hash-algorithm, a message digest is generated and encrypted with a private key to produce a digital signature. This operation makes sure that any changes made to the data will be detected during the upgrade process. It helps ensure that no malware can be installed on System z products during firmware updates. It enables, with other existing security functions, System z10 CPACF functions to comply with Federal Information Processing Standard (FIPS) 140-2 Level 1 for Cryptographic Licensed Internal Code (LIC) changes. The enhancement follows the System z focus of security for the HMC and the SE.

- ▶ **Optional user password on disruptive confirmation:** The requirement to supply a user password on disruptive confirmation is optional. The general recommendation remains to require a password.

- ▶ Improved consistency of confirmation panels on the HMC and the SE: Attention indicators are on the top of panels, and there will be a list of objects affected by the action (target as well as secondary objects), for example, LPARs if the target is CPC.
- ▶ Serviceability enhancements for FICON channels:
 - Simplified problem determination to more quickly detect fiber optic cabling problems in a storage area network.
 - All FICON channel error information is forwarded to the HMC, thus facilitating detection and reporting trends and threshold for the channels including aggregate views including data from multiple systems.

3.3 Common time functions

Each server must have an accurate time source to maintain a time-of-day value. Logical partitions use its server time. When servers participate in Parallel Sysplex, coordinating the time across all the systems in a complex is critical to its operation.

3.3.1 Sysplex Timer

A Sysplex Timer is a device that provides the synchronization for the time-of-day (TOD) clocks of multiple servers, and thereby allows events started by different servers to be properly sequenced in time. For instance, when multiple servers update the same database, all updates are required to be time stamped in proper sequence.

The z10 EC and the other System z servers can attach to the IBM Sysplex Timer units. More information can be found in the Redbooks publication *S/390® Timer Management and IBM 9037 Sysplex Timer*, SG24-2070.

ETR attachment

New with the z10 EC server is the shipment as a standard feature of two External Time Reference (ETR) cards. Each card contains an ETR port for Sysplex Timer connection.

Thus, a redundant dual-path interface to IBM Sysplex Timers can be created, which may be used for timing synchronization between systems. This redundant design allows continued operation even if a single ETR card fails, and also allows concurrent maintenance. The two z10 EC server ETR cards are located in the processor cage of the z10 EC server.

3.3.2 Server Time Protocol (STP)

Server Time Protocol is a message-based protocol in which timekeeping information is passed over data links between servers. The timekeeping information is transmitted over externally defined coupling links.

The STP feature is intended to be the supported method for maintaining time synchronization between System z servers and CFs.

The STP design uses a new concept called Coordinated Timing Network (CTN). A CTN is a collection of servers and CFs that are time synchronized to a time value called Coordinated Server Time (CST). Each server and CF planned to be configured in a CTN must be STP-enabled. STP is intended for servers that are configured to participate in a Parallel Sysplex or in a sysplex (without a CF), as well as servers that are not in a sysplex, but must be time synchronized.

STP is implemented in LIC as a server-wide facility of z10 EC (and other System z servers and CFs). STP presents a single view of time to PR/SM and provides the capability for multiple servers and CFs to maintain time synchronization with each other. A System z server or CF may be enabled for STP by installing the STP feature.

STP provides the following additional value over the Sysplex Timer:

- ▶ STP supports a multi-site timing network of up to 100 km (62 miles) over fiber optic cabling, without requiring an intermediate site. This allows a Parallel Sysplex to span these distances and reduces the cross-site connectivity required for a multi-site Parallel Sysplex.
- ▶ The STP design allows more stringent synchronization between servers and CFs using short communication links, such as PSIFB or ICB-4 links, compared with servers and CFs using long ISC-3 links across sites. With the z10 EC server, STP will support coupling links over InfiniBand.
- ▶ STP helps eliminate infrastructure requirements, such as power and space, needed to support the Sysplex Timers.
- ▶ STP helps eliminate maintenance costs associated with the Sysplex Timers.
- ▶ STP may reduce the fiber optic infrastructure requirements in a multi-site configuration. Dedicated links may not be required to transmit timing information.

The CTN concept is used to help meet two key goals of z10 EC and System z customers:

- ▶ Concurrent migration from an existing External Time Reference (ETR) network to a timing network using STP.
- ▶ Capability of servers to be synchronized in the timing network that contains a collection of servers and has at least one STP-configured server stepping to timing signals provided by the Sysplex Timer. Such a network is called a Mixed CTN.

STP supports dial-out time services to set the time to an international time standard, such as Coordinated Universal Time (UTC), as well as to adjust to the time standard on a periodic basis. In addition, setting local time parameters, such as time zone and Daylight Saving Time (DST), and automatic updates of DST are supported.

STP is available as a chargeable feature on the System z servers and is supported by z/OS starting with V1.7, which requires PTFS to enable STP support. This support is included with z/OS V1.8.

Network Time Protocol (NTP) client support

The use of Network Time Protocol servers as an external time source (ETS) usually fulfills a requirement for a time source or common time reference across heterogeneous platforms. In most cases, this fulfillment is an NTP server that obtains the exact time through satellite.

NTP client support is available on z10 servers and has been available on z9 servers since October 2007. With this implementation, the z10 and z9 servers support the use of NTP servers as time sources.

NTP client support is added to the support element (SE) code of the z10 and z9 servers. The code interfaces with the NTP servers. This allows an NTP server to become the single time source for z10 and z9 servers, and for other servers that have NTP clients. NTP can be used only for an STP-only CTN environment.

This support satisfies the following 2007 IBM Statement of Direction related to STP:

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.

Pulse per second (PPS) support

Some NTP servers also provide a PPS output signal. The PPS output signal is more accurate (within 10 microseconds) than that from the HMC dial-out function or an NTP server without PPS (within 100 milliseconds).

Each of the two standard ETR cards on z10 servers also have a PPS port (for a coaxial cable connection) that can be used by STP in conjunction with the NTP client.

Continuous availability of NTP servers used as an External Time Source

If the Preferred Time Server or Current Time Server (CTS) cannot access the NTP Server or the pulse per second signal from the NTP server, the Backup Time Server (BTS), if configured to a different NTP server, can calculate the External Time Source (ETS) adjustment required and propagate it to the PTS/CTS. The PTS/CTS will continue to perform the necessary time adjustment steering.

NTP server on HMC

NTP server capability on the HMC addresses the potential security concerns that users may have for attaching NTP servers directly to the HMC/SE LAN. Note that when using the HMC as the NTP server, there is no pulse per second capability available.

Enhanced STP recovery when Internal Battery Feature is in use

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

STP configuration and time information

STP configuration and time information are restored across Power on Resets (POR) or after power outages for STP-only CTNs. The user will not have to reinitialize the time or reassign the roles of PTS/CTS and BTS across POR or after power outage events. This is valid for a single server CTN (PTS/CTS) as well as for a dual server CTN (PTS/CTS and BTS).

Application programming interface (API) to automate STP CTN reconfiguration

If the PTS fails and the BTS takes over as CTS, an API is available on the HMC so that you can automate the reassignment of the PTS, BTS, and Arbiter roles. For additional details about the API, refer to *System z Application Programming Interfaces*, SB10-7030.

This fulfills a Statement of Direction made in October 2006.

For a more in-depth discussion of STP, refer to the *Server Time Protocol Planning Guide*, SG24-7280, and the *Server Time Protocol Implementation Guide*, SG24-7281.

3.4 Capacity on Demand (CoD) enhancements

Based on customer demand and changes in the market requirements, the System z10 servers introduce a number of enhancements to the on demand offerings. These changes provide more flexibility and control to the customer in order to ease the administrative burden in the handling of the offerings and to give the customer finer control over resources needed to meet the resource requirements in various situations.

The System z10 servers have the capability of *concurrent* upgrades, providing additional capacity with no *server* outage. In most cases, with prior planning and operating system support, a concurrent upgrade can also be *nondisruptive* to the operating system.

It is important to note that these upgrades are based on the enablement of resources already physically present in the z10 servers.

Capacity upgrades cover both permanent and temporary changes to the installed capacity. This can be done using the Customer Initiated Upgrade (CIU) facility, without requiring IBM service personnel involvement. Such upgrades are initiated through the Web, using IBM Resource Link™. Use of the CIU facility requires a special contract between the customer and IBM, through which terms and conditions for online CoD buying of upgrades and other types of CoD upgrades are accepted. For more information consult the IBM Resource Link site:

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

With CoD, the z10 introduces the possibility of having more than one temporary capacity upgrade active at any point in time. Eight different temporary upgrades can be active at the same time, with one of them being an On/Off Capacity on Demand (On/Off CoD) upgrade. The others can be a combination of upgrade offerings. Furthermore, upgrades can be performed concurrently, and they can be replenished even when active. It is also possible to do permanent upgrades while temporary upgrades are active.

The content of the on demand upgrade records can be used in such a way that subsets of the capacity can be activated, and additional resources in the upgrade can be added or taken away without having to go back to the base configuration. This removes the requirement for a customer to have several On/Off CoD upgrade records installed and to have to switch between them, potentially impacting availability, to meet varying workload demands. Thus, we recommend that a single OOCOD with the largest possible configuration be used.

The following sections discuss permanent and temporary upgrades and provisioning.

For more information regarding the Capacity on Demand offerings, refer to the *IBM System z10 Enterprise Class Technical Guide*, SG24-7516, and *IBM System z10 Enterprise Class Capacity on Demand*, SG24-7504.

3.4.1 Permanent and temporary upgrades

Table 3-1 summarizes the CoD offerings that are available for the z10 servers.

Table 3-1 Capacity on Demand summary

		Upgrades	Process
Permanent	Online permanent upgrade	CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs, and memory	Performed through the Resource Link application
Temporary	On/Off CoD CBU CPE	CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs	Performed through the Resource Link application

Permanent upgrades

Permanent upgrades of processors (CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs) and memory, or changes to a server's Model-Capacity Identifier, up to the limits of the installed books on an existing z10 server, can be performed by the customer through the IBM On-line Permanent Upgrade offering, using the CIU facility. These permanent upgrades require a special contract

between the customer and IBM, through which the terms and conditions of the offering are accepted.

Temporary upgrades

Temporary upgrades of a System z10 server can be done by On/Off CoD, Capacity Backup (CBU) or Capacity for Planned Event (CPE) ordered from the CIU facility. These temporary upgrades require a special contract between the customer and IBM, through which the terms and conditions of the offering are accepted.

On/Off Capacity on Demand

On/Off CoD is a function available on the z10 server that enables *concurrent* and *temporary* capacity growth of the server. On/Off CoD *can* be used for customer peak workload requirements, for any length of time, and has a daily hardware charge and may have an associated SW charge. On/Off CoD offerings can be pre-paid or post-paid. Capacity tokens are introduced on the System z10 servers. Capacity tokens are always present in pre-paid offerings and can be present in post-paid if the customer so desires. In both cases capacity tokens are being used to control the maximum resource and financial consumption.

The customer's charges for software may vary according to the license agreement for the individual products. The IBM Software Group representative should be contacted for exact details of an impact to charges for IBM program products.

Using On/Off CoD, the customer can concurrently add processors (CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs), increase the CP capacity level, or both.

Capacity Backup (CBU)

CBU allows the customer to perform a *concurrent* and *temporary* activation of additional CPs, ICFs, IFLs, zAAPs, zIIPs, and SAPs, an increase of the CP capacity level, or both, in the event of an unforeseen loss of System z capacity within the customer's enterprise, or to perform a test of the customer's disaster recovery procedures. The capacity of a CBU upgrade cannot be used for peak workload management.

CBU features are optional and require unused capacity to be available on installed books of the backup server, either as unused PUs or as a possibility to increase the CP capacity level on a sub-capacity server, or both. A CBU contract must be in place before the LIC-CC code that enables this capability can be loaded on the server. An initial CBU record provides for at least five tests (each up to 10 days in duration) and one disaster activation (up to 90 days in duration) and can be configured to be valid for up to five years.

Capacity for Planned Event (CPE)

Capacity for Planned Event allows the customer to perform a *concurrent* and *temporary* activation of additional CPs, ICFs, IFLs, zAAPs, zIIPs, and SAPs, an increase of the CP capacity level, or both, in the event of a planned outage of System z capacity within the customer's enterprise (for example, data center changes or system maintenance). CPE cannot be used for peak workload management and is available for up to a maximum of three days.

The CPE feature is optional and requires unused capacity to be available on installed books of the back-up server, either as unused PUs or as a possibility to increase the CP capacity level on a sub capacity server, or both. A CPE contract must be in place before the LIC-CC that enables this capability can be loaded on the server.

3.4.2 z/OS capacity provisioning

Capacity provisioning helps customers manage the CP, zAAP, and zIIP capacity of z10 EC servers that are running one or more instances of the z/OS operating system. Based on On/Off CoD, temporary capacity may be activated and deactivated under control of a defined policy. Combined with functions in z/OS, the z10 provisioning capability gives the customer a flexible, automated process to control the configuration and activation of On/Off CoD offerings.

Provisioning architecture overview

The provisioning architecture enhances the already rich on demand environment by opening up interfaces to the z/OS operating system. The z/OS operating system can interrogate the on demand environment and query which resources are in the On/Off CoD offerings and the status of the resources.

z/OS capacity provisioning simplifies the monitoring of critical workloads, and its automation features can help activate additional resources faster than manual operation. When using capacity provisioning, you can select different levels of automation to provide you with an appropriate level of control. For example, you can:

- ▶ Activate and deactivate temporary capacity through operator commands (manual mode).
- ▶ Activate and deactivate temporary capacity based on a defined schedule, without considering the actual workload performance.
- ▶ Instruct the Provisioning Manager (described in “Capacity Provisioning Manager” on page 67) to suggest changes to the capacity of the z10 server based on the observation of defined workloads. In this case the operator will have to confirm the suggested changes.
- ▶ Instruct the Provisioning Manager to automatically implement changes to the capacity of the z10 server based on the observation of defined workloads.

You may also run capacity provisioning in analysis mode. In this mode, the operator will be informed when an action would have occurred according to the defined rules. However, no action will be taken unless the operator manually enters the necessary commands.

The capacity provisioning function included in the z/OS operating system is part of the z/OS MVS Base Control Program (BCP). It includes the following items:

- ▶ Capacity Provisioning Manager, which is the server program
- ▶ Capacity Provisioning Control Center, which is the workstation code
- ▶ Sample datasets and files

Capacity Provisioning Manager

The Capacity Provisioning Manager monitors the workload on a set of z/OS instances and organizes the allocation of additional capacity when required. The systems to be observed are defined in a domain configuration file. Details of additional capacity and the rules for its allocation are held in a policy file. These two files are created and maintained using the Capacity Provisioning Control Center.

Capacity Provisioning Control Center

The Capacity Provisioning Control Center is installed on a workstation and is the graphical user interface to capacity provisioning. Through this interface the administrators work with provisioning policies and domain configurations and can transfer these to the Capacity Provisioning Manager.

Sample datasets and files

The capacity provisioning component includes several samples to simplify customization and help the definition of your provisioning policies.

Workload management

z/OS can apply its workload management disciplines to the on demand environment using the same mechanisms and controls already used inside z/OS. The provisioning of resources can be placed under control of automation functions using well-known parameters in the z/OS Workload Manager (WLM). Automation processes can activate resources in the OOCOD offerings as dictated by automation policies and return the resources when they are no longer needed according to the automation policies. Manual and supervised modes (where changes need authorization) are also available.

The CPM can access the offering records through standard APIs. Resource and capacity token information can be interrogated by the CPM, and CPM will check all elements of the records before provisioning any resources.

3.5 Throughput optimization enhancements

The z990 was the first server to use the concept of books. Despite the memory being distributed through the books and books having individual Level 2 caches, all processors have access to all the Level 2 caches and memory. Thus, the server is managed as a memory coherent symmetric multi-processor (SMP).

Processors within the z10 EC book structure have different *distance to memory* attributes. As described on 2.4, “CEC cage and books” on page 22, books are connected in a star configuration, which helps to minimize the distance.

Other non-negligible effects result from data latency when grouping and dispatching work on a set of available logical processors. In order to minimize latency, one can aim to dispatch and later re-dispatch work to a group of physical CPUs that share the same Level 2 (L2) cache.

PR/SM manages the utilization of physical processors by logical partitions by dispatching the logical processors on the physical processors. But PR/SM is not aware of which workloads are being dispatched by the operating system in which logical processors. The Workload Manager (WLM) component of z/OS has the information at the task level, but is unaware of physical processors. This disconnect is solved by enhancements on z10 EC that allow PR/SM and WLM to work more closely together. They can cooperate to create an affinity between task and physical processor rather than between logical partition and physical processor. This is known as HiperDispatch.

HiperDispatch

HiperDispatch, exclusive on z10, combines two functional enhancements, one in the z/OS dispatcher and one in PR/SM. This is intended to improve efficiency both in the hardware and in z/OS.

In general, the PR/SM dispatcher assigns work to a minimum number of logical processors needed for the priority (weight) of the LPAR. The end result is to reduce the multi-processor effects and lower the interference among multiple partitions.

The z/OS dispatcher is enhanced to operate with multiple dispatching queues, and tasks are distributed among these queues. The current implementation operates with an average of four logical processors per queue. Specific z/OS tasks may then be dispatched to a small

subset of logical processors, which PR/SM will tie to the same physical processors, thus improving the hardware cache re-use and locality of reference characteristics such as reducing the rate of cross-book communication.

To use the correct logical processors, the z/OS dispatcher obtains the necessary information from PR/SM through new interfaces implemented on the z10 EC. The entire z10 EC stack (hardware, firmware, and software) now tightly collaborates to obtain the hardware's full potential.

The HiperDispatch function can be switched on and off dynamically without requiring an IPL.

3.6 Reliability, availability, and serviceability improvements

The System z10 EC server presents numerous enhancements in the reliability, availability, and serviceability areas. In the availability area focus was given to reduce the planning requirements, while continuing to improve the elimination of planned, scheduled, and unscheduled outages.

Enhanced driver maintenance (EDM) helps reduce the necessity and the eventual duration of a planned outage. One of the contributors to planned outages is LIC updates performed in support of new features and functions. When properly configured, the System z10 can concurrently activate a new LIC level. Concurrent activation of the select new LIC level used to be supported only at specific synchronization points. It is possible to concurrently activate a select new LIC level anywhere in the maintenance stream. However, certain LIC updates are still not supported this way.

The design and packaging of the z10 EC represents a reduction in the number of chips necessary to implement the processor, cache, and infrastructure functions. Although the z9 was implemented using 16 chips on the MCM, the z10 EC uses only seven chips on the MCM. This reduction helps to improve the availability characteristics of the new server.

Availability enhancements include single processor core checkstop and sparing, point-to-point fabric for SMP, and fixed size HSA.

The ICB-4 connections used for Parallel Sysplex connectivity can now be reconfigured concurrently, and no longer need a server outage.

If an additional system assist processor (SAP) is required on a z10 EC server (for example, as a result of a disaster recovery situation), the SAPs can be concurrently added to the server configuration.

It is possible to concurrently add CPs, zAAPs, zIIPs, IFLs, and ICFs processors to an LPAR. This is supported by z/VM V5R3 and later with appropriate PTFs and by z/OS. Previously, proper planning was required in order to concurrently add CPs, zAAPs, and zIIPs to a z/OS LPAR.

Concurrently adding memory to an LPAR is also possible and is supported by z/OS and z/VM.

z10 EC supports dynamically adding Crypto Express features to an LPAR by providing the ability to change the cryptographic information in the image profiles without outage to the LPAR. Users can also dynamically delete or move Crypto Express features. This enhancement is supported by z/OS, z/VM, and Linux on System z.

The System Activity Display (SAD) screens now include energy efficiency displays.

These and additional features are further described *IBM System z10 Enterprise Class Technical Guide*, SG24-7516.

3.7 Parallel Sysplex technology

Parallel Sysplex technology is a clustering technology for logical and physical servers, allowing the highly reliable, redundant, and robust System z technology to achieve near-continuous availability. Both hardware and software tightly cooperate to achieve this result. The hardware components comprise:

- ▶ **Coupling Facility (CF):** This is the cluster center. It can be implemented either as an LPAR of a stand-alone System z server or as an additional LPAR of a System z server where other loads are running. Processor units characterized as either CPs or ICFs can be configured to this LPAR. ICFs are often used because they do not incur any software license charges. Two CFs are recommended for availability.
- ▶ **Coupling Facility Control Code (CFCC):** This IBM Licensed Internal Code is both the *operating system* and the *application* that executes in the CF. No other code executes in the CF.¹
- ▶ **Coupling links:** These are high-speed links connecting the several system images (each running in its own logical partition) that participate in the Parallel Sysplex. At least two connections between each physical server and the CF should exist. When all of the system images belong to the same physical server, internal coupling links are used.

On the software side, the z/OS operating system exploits the hardware components to create a Parallel Sysplex². Normally, two or more z/OS images are clustered to create a Parallel Sysplex, although it is possible to have a configuration setting with a single image, called a *monoplex*. Multiple clusters can span several System z servers although a specific image (logical partition) can belong to only one Parallel Sysplex.

A z/OS Parallel Sysplex implements a shared-all access to data. This is facilitated by System z I/O virtualization capabilities such as Multiple Image Facility (MIF). MIF allows several logical partitions to share I/O paths in a totally secure way, maximizing utilization and greatly simplifying the configuration and connectivity.

In short, a Parallel Sysplex comprises one or more z/OS operating system images coupled through one or more Coupling Facilities. A properly configured Parallel Sysplex cluster is designed to maximize availability at the application level.

¹ CFCC can also execute in a z/VM Virtual Machine (as a z/VM guest system). In fact, a complete Parallel Sysplex can be set up under z/VM allowing, for instance, testing and operations training. This setup is not recommended for production environments.

² TPF and z/TPF also exploit the CF hardware components. However, the term Parallel Sysplex exclusively applies to z/OS exploitation of CF.

The major characteristics of a Parallel Sysplex are:

- ▶ Data sharing with integrity: The CF is key to the implementation of a share-all access to data. Every z/OS system image has access to all the data. Subsystems in z/OS declare resources to the CF. The CF accepts and manages lock and unlock requests on those resources, guaranteeing data integrity. A duplicate CF further enhances the availability. Key exploiters of this capability are DB2, WebSphere MQ, WebSphere ESB, IMS, and CICS.
- ▶ Continuous (application) availability: Changes, such as software upgrades and patches, can be introduced one image at a time, while the remaining images continue to process work. For additional details see the manual *Parallel Sysplex Application Considerations*, SG24-6523.
- ▶ High capacity: scales from two to 32 images. Remember that each image can have from 1 to 64 processor units. CF scalability is near linear. This contrasts with other forms of clustering that employ n-to-n messaging, leading to rapidly degrading performance with growth of the number of nodes.
- ▶ Dynamic workload balancing: Viewed as a single logical resource, work can be directed to any of the Parallel Sysplex cluster operating system images where capacity is available.
- ▶ Systems management: The architecture provides the infrastructure to satisfy a customer requirement for continuous availability, while enabling techniques for achieving simplified systems management consistent with this requirement.
- ▶ Resource sharing: A number of base z/OS components exploit Coupling Facility shared storage. This exploitation enables sharing of physical resources with significant improvements in cost, performance, and simplified systems management.
- ▶ Single system image: The collection of system images in the Parallel Sysplex appears as a single entity to the operator, the user, the database administrator, and so on. A single system image ensures reduced complexity from both operational and definition perspectives.

Figure 3-3 illustrates the components of a Parallel Sysplex as implemented within the System z architecture. Figure 3-3 is intended only as an example. It shows one of many possible Parallel Sysplex configurations. Many other possibilities exist.

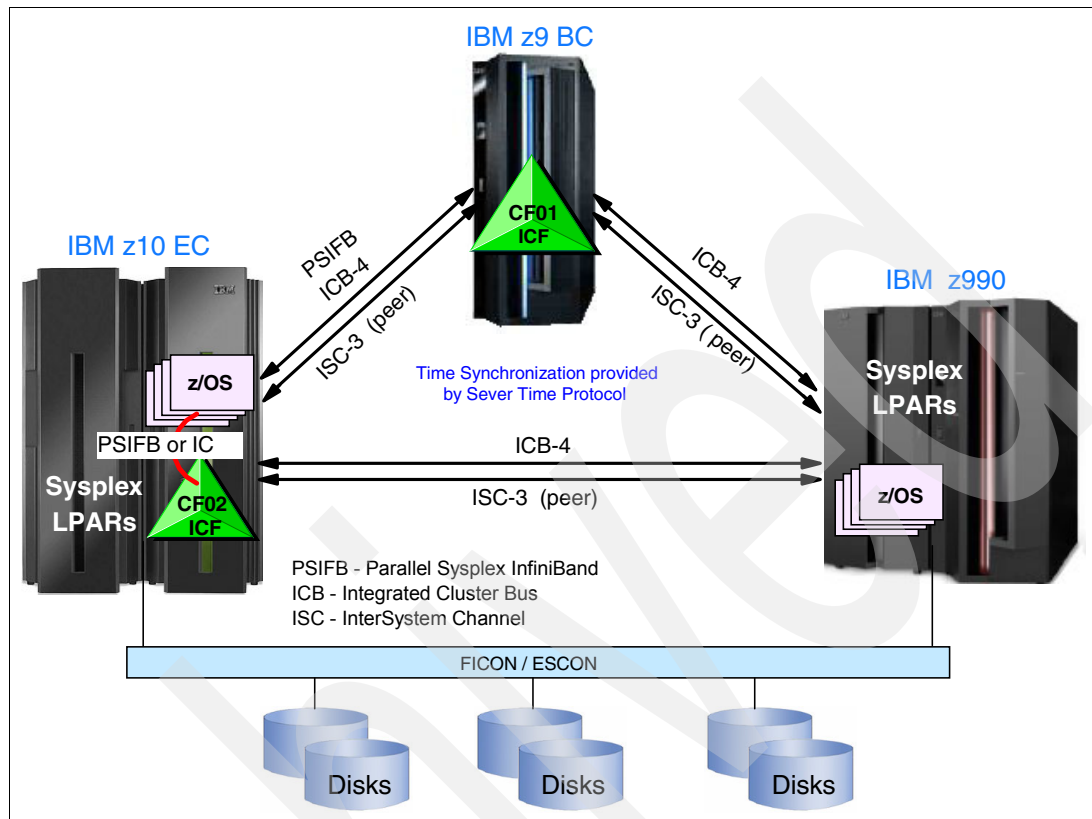


Figure 3-3 Sysplex hardware overview

Figure 3-3 shows a z10 EC server containing multiple z/OS sysplex partitions and an internal coupling facility (CF02), a z9 Business Class server containing a stand-alone CF (CF01), and a z990 containing multiple z/OS sysplex partitions. STP over coupling links provides time synchronization to all servers. Appropriate CF link technology (PSIFB, ICB-4, ISC-3) selection depends on server configuration.

Through this state-of-the-art cluster technology, the power of multiple z/OS images can be harnessed to work in concert on shared workloads and data. The System z Parallel Sysplex cluster takes the commercial strengths of the z/OS platform to improved levels of system management, competitive price/performance, scalable growth, and continuous availability.

3.8 Summary

Multiple forces are driving a transformation of the data center, such as demands to IT to improve cost and service delivery, manage escalating complexity, and better secure the enterprise. Aligning IT more closely with the business has become a primary goal.

The Dynamic Infrastructure, the IBM vision for the future of the IT infrastructure, and with it the data center, is designed to optimize service delivery and provide exceptional efficiency. The z10 EC has a key role at the core of this infrastructure, helping to simplify it and providing powerful shared virtual resources while lowering the cost of ownership.

The z10 EC is a big step forward in the mainframe evolution. It is, at the same time, a revolution evidenced in its new microprocessor design, use of high frequency, and additional exploitation of open technologies such as InfiniBand.

The z10 EC extreme virtualization capabilities, very large capacity range (1:140), memory growth (16 GB to 1.5 TB), and I/O bandwidth are provided *in the same footprint*. All this distinctly shows the server's readiness for IT infrastructure simplification through image consolidation and application integration.

When considering the complementarity of supported environments, with emphasis on z/OS and Linux on System z, and the wide scope of technologies exploited by the thousands of supported applications, the z10 EC is the place for both new applications and traditional ones. Furthermore, the investment value locked in traditional applications can be exploited through novel interfaces such as Web services.

z10 EC advanced security features, including key generation and management, make it the choice for an enterprise security hub. Its high availability (up to 99.999% at the *application* level with Parallel Sysplex) and resilience naturally recommend it as the platform to run the mission-critical middleware and applications and host, manage, and care for the enterprise data.

The granular capacity increments and capacity on demand options allow the server to grow with the enterprise. In fact, when also considering the advantages of server consolidations, z10 EC may well help grow the enterprise.

Archived

Software support

This chapter focuses on operating system requirements and support considerations for the z10 EC and its features.

This chapter discusses the following topics:

- ▶ 4.1, “Software support summary” on page 76
- ▶ 4.2, “Support by operating system” on page 78
- ▶ 4.3, “Support for selected functions” on page 85
- ▶ 4.4, “z/OS considerations” on page 93
- ▶ 4.5, “Coupling Facility and CFCC considerations” on page 95
- ▶ 4.6, “IOCP” on page 95
- ▶ 4.8, “ICKDSF” on page 96
- ▶ 4.9, “Software licensing considerations” on page 97

Support of the System z10 EC functions is dependent on the operating system version and release. This information is subject to change. Therefore, for the most current information, refer to the Preventive Service Planning (PSP) bucket for 2097DEVICE.

4.1 Software support summary

The software portfolio for the System z10 EC server includes a large variety of operating systems and middleware that support the most recent and significant technologies. Continuing the mainframe-rich tradition, five major operating systems are supported on the z10 EC:

- ▶ z/OS
- ▶ z/VM
- ▶ z/VSE
- ▶ z/TPF
- ▶ Linux on System z

Operating systems summary

Table 4-1 summarizes the current and minimum operating system levels required to support the z10 EC. Note that operating system levels that are no longer in service are not covered in this publication. These older levels may provide support for some features.

Table 4-1 z10 EC operating system requirements

Operating system	ESA/390 (31-bit mode)	z/Architecture (64-bit mode)	End of service	Notes
z/OS V1R11	No	Yes	Not announced	Refer to the z/OS, z/VM, z/VSE, and z/TPF subsets of the 2097DEVICE Preventative Service Planning (PSP) bucket prior to installing the IBM System z10 EC.
z/OS V1R10	No	Yes	September 2011 ^a	
z/OS V1R9	No	Yes	September 2010 ^a	
z/OS V1R8 ^b	No	Yes	September 2009 ^b	
z/OS V1R7 ^c	No	Yes	September 2008 ^c	
z/VM V6R1 ^d	No ^e	Yes	April 2013	
z/VM V5R4	No ^e	Yes	September 2013 ^a	
z/VM V5R3	No ^e	Yes	September 2010 ^a	
z/VSE V4R2	No ^f	Yes ^g	Not announced	
z/VSE V4R1	No ^f	Yes ^g	Not announced	
z/TPF V1R1	No	Yes	Not announced	
TPF V4R1	Yes	No	December 2010	
Linux on System z	See Table 4-5 on page 83.	See Table 4-5 on page 83.	See footnote ^h	Novell SUSE SLES 11 Novell SUSE SLES 10 Novell SUSE SLES 9 Red Hat RHEL 5 Red Hat RHEL 4

- a. Planned date. All statements regarding IBM plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.
- b. With the announcement of IBM Lifecycle Extension for z/OS V1.8 fee-based corrective service can be ordered for up to two years after the withdrawal of service for z/OS V1R8.
- c. With the announcement of IBM Lifecycle Extension for z/OS V1.7 fee-based corrective service can be ordered for up to two years after the withdrawal of service for z/OS V1R7.
- d. z/VM V6R1 requires an architectural level set exclusive to z10.

- e. z/VM supports both ESA/390 mode and z/Architecture mode virtual machines.
- f. ESA/390 is not supported. However, 31-bit mode is supported.
- g. z/VSE V4R1 and later support 64-bit real addressing only. They do not support 64-bit addressing for user, system, or vendor applications.
- h. For information about support-availability of Linux on System z distributions, see:
Novell SUSE:
<http://support.novell.com/lifecycle/lcSearchResults.jsp?st=Linux+Enterprise+Server&x=32&y=11&s1=-1&sg=-1&pid=1000>
Red Hat:
<http://www.redhat.com/security/updates/errata/>

Note: Exploitation of several features depends on a particular operating system. In all cases, PTFs might be necessary with the operating system level indicated. PSP buckets are continuously updated and should be reviewed regularly when planning for installation of a new server. They contain the latest information about maintenance.

PSP buckets contain installation information, hardware and software service levels, service recommendations, and cross-product dependencies.

Middleware

Middleware offerings for the z10 EC environments include:

- ▶ Transaction processing
 - WebSphere Application Server and WebSphere Extended Deployment
 - CICS Transaction Server
 - CICS Transaction Gateway
 - IMS DB and IMS DC
 - IMS Connect
- ▶ Application integration and connectivity
 - WebSphere Message Broker
 - WebSphere MQ
 - WebSphere ESB
- ▶ Process integration
- ▶ WebSphere Process Server
- ▶ WebSphere MQ Workflow
- ▶ WebSphere Business Integration Server
- ▶ Database
 - ▶ DB2 for z/OS
 - ▶ DB2 for Linux
 - ▶ DB2 Connect™

Operations

The Tivoli brand has a large product set that includes:

- ▶ Tivoli Service Management Center
- ▶ Tivoli Information Management for z/OS
- ▶ Tivoli Workload Scheduler
- ▶ Tivoli OMEGAMON® XE
- ▶ Tivoli System Automation

Security

A highly secure System z environment can be implemented at various levels using the following products:

- ▶ Security Server feature of z/OS (includes Resource Access Control Facility (RACF®) and LDAP server)
- ▶ Tivoli Access Manager
- ▶ Tivoli Federated Identity Manager
- ▶ z/OS Communications Server and Policy Agent (for policy-based network security)

Application development and languages

Many languages are available for the z10 EC environments. Because the Linux environment is similar to Linux on other servers, we focus on the z/OS environment.

In addition to the traditional COBOL, PL/I, FORTRAN, and Assembler languages, C, C++, and Java, including J2EE and batch environments, are available.

Development can be conducted using the latest software engineering technologies and advanced IDEs. The extensive tool set uses a workstation environment for development and testing, with final testing and deployment performed on z/OS. Application development tools, many of which have components based on the Eclipse platform, include:

- ▶ Rational® Application Developer for WebSphere
- ▶ Rational Developer for System z
- ▶ WebSphere developer for System z
- ▶ Rational Rose® product line
- ▶ Rational Software Architect and Software Modeler

The following Web site is organized by category and has an extensive set of links to information about software for System z:

<http://www-306.ibm.com/software/sw-bycategory/systemz>

4.2 Support by operating system

In this section we list the support of new System z10 EC functions by the current operating systems. See the companion manual *z10 EC IBM System z10 Enterprise Class Technical Guide*, SG24-7516, for a detailed description of z10 EC and its features. For an in-depth description of all I/O features refer to the *IBM System z Connectivity Handbook*, SG24-5444.

4.2.1 z/OS

z/OS Version 1 Release 9 is the earliest service release supporting the z10 EC. Although service support for z/OS Version 1 Release 8 ended in September of 2009, a fee-based extension for defect support (for up to two years) can be obtained by ordering the IBM Lifecycle Extension for z/OS V1.8. Similarly, IBM Lifecycle Extension for z/OS V1.7 provides fee-based support for z/OS V1.7 up to September 2010. Service support for z/OS Version 1 Release 6 ended on September 30, 2007. Also note that z/OS.e is not supported on z10 EC and that the last release of z/OS.e was z/OS.e Version 1 Release 8.

Table 4-2 summarizes the z10 EC function support requirements for the currently supported z/OS releases. It uses the following conventions:

- Y** The function is supported.
- N** The function is not supported.

Table 4-2 z/OS support summary

Function	V1R11	V1R10	V1R9	V1R8 ^a	V1R7 ^a
z10 EC	Y	Y	Y	Y	Y
Greater than 54 PUs single system image	Y	Y	Y	N	N
Dynamic add of logical CPs	Y	Y	N	N	N
zAAP on zIIP	Y	Y ^d	Y ^d	N	N
Large memory > 128 GB	Y	Y	Y	Y ^d	N
Large page support	Y	Y ^d	Y ^d	N	N
Hardware decimal floating point	Y ^b	Y ^b	Y ^{bd}	Y ^{bd}	Y ^{bd}
CPACF protected public key	Y ^c	Y ^c	Y ^c	N	N
Enhanced CPACF	Y	Y	Y ^c	Y ^c	Y ^c
Personal account numbers of 13 to 19 digits	Y ^c	Y ^c	Y ^c	Y ^c	Y ^c
Crypto Express3	Y ^c	Y ^c	Y ^c	N	N
Capacity Provisioning Manager	Y	Y ^d	Y ^d	N	N
HiperDispatch	Y	Y ^d	Y ^d	Y ^d	Y ^{de}
HiperSockets multiple write facility	Y	Y	Y ^d	N	N
High Performance FICON	Y	Y ^d	Y ^d	Y ^d	Y ^d
FICON Express8	Y ^f	Y ^f	Y ^f	Y ^f	Y ^f
OSA-Express3 10 Gigabit Ethernet LR CHPID type OSD	Y	Y	Y	Y	Y
OSA-Express3 10 Gigabit Ethernet SR CHPID type OSD	Y	Y	Y	Y	Y
OSA-Express3 Gigabit Ethernet LX using four ports CHPID types OSD and OSN	Y	Y	Y ^d	Y ^d	N
OSA-Express3 Gigabit Ethernet LX using two ports CHPID types OSD and OSN	Y	Y	Y	Y	Y
OSA-Express3 Gigabit Ethernet SX using four ports CHPID types OSD and OSN	Y	Y	Y ^d	Y ^d	N

Function	V1R11	V1R10	V1R9	V1R8 ^a	V1R7 ^a
OSA-Express3 Gigabit Ethernet SX using two ports CHPID types OSD and OSN	Y	Y	Y	Y	Y
OSA-Express3 1000BASE-T Ethernet using four ports CHPID types OSC, OSD, and OSN ^g	Y	Y	Y ^d	Y ^d	N
OSA-Express3 1000BASE-T Ethernet using two ports CHPID types OSC, OSD, OSE and OSN ^g	Y	Y	Y	Y	Y ^d
Coupling using InfiniBand CHPID type CIB	Y	Y	Y	Y	Y
InfiniBand coupling links (12x IB-SDR or 12x IB-DDR) at a distance of 150 m	Y	Y	Y	Y	Y ^d
InfiniBand coupling links (1x IB-SDR or 1x IB-DDR) at an unrepeated distance of 10 km	Y	Y ^d	Y ^d	Y ^d	Y ^d
CFCC Level 16	Y	Y ^d	Y ^d	Y ^d	Y ^d
Assembler instruction mnemonics	Y	Y	Y ^d	Y ^d	Y ^d
C/C++ exploitation of z10 hardware instructions	Y	Y	Y ^d	Y ^d	N
Layer 3 VMAC	Y	Y	Y	Y ^d	N
Large dumps	Y	Y	Y ^d	Y ^d	N
CPU measurement facility	Y	Y	Y ^d	Y ^d	N

- a. With the announcement of IBM Lifecycle Extension for z/OS V1.8, fee-based corrective service can be ordered for up to two years after the withdrawal of service for z/OS V1R8. Similarly, IBM Lifecycle Extension for z/OS V1.7 provides fee-based support for z/OS V1.7 up to September 2010.
- b. The level of decimal floating-point exploitation varies with z/OS release and PTF level.
- c. FMIDs are shipped in a Web deliverable.
- d. PTFs are required.
- e. Requires Web deliverable support for zIIP.
- f. Support varies with operating system and level. See "FICON Express8" on page 89 for details.
- g. CHPID type OSN does not use ports. LPAR-to-LPAR communication is used.

4.2.2 z/VM

At general availability, z/VM V5R4 and later provide exploitation support and z/VM V5R3 provides compatibility support only. Table 4-3 lists the z10 EC functions currently supported for z/VM releases. It uses the following conventions:

- Y** The function is supported.
- N** The function is not supported.

Table 4-3 z/VM support summary

Function	V6R1	V5R4	V5R3
z10 EC	Y	Y	Y ^a
Greater than 54 PUs for single system image	N ^b	N ^b	N ^b
Dynamic add of logical CPs	Y	Y	Y ^g
zAAP on zIIP	Y ^c	Y ^c	N
Large memory > 128 GB	Y ^d	Y ^d	Y ^d

Function	V6R1	V5R4	V5R3
Large page support	N ^e	N ^e	N ^e
Hardware decimal floating point	Y ^f	Y ^f	Y ^f
CPACF protected public key	N ^e	N ^e	N ^e
Enhanced CPACF	Y	Y ^g	Y ^g
Personal account numbers of 13 to 19 digits	Y ^f	Y ^f	Y ^f
Crypto Express3	Y ^f	Y ^f	Y ^f
Execute relative guest exploitation	Y ^f	Y ^f	Y ^f
Capacity provisioning	N ^e	N ^e	N ^e
HiperDispatch	N ^e	N ^e	N ^e
Restore subchannel facility	Y	Y	Y
HiperSockets multiple write facility	N ^e	N ^e	N ^e
High Performance FICON	N ^e	N ^e	N ^e
FICON Express8	Y ^h	Y ^h	Y ^h
OSA-Express QDIO data connection isolation for z/VM environments	Y	Y ^g	Y ^g
OSA-Express3 10 Gigabit Ethernet LR CHPID type OSD	Y	Y	Y
OSA-Express3 10 Gigabit Ethernet SR CHPID type OSD	Y	Y	Y
OSA-Express3 Gigabit Ethernet LX using four ports CHPID types OSD and OSN	Y	Y	Y ^g
OSA-Express3 Gigabit Ethernet LX using two ports CHPID types OSD and OSN	Y	Y	Y
OSA-Express3 Gigabit Ethernet SX using four ports CHPID types OSD and OSN	Y	Y	Y ^g
OSA-Express3 Gigabit Ethernet SX using two ports CHPID types OSD and OSN	Y	Y	Y
OSA-Express3 1000BASE-T Ethernet using four ports CHPID types OSC, OSD ⁱ , OSE, and OSN ^j	Y	Y	Y ^g
OSA-Express3 1000BASE-T Ethernet using two ports CHPID types OSC, OSD, OSE, and OSN ^j	Y	Y	Y
Dynamic I/O support for InfiniBand CHPIDs	Y	Y	Y
InfiniBand coupling links (1x IB-SDR or 1x IB-DDR) at an unrepeated distance of 10 km	N	N	N
CFCC Level 16	Y ^f	Y ^f	Y ^f

a. Compatibility support only: z/VM and guests are supported at the System z9 functionality level. No exploitation of new hardware unless otherwise noted.

b. A maximum of 32 PUs per system is supported by z/VM V5R3 and later. Guests can be configured with up to 64 virtual PUs.

- c. Available for z/OS on virtual machines without virtual zAAPs defined when the z/VM LPAR does not have zAAPs defined.
- d. 256 GB of central memory are supported by z/VM V5R3 and later. z/VM V5R3 and later support more than 1 TB of virtual memory in use for guests.
- e. Not available to guests.
- f. Supported for guest use only.
- g. PTFs are required.
- h. Support varies with operating system and level. See “FICON Express8” on page 89 for details.
- i. PTFs are required for CHPID type OSD.
- j. CHPID type OSN does not use ports, it uses LPAR-to-LPAR communication.

Notes: We recommend that the capacity of z/VM logical partitions and any guests, in terms of the number of IFLs and CPs, real or virtual, be adjusted in face of the PU capacity of the z10 EC.

4.2.3 z/VSE

Table 4-4 lists z10 EC support requirements for the currently supported z/VSE releases. It uses the following conventions:

- Y** The function is supported.
- N** The function is not supported.

Table 4-4 z/VSE support summary

Function	V4R2	V4R1
z10 EC	Y ^a	Y ^a
CPACF protected public key	N	N
Enhanced CPACF	Y	Y
Crypto Express3	Y ^b	N
FICON Express8	Y ^c	Y ^c
OSA-Express3 10 Gigabit Ethernet LR CHPID type OSD	Y	Y
OSA-Express3 10 Gigabit Ethernet SR CHPID type OSD	Y	Y
OSA-Express3 Gigabit Ethernet LX using four ports CHPID types OSD	Y ^e	Y ^e
OSA-Express3 Gigabit Ethernet LX using two ports CHPID types OSD and OSN	Y	Y
OSA-Express3 Gigabit Ethernet SX using four ports CHPID types OSD	Y ^e	Y ^e
OSA-Express3 Gigabit Ethernet SX using two ports CHPID types OSD and OSN	Y	Y
OSA-Express3 1000BASE-T Ethernet using four ports CHPID types OSC, OSD ^e , OSE, and OSN ^d	Y	Y ^e
OSA-Express3 1000BASE-T Ethernet using two ports CHPID types OSC, OSD ^e , OSE, and OSN ^d	Y	Y

- a. z/VSE V4 is designed to exploit z/Architecture, specifically 64-bit real-memory addressing, but does not support 64-bit virtual memory addressing.
- b. PTFs are required.
- c. Support varies with operating system and level. See “FICON Express8” on page 89 for details.
- d. CHPID type OSN does not use ports. All communication is LPAR to LPAR.
- e. Exploitation of two ports per CPHID type OSD requires a minimum of z/VSE V4R1 with PTFs.

4.2.4 Linux on System z

Linux on System z distributions are built separately for the 31-bit and 64-bit addressing modes of the z/Architecture. The newer distribution versions are built for 64-bit only. You can run 31-bit applications in the 31-bit emulation layer on a 64-bit Linux on System z distribution.

None of the current versions of Linux on System z distributions (SLES 9, SLES 10, SLES 11, RHEL 4, RHEL 5) require z10 toleration support, so that any release of these distributions can run on System z10 servers.

Table 4-5 lists the most recent service levels of the current SUSE and Red Hat releases at the time of writing.

Table 4-5 Current Linux on System z distributions as of October 2009, by z/Architecture mode

Linux distribution	ESA/390 (31-bit mode)	z/Architecture (64-bit mode)
Novell SUSE SLES 11	No	Yes
Novell SUSE SLES 10 SP3	No	Yes
Novell SUSE SLES 9 SP4	Yes	Yes
Red Hat RHEL 5.4	No	Yes
Red Hat RHEL 4.8	Yes	Yes

Table 4-6 lists selected System z10 features, showing the minimum level of Novell SUSE and Red Hat distributions that support each feature.

Table 4-6 Linux on System z support summary

Function	Novell SUSE	Red Hat
z10 EC	SLES 9, SLES 10	RHEL 4, RHEL 5
Large page support	SLES 10 SP2	RHEL 5.3
Hardware decimal floating point	SLES 11	N ^b
CPACF protect public key	N	N
Enhanced CPACF	SLES 10 SP2	RHEL 5.3
Crypto Express3	SLES 10 SP3 ^a	RHEL 5.4 ^a
HiperSockets Layer 2 support	SLES 10 SP2	RHEL 5.3
FICON Express8	SLES 9, SLES 10	RHEL 4, RHEL 5
High Performance FICON	Note ^b	Note ^b

Function	Novell SUSE	Red Hat
FICON Express4 ^b CHPID type FCP	SLES 9, SLES 10	RHEL 4, RHEL 5
OSA-Express3 using four ports CHPID type OSD	SLES 10	RHEL 4, RHEL 5.2
OSA-Express3 using two ports CHPID type OSD	SLES 9, SLES 10 SP2	RHEL 4, RHEL 5
OSA-Express3 CHPID type OSN	SLES 9 SP2, SLES 10	RHEL 4.3, RHEL 5

a. Toleration support only.

b. FICON Express4 10KM LX, 4KM LX, and SX features are withdrawn from marketing. All FICON Express2 and FICON features are withdrawn from marketing

IBM is working with its Linux distribution partners so that exploitation of further z10 EC functions will be provided in future Linux on System z distribution releases. We recommend that:

- ▶ SUSE SLES 11 or Red Hat RHEL 5 be used in any new projects for the z10 EC
- ▶ Any Linux distributions be updated to their latest service level before migration to z10 EC
- ▶ The capacity of any z/VM and Linux logical partitions, as well as any z/VM guests, in terms of the number of IFLs and CPs, real or virtual, be adjusted in face of the PU capacity of the z10 EC

4.2.5 TPF and z/TPF

Table 4-7 lists the z10 EC function' support requirements for the currently supported TPF and z/TPF releases. It uses the following conventions:

Y The function is supported.

N The function is not supported.

Table 4-7 TPF and z/TPF support summary

Function	z/TPF V1R1	TPF V4R1
z10 EC	Y	Y
Greater than 54 PUs for single system image	Y	N
Large memory > 128 GB (4 TB)	Y	N
CPACF protected public key	N	N
Enhanced CPACF	Y	N
Crypto Express3 (accelerator mode only)	Y	N
HiperDispatch	N	N
FICON Express8	Y ^a	Y
OSA-Express3 10 Gigabit Ethernet LR CHPID type OSD	Y	Y ^b
OSA-Express3 10 Gigabit Ethernet SR CHPID type OSD	Y	Y ^b

Function	z/TPF V1R1	TPF V4R1
OSA-Express3 Gigabit Ethernet LX using four ports CHPID types OSD and OSN ^c	Y ^b	N
OSA-Express3 Gigabit Ethernet LX using two ports CHPID types OSD and OSN ^c	Y ^b	Y ^b
OSA-Express3 Gigabit Ethernet SX using four ports CHPID types OSD and OSN ^c	Y ^b	N
OSA-Express3 Gigabit Ethernet SX using two ports CHPID types OSD and OSN ^c	Y	Y ^b
OSA-Express3 1000BASE-T Ethernet using four ports CHPID types OSC, OSD, and OSN ^c	Y ^b	N
OSA-Express3 1000BASE-T Ethernet using two ports CHPID types OSC, OSD, and OSN ^c	Y ^b	Y ^b
Coupling over InfiniBand CHPID type CIB	Y ^d	Y ^d
CFCC Level 16	Y	Y

- a. See "FICON Express8" on page 89 for details.
- b. PTFs are required.
- c. CHPID type OSN does not use ports. It uses LPAR-to-LPAR communication.
- d. Compatibility is supported.

4.3 Support for selected functions

In this section we review the operating system support of a very small set of functions introduced by the z10 EC. They were selected due to their high importance.

4.3.1 Single system image

A single system image can control several processing units (PUs) such as central processors CPs, zIIPs, zAAPs, and IFLs, as appropriate. See "PU characterization" on page 10 for a description.

Table 4-8 shows the maximum number of PUs supported for each operating system image.

Table 4-8 Single system image software support

Operating system	Maximum number of (CPs+zIIPs+zAAPs) ^a or IFLs per system image
z/OS V1R11	64
z/OS V1R10	64
z/OS V1R9	64
z/OS V1R8	32
z/OS V1R7	32
z/VM V6R1	24 ^{b,c}
z/VM V5R4	32 ^{b,c}

Operating system	Maximum number of (CPs+zIIPs+zAAPs) ^a or IFLs per system image
z/VM V5R3	32 ^b
z/VSE V4	z/VSE Turbo Dispatcher can exploit up to four CPs and tolerates up to 10-way LPARs
Linux on System z	Novell SUSE SLES 11: 64 CPs or IFLs Novell SUSE SLES 10: 64 CPs or IFLs Novell SUSE SLES 9: 64 CPs or IFLs Red Hat RHEL 5: 64 CPs or IFLs Red Hat RHEL 4: 8 CPs or IFLs
z/TPF V1R1	64 CPs
TPF V4R1	16 CPs

a. The number of purchased zAAPs and the number of purchased zIIPs cannot each exceed the number of purchased CPs. A logical partition can be defined with any number of the available zAAPs and zIIPs. The total refers to the sum of these PU characterizations.

b. z/VM guests can be configured with up to 64 virtual PUs.

c. The z/VM-mode LPAR supports CPs, zAAPs, zIIPs, IFLs, and ICFs.

4.3.2 z/VM-mode LPAR

System z10 supports a new logical partition (LPAR) mode, named z/VM-mode, exclusively for running z/VM. This new LPAR mode requires z/VM V5R4 or later and allows z/VM to utilize a wider variety of specialty processors in a single LPAR. For instance, in a z/VM-mode LPAR, z/VM can manage Linux on System z guests running on IFL processors while also managing z/VSE and z/OS on CPs and allowing z/OS to fully exploit zIIPs and zAAPs.

4.3.3 Dynamic PU exploitation

z/OS has long been able to define reserved PUs to an LPAR for the purpose of nondisruptively bringing online the additional computing resources when needed. This no longer requires pre-planning.

z/OS V1R10 and z/VM V5R4 offer a similar capability. No pre-planning is required. The ability to dynamically define and change the number and type of reserved PUs in an LPAR profile can be used for that purpose. The new resources are immediately made available to the operating systems and, in the z/VM case, to its guests.

4.3.4 zAAP on zIIP capability

This new capability, exclusive to System z10 and System z9 servers under defined circumstances, enables workloads eligible to run on Application Assist Processors (zAAPs) to run on Integrated Information Processors (zIIP). It is intended as a means to optimize the investment on existing zIIPs, not as a replacement for zAAPs. The rule of at least one CP installed per zAAP and zIIP installed still applies. Exploitation of this capability is by z/OS only, and is only available in these situations:

- ▶ When there are zIIPs but no zAAPs installed in the server.
- ▶ When z/OS is running as a guest of z/VM V5R4 or later, and there are no zAAPs defined to the z/VM LPAR. The server may have zAAPs installed. Because z/VM can dispatch both virtual zAAPs and virtual zIIPs on real CPs¹, the z/VM partition does not require any real zIIPs defined to it, although we recommend the use of real zIIPs due to software licensing reasons.

Table 4-9 summarizes this support.

Table 4-9 Availability of zAAP on zIIP support

		z/OS is running on an LPAR ^a	z/OS is running as a z/VM guest		
			z/VM LPAR has zAAPs defined	No zAAPs defined to z/VM LPAR	
				Virtual zAAPs defined for z/OS guest	No virtual zAAPs for z/OS guest ^b
zAAPs installed on the server	Yes	No	No	No	Yes
	No	Yes	Not valid	No	Yes

a. zIIPs must be defined to the z/OS LPAR.

b. Virtual zIIPs must be defined to the z/OS virtual machine.

4.3.5 Large memory

Table 4-10 lists the maximum amount of main storage supported by current operating systems. Expanded storage, although part of the z/Architecture, is currently exploited only by z/VM. On z10 EC a maximum of 1 TB of main storage can be defined to a logical partition.

Table 4-10 Maximum memory supported by operating system

Operating system	Maximum supported main storage
z/OS	z/OS V1R11 supports 4 TB and up to 1.5 TB per server ^a z/OS V1R10 supports 4 TB and up to 1.5 TB per server ^a z/OS V1R9 supports 4 TB and up to 1.5 TB per server ^a z/OS V1R8 supports 4 TB and up to 1.5 TB per server ^a z/OS V1R7 supports 128 GB
z/VM	z/VM V6R1 supports 256 GB z/VM V5R4 supports 256 GB z/VM V5R3 supports 256 GB
z/VSE	z/VSE V4R2 supports 32 GB z/VSE V4R1 supports 8 GB
Linux on System z (64-bit)	Novell SUSE SLES 11 supports 4 TB Novell SUSE SLES 10 supports 4 TB Novell SUSE SLES 9 supports 4 TB Red Hat RHEL 5 supports 64 GB Red Hat RHEL 4 supports 64 GB
TPF and z/TPF	z/TPF supports 4 TB ^a TPF runs in ESA/390 mode and supports 2 GB

a. System z10 EC restricts the LPAR memory size to 1 TB.

4.3.6 Dynamic LPAR memory upgrade

A logical partition can be defined with both an initial and a reserved amount of memory. At activation time the initial amount is made available to the partition and the reserved amount can later be added, partially or totally. Those two memory zones do not have to be contiguous

¹ The z/VM system administrator can use the SET CPUAFFINITY command to influence the dispatching of virtual specialty engines on CPs or real specialty engines.

in real memory but appear as logically contiguous to the operating system running in the LPAR.

Until now, only z/OS was able to take advantage of this support by nondisruptively acquiring and releasing memory from the reserved area. z/VM V5R4 and later are able to acquire memory nondisruptively and immediately make it available to guests. z/VM virtualizes this support to its guests, which can also increase their memory nondisruptively. Releasing memory is still a disruptive operation.

4.3.7 Hardware decimal floating point

Industry support for decimal floating point is growing, with IBM leading the open standard definition. Examples of support for the draft standard IEEE 754r include Java BigDecimal, C#, XML, C/C++, GCC, COBOL, and other key software vendors such as Microsoft® and SAP.

Hardware decimal floating point support was introduced with the z9 EC. The z10 EC, however, has a new decimal floating point accelerator feature, described in the *z10 EC IBM System z10 Enterprise Class Technical Guide*, SG24-7516.

Table 4-11 lists the operating system support for decimal floating point. See also “Decimal floating point (z/OS XL C/C++ considerations)” on page 95.

Table 4-11 Minimum support requirements for decimal floating point

Operating system	Support requirements
z/OS	<ul style="list-style-type: none"> ▶ z/OS V1R9: Support includes XL, C/C++, HLASM, Language Environment®, DBX, and CDA RTLE. ▶ z/OS V1R8: Support includes HL ASM, Language Environment, DBX, and CDA RTLE. ▶ z/OS V1R7: Support of the High Level Assembler (HLASM) only.
z/VM	z/VM V5R3: Support for guest use.
Linux	Novel SUSE SLES 11.

4.3.8 High Performance FICON for System z10

High Performance FICON for System z10 (zHPF) is a new FICON architecture for protocol simplification and efficiency, reducing the number of information units (IUs) processed. Enhancements have been made to the z/Architecture and the FICON interface architecture to provide optimizations for online transaction processing (OLTP) workloads.

When exploited by the FICON channel, the z/OS operating system, and the control unit (new levels of Licensed Internal Code are required), the FICON channel overhead can be reduced and performance can be improved. Additionally, the changes to the architectures provide end-to-end system enhancements to improve reliability, availability, and serviceability (RAS). Table 4-12 lists the minimum support requirements for zHPF.

Table 4-12 Minimum support requirements for zHPF

Operating system	Support requirements
z/OS	z/OS V1R8 with PTFs.
z/VM	Not supported. Not available to guests.
Linux	IBM is working with its Linux distribution partners so that exploitation of appropriate z10 BC functions be provided in future Linux on System z distribution releases.

zHPF channel programs can be exploited by z/OS OLTP I/O workloads DB2, VSAM, PDSE, and zFS, which transfer small blocks of fixed size data (4 K blocks). zHPF implementation, along with matching support by the DS8000 series, provides support for I/Os that transfer less than a single track of data as well as multitrack operations.

For more information about FICON channel performance, see the performance technical papers on the System z I/O Connectivity Web site at:

http://www-03.ibm.com/systems/z/hardware/connectivity/ficon_performance.html

zHPF is exclusive to System z10. The FICON Express8, FICON Express4², and FICON Express2 features (CHPID type FC) support both the existing FICON protocol and the zHPF protocol concurrently in the server Licensed Internal Code.

FICON Express8

FICON Express8 is the newest generation of FICON features. They provide a link rate of 8 Gbps, with autonegotiation to 4 or 2 Gbps, for compatibility with previous devices and investment protection. Both 10KM LX and SX connections are offered (in a given feature all connections must have the same type).

With FICON Express8 customers may be able to consolidate existing FICON, FICON Express2, and FICON Express4 channels while maintaining and enhancing performance.

² FICON Express4 10KM LX, 4KM LX, and SX features are withdrawn from marketing. All FICON Express2 and FICON features are withdrawn from marketing.

Table 4-13 lists the minimum support requirements for FICON Express8.

Table 4-13 Minimum support requirements for FICON Express8

Operating system	z/OS	z/VM	z/VSE	Linux on System z	z/TPF	TPF
Native FICON and Channel-to-Channel (CTC) CHPID type FC	V1R7	V5R3	V4R1	SUSE SLES 9 RHEL 4	V1R1	V4R1 PUT 16
zHPF single track operations CHPID type FC	V1R7 ^a	NA	NA	NA	NA	NA
zHPF multitrack operations CHPID type FC	V1R9 ^a	NA	NA	NA	NA	NA
Support of SCSI devices CHPID type FCP	NA	V5R3	V4R1	SUSE SLES 9 RHEL 4	NA	NA

a. PTFs required.

MIDAW facility

The Modified Indirect Data Address Word (MIDAW) facility is a system architecture and software exploitation designed to improve FICON performance. This facility is available only on System z9 and System z10 servers and is exploited by the media manager in z/OS. Table 4-14 lists the minimum support requirements for the MIDAW facility.

Table 4-14 Minimum support requirements for MIDAW

Operating system	Support requirements
z/OS	z/OS V1R7
z/VM	z/VM V5R3 ^a

a. Supported for guest exploitation

The MIDAW facility provides a more efficient structure for certain categories of data-chaining I/O operations:

- ▶ MIDAW can significantly improve FICON performance for extended format (EF) data sets. Non-extended data sets can also benefit from MIDAW.
- ▶ MIDAW can improve channel utilization and can significantly improve I/O response time. This reduces FICON channel connect time, director ports, and control unit overhead.

From IBM laboratory tests it is expected that applications that use EF data sets (such as DB2, or long chains of small blocks) gain significant performance benefits using the MIDAW facility.

For more information about FICON, FICON channel performance, and MIDAW, see the I/O Connectivity Web page:

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

An excellent paper called *How does the MIDAW Facility Improve the Performance of FICON Channels Using DB2 and other workloads?*, REDP-4201, is available at:

<http://www.redbooks.ibm.com/redpapers/pdfs/redp4201.pdf>

Also see *IBM TotalStorage DS8000 Series: Performance Monitoring and Tuning*, SG24-7146.

4.3.9 Cryptographic support

z10 EC provides two major groups of cryptographic functions, which, from an application program perspective, are both synchronous and asynchronous cryptographic functions:

- ▶ Synchronous cryptographic functions are provided by the CP Assist for Cryptographic Function (CPACF).
- ▶ Asynchronous cryptographic functions are provided by the Crypto Express features.

The minimum software support levels are listed in the following sections. The latest PSP buckets should be obtained and reviewed to ensure that the latest support levels are known and included as part of the implementation plan.

CPACF

In z10 EC, the CP Assist for Cryptographic Function was extended to support the full standard for AES (symmetric encryption) and SHA (hashing). For a full description refer to the *z10 EC IBM System z10 Enterprise Class Technical Guide*, SG24-7516. Support for this function is provided through Web-delivered code. Table 4-15 lists the support requirements for CPACF enhancements.

Table 4-15 Support requirements for CPACF enhancements

Operating system	Support requirements
z/OS ^a	z/OS V1R7: The function varies by release. Protected public key requires z/OS V1R9 and later plus PTFs.
z/VM	z/VM V5R3 and later: Supported for guest use, but protected key is not supported.
z/VSE	z/VSE V4R1 and later, and IBM TCP/IP for VSE/ESA V1R5 with PTFs.
Linux on System z	Novell SUSE SLES 9 SP3, SLES 10. and SLES 11. Red Hat RHEL 4.3 and RHEL 5. The z10 EC CPACF enhancements can be used with: <ul style="list-style-type: none"> ▶ Novell SUSE SLES 10 SP2 and SLES 11 ▶ Red Hat RHEL 5.2
TPF and z/TPF	TPF V4R1 and z/TPF V1R1

a. CPACF is also exploited by several IBM software product offerings for z/OS, such as IBM WebSphere Application Server for z/OS.

Crypto Express3 and Crypto Express2

Support of Crypto Express3 and Crypto Express2 functions varies by operating system and release. Table 4-16 lists the software requirements for the Crypto Express3 and Crypto Express2 features when configured as a coprocessor or an accelerator and support for the base and enhanced functions. Several functions require software support, which can be downloaded from the Web (see “Web deliverables” on page 92). For integrated cryptographic information see 4.3.10, “z/OS ICSF” on page 92.

Table 4-16 *Crypto Express2 and Crypto Express3 support on z10 EC*

Operating system	Crypto Express3	Crypto Express2
z/OS	V1R11: Web deliverable V1R10: Web deliverable V1R9: Web deliverable V1R8: not supported V1R7: not supported	V1R11: included in base V1R10: included in base V1R9: included in base V1R8: included in base V1R7: Web deliverable
z/VM	V5R3: service required; supported for guest use only	V5R3: supported for guest use only.
z/VSE	V4R2 with IBM TCP/IP for VSE/ESA V1R5; service required	V4R1 with IBM TCP/IP for VSE/ESA V1R5; service required
Linux on System z	Note ^a Novell SUSE SLES 11 Novell SUSE SLES 10 SP3 Red Hat RHEL 5.4	Novell SUSE SLES 11 Novell SUSE SLES 10 Novell SUSE SLES 9 SP3 Red Hat RHEL 5.1 Red Hat RHEL 4.4
TPF V4R1	Not supported	Not supported
z/TPF V1R1	Service required (accelerator mode only)	Service required (accelerator mode only)

a. Support for Crypto Express3 is provided at the same functional level as for Crypto Express2.

Web deliverables

For z/OS downloads see the z/OS Web site:

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

4.3.10 z/OS ICSF

Integrated Cryptographic Service Facility (ICSF) is a base component of z/OS and is designed to transparently use the available cryptographic functions, whether CPACF or Crypto Express features, to balance the workload and help address the bandwidth requirements of the customer's applications.

Specific support is available as a Web download for z/OS V1R7 (FMID HCR7730) and for z/OS V1R8 (FMID HCR7731) in support of the cryptographic coprocessor and accelerator functions as well as the CPACF AES, PRNG, and SHA support. The z/OS V1R9 has this support (FMID HCR7740) integrated in the base, so no download is necessary.

For support of the SHA-384 and SHA-512 function on z/OS V1R7 and later, download and installation of FMID HCR7750 is required.

Support for the most recent functions, which include Secure Key AES, new Crypto Query Service, enhanced IPv6 support, and enhanced SAF Checking and Personal Account

Numbers with 13 to 19 digits, is provided by FMID HCR7751, which is available for z/OS V1R8 and later.

Support for the Crypto Express3, Crypto Express3-1P, and CPACF protected key is provided for z/OS V1R9 and later by FMID HCR7770. Planned availability for FMID HCR7770 is November 2009.

ICSF considerations

Consider the following points regarding the version of Web-delivered ICSF code:

- ▶ Increased size of the PKDS file: This is required to allow 4096-bit RSA keys to be stored. If you use the PKDS for asymmetric keys you must copy your PKDS to a larger VSAM data set before using the new version of ICSF. The ICSF options file must be updated with the name of the new data set. ICSF can then be started.

A toleration PTF must be installed on any system that is sharing the PKDS with a system running HCR7750 ICSF. The PTF allows the PKDS to be larger and prevents any service from accessing 4096-bit keys stored in a HCR7750 PKDS.

- ▶ Reduced support for retained private keys. Applications that make use of the retained private key capability for key management will no longer be able to store the private key in the crypto coprocessor card. The applications will continue to be able to list the retained keys and to delete them from the crypto coprocessor cards.

4.4 z/OS considerations

z10 EC base processor support is required in z/OS. With that exception, software changes do not require the new z10 EC functions and, equally, the new functions do not require functional software. The approach has been to, where applicable, automatically decide to enable or disable a function based on, respectively, the presence or absence of the required hardware and software.

General recommendations

The new System z10 EC introduces the latest System z technology. Although support is provided by z/OS starting with z/OS V1R7, exploitation of z10 EC is dependent on the z/OS release. z/OS.e is *not* supported on z10 EC.

In general, we recommend that you:

- ▶ Do not migrate software releases and hardware at the same time.
- ▶ Keep members of the sysplex at the same software level other than during brief migration periods.
- ▶ Review z10 EC restrictions and considerations prior to creating an upgrade plan.

HCD

When using HCD on z/OS V1R6 to create a definition for z10 EC, *all* subchannel sets must be defined or the VALIDATE will fail. On z/OS V1R7, HCD or HCM will assist in the definitions.

InfiniBand coupling links

Each system can use, or not use, InfiniBand coupling links independently of what other systems are doing, and do so in conjunction with other link types.

InfiniBand coupling connectivity is only available when other systems also support InfiniBand coupling. We recommend that you consult the *Coupling Facility Configuration Options* white paper when planning to exploit the InfiniBand coupling technology, available at:

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

Large page support

The large page support function is not enabled without the required software support. Without the large page support, page frames are allocated at the current 4 KB size.

Memory reserved for large page support can be defined with the following new parameter in the IEASYSxx member of SYS1.PARMLIB:

```
LFAREA=xx%|xxxxxxM|xxxxxxG
```

This parameter *cannot* be changed dynamically.

HiperDispatch

There is a new HIPERDISPATCH=YES/NO parameter in the IEAOPTxx member of SYS1.PARMLIB and on the SET OPT=xx command to control whether HiperDispatch is enabled or disabled for a z/OS image. It can be changed dynamically (without an IPL or any outage).

To effectively exploit HiperDispatch, adjustment of defined WLM goals and policies may be required. We recommend that you review WLM policies and goals and update them as necessary. You may want to run with the new policies and HiperDispatch on for a period, turn it off and use the older WLM policies while analyzing the results of using HiperDispatch, re-adjust the new policies, and repeat the cycle as needed. In order to change WLM policies, turning HiperDispatch off then on is not necessary.

A health check is provided to verify whether HiperDispatch is enabled on a z10 EC system.

Capacity provisioning

Installation of the capacity provision function on z/OS requires:

- ▶ Setting up and customizing z/OS RMF, including the Distributed Data Server (DDS)
- ▶ Setting up the z/OS CIM Server (a z/OS base element with z/OS V1R9)
- ▶ Performing capacity provisioning customization as described in the *z/OS MVS Capacity Provisioning User's Guide*, SA33-8299

Exploitation of the capacity provisioning function requires:

- ▶ TCP/IP connectivity to observed systems
- ▶ TCP/IP connectivity from the observing system to the HMC of observed systems
- ▶ RMF Distributed Data Server must be active
- ▶ CIM Server must be active
- ▶ Security and CIM customization
- ▶ Capacity Provisioning Manager customization

In addition, the Capacity Provisioning Control Center must be downloaded from the host and installed on a PC workstation. This application is only used to define policies. It is not required to manage operations.

Customization of the capacity provisioning function is required on the operating system that will observe other z/OS systems in one or multiple sysplexes. For a description of the capacity provisioning domain refer to the *z/OS MVS Capacity Provisioning User's Guide*, SA33-8299.

Also see *IBM System z10 Enterprise Class Capacity on Demand*, SG24-7504, which discusses capacity provisioning in more detail.

Decimal floating point (z/OS XL C/C++ considerations)

The two new options for the C/C++ compiler are ARCHITECTURE and TUNE. They require z/OS V1R9.

The ARCHITECTURE C/C++ compiler option selects the minimum level of machine architecture on which your program will run. Note that certain features provided by the compiler require a minimum architecture level. ARCH(8) exploits instructions available on the z10 EC.

The TUNE compiler option allows optimization of the application for a specific machine architecture within the constraints imposed by the ARCHITECTURE option. The TUNE level must not be lower than the setting in the ARCHITECTURE option.

For more information about the ARCHITECTURE and TUNE compiler options refer to the *z/OS V1R9.0 XL C/C++ User's Guide*, SC09-4767. See the Authorized Program Analysis Report, APAR PK60051, which provides guidance to installation of the z/OS V1.9 XL C/C++ compiler on a z/OS V1.8 system.

Note: A C/C++ program compiled with the ARCH() or TUNE() options run only on z10 EC servers, otherwise an operation exception will result. This is a consideration for programs that might have to run on different level servers during development, test, production, and fallback or DR.

4.5 Coupling Facility and CFCC considerations

Coupling Facility connectivity to a z10 EC server is supported on the z10 BC, z9 EC, z9 BC, z990, z890, or another z10 EC server. The logical partition running the Coupling Facility Control Code (CFCC) can reside on any of the supported servers previously listed.

Because coupling link connectivity to z800 and z900 is *not* supported, this might affect the introduction of z10 EC into existing installations and require additional planning. For more information refer to the *IBM System z10 Enterprise Class Technical Guide*, SG24-7516.

System z servers support CFCC Level 16. To support migration from one CFCC level to the next, different levels of CFCC can be run concurrently as long as the Coupling Facility logical partitions are running on different servers. (CF logical partitions running on the same server share the same CFCC level.)

For additional details about CFCC code levels, see the Parallel Sysplex Web site at:

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

4.6 IOCP

All System z servers require a description of their I/O configuration. This description is stored in input/output configuration data set (IOCDs) files. The input/output configuration program (IOCP) allows creation of the IOCDs file from a source file known as the input/output configuration source (IOCS).

The IOCS file contains detailed information for each channel and path assignment, each control unit, and each device in the configuration.

The required level of IOCP for the z10 EC is V2 R1 L0 (IOCP 2.1.0). See the *Input/Output Configuration Program User's Guide*, SB10-7037, for details.

4.7 Worldwide portname (WWPN) prediction tool

A part of the installation of your IBM System z10 server is the pre-planning of the Storage Area Network (SAN) environment. IBM has made available a standalone tool to assist with this planning prior to the installation.

The tool, known as the worldwide port name prediction tool, assigns WWPNs to each virtual Fibre Channel Protocol (FCP) channel/port using the same WWPN assignment algorithms that a system uses when assigning WWPNs for channels utilizing N_Port Identifier Virtualization (NPIV). Thus, the SAN can be set up in advance, allowing operations to proceed much faster once the server is installed.

The WWPN prediction tool takes a .csv file containing the FCP-specific I/O device definitions and creates the WWPN assignments that are required to set up the SAN. A binary configuration file that can be imported later by the system is also created. The .csv file can either be created manually or exported from the Hardware Configuration Definition/Hardware Configuration Manager (HCD/HCM).

The WWPN prediction tool on System z10 (CHPID type FCP) requires at a minimum:

- ▶ z/OS V1R8, V1R9, and V1R10 with PTFs
- ▶ z/VM V5R3, V5R4 with PTFs and V6R1

The WWPN prediction tool is available for download at the Resource Link and is applicable to all FICON channels defined as CHPID type FCP (for communication with SCSI devices) on System z10. See:

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

4.8 ICKDSF

The ICKDSF Release 17 device support facility is required on all systems that share disk subsystems with a z10 EC server.

ICKDSF supports a modified format of the CPU information field, which contains a 2-digit logical partition identifier. ICKDSF uses the CPU information field instead of CCW reserve/release for concurrent media maintenance. It prevents multiple systems from running ICKDSF on the same volume, and at the same time allows user applications to run while ICKDSF is processing. In order to prevent any possible data corruption, ICKDSF must be able to determine all sharing systems that can potentially run ICKDSF. Therefore, this support is required for the z10 EC.

Important: The need for ICKDSF Release 17 applies even to systems that are not part of the same sysplex or that are running other than the z/OS operating system, such as z/VM.

4.9 Software licensing considerations

The System z10 EC mainframe IBM software portfolio includes operating system software (that is, z/OS, z/VM, z/VSE, and z/TPF) and middleware that runs on these operating systems. It also includes middleware for Linux on System z environments.

Two major metrics for software licensing are available from IBM, depending on the software product:

- ▶ Monthly License Charge (MLC)
- ▶ International Program License Agreement (IPLA)

MLC pricing metrics have a recurring charge that applies each month. In addition to the right to use the product, the charge includes access to IBM product support during the support period. MLC metrics, in turn, include a variety of offerings. Those applicable to the System z10 EC are:

- ▶ Workload License Charges (WLC)
- ▶ System z New Application License Charges (zNALC)
- ▶ Parallel Sysplex License Charges (PSLC)
- ▶ Midrange Workload License Charges (MWLC)

IPLA metrics have a single, up-front charge for an entitlement to use the product. An optional and separate annual charge called *subscription and support* entitles customers to access IBM product support during the support period and also receive future releases and versions at no additional charge. For details, consult the *IBM System z Software Pricing Reference Guide*, Web page:

http://www-03.ibm.com/servers/eserver/zseries/library/refguides/sw_pricing.html

4.9.1 Workload License Charges (WLC)

Workload License Charges require z/OS or z/TPF operating systems in 64-bit mode. Any mix of z/OS, z/VM, Linux, z/VSE, TPF, and z/TPF images is allowed.

The two WLC license types are:

- ▶ Flat WLC (FWLC): Software products licensed under FWLC are charged at the same flat rate, no matter what is the capacity of (MSUs) the server.
- ▶ Variable WLC (VWLC): This type applies to products such as z/OS, DB2, IMS, CICS, MQSeries®, and Lotus® Domino®. VWLC software products can be charged as:
 - Full-capacity: The server's total number of MSUs is used for charging. Full-capacity is applicable when the server is not eligible for sub-capacity.
 - Sub-capacity: Software charges are based on the logical partition's usage where the product is running.

WLC sub-capacity allows software charges based on logical partition usage instead of the server's total number of MSUs. Sub-capacity removes the dependency between software charges and server (hardware) installed capacity.

Sub-capacity is based on the logical partition's rolling 4-hour average usage. It is *not* based on the usage of each product,³ but on the usage of the logical partitions where it runs. The VWLC licensed products running on a logical partition will be charged by the maximum value of this partition's rolling 4-hour average usage within a month.

³ With the exception of products licensed using the SALC pricing metric

The logical partition's rolling 4-hour average usage can be limited by a *defined capacity* definition on the partition's image profiles. This activates the *soft capping* function of PR/SM, limiting 4-hour average partition usages above the defined capacity value. Soft capping controls the maximum rolling 4-hour average usage (the *last* 4-hour average value at every 5-minute interval), but does *not* control the maximum *instantaneous* partition use.

Also available is an LPAR group capacity limit, which allows you to set soft capping of PR/SM for a group of logical partitions running z/OS.

Even using the soft capping option, the partition's use can reach up to its maximum share based on the number of logical processors and weights in the image profile. Only the rolling 4-hour average use is tracked, allowing usage peaks above the defined capacity value.

As with the Parallel Sysplex License Charges (PSLC) software license charge type, the aggregation of the servers' capacities within the same Parallel Sysplex is also possible in WLC following the same prerequisites.

The Entry Workload License Charges (EWLC) charge type is not offered for IBM System z10 EC.

For further information about WLC and details about how to combine logical partitions usage, see the publication *z/OS Planning for Workload License Charges*, SA22-7506, available from:

http://www-03.ibm.com/systems/z/os/zos/bkserv/find_books.html

4.9.2 System z New Application License Charges (zNALC)

System z New Application License Charges offers a reduced price for the z/OS operating system on logical partitions running a qualified *new workload* application such as Java language business applications running under WebSphere Application Server for z/OS, Domino, SAP, PeopleSoft, and Siebel.

z/OS with zNALC provides a strategic pricing model available on the full range of System z servers for simplified application planning and deployment. zNALC allows for aggregation across a qualified Parallel Sysplex, which can provide a lower cost for incremental growth across new workloads that span a Parallel Sysplex.

For additional information see the zNALC Web page:

<http://www-03.ibm.com/servers/eserver/zseries/swprice/znalc.html>

4.9.3 Select Application License Charges (SALC)

Select Application License Charges applies to WebSphere MQ for System z only. It allows a WLC customer to license MQ under product utilization rather than the sub-capacity pricing provided under WLC.

WebSphere MQ is typically a low-usage product that runs pervasively throughout the environment. Clients who run WebSphere MQ at a very low usage may benefit from SALC. Alternatively, you can still choose to license WebSphere MQ under WLC.

A reporting function, which IBM provides in the operating system IBM Software Usage Report Program, is used to calculate the daily MSU number. The rules to determine the billable SALC MSUs for WebSphere MQ use the following algorithm:

1. Determines the highest daily usage of a program⁴ family, which is the highest of 24 hourly measurements recorded each day
2. Determines the monthly usage of a program² family, which is the fourth highest daily measurement recorded for a month
3. Uses the highest monthly usage determined for the next billing period

For additional information about SALC, see the Other MLC Metrics Web page:

<http://www.ibm.com/servers/eserver/zseries/swprice/other.html>

4.9.4 Midrange Workload Licence Charges

Midrange Workload Licence Charges (MWLC) applies to z/VSE V4 when running on System z10 and System z9 servers. The exceptions are the z10 BC and z9 BC servers at capacity setting A01 to which zELC applies.

Similarly to Workload Licence Charges, MWLC can be implemented in full-capacity or sub-capacity mode. MWLC applies to z/VSE V4 and several IBM middleware products for z/VSE. All other z/VSE programs continue to be priced as before.

The z/VSE pricing metric is independent of the pricing metric for other systems (for instance, z/OS) that might be running on the same server. When z/VSE is running as a guest of z/VM, z/VM V5R3 or later is required.

To report usage, the sub-capacity report tool is used. One SCRT report per server is required.

For additional information see the MWLC Web page:

<http://www.ibm.com/servers/eserver/zseries/swprice/mwlc.html>

4.9.5 System z International Program License Agreement (IPLA)

On the mainframe, the following types of products are generally in the IPLA category:

- ▶ Data management tools
- ▶ CICS tools
- ▶ Application development tools
- ▶ Certain WebSphere for z/OS products
- ▶ Linux middleware products
- ▶ z/VM Versions 5 and 6

Generally, three pricing metrics apply to IPLA products for System z:

VU Value unit pricing, which applies to the IPLA products that run on z/OS. Value unit pricing is typically based on the number of MSUs and allows for lower cost of incremental growth. Examples of eligible products are IMS tools, CICS tools, DB2 tools, application development tools, and WebSphere products for z/OS and OS/390®.

EBVU Engine-based value unit pricing enables a lower cost of incremental growth with additional engine-based licenses purchased. Examples of eligible products include

⁴ Program refers to all active versions of MQ

z/VM V5 and certain z/VM middleware, which are priced based on the number of engines.

PVU Processor value units. Here the number of engines is converted into processor value units under the Passport Advantage® terms and conditions. Most Linux middleware is also priced based on the number of engines.

For additional information see the System z IPLA Web page at:

<http://www.ibm.com/servers/eserver/zseries/swprice/zipla/>

4.10 References

Planning information for each operating system is available on the following support Web pages:

- ▶ z/OS
<http://www.ibm.com/systems/support/z/zos>
- ▶ z/VM
<http://www.ibm.com/systems/support/z/zvm>
- ▶ z/TPF
<http://www.ibm.com/tpf/maint/supportgeneral.htm>
- ▶ z/VSE
<http://www.ibm.com/servers/eserver/zseries/zvse/>
- ▶ Linux on System z
<http://www.ibm.com/systems/z/os/linux>

Frequently asked questions

Q: What is System z?

A: IBM System z is a brand name for IBM mainframe computers. It is the line of computers that started in 1964 with S/360 and evolved over the decades. It still preserves backward compatibility with previous systems while bringing new features and technologies.

Q: How does System z fit into the IBM Dynamic Infrastructure initiative?

A: In the IBM vision, a Dynamic Infrastructure drives a new scale of efficiency and service excellence for businesses, helping to align IT with business goals.

The Dynamic Infrastructure, being an evolutionary model for efficient IT delivery, provides a highly dynamic, efficient and shared environment which allows IT to better manage costs, improve service levels, improve operational performance and resiliency, and more quickly respond to business needs. Operational issues are addressed through consolidation, virtualization, energy offerings, and service management. The existence of virtualized resource pools for server platforms, storage systems, networks, and applications enables delivery IT to users in a more fluid way.

System z is considered to be the most robust, secure, and virtualized platform in the industry. The z10 server introduces just-in-time deployment of additional resources, known as Capacity on Demand (CoD). CoD provides flexibility, granularity, and responsiveness by allowing the user to dynamically change capacity when business requirements change. Considering the other elements of the stack (software and services), System z is an up-to-date evolutionary platform that truly is the cornerstone for implementation of a Dynamic Infrastructure.

Q: What security classifications do the System z servers have?

A: System z servers are certified at the highest security level in the industry: EAL 5 Common Criteria for Logical Partitions. System z10 EC received its certification on October 29, 2008, and System z10 BC on May 4, 2009. In addition, the following operating systems have received the EAL 4+ with CAPP and LSPP certification: z/OS 1.8 and later with RACF and z/VM 5.3 with RACF. Novell SUSE SLES 9, Novell SUSE SLES 10, and Red Hat RHEL 4 are certified for EAL 4+ with CAPP and Red Hat RHEL 5 has a EAL 4+ with CAPP e LSPP certification.

Q: Does z10 use the Power processor?

A: The processor in z10 was developed in cooperation with the Power6 team. Both processors share many components, such as IBM 65nm SOI technology, design building blocks, large portions of execution units, floating point units, EI3 interface technology, core pipeline design style, high frequency, and low latency, mostly in order instruction execution. However, they have different personalities, different cache hierarchy, SMP topology and protocol, and different chip organization. The z10 implements the Complex Instruction Set (CISC) of the z/Architecture and POWER® implements a RISC architecture. They are siblings, not identical twins.

Q: Processor clock speed has more than doubled when compared with System z9 EC. Why does the capacity in MIPS not show double growth?

A: The published processor capacity indexes (not MIPS) are comparisons of the capacity of a processor running a standard workload mix with a defined processor taken as the base. Clock speed is and has always been only one of many factors influencing the performance characteristics of a microprocessor in a given operating environment. Penalties to be paid by increasing clock speeds are partially offset by architecture, hardware, and operating system design improvements. The combination of these has led to the performance figures published for specific workloads. Variability of workload characteristics will show similar variability in performance characteristics.

Q: It is well known that problems are associated with high-frequency processors. Why do you still use them?

A: The classic scaling of transistor switching speeds as defined by Moore's law has slowed down because at ultra-high frequencies power leakage must be overcome. In the coming years slower clock speed increases must be expected while the industry as a whole is searching for and implementing techniques based on complex chip, processor, and system designs, while software development must concentrate on introducing more parallelism in their designs.

Q: What is new about PU sparing?

A: Sparing works the same way as in System z9 servers. There are two spare PUs per server, regardless of number of books. Sparing of a single core is now possible.

Q: What is HiperDispatch?

A: HiperDispatch is a name for several improvements in interaction between PR/SM and z/OS. It is a mechanism that recognizes the physical processor where the work was started and then dispatches subsequent work to the same physical processor. This helps to reduce the movement of cache and data and improves overall system throughput. HiperDispatch is available only with z10 PR/SM and z/OS functions.

Q: What are the consequences when I switch off HiperDispatch in z/OS?

A: For some workloads that are dispatched on many processors across multiple books it may decrease the performance because processor caches must be reloaded more often.

Q: Can I run Linux on z10?

A: Yes. Major Linux on System z distributions include Novell SUSE and Red Hat. IBM is working with these Linux distribution partners to provide Linux with appropriate functionality on all of its hardware platforms.

Q: Can I run AIX® on z10?

A: No, because no AIX version is available that can run on z10. AIX is designed for IBM Power Systems™ (System p and System i®).

Q: Can I run MS Windows on z10?

A: No, because no Windows operating system is available for z10 EC.

Q: Can I run Sun Solaris on z10?

A: No. There is no Solaris on z10 offering. However, as a result of a collaboration project between IBM and the Open Source community to investigate the feasibility of bringing OpenSolaris to System z, a distribution of OpenSolaris for System z is now available to run as a guest of z/VM. Note that IBM does not warrant, and is not responsible for support of, a non-IBM operating system. For the announcement details see:

http://www-01.ibm.com/common/ssi/rep_ca/5/897/ENUS108-875/ENUS108-875.PDF

Q: Can any z990 or z9 EC model be upgraded to a z10 EC?

A: Yes.

Q: Can I do all upgrades within the z10 concurrently?

A: Most upgrades are concurrent. For example, if memory is already installed in the book, enabling it with Licensed Internal Code (LIC) is a concurrent action. But if the memory must be installed physically first, the action may not be concurrent. It also depends on the number of available books and their configuration. With proper planning, the user may be able to avoid planned outages.

Q: What is the difference between concurrent and nondisruptive upgrade?

A: In general, concurrency addresses the continuity of operations of just the hardware part of an upgrade, for instance, whether a server (as a box) is required to be switched off during the upgrade. Disruptive versus nondisruptive refers to whether the running software or operating system must be restarted for the upgrade to take effect. Thus, even concurrent upgrades can be disruptive to those operating systems or programs that do not support them while at the same time being nondisruptive to others.

Q: What is meant by the plan-ahead memory function on the System z10?

A: Memory can be upgraded concurrently using LIC-CC if physical memory is available on the z10 server. The plan-ahead memory function provides the ability to plan for nondisruptive memory upgrades by having the system pre-plugged based on a target configuration. Pre-plugged memory will be enabled through a LIC-CC order placed by the customer.

Q: What is the difference between the plan-ahead memory and the flexible memory option?

A: Plan-ahead memory should not be confused with flexible memory support. Plan-ahead memory is for a permanent increase of installed memory, while flexible memory provides a temporary replacement of a part of memory that becomes unavailable.

Q: What is the benefit of 1 MB page size? Should I switch from 4 KB to 1 MB?

A: For workloads with large memory requirements, large pages cause the Translation Lookaside Buffer (TLB) to better represent the working set and suffer fewer misuses by allowing a single TLB entry to cover more address translations. Exploiters of large pages are better represented in the TLB and are expected to perform better. Long-running memory access-intensive applications especially benefit. Short processes with small working sets see little or no improvement. The use of large pages must be decided based on knowledge obtained from measurement of memory usage and page translation overhead for a specific workload.

Under z/OS, large pages are treated as fixed pages and are never paged out. They are only available for 64-bit virtual private storage such as virtual memory located above 2 GB.

IBM is working with its Linux distribution partners to include support in future Linux on System z distribution releases, where the benefits will be similar to those described above.

Q: How can I control and limit the costs related to On/Off CoD?

A: Customers can limit the financial exposure towards On/Off CoD and order an exact

amount of temporary On/Off CoD processing capacity in the form of capacity tokens. The actual balance of capacity tokens can be checked anytime.

Q: What is a capacity token?

A: A capacity token is a representation of resources available for a given period of time. The measurement units used are MSU days for CP capacity and specialty engine days for specialty engines. One MSU token is worth one MSU day and one specialty engine token is worth one specialty engine day.

Q: What is a pre-paid capacity?

A: Pre-paid capacity is temporary processing capacity that can be ordered, paid for, and kept in reserve for future consumption.

Q: What is CPE?

A: Capacity for Planned Event is a Capacity on Demand offering. It delivers a replacement capacity for a planned event-like data center planned outage, server move, and so on. When it is activated, a specific pre-ordered configuration can be set online and used for up to three days. Upon activation, it does not incur additional charges from IBM.

Q: Can I use NTP instead of STP?

A: No. However, z10 server can be synchronized against an external time sources through NTP.

Q: Mainframes are known for EBCDIC code pages. How can Linux, which is ASCII, run on z10?

A: There is no requirement for EBCDIC defined in the z/Architecture. For example, z/OS supports EBCDIC, ASCII, and Unicode. Linux on System z uses ASCII and Unicode.

Q: How many new machine instructions were added to z10?

A: More than 50 instructions were added.

Q: What are the new machine instructions used for?

A: The instructions added to z/Architecture are the result of active collaboration between hardware and software designers, specifically with the compiler teams. Hardware and software are being co-optimized, while maintaining full upward compatibility. Most of these instructions are specifically targeted to be used by compilers to improve the efficiency of generated code. Examples are combining two simple functions into a single instruction or reducing the number of active general registers needed for a sequence. Another group of new instructions improves software/hardware synergy by enabling software to give hints to the hardware on caching of specific blocks of memory and by communicating the effective SMP topology so that processes can be kept close to their cached data. The remaining instructions provide minor extensions to existing functions.

Q: When using a sub-capacity model, is it possible to mix different CP feature codes?

A: No, all CP feature codes must be the same. In other words, all CPs must be at the same capacity level.

Q: Are subcapacity versions of specialty engines available?

A: No, specialty engines are always at maximum capacity.

Q: Is it possible to have more than one book although purchased processors would fit into fewer books?

A: Yes, there is no restriction that prevents this. Enhanced book availability uses this approach in order to avoid planned outages.

Q: I cannot have more zAAPs than CPs and I cannot have more zIIPs than CPs. Combined, can I have more zAAPs plus zIIPs than CPs?

A: Yes. For each CP you can have one zAAP and one zIIP.

Q: Did I read that right? I will be authorized to run zAAP-eligible workloads on zIIPs?

A: Yes, that is correct, but several conditions apply. z/OS V1.11 is enhanced with a new function that can enable System z Application Assist Processor (zAAP) eligible workloads to run on System z Integrated Information Processors (zIIPs). This function can enable you to run zIIP-eligible and zAAP-eligible workloads on the zIIP. This capability is ideal for customers without enough zAAP-eligible or zIIP-eligible workload to justify a speciality engine today. The combined eligible workloads may make the acquisition of a zIIP cost effective. The capability is also intended to provide more value for customers having only zIIP processors installed by making Java and XML-based workloads run on existing zIIPs. Customers who have already invested in zAAPs, or both zAAPs and zIIPs, should continue to use these, as this maximizes the new workload potential for the platform. The capability is available with z/OS V1.11 (and with z/OS V1.9 and V1.10 with PTF for APAR OA27495 installed) on all z9 and z10 servers. The capability does *not* provide for overflow so that additional zAAP-eligible workload can spill over and run on the zIIP. It enables the zAAP-eligible workload to run on the zIIP *only* when no zAAP is installed on the server. So the capability is not available on a server with zAAPs installed.

Q: Can I mix dedicated and shared processors in one logical partition?

A: No. Dedicated and shared processors cannot be mixed in one logical partition, regardless of their type.

Q: Is it possible to order an IFL-only server?

A: Yes.

Q: Why can I not use zAAP for Java workload in Linux?

A: zAAP is designed to offload Java workload from CPs in z/OS to keep MSU values lower while providing more processing capacity for Java workload. Because there is no MSU measurement for Linux, it makes no sense to use zAAP there.

Q: What is a z/VM-mode LPAR?

A: A z/VM-mode LPAR is a special type of partition designed to allow z/VM guests to utilize a broader range of specialty processors. This new LPAR mode allows z/VM and its guests to utilize CPs, IFLs, zAAPs, zIIPs, and ICFs in the same logical partition.

Q: What is the Capacity Provisioning Manager?

A: The Capacity Provisioning Manager is software delivered with the z/OS BCP feature. It is a component that allows you to switch OOCoD records on and off automatically according to defined policies. It monitors RMF metrics to decide when to activate OOCoD and when to deactivate it.

Q: What is the watts per square foot ratio for z10 EC?

A: It is approximately 1 163. For a comparison, it is about 709 for System z9 EC.

Q: How much memory should I plan for HSA?

A: With z10 EC no planning of memory for HSA is required. Each z10 EC server by default contains a 16 GB memory that is fixed and used for HSA. This memory is *not* part of the memory purchased by the customer. HSA never occupies additional memory outside of its 16 GB.

Q: Can I run out of HSA space?

A: No. The HSA size is large enough to hold all possible definitions.

Q: What are the increments for ordering the memory?

A: Up to 256 GB, the increment is 16 GB. From 256 GB to 512 GB it is 32 GB. From 512 GB to 944 GB it is 48 GB. From 944 GB it is 64 GB.

Q: Can I connect z10 to a SAN?

A: Yes. FICON cards support both FICON and FC protocols. If the operating system supports the FC protocol, it can participate in the same environment as any other operating system supporting the FC protocol.

Q: Can I use the 2-port FICON Express4 adapter with a z10 EC server?

A: No. It is unique to the z10 BC and z9 BC.

Q: Can I carry forward my older OSA-Express adapters?

A: In general, yes. Check Table B-1 on page 109 for details.

Q: What is the High Performance FICON for System z (zHPF)?

A: zHPF is an extension of the FICON channel architecture compatible with FC-FS, FC-SW, FC-SB-2, FC-SB-3, and FC-SB-4 Fibre Channel standards. If fully exploited by the FICON channel, z/OS, and control unit, it reduces the FICON channel overhead between z/OS and the control unit, thus improving the channel performance.

Q: What is InfiniBand?

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

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

A: The goal of System z was to introduce an industry-standard, high-speed host bus physical interface to replace the self-timed interconnect (STI) proprietary host bus interface. At the same time, System z was looking for an industry-standard protocol to provide Parallel Sysplex coupling link connectivity. The InfiniBand host bus physical interface supports 12x Double Data Rate (12x IB-DDR) with a link speed of up to 6 GBps when attached to a z10 and supports 12x Single Data Rate (12x IB-SDR) with a link speed of up to 3 Gbps when a z10 is attached to a z9.

Q: How can I use InfiniBand on z10?

A: The z10 takes advantage of the InfiniBand I/O bus that includes the InfiniBand Double Data Rate (IB-DDR) infrastructure, which replaces the self-timed interconnect features found in prior System z servers. z10 also uses Parallel Sysplex InfiniBand (PSIFB) links.

Q: Why is the InfiniBand connection speed to z9 only up to 3 GB per second while it is up to 6 GB per second to z10 EC?

A: The System z9 internal bus cannot handle more than 3 GBps. The connection speed automatically adjusts to the slower server.

Q: What are the maximum supported distances of InfiniBand coupling links?

A: The HCA2-O fanout (FC 0163) provides the distance up to 150 m, while HCA2-O LR fanout (FC0168) supports up to 10 km unrepeated (up to 100 km with repeaters).

Q: InfiniBand coupling links are faster than ISC and ICB links. Should I convert all my ISC and ICB links to InfiniBand?

A: The InfiniBand link data rates of 6 GBps, 3 GBps, 2.5 Gbps, and 5 Gbps do *not* represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

When comparing coupling links data rates, InfiniBand (12x IB-SDR or 12x IB-DDR) may be higher than ICB-4 and InfiniBand (1x IB-SDR or 1x IB-DDR) may be higher than that of ISC-3, but with InfiniBand the service times of coupling operations are greater, and the actual throughput may be less than with ICB-4 links or ISC-3 links.

Refer to the *Coupling Facility Configuration Options* white paper for a more specific explanation of when to continue using the current ICB or ISC-3 technology versus migrating to InfiniBand coupling links. The white paper is available from:

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

Q: What CFCC level is supplied with the z10 server?

A: The current level is CFCC 16.

Q: CFCC level 16 is a new level. What new functionality is provided by it?

A: CFCC level 16 contains the following improvements:

- System-Managed CF Structure Duplexing enhancements

Prior to CFCC level 16, Systems-Managed CF Structure Duplexing requires two protocol exchanges to occur synchronously to CF processing of the duplexed structure request. CFCC level 16 allows one of these requests to be asynchronous to CF processing. This implies that the CF-to-CF exchange will occur without z/OS waiting for acknowledgement. This allows faster service time, with more benefits as the Coupling Facilities are further apart, such as in a multi-site Parallel Sysplex. Both Coupling Facilities must be at CFCC level 16 for these enhancements to occur.

- List Notification improvements

Today when a list changes its state from empty to non-empty all its connectors are notified. The first connector notified reads the new message, but subsequent readers will find nothing. CFCC Level 16 approaches this differently to improve CPU utilization. It only notifies one connector in a round-robin fashion, and if the shared queue (as in IMS Shared Queue and WebSphere MQ Shared Queue) is read within a fixed period of time, the other connectors do not need to be notified. If the list is not read again within the time limit the other connectors are informed.

Q: With all those changes what will the CF memory requirements be?

A: No significant CF structure sizing changes are expected when going from CFCC level 15 to CFCC level 16. However, we strongly recommend using the CF Sizer tool available at:

<http://www.ibm.com/systems/z/cfsizer/>

Archived

Channel options

Table B-1 lists the attributes of the channel options supported on z10 EC servers with the required connector and cable types, the maximum unrepeated distance, and the bit rate.

The z10 EC Model E12 supports up to 64 I/O slots (960 CHPIDs in three I/O cages), while models E26, E40, E56, and E64 support up to 84 I/O slots (1024 CHPIDs in three I/O cages). At least one ESCON, FICON, ICB, or ISC feature is required.

Table B-1 System z10 EC channel feature support

Channel feature	Feature codes	Bit rate	Connector	Cable type	Maximum unrepeated distance ^a
Enterprise Systems CONnection (ESCON)					
16-port ESCON	2323	200 Mbps	MT-RJ	MM 62.5 μ m	3 km (800)
Fiber Connection (FICON)					
FICON Express ^b LX	2319	1 Gbps	LC Duplex	SM 9 μ m	10 km
FICON Express ^b SX	2320	1 Gbps	LC Duplex	MM 62.5 μ m MM 50 μ m	300 m (984) 860 m (2822) 500 m (1640)
FICON Express2 ^b LX	3319	2 Gbps	LC Duplex	SM 9 μ m	10 km
FICON Express2 ^b SX	3320	2 Gbps	LC Duplex	MM 62.5 μ m MM 50 μ m	150 m (492) 500 m (1640) 300 m (984)

Channel feature	Feature codes	Bit rate	Connector	Cable type	Maximum unrepeat distance ^a
FICON Express4 SX ^b	3322	4 Gbps	LC Duplex	MM 62.5 µm MM 50 µm	70 m (230) 380 m (1247) 150 m (492)
		2 Gbps	LC Duplex	MM 62.5 µm MM 50 µm	150 m (492) 500 m (1640) 300 m (984)
		1 Gbps	LC Duplex	MM 62.5 µm MM 50 µm	300 m (984) 860 m (2822) 500 m (1640)
FICON Express4 ^b 4KM LX	3324	1, 2, or 4 Gbps	LC Duplex	SM 9 µm	4 km
FICON Express4 ^b 10KM LX	3321	1, 2, or 4 Gbps	LC Duplex	SM 9 µm	10 km/20 km ^c
FICON Express8 SX	3326	8 Gbps	LC Duplex	MM 62.5 µm MM 50 µm	21 m (69) 150 m (492) 50 m (164)
		4 Gbps	LC Duplex	MM 62.5 µm MM 50 µm	70 m (230) 380 m (1247) 150 m (492)
		2 Gbps	LC Duplex	MM 62.5 µm MM 50 µm	150 m (492) 500 m (1640) 300 m (984)
FICON Express8 10KM LX	3325	2, 4, or 8 Gbps	LC Duplex	SM 9 µm	10 km
Open Systems Adapter (OSA)					
OSA-Express2 GbE LX	3364	1 Gbps	LC Duplex	SM 9 µm	5 km
				MCP	550 m (500)
OSA-Express2 GbE SX	3365	1 Gbps	LC Duplex	MM 62.5 µm	220 m (166) 275 m (200)
				MM 50 µm	550 m (500)
OSA-Express2 1000BASE-T Ethernet	3366	10/100/ 1000	RJ45	UTP Cat5	100 m
OSA-Express2 10 GbE LR	3368	10 Gbps	SC Duplex	SM 9 µm	10 km
OSA-Express3 GbE LX	3362	1 Gbps	LC Duplex	SM 9 µm	5 km
				MCP	550 m (500)
OSA-Express3 GbE SX	3363	1 Gbps	LC Duplex	MM 62.5 µm	220 m (166) 275 m (200)
				MM 50 µm	550 m (500)
OSA-Express3 1000BASE-T Ethernet	3367	10/100/ 1000	RJ45	UTP Cat5	100 m

Channel feature	Feature codes	Bit rate	Connector	Cable type	Maximum unrepeated distance ^a
OSA-Express3 10 GbE LR	3370	10 Gbps	LC Duplex	SM 9 μm	10 km
OSA-Express3 10 GbE SR	3371	10 Gbps	LC Duplex	MM 62.5 μm	33 m (200)
				MM 50 μm	300 m (2000) 82 m (500)
Parallel Sysplex					
IC	n/a		N/A	N/A	N/A
ICB-4	3393	2 GBps	0229/0230 ^c		10 m
ISC-3 (peer mode)	0217 0218 0219	2 Gbps	LC Duplex	SM 9 μm MCP 50 μm	10 km/20 km 550 m (400)
ISC-3 (RPQ 8P2197 Peer mode at 1 Gbps) ^d		1 Gbps		SM 9 μm	20 km ^c
PSIFB	0163	6 GBps	MPO	OM3 MM 50 μm	150 m
PSIFB LR	0168	5 Gbps	LC Duplex	SM 9 μm	10 km/100 km ^e
ETR ^f	n/a	8 Mbps	MT-RJ	MM 62.5 μm MM 50 μm	3 km (26 km) 2 km (24 km)
Cryptography					
Crypto Express2	0863	N/A	N/A	N/A	N/A
Crypto Express3	0864	N/A	N/A	N/A	N/A

- a. Minimum fiber bandwidth in MHz/km for multi-mode fiber optic links are included in parentheses were applicable.
- b. Feature is only available if carried forward by an upgrade from a previous server.
- c. FC 0229 cable: z10 EC to z9, z990, or z890.
FC 0230 cable: z10 EC to z10.
- d. RPQ 8P2197 enables the ordering of a different daughter card supporting 20 km unrepeated distance for 1 Gbps peer mode. RPQ 8P2262 is a requirement for that option, and other than the normal mode the channel increment is two, that is, both ports (FC 0219) at the card must be activated.
- e. Up to 100 km at 2.5 Gbps, with repeater (System z qualified DWDM vendor product that supports 1x IB-SDR)
- f. The External Time Reference (ETR) replaces the ETR feature available in previous servers. Two ETR cards are a standard feature in the z10 EC servers.

Archived

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks publications

For information about ordering these publications, see “How to get Redbooks publications” on page 114. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM System z10 Enterprise Class Technical Guide*, SG24-7516
- ▶ *IBM System z Strengths and Values*, SG24-7333
- ▶ *Getting Started with InfiniBand on System z10 and System z9*, SG24-7539
- ▶ *IBM System z Connectivity Handbook*, SG24-5444
- ▶ *Server Time Protocol Planning Guide*, SG24-7280
- ▶ *Server Time Protocol Implementation Guide*, SG24-7281
- ▶ *IBM System z9 Enterprise Class Configuration Setup*, SG24-7571
- ▶ *IBM System z10 Enterprise Class Capacity On Demand*, SG24-7504
- ▶ *IBM TotalStorage DS8000 Series: Performance Monitoring and Tuning*, SG24-7146
- ▶ *How does the MIDAW Facility Improve the Performance of FICON Channels Using DB2 and other workloads?*, REDP-4201
- ▶ *Introduction to the New Mainframe: z/VM Basics*, SG24-7316

Online resources

These Web sites are also relevant as further information sources:

- ▶ ResourceLink Web site
<http://www.ibm.com/servers/resourcelink>
- ▶ Large Systems Performance Reference (LSPR)
<http://www-03.ibm.com/servers/eserver/zseries/lspr/>
- ▶ MSU ratings
<http://www-03.ibm.com/servers/eserver/zseries/library/swpriceinfo/hardware.html>

Other publications

These publications are also relevant as further information sources:

- ▶ *Hardware Management Console Operations Guide Version 2.10.0*, SC28-6867
- ▶ *Support Element Operations Guide V2.10.0*, SC28-6868

- ▶ *IOCP User's Guide*, SB10-7037
- ▶ *Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152
- ▶ *Planning for Fiber Optic Links*, GA23-0367
- ▶ *System z10 Enterprise Class Capacity on Demand User's Guide*, SC28-6871
- ▶ *CHPID Mapping Tool User's Guide*, GC28-6825
- ▶ *Common Information Model (CIM) Management Interfaces*, SB10-7154
- ▶ *System z10 Enterprise Class Installation Manual*, GC28-6865
- ▶ *System z10 Enterprise Class Installation Manual for Physical Planning*, GC28-6864
- ▶ *System z10 Enterprise Class Processor Resource/Systems Manager Planning Guide*, SB10-7153
- ▶ *System z10 Enterprise Class System Overview*, SA22-1084
- ▶ *System z10 Enterprise Class Service Guide*, GC28-6866
- ▶ *IBM System z Functional Matrix*, ZSW0-1335
- ▶ *z/Architecture Principles of Operation*, SA22-7832
- ▶ *z/OS Cryptographic Services Integrated Cryptographic Service Facility Administrator's Guide*, SA22-7521
- ▶ *z/OS Cryptographic Services Integrated Cryptographic Service Facility Application Programmer's Guide*, SA22-7522
- ▶ *z/OS Cryptographic Services Integrated Cryptographic Service Facility Messages*, SA22-7523
- ▶ *z/OS Cryptographic Services Integrated Cryptographic Service Facility Overview*, SA22-7519
- ▶ *z/OS Cryptographic Services Integrated Cryptographic Service Facility System Programmer's Guide*, SA22-7520

How to get Redbooks publications

You can search for, view, or download Redbooks publications, Redpapers, publications, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

50.0 μm 35
62.5 μm 35

A

Active Energy Manager (AEM) 44
Advanced Encryption Standard (AES) 11, 37
AES 11, 37

C

cage, CEC and I/O 12
Capacity Backup (CBU) 66
Capacity for Planned Event 104
Capacity for Planned Events (CPE) 66
Capacity provisioning 67
Capacity Provisioning Control Center 67
Capacity Provisioning Manager 67
capacity ratios 15
capacity token 104
CBU 65–66
CF 40
CFCC level 107
Channel path 11
chip lithography 24
CHPID 40
Commercial Batch Short (CB-S) 15
Compression Unit 37
cooling 43
Coupling Facility (CF) 40
Coupling Facility Control Code level 107
Coupling Link 29
coupling links 14
CP characterization 10
CP Cryptographic Assist Facility 37
CPACF
 cryptographic capabilities 6, 10
 description 37
 PU design 37
CPACF enhancements 61
CPE 65
Crypto Express2 11–12
 accelerator 13, 38
 coprocessor 13, 38
 support 92
Crypto Express3
 accelerator 38
 coprocessor 38
Cryptographic Accelerator (CA) 38
Cryptographic Coprocessor (CC) 38
cryptographic hardware 25
Customer Initiated Upgrade (CIU) facility 65

D

data connection isolation 59
Data Encryption Standard (DES) 11, 37
Decimal Floating Point 25
DES 11, 37
Dynamic Infrastructure 101
dynamic infrastructure 6

E

EAL 5 101
EAL5 46
Enhanced driver maintenance (EDM) 69
enterprise service bus (ESB) 6
ESA/390 25
ETR attachment 62
Extended Address Volumes (EAV) 48

F

FCP 13
 enhancements for small block sizes 57
 switch 13
Federal Information Processing Standard (FIPS) 13
Fibre Channel Protocol (FCP) 13, 57
FICON
 extended distance 56
 name server registration 56
FICON Express 34
FICON Express2 33
FICON Express4 13, 32
FICON Express8 32
FICON to ESCON
 conversion function 32
flexible memory 53
footprint 43

H

Hardware system area (HSA) 11
HiperDispatch 16, 102
HiperSockets 13–14
 IPv6 13
 multiple write facility 60
 zIIP-Assisted 60
HiperSockets Layer 2 support 60
HiperSockets Multiple Write Facility 60
HMC 42
HMC applications 61
HyperPAV 48

I

I/O cage 12
 I/O slot 12
I/O device 28

- I/O operation 90
- I/O virtualization 48
- IC3 40
- ICB-4 link 40
- ICF characterization 10
- ICSF 92
- IFL characterization 10
- InfiniBand 12, 54–55, 106
- Input/Output Configuration Dataset (IOCDS) 40
- Internal Battery Feature (IBF) 44, 64
- ISC-3 39
- ITRR 15

L

- L1 cache 25
- LAN 39
- large page support 53
- Licensed Internal Code (LIC) 27
- Linux for System z 76, 83
- Local area network (LAN) 39
- logical processors 47

M

- MCM 9
- memory
 - card 27
 - size 26
- memory bus adapter (MBA) 28
- MIDAW facility 95–96
- mode conditioner patch (MCP) 35
- Model S08 19
- Model S54 19
- MSU value 14

N

- Network Control Program 58
- Network Time Protocol (NTP) 63
- NPIV 57
- NTP server 64

O

- On/Off CoD 65–66
- On-line Permanent Upgrade 65
- operating system 65, 76–78
 - support 78
 - support Web page 100
- OSA for NCP 58
- OSA-Express 13
- OSA-Express2 35
 - Gigabit Ethernet 58
- OSN 58
- OSA-Express3 34

P

- Parallel Access Volume (PAV) 48
- Parallel Sysplex 40, 98
 - License Charge 98

- Web site 95
- PCI-e
 - cryptographic adapter 38
 - cryptographic coprocessor 38
- Peer-to-Peer Remote Copy (PPRC) 48
- Performance 14
- permanent upgrade 65
- personal identification number (PIN) 39
- physical memory 26–27
- plan-ahead memory 53
- power-on reset (POR) 25
- PR/SM 46
- prediction tool 57
- pre-paid capacity 104
- processor unit (PU) 25
- Pseudo Random Number Generation (PRNG) 11, 37, 92
- PSIFB 40
- PSP buckets 77
- PU characterization 25
- pulse per second (PPS) support 64

Q

- QDIO interface isolation 59
- QDIO optimized latency mode 60

R

- Redbooks Web site 114
 - Contact us xi
- refrigeration 43
- Resource Link 65

S

- SALC 98
- SAP characterization 10
- SE 42
- Secure Hash Algorithm (SHA-1) 11
- Secure Sockets Layer (SSL) 13, 37–38
- Select Application License Charges 98
- Self-Timed Interconnect (STI) 40
- SHA-1 37
- SHA-2 37
- SHA-256 11
- single system image 85
- soft capping 98
- software licensing 95–97
- software support 85
- STP configuration and time information 64
- STP CTN reconfiguration 64
- Subchannels 11
- support requirement
 - z/OS 79, 84
- Sysplex Timer, time-of-day synchronization 62
- System Assist Processor (SAP) 11
- system assist processor (SAP) 19
- system image 85
- System z9 Integrated Information Processor 10

T

temporary upgrades 66
TKE workstation 14
Triple Data Encryption Standard (TDES) 11, 37
Trusted Key Entry (TKE) 39

U

unplanned upgrades 65
user 13
User Defined Extensions (UDX) 13, 38

V

VMAC support 59

W

Web 2.0 5
WebSphere 6
WebSphere MQ 98
Workload License Charge (WLC) 97–98
 Flat WLC (FWLC) 97
 sub-capacity 97
 Variable WLC (VWLC) 97
Workload Manager (WLM) 68

Z

z/Architecture 6, 25, 76, 83
z/OS 80, 92
z/TPF 16
z/VM 49
z/VM-mode 46
zAAP characterization 10
zHPF 55, 106
zIIP 10

Archived

Archived



IBM System z10 Enterprise Class Technical Introduction



The server's role in a dynamic infrastructure

Key functional elements and enhancements

Hardware and software capabilities

This IBM Redbooks publication introduces the IBM System z10 Enterprise Class server, which is based on z/Architecture. It builds on the inherent strengths of the System z platform, delivering new technologies and virtualization that are designed to offer improvements in price and performance for key workloads, as well as enabling a new range of solutions. The z10 EC further extends System z's leadership in key capabilities with the delivery of expanded scalability for growth and large-scale consolidation, availability to help reduce risk and improve flexibility to respond to changing business requirements, and improved security. The z10 EC is at the core of the enhanced System z platform that is designed to deliver technologies that businesses need today along with a foundation to drive future business growth.

This document provides basic information about z10 EC capabilities, hardware functions and features, and associated software support. It is intended for IT managers, architects, consultants, and anyone else who wants to understand the new elements of the z10 EC. The changes in this third edition are based on the IBM Hardware Announcement, dated October 20, 2009.

This book is intended as an introduction to the z10 EC mainframe. Readers are not expected to be generally familiar with current IBM System z technology and terminology.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks