# IBM z13 and IBM z13s Technical Introduction

Bill White

Cecilia A De Leon

Edzard Hoogerbrug

Ewerson Palacio

Franco Pinto

Barbara Sannerud

Martin Soellig

John Troy

Jin J Yang

z Systems

IBM

**IBM**  International Technical Support Organization

**IBM z13 and IBM z13s Technical Introduction**

March 2016

> **Note:** Before using this information and the product it supports, read the information in "Notices" on page vii.

**Second Edition (March 2016)**

This edition applies to the following IBM z Systems™ platforms: IBM z13™ and z13s™.

# Contents

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

| | | |
|---|---|---|
| AIX® | Language Environment® | System z® |
| CICS® | Lotus® | System z10® |
| DataPower® | Parallel Sysplex® | System z9® |
| DB2® | Passport Advantage® | WebSphere® |
| Domino® | POWER7® | Worklight® |
| DRDA® | PowerVM® | z Systems™ |
| DS8000® | PR/SM™ | z/Architecture® |
| FICON® | Processor Resource/Systems | z/OS® |
| GDPS® | Manager™ | z/VM® |
| Guardium® | RACF® | z/VSE® |
| HiperSockets™ | Rational® | z10™ |
| HyperSwap® | Rational Rose® | z13™ |
| IBM® | Redbooks® | z9® |
| IBM MobileFirst™ | Redbooks (logo) ® | zEnterprise® |
| IBM z Systems™ | Resource Link® | zSecure™ |
| IBM z13™ | RMF™ | |
| IMS™ | System Storage® | |

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds 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.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

# New Availability Intelligence About Your z/OS Infrastructure

**IntelliMagic**

**You already know z13 enables new "Systems of Insight" for your business users through how it processes your digital transactions.**

**But did you know your IT staff can have new "Systems of Insight" about the z/OS infrastructure itself through the "Availability Intelligence" that IntelliMagic Vision automatically produces?**

## What is Availability Intelligence?

- It is what enables you to proactively protect availability instead of simply recovering it after it is lost.
- It is intelligence that goes significantly beyond what your current online monitors and SAS/MXG reports are producing today.
- It is the result of mining valuable information "trapped" in structured data, such as RMF and SMF, to better manage the risk of z/OS resources during peak workloads.
- It is produced by analytics together with built-in knowledge about the infrastructure resources and includes scoring (risk quantification)
- It provides early warning about the most common causes of disruption to your continuous availability status, and it also reduces your MTTR for unpredictable disruptions.

## What is IntelliMagic Vision?

- A software solution (also available as Software as a Service) that produces Availability Intelligence about your z13 or any other z/OS environment.
- It supports detailed intelligence about resources such as Disk, Replication, Tape/Virtual Tape, Coupling Facility, XCF, Processors, WLM, and more.
- It's a solution that is easily demonstrated with your historical RMF/SMF data, showing you the early warning that would have helped you avoid your last performance/capacity incident.

**Request a Brief Demo »»**
info.intellimagic.com/z13RB

**IntelliMagic software and services are used daily in some of the largest, most demanding z/OS sites in the world.**

IntelliMagic • www.intellimagic.com

THIS PAGE INTENTIONALLY LEFT BLANK

# IBM Redbooks promotions

## Find and read thousands of IBM Redbooks publications

▶ Search, bookmark, save and organize favorites

▶ Get personalized notifications of new content

▶ Link to the latest Redbooks blogs and videos

**Get the latest version of the Redbooks Mobile App**

iOS
**Download Now**
Android

Redbooks

**Creating Hybrid Clouds with IBM Bluemix Integration Services**

David Kwock
Rahul Gupta
Vasfi Gucer

☁ Cloud

Analytics

IBM                     Solution Guide

## Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

It's good to be **noticed**.

**ibm.com/Redbooks**
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

# Preface

This IBM® Redbooks® publication introduces the latest IBM z Systems™ platforms, the IBM z13™ and IBM z13s. It includes information about the z Systems environment and how it can help integrate data, transactions, and insight for faster and more accurate business decisions.

The z13 and z13s are state-of-the-art data and transaction systems that deliver advanced capabilities that are vital to modern IT infrastructures. These capabilities include:

► Accelerated data and transaction serving
► Integrated analytics
► Access to the API economy
► Agile development and operations
► Efficient, scalable, and secure cloud services
► End-to-end security for data and transactions

This book explains how these systems use both new innovations and traditional z Systems strengths to satisfy growing demand for cloud, analytics, and mobile applications. With one of these z Systems platforms as the base, applications can run in a trusted, reliable, and secure environment that both improves operations and lessens business risk.

## Authors

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

**Bill White** is a Project Leader and Senior z Systems Networking and Connectivity Specialist at IBM Redbooks, Poughkeepsie Center.

**Cecilia A De Leon** is a Certified IT Specialist in the Philippines. She has 15 years of experience in z Systems field. She has worked at IBM for 7 years. She holds a degree in Computer Engineering from Mapua Institute of Technology. Her areas of expertise include z Systems servers and operating systems. In her current role as Client Technical Specialist, she supports mainframe clients and IBM sales representatives on technical sales engagements. She has also worked as a systems programmer for large banks in the Philippines.

**Edzard Hoogerbrug** is a System Support Representative in The Netherlands. During the past 26 years, he has worked in various roles within IBM mainly in the mainframe world. He has work 14 years for EMEA L2 support for z Systems after a 2.5 years assignment in Montpellier France. He holds a degree in electrotechnology. His areas of expertise include failure analysis on hardware for z Systems.

**Ewerson Palacio** is an IBM Distinguished Engineer and a Certified Consulting IT Specialist for Large Systems in Brazil. He has more than 40 years of experience in IBM large systems. Ewerson holds a Computer Science degree from Sao Paulo University. His areas of expertise include z Systems client technical support, mainframe architecture, infrastructure implementation, and design. He is an ITSO z Systems hardware official speaker who has presented technical ITSO seminars, workshops, and private sessions to IBM clients, IBM IT Architects, IBM IT Specialists, and IBM Business Partners around the globe. He has also been a z Systems Hardware Top Gun training designer, developer, and instructor for the last generations of the IBM high-end servers. Ewerson leads the Mainframe Specialty Services

Area (MF-SSA), which is part of GTS Delivery, Technology and Engineering (DT&E). He is a member of the IBM Academy of Technology.

**Franco Pinto** is a Client Technical Specialist in IBM Switzerland. He has 20 years of experience in the mainframe and IBM z/OS® field. His areas of expertise include z Systems technical pre-sales covering mainframe sizing and installation planning, and providing support on existing and new z Systems functions.

**Barbara Sannerud** is a Worldwide Technical Enablement Manager. She has 30 years of experience in services, strategy, marketing, and product management. She recently was an Offering Manager for z/OS, and also held competitive marketing roles. She holds a math and MBA degree and joined IBM from the software and professional services industries where she specialized in performance, systems management, and security.

**Martin Soellig** is a Consulting IT Specialist in Germany. He has 26 years of experience working in the z Systems field. He holds a degree in Mathematics from University of Hamburg. His areas of expertise include z/OS and z Systems hardware, specifically in IBM Parallel Sysplex® and GDPS® environments.

**John Troy** is a z Systems and Storage hardware National Top Gun in the northeast area of the United States. He has 35 years of experience in the service field. His areas of expertise include z Systems server and high-end storage systems technical and customer support. John has been a z Systems hardware technical support course designer, developer, and instructor for the last six generations of IBM high-end servers.

**Jin J Yang** is a Senior System Service Representative in China. He joined IBM in 1999 to support z Systems products maintenance for clients in China. He has been working in the Technical Support Group to provide second-level support for z Systems clients as a country Top Gun since 2009. His areas of expertise include z Systems hardware, channel connectivity, z/VM and Linux on z Systems.

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

► Octavian Lascu, Shawn Tooley
  IBM Redbooks

► Vicente Ranieri Jr.
  IBM Brazil

► Chris Brooker, Rick Brown, Anuja Deedwaniya, Romney White
  IBM US

Thanks to the authors of the first edition of this book, *IBM z13 Technical Introduction*, published in March 2015:

Octavian Lascu
Erik Bakker
Hans-Peter Eckam
Parwez Hamid
Rakesh Krishnakumar
Steven LaFalce
Lourenço Luitgards Moura Neto
Mauricio Andozia Nogueira
Frank Packheiser
Ewerson Palacio
Giancarlo Rodolfi
André Spahni

# Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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 publications in one of the following ways:

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

**ibm.com**/redbooks

► Send your comments in an email 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

# Stay connected to IBM Redbooks

► Find us on Facebook:

http://www.facebook.com/IBMRedbooks

► Follow us on Twitter:

http://twitter.com/ibmredbooks

► Look for us on LinkedIn:

http://www.linkedin.com/groups?home=&gid=2130806

► Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

► Stay current on recent Redbooks publications with RSS Feeds:

http://www.redbooks.ibm.com/rss.html

**1**

# Platforms matter

It is no surprise that today's massive data growth rates and increasing consumer expectations are changing the way businesses and organizations assess their IT infrastructure. Implementing a truly effective IT infrastructure now requires *each* of these key capabilities:

► Deliver lightning-fast response times

► Receive and process requests from anyone at any time, regardless of location or choice of device

► Secure all data, communications, identities, and business processes

► Predict and adapt to activity peaks and shifts in consumer interests

The latest members of the IBM z Systems family, the IBM z13 and IBM z13s, can provide the capacity and flexibility you need to shape your IT infrastructure to take advantage of the opportunities of the digital age.

Innovations such as large memory, compression technologies, vector processing, and improved chip design give these mainframes an operational edge that can extend any organization's ability to store vast amounts of information to harvest real-time insight. The z13 and z13s also have unmatched virtualization capabilities, support for enterprise-grade Linux platforms and open source applications, and numerous feature enhancements to better handle both new and current workloads with speed, security, and resiliency.

This chapter introduces the technical capabilities of both the z13 and z13s and includes these topics:

► 1.1, "The z13 and z13s: Built for today and tomorrow" on page 2
► 1.2, "z13 and z13s technical description" on page 4
► 1.3, "Software support" on page 11

**1**

# 1.1  The z13 and z13s: Built for today and tomorrow

The z13 offers a fast, scalable, and securable enterprise system. Compared to its predecessor platforms, the z13 provides more of what you need to satisfy today's growing IT demands:

► Compute power for increased throughput
► Large-scale memory to process data faster
► Industry-unique cache design to optimize performance
► Accelerated I/O bandwidth to move massive amounts of data
► Data compression to economically store and process information
► High-speed cryptographic operations to help secure transactions

The z13s offers a smart alternative for businesses and organizations that need the advanced capabilities and qualities of service of z13 at a more economical entry point.

Figure 1-1 shows the IBM z13 on the left and the single-frame IBM z13s on the right.



*Figure 1-1    IBM z13 and IBM z13s platforms*

From the microprocessors to the software, the z13 and z13s are designed to meet changing market demands and are optimized for current and future IT infrastructure needs. They meet these needs by supplying:

► Accelerated data and transaction serving
► Integrated analytics for insight
► Access to the API economy
► An agile application development and operations environment
► Efficient, scalable, and secure cloud services
► End-to-end security for data and transactions

### 1.1.1  Accelerated data and transaction serving

The z13 and z13s are built with capabilities that process data extremely fast. These capabilities include I/O features that improve the latency of data access (while reducing the elapsed time) and shared memory communications that are transparent to applications and offer direct, high-speed data transfer within the z Systems platform or to other z Systems platforms.

The z13 and z13s also support a larger memory pool for faster analysis of large volumes of data. The cache structure is optimized for improved proximity of processing to data. In addition, the simultaneous multithreading (SMT) feature delivers more throughput for efficient and competitive operational analytics of Java, Linux, and many other eligible workloads.

The high-performance on-chip hardware compression acceleration facility, introduced with z13, helps to compress more data, thereby saving disk space and reducing data transfer time, while lowering I/O bandwidth requirements and providing better time-to-value.

Speed and access to data define a competitive advantage for the enterprise. The quicker meaningful information is detected from data, the better prepared you are to react to changing market dynamics.

### 1.1.2  Integrated analytics for insight

In the digital age, data has become the new currency that organizations use to capture business advantages, particularly with the help of advanced analytics. Sophisticated analytic algorithms can be applied during transactions to analyze them in real time (for example, fraud detection). Companies in banking, finance, the public sector, and healthcare must be able to use the core transactional systems and data repositories on z Systems and extend them using modern APIs to create new applications.

In this world of constant information exchange, a highly responsive, secure, and available infrastructure is mandatory. Advanced analytics depend on having a highly capable infrastructure that can scale to meet both current service requirements and new needs as they emerge. The z13 and z13s are built with these requirements in mind.

### 1.1.3  Access to the API economy

Mobile, web, and cloud technologies are enabling new business models where innovative services can be created in just hours using a palette of easy-to-consume APIs. Technologies are introduced that make z Systems platforms fully conversant with REST APIs, putting z13 and z13s at the heart of the API economy. Thanks to IBM products such as z/OS Connect Enterprise Edition, contemporary mobile, digital, and cloud applications can directly and securely use z Systems business applications as an equal partner in the API economy.

### 1.1.4  An agile application development and operations environment

The z13 and z13s deliver an approach for enterprise-grade Linux with features that are designed for availability and virtualization with IBM z/VM and that have a focus on open standards and architectures, including support of kernel-based virtual machines (KVMs).

KVM for IBM z Systems is a software distribution that can coexist with all other supported operating systems on the z13 and z13s. It is optimized for the z Systems architecture and provides standard Linux and KVM interfaces for operational control of the environment. In addition, KVM for IBM z Systems integrates with standard OpenStack virtualization

management, allowing businesses and organizations to incorporate Linux images into existing infrastructure and cloud offerings. In addition, provisioning capabilities, such as IBM Dynamic Partition Manager, provide (in conjunction with KVM for IBM z Systems) a catalyst for building open source-based solutions in an agile way and rapid creation of new services.

### 1.1.5 Efficient, scalable, and secure cloud services

Growing digital demands necessitate that IT infrastructure be aligned with new and changing operational and business processes. The infrastructure must serve consumers through a secure, flexible "as a service" model that can help meet needs with acceptable risk to the business. To accommodate these demands, businesses and organizations are looking to deploy a cloud solution to extend their current environment with new cloud services. The z13 and z13s can participate in any type of cloud environment. In addition, they offer the modern APIs, the workload isolation, and the scale and the resiliency needed. Also, enhanced cryptography creates a highly secure environment that protects the privacy of customer information.

### 1.1.6 End-to-end security for data and transactions

Analytics, mobile, and cloud all have one aspect in common: the need for a platform with deeply integrated security. The z13 and z13s excel with security features that are *built into* the hardware, firmware, and operating systems. The built-in features range from storage protection keys and workload isolation, to granular audit capabilities, and more. The z13 and z13s also offer processor-based encryption and advanced cryptographic capabilities with secure key management.

In addition, IBM has surrounded the z Systems platforms with tools to aid with analytics, user authentication, application security, and regulatory compliance. Examples include:

▶ Multi-Factor Authentication helps to prevent unauthorized access to sensitive data.

▶ IBM zSecure™ and IBM Security Guardium® can automate and speed up security intelligence in the cloud.

▶ z Appliance Container Infrastructure enables the secure deployment of software appliances.

Together, these advanced security options can help to ensure end-to-end privacy and protection of data and transactions.

## 1.2 z13 and z13s technical description

The z13s shares many characteristics with the z13 but on a smaller scale. Both use several leading-edge technologies, including Silicon-On-Insulator 22nm (CMOS 14S0 process), storage-class memory, and Peripheral Component Interconnect Express (PCIe) Generation 3 (PCIe Gen3).

The z13 and z13s, when compared to their predecessors (IBM zEC12 and zBC12), offer improvements in areas such as faster, more efficient, and redesigned high-frequency chips, additional granularity options, better availability, faster encryption, and enhanced on-demand options.

## 1.2.1  Technical highlights

The z13 and z13s are highly scalable symmetric multiprocessor (SMP) systems.

The z13 that can be configured with up to 141 characterizable Processor Units, and an architecture that ensures continuity and upgradeability from the previous zEC12 and z196. Five z13 models are offered: N30, N63, N96, NC9, and NE1.

The z13s offers up to 20 characterizable Processor Units. Like the z13, its architecture ensures continuity and upgradeability from its predecessor platforms, the zBC12 and z114. Two z13s models are offered: N10 and N20.

Table 1-1 shows the major technical enhancements in the z13 and z13s over predecessor systems.

*Table 1-1   Technical enhancements in the z13 and z13s*

| Feature | z13<br>(N30, N63, N96, NC9, NE1) | z13s<br>(N10, N20) |
|---|---|---|
| Greater total system capacity and more subcapacity settings. Architecture ensures continuity and upgradeability from previous models. | Up to 141 characterizable Processor Units | Up to 20 characterizable Processor Units |
| Multi-core, single-chip modules running to help improve the execution of processor-intensive workloads. | 5.0 GHz | 4.3 GHz |
| More real memory per system, ensuring high availability in the memory subsystem through use of proven redundant array of independent memory (RAIM) technology. | Up to 10 TB of addressable real memory per system | Up to 4 TB of addressable real memory per system |
| A large fixed hardware system area (HSA) that is managed separately from client-purchased memory. | 96 GB | 40 GB |
| Proven technology (fourth-generation high frequency and second-generation out-of-order design) with a single instruction, multiple data (SIMD) processor that increases parallelism to accelerate analytics processing. In addition, simultaneous multithreading (SMT) increases processing efficiency and throughput and raises the number of instructions in flight by more than 40%. | | |
| Processor cache structure improvements and larger cache sizes to help with more of today's demanding production workloads. Both z13 and z13s offer:<br>► First-level cache (L1 private): 96 KB for instructions, 128 KB for data<br>► Second-level cache (L2): 2 MB for data and instructions<br>► Third-level cache (L3): 64 MB (plus an additional 224 MB for the non-Data Integrated Coherent (NIC) Directory)<br>► Fourth-level cache (L4): 480 MB | | |
| Improved cryptographic functionality and performance, achieved by having one dedicated cryptographic coprocessor per core. | | |
| Channel subsystem enhancements for I/O resilience. The number of logical channel subsystems (LCSS) has grown compared to predecessor systems, as has the number of logical partitions (LPAR). Additional subchannel sets are implemented to improve addressability, and the number of I/O devices has jumped from 24,000 devices to 32,000 devices. | ► Six LCSS<br>  (up from four)<br>► 85 LPAR<br>  (up from 60)<br>► Four subchannel sets<br>  (up from three) | ► Three LCSS<br>  (up from two)<br>► 40 LPAR<br>  (up from 30)<br>► Three subchannel sets<br>  (up from two) |

To ensure a completely balanced and highly available system, both the z13 and z13s include these additional features and functions:

- ► Enhanced LPAR resource allocation algorithms for Processor Units and memory.
- ► Industry-standard flash solid-state drives (SSD) that are mounted in Flash Express feature cards and can be used to handle paging workload spikes and improve availability.
- ► Crypto Express5S with enhanced support of cryptographic functions for up to 85 domains on z13 (up to 40 domains on the z13s).
- ► PCIe Gen3 (Generation 3) fanouts provide 16 GBps bus connections to the PCIe I/O features.
- ► Functionality in RoCE Express, with a second port now available, along with the ability to share adapters between LPARs.
- ► Integrated Coupling Adapter-Short Reach (ICA-SR) for coupling links.
- ► 1U Support Elements (SEs) to replace the SE notebooks.
- ► Redundant System Control Hubs (SCHs) that replace the Bulk Power Hubs (BPHs)
- ► An optional, rack-mountable 1U Hardware Management Console (HMC) for installation in a customer supplied rack.
- ► Air cooled systems such as cooling radiators with N+2 redundant design.
- ► The IBM z Advanced Workload Analysis Reporter (IBM zAware), which provides a smart solution for detecting and diagnosing anomalies in z/OS systems and Linux on z Systems.

Figure 1-2 and Figure 1-3 on page 7 compare the z13 and z13s with previous z Systems platforms in the following major areas:

- ► Single-engine processing capacity[1]
- ► Number of Processor Units
- ► Memory
- ► I/O bandwidth
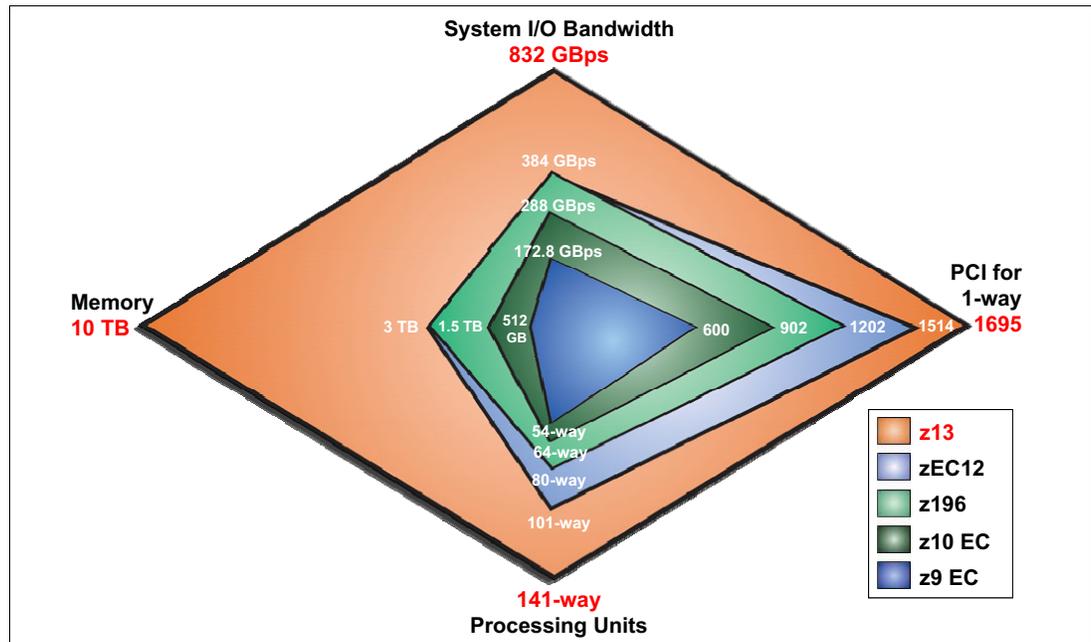


Figure 1-2   Platform design: The z13 versus its predecessors

---

[1] Based on the processor capacity index (PCI). PCI values can be obtained from *Large Systems Performance Reference*, SC28-1187.

You can find additional details about the z13 in Chapter 2, "IBM z13 hardware overview" on page 13. For an in-depth description of IBM z13 functions and features, see also *IBM z13 Technical Guide,* SG24-8251.
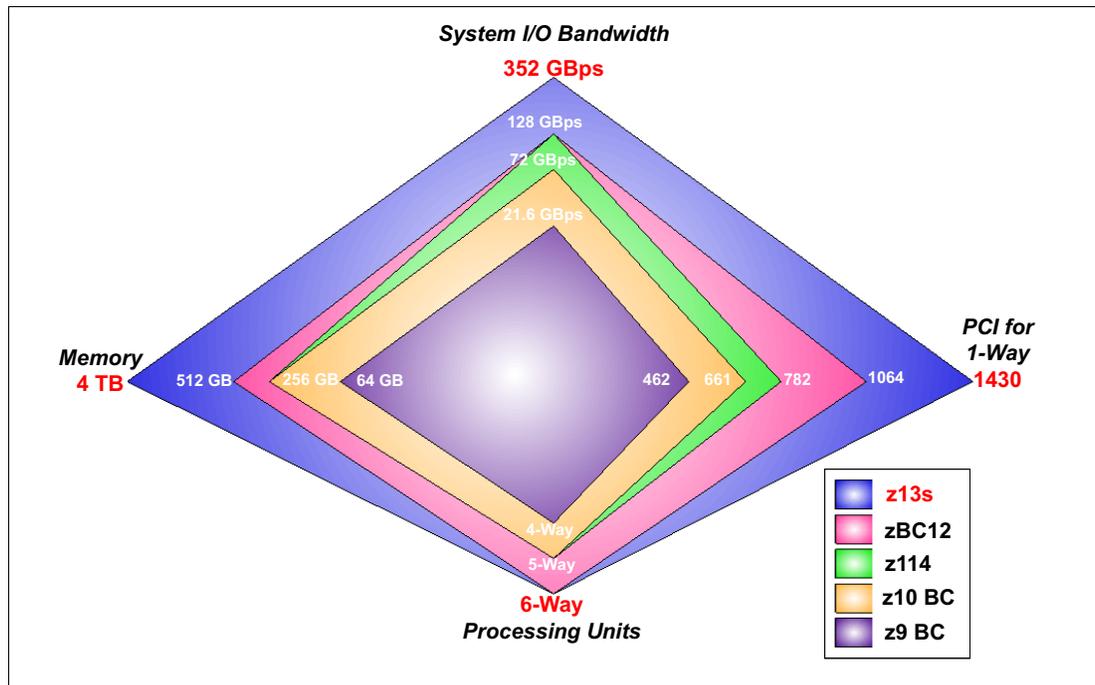


*Figure 1-3   Platform design: The z13s versus its predecessors*

You can find additional details about the z13s in Chapter 3, "IBM z13s hardware overview" on page 29. For an in-depth description of IBM z13s, see *IBM z13s Technical Guide*, SG24-8294.

## 1.2.2  Storage connectivity

Storage connectivity is provided on the z13 and z13s by Fibre Channel connection features. IBM Fibre Connection (FICON®) features follow the established Fibre Channel (FC) standards to support data storage and access requirements, along with the latest FC technology in storage and access devices. FICON Express features support these protocols:

► Native FICON

   This enhanced protocol (over FC) provides for communication across channels, channel-to-channel (CTC) connectivity, and with FICON devices such as disks, tapes, and printers. It is used in z/OS, z/VM, IBM z/VSE®, z/TPF, and Linux on z Systems environments.

► Fibre Channel Protocol (FCP)

   This is a standard protocol for communicating with disk and tape devices through Fibre Channel switches and directors. The FCP channel can connect to FCP SAN fabrics and access FCP/SCSI devices. FCP is used by z/VM, KVM for IBM z Systems, z/VSE, and Linux on z Systems environments.

FICON Express16S features provide significant improvements in throughput and response time for performance-critical middleware and to shrink the batch window required to accommodate I/O-bound batch work. FICON Express16S features are implemented using PCIe cards and offers better port granularity and improved capabilities over the previous FICON Express features. FICON Express16S features support a link data rate of 16 Gbps (4, 8, or 16 Gbps auto-negotiate) and is the preferred technology for new systems.

For details about all of the available FICON Express features, see 4.1, "I/O features at a glance" on page 44.

## 1.2.3  Network connectivity

The z13 and z13s are fully virtualized platforms that can support many system images at once. Therefore, network connectivity covers not only the connections between the platform and external networks but also supports specialized internal connections for intra-system communication.

The Open Systems Adapter-Express (OSA-Express) features provide local networking (LAN) connectivity and comply with IEEE standards. In addition, OSA-Express features assume several functions of the TCP/IP stack that normally are performed by the Processor Unit, allowing significant performance benefits by offloading processing from the operating system.

OSA-Express5S features provide significant improvements in throughput and response time for network traffic via a LAN environment. OSA-Express5S features are implemented using PCIe cards and are the preferred features for network connectivity.

### HiperSockets
IBM HiperSockets™ is an integrated function of the z Systems platforms that supplies attachments to up to 32 high-speed virtual local area networks with minimal system and network overhead.

HiperSockets is a function of the Licensed Internal Code (LIC) and provides LAN connectivity across multiple system images on the same z Systems platform by performing memory-to-memory data transfers in a secure way. The HiperSockets function eliminates the use of I/O subsystem operations and having to traverse an external network connection to communicate between logical partitions in the same z Systems platform. In this way, HiperSockets can help with server consolidation by connecting virtual servers and simplifying the enterprise network.

### 10GbE RoCE Express
The 10 Gigabit Ethernet (10GbE) RoCE Express feature uses Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) to provide fast memory-to-memory communications between two z Systems platforms or within a z Systems platform.

The feature is designed to help reduce consumption of CPU resources for applications that use the TCP/IP stack (such as IBM WebSphere® accessing an IBM DB2® database). It can also help reduce network latency with memory-to-memory transfers by using Shared Memory Communications–Remote Direct Memory Access (SMC-R) in z/OS V2R1 or later.

With SMC-R, you can transfer huge amounts of data quickly, at low latency. SMC-R is totally transparent to the application, requiring no code changes and thus enabling rapid time to value.

**Shared Memory Communications–Direct Memory Access (SMC-D)**

The z13s and z13 also use a recently introduced communications protocol called Shared Memory Communications–Direct Memory Access (SMC-D). SMC-D is similar to SMC-R but is intended for communications *within* the same z Systems platform, optimizing operating systems communications in a way that is transparent to socket applications. It reduces the CPU cost of TCP/IP processing in the data path, enabling highly efficient and application-transparent communications. Notably, SMC-D requires no additional physical resources (such as RoCE adapters, PCI bandwidth, ports, I/O slots, network resources, 10GbE switches, and so on). Instead, SMC-D uses either HiperSockets or an OSA-Express feature for establishing the initial connection.

For more information about the available OSA-Express features see 4.1, "I/O features at a glance" on page 44.

## 1.2.4 Special-purpose features and functions

When it comes to z Systems development, IBM takes a *total systems* view. The z Systems stack is built around digital services, agile application development, connectivity, and systems management, creating an integrated, diverse platform with specialized hardware and dedicated computing capabilities.

The z13 and z13s deliver a range of features and functions, allowing Processor Units to concentrate on computational tasks, while other, specialized features take care of the rest. Some of the special-purpose features and functions that are offered with the z13 and z13s include:

► Cryptography

The tamper-sensing and tamper-responding Crypto Express5S features provide high-performance cryptographic operations at twice the rate of the earlier Crypto Express4s. This specialized hardware performs DES, TDES, AES, RSA, SHA-1, and SHA-2 cryptographic operations. Crypto Express5S features are designed to meet the FIPS 140-2 Level 4 security requirements for hardware security modules.

► Flash Express

This feature is enhanced to improve availability and handling of paging workload spikes. It also acts as an overflow area for large data structures.

► IBM zEnterprise® Data Compression (zEDC) Express

This optional feature delivers an integrated solution to help reduce CPU consumption, optimize performance of compression-related tasks, and enable more efficient use of storage resources.

► z Appliance Container Infrastructure

This is a special-purpose firmware partition that is isolated from production and enables the secure deployment of software appliances.

► GDPS Virtual Appliance

The GDPS Virtual Appliance is a fully integrated, continuous availability, and disaster recovery solution for Linux on z Systems that can help improve availability and time-to-value.

► Dynamic Partition Manager (DPM)

DPM is a guided management interface used to define the z Systems hardware and virtual infrastructure, including integrated dynamic I/O management that runs KVM for IBM z Systems environments.

- ► z Advanced Workload Analysis Reporter (zAware)

  zAware is an integrated expert solution that uses sophisticated analytics to help identify potential workload problems and improve overall service levels. zAware runs analytics in z Appliance Container Infrastructure, enabling more effective detection of potential production problems.

- ► Simultaneous multithreading (SMT)

  With SMT, you can process up two simultaneous threads in a single core to optimize throughput. An operating system with SMT support can be configured to dispatch work to a thread on a zIIP[2] or an IFL[3].

- ► Single-instruction, multiple-data (SIMD)

  SIMD is set of instructions that allows optimization of code to complex mathematical models and business analytics vector processing. The set of SIMD instructions are a type of data-parallel computing and vector processing that can decrease the amount of code and accelerate mathematical computations with integer, string, character, and floating point data types.

- ► Dynamic memory reassignment

  With this technology, an algorithm is used to dynamically move memory between CPC drawers to improve performance without impacting the operating system.

You can read more about these and other z13 and z13s features in Chapter 5, "Strengths of the z13 and z13s" on page 71.

## 1.2.5 Capacity on Demand and performance

The z13 and z13s enable just-in-time deployment of processor resources. The Capacity on Demand (CoD) function allows users to dynamically change available system capacity. This function helps companies respond to new business requirements with flexibility and precise granularity.

Also contributing to the additional capacity on the z13 and z13s are numerous improvements in processor chip design, including new instructions, multithreading, and redesigned and larger caches.

In the same footprint, the z13 can deliver up to 40%[4] more capacity than the largest 101-way zEC12. The z13 1-way system has approximately 10% more capacity than the zEC12 1-way system. Similarly the z13s one-way system has approximately 34% more capacity than the zBC12 one-way, and 43% for the zBC12 6-way. The 20-way system has approximately 109% more capacity than the zBC12 H13 (13-way).

---

[2] IBM z Integrated Information Processor (zIIP) is used under z/OS for designated workloads, which include IBM Java Virtual Machine (JVM), various XML System Services, and others.

[3] An Integrated Facility for Linux (IFL) is exclusively used with Linux on z Systems and for running the z/VM or KVM hypervisor in support of Linux.

[4] Variations on all the observed increased performance depend on the workload type.

### 1.2.6  Reliability, availability, and serviceability

The z13 and z13s offer the same high quality of service and *reliability, availability,* and *serviceability* (RAS) that is traditional in z Systems platforms. The RAS strategy employs a building-block approach that is designed to meet stringent client requirements for achieving continuous, reliable operation. These are the RAS building blocks:

- ► Error prevention
- ► Error detection
- ► Recovery
- ► Problem determination
- ► Service structure
- ► Change management
- ► Measurement
- ► Analysis

The RAS design objective is to manage change by learning from previous product releases and investing in new RAS functionality to eliminate or minimize all sources of outages.

## 1.3  Software support

The z13 and z13s support a wide range of IBM and independent software vendor (ISV) software solutions. This includes traditional batch and online transaction processing (OLTP) environments, such as IBM Customer Information Control System (CICS®), IBM Information Management System (IMS™), and IBM DB2 It also includes these web services (in addition others that are not listed):

- ► Java platform

- ► Linux and open standards applications

- ► WebSphere

- ► IBM MobileFirst™ Platform Foundation (formerly IBM Worklight®) for mobile application development

The following operating systems are supported by IBM z13 and z13s:

- ► z/OS Version 2 Release 2 with PTFs (exploitation)
- ► z/OS Version 2 Release 1 with PTFs (exploitation)
- ► z/OS Version 1 Release 13 with PTFs (limited exploitation)
- ► z/OS Version 1 Release 12 with PTFs (limited exploitation)
- ► z/VM Version 6 Release 4 (previewed)
- ► z/VM Version 6 Release 3 with PTFs (exploitation)
- ► z/VM Version 6 Release 2 with PTFs (compatibility support)
- ► z/VSE Version 6 Release 1
- ► z/VSE Version 5 Release 1 with PTFs (compatibility support)
- ► z/VSE Version 5 Release 2 with PTFs (compatibility support)
- ► z/TPF Version 1 Release 1 (compatibility support)
- ► Linux on z Systems:
  - – SUSE: SUSE Linux Enterprise Server (SLES) 11 and SLES 12
  - – Red Hat: Red Hat Enterprise Linux (RHEL) 5, RHEL 6, and RHEL7
  - – Ubuntu: 16.04 LTS
- ► KVM for IBM z Systems 1.1.1

For more information about the z13 and z13s software support, see Chapter 6, "Operating system support" on page 95.

### 1.3.1  IBM compilers

Compilers are built with specific knowledge of the system architecture, which is used during code generation. Therefore, using the latest compilers is essential to extract the maximum benefit of a platform's capabilities. IBM compilers make use of the latest architecture enhancements and new instruction sets to deliver additional value.

With IBM Enterprise COBOL for z/OS and IBM Enterprise PL/I for z/OS, decades of IBM experience in application development can be used to integrate COBOL and PL/I with web services, XML, and Java. Such interoperability makes it possible to capitalize on existing IT investments, while smoothly incorporating new, web-based applications into the infrastructure.

z/OS, XL C/C++, and XL C/C++ for Linux on z Systems help with creating and maintaining critical business applications that are written in C or C++ to maximize application performance and improve developer productivity. These compilers transform C or C++ source code into executable code that fully leverages the z Systems architecture. This is possible thanks to hardware-tailored optimizations, built-in functions, performance-tuned libraries, and language constructs that simplify system programming and boost application runtime performance.

**2**

# IBM z13 hardware overview

This chapter expands on the description of the key hardware elements of the IBM z13 that was presented in Chapter 1, "Platforms matter" on page 1. It includes the following topics:

# 2.1 z13 models and upgrade paths

The z13 has an assigned machine type (MT) of 2964, which uniquely identifies the central processor complex (CPC). The z13 is offered in the following models:

- ► z13 N30: One CPC drawer and a maximum of 30 characterizable Processor Units
- ► z13 N63: Two CPC drawers and a maximum of 63 characterizable Processor Units
- ► z13 N96: Three CPC drawers and a maximum of 96 characterizable Processor Units
- ► z13 NC9: Four CPC drawers and a maximum of 129 characterizable Processor Units
- ► z13 NE1: Four CPC drawers and a maximum of 141 characterizable Processor Units

The last two digits of the model ID define the maximum number of Processor Units that are available for characterization. Processor Units are delivered in single-engine (core) increments. All z13 models use six, seven, or eight Processor Unit cores on six Processor Unit single chip modules in one CPC drawer.

Spare Processor Units, system assist processors (SAPs), and one integrated firmware processor (IFP) are integral to the system. Table 2-1 provides a summary that includes SAPs and spare Processor Units for the various models. For an explanation of Processor Unit characterization, see "Processor Unit characterization" on page 20.

*Table 2-1   z13 model summary (machine type 2964)*

| Model | Drawers/ Processor Units | Characterizable Processor Unit | Standard SAPs | Spares | Integrated firmware processor |
|-------|--------------------------|-------------------------------|---------------|--------|-------------------------------|
| N30 | 1/39 | 0–30 | 6 | 2 | 1 |
| N63 | 2/78 | 0–63 | 12 | 2 | 1 |
| N96 | 3/117 | 0–96 | 18 | 2 | 1 |
| NC9 | 4/156 | 0–129 | 24 | 2 | 1 |
| NE1 | 4/168 | 0–141 | 24 | 2 | 1 |

The z13 offers 231 capacity levels. There are 141 capacity levels based on the number of physically used *central processors* (CPs), plus up to 90 additional subcapacity models for the first 30 CPs. There is also one model for all *Integrated Facility for Linux* (IFL) or all *Internal Coupling Facility* (ICF) configurations. This topic is described in more detail in "Processor Unit characterization" on page 20.

Figure 2-1 summarizes the upgrade paths to the z13.
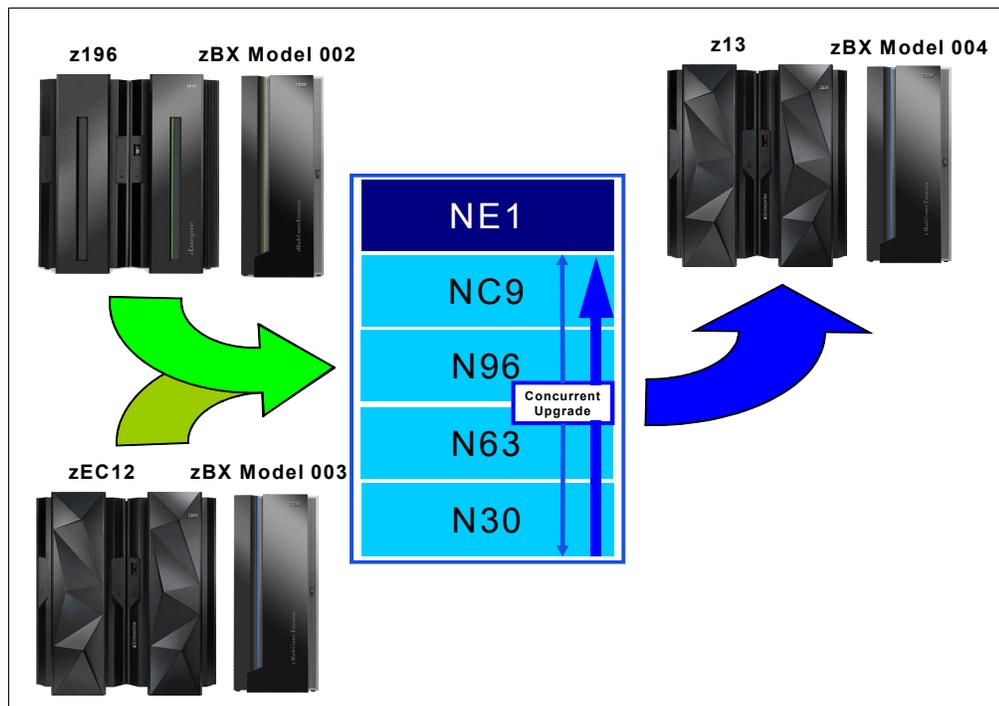


*Figure 2-1   z13 upgrade paths*

On the z13, concurrent upgrades (explained in Chapter 5, "Strengths of the z13 and z13s" on page 71) are available for CPs, IFLs, ICFs, z Systems Integrated Information Processors (zIIPs), and SAPs. However, concurrent Processor Unit upgrades require that additional Processor Units were physically installed, but not activated, at a previous time.

If an upgrade request cannot be accomplished in the customer's existing configuration, a hardware upgrade is required in which one or more CPC drawers is added to accommodate the desired capacity. On the z13, additional CPC drawers can be installed concurrently, but upgrading from any z13 model to model NE1 is disruptive because this upgrade requires the replacement of all installed CPC drawers.

Spare Processor Units are used to replace defective Processor Units and there are always two spare Processor Units on a z13. In the rare event of a Processor Unit failure, one of the spare Processor Units is immediately and transparently activated and assigned the characteristics of the failing Processor Unit.

When a z196 with a zBX Model 002 is upgraded to z13, the zBX is converted to a Model 004. When a zEC12 with a zBX Model 003 is upgraded to z13, the zBX is converted to a Model 004. The virtualization and configuration data are preserved, but the upgrade process requires downtime.

## 2.2 z13 frames and cabling

The z13 is always a two-frame system, the *A Frame* and the *Z Frame,* and can be delivered as an air-cooled system or as a water-cooled one.

The two frames form the z13. The type of I/O drawer (and the number of drawers) can vary based on the number of I/O features. For a new build system, a combination of up to five PCIe I/O drawers can be installed. A miscellaneous equipment specification (MES) can carry forward up to two I/O drawers (8-slot). This configuration provides for a maximum of 16 non-PCIe features that can be carried forward.

In addition, the z13 (both new builds and MES orders) offers top-exit options for the fiber optic and copper cables used for I/O and power. These options (*Top Exit Power* and *Top Exit I/O Cabling*) give you more flexibility in planning where the system will reside, potentially freeing you from running cables under a raised floor and increasing air flow over the system.

The radiator-cooled z13 models support installation on raised floor and non-raised floor environments. For water-cooled models, only the raised floor option is available.

Figure 2-2 shows an internal, front view of the two frames of an air-cooled z13 system with the maximum five PCIe I/O drawers, including the top-exit I/O and power cable options.
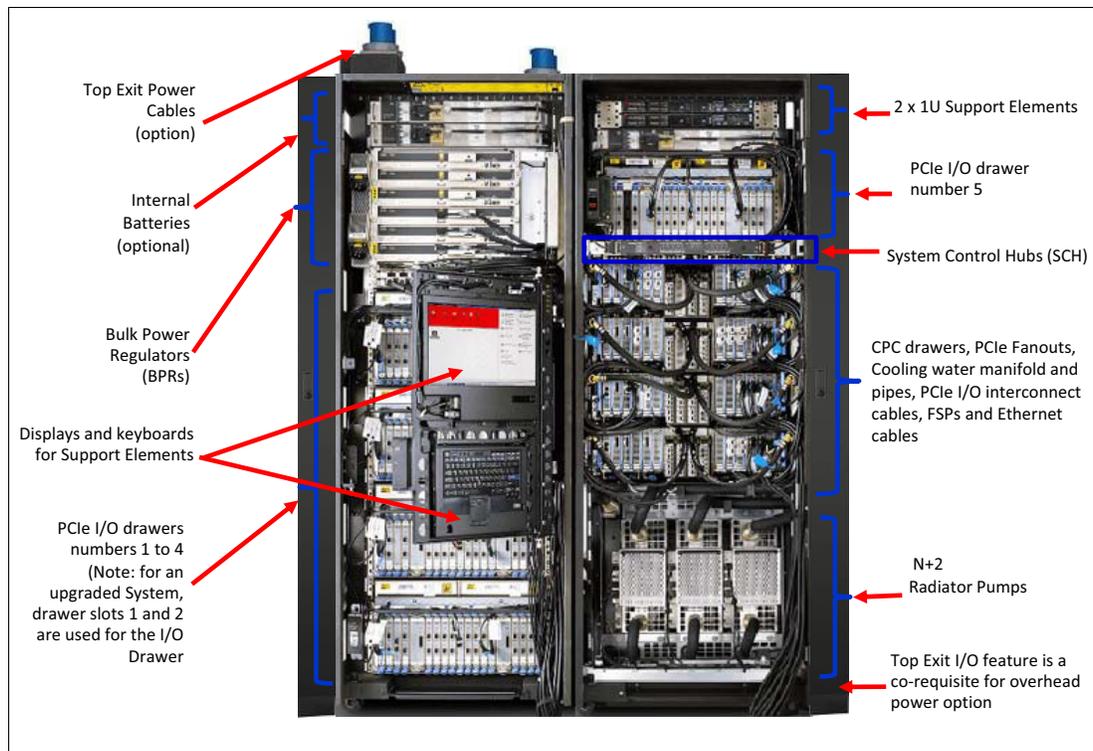


*Figure 2-2   z13 internal, front view: Air-cooled platform with five PCIe I/O drawers*

Figure 2-3 shows an internal, front view of the two frames of a water-cooled platform without the top exit I/O and power cable options.
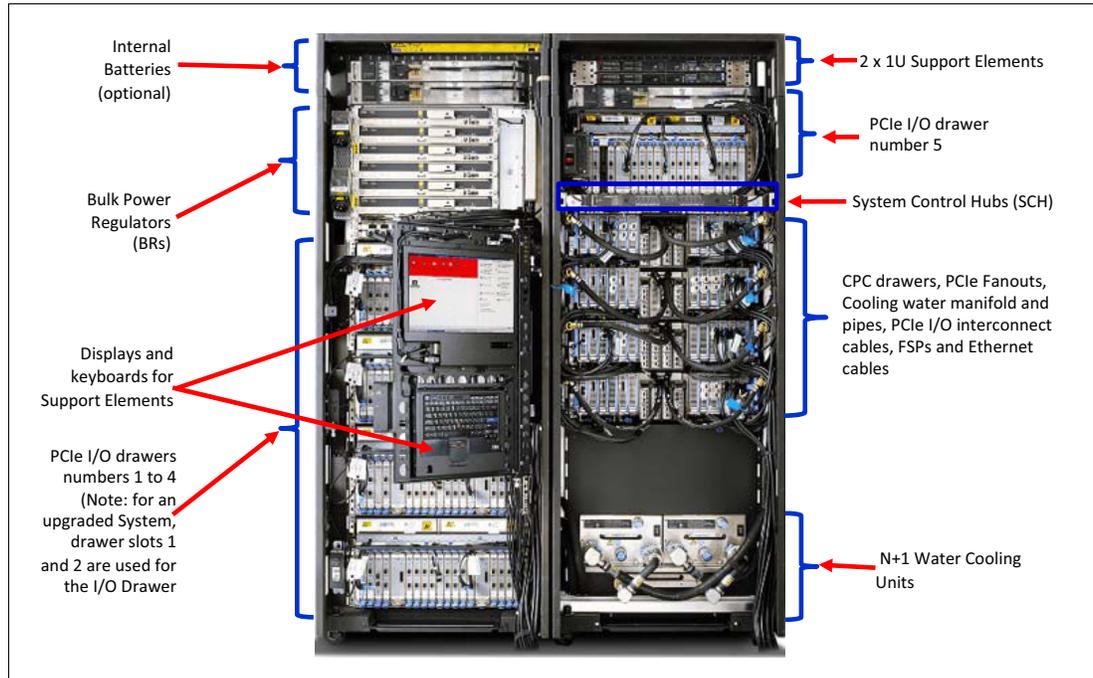


*Figure 2-3   z13 internal, front view: Water-cooled platform with five PCIe I/O drawers*

## 2.3  z13 CPC drawers

The z13 is a multiple CPC drawer system that is capable of holding up to four drawers in the A Frame.

Each CPC drawer contains the following elements:

► Single chip modules

– Six Processor Unit single chip modules, each containing eight Processor Unit cores (water-cooled)

– Two storage controller single chip modules, with a total of 960 MB L4 cache

Single chip modules are described in 2.3.1, "Single chip modules" on page 18. Also refer to Table 2-1 on page 14 for the model summary and the relation between the number of CPC drawers and number of available Processor Units.

► Memory

– A minimum of 256 GB and a maximum of 2.5 TB of memory (excluding 96 GB HSA) is available for client use. See Table 2-2 on page 21 for details.

– Either 20 or 25 memory DIMMs are plugged in a CPC drawer.

► Fanouts

Fanouts provide the connectivity from the CPC drawer to the I/O features via the I/O drawers. A combination of up to four InfiniBand host channel adapter fanouts (HCA3-Optical, HCA2-Copper) and up to 10 PCIe third generation fanouts (PCIe Gen3).

Each fanout has one, two, or four ports, so up to 40 connections are supported:

– One-port PCIe 16 GBps I/O fanout, each supporting one domain in 32-slot PCIe I/O drawers)

– ICA SR two-port fanout for coupling links (two links, 8 GBps each)

– HCA3-O 12x InfiniBand fanout for coupling links (two ports at 6 GBps each)

– HCA3-O LR 1x InfiniBand fanout for coupling links (four ports, 5 Gbps each)

– HCA2-Copper fanouts (two ports per fanout, 6 Gbps each) supported only for I/O drawers that are carried forward, for a maximum of two features

► Two Distributed Converter Assemblies (DCAs) that provide power to the CPC drawer

Loss of one DCA leaves enough power to satisfy the power requirements of the entire drawer. The DCAs can be concurrently maintained.

► Two Flexible Support Processors (FSP) that provide redundant interfaces to the internal management network

As shown in Figure 2-4, all CPC drawers are interconnected with high-speed communications links through the L4 shared caches. The z13 has 960 MB of L4 cache per CPC drawer, which is 2.5 times larger than its predecessor system.
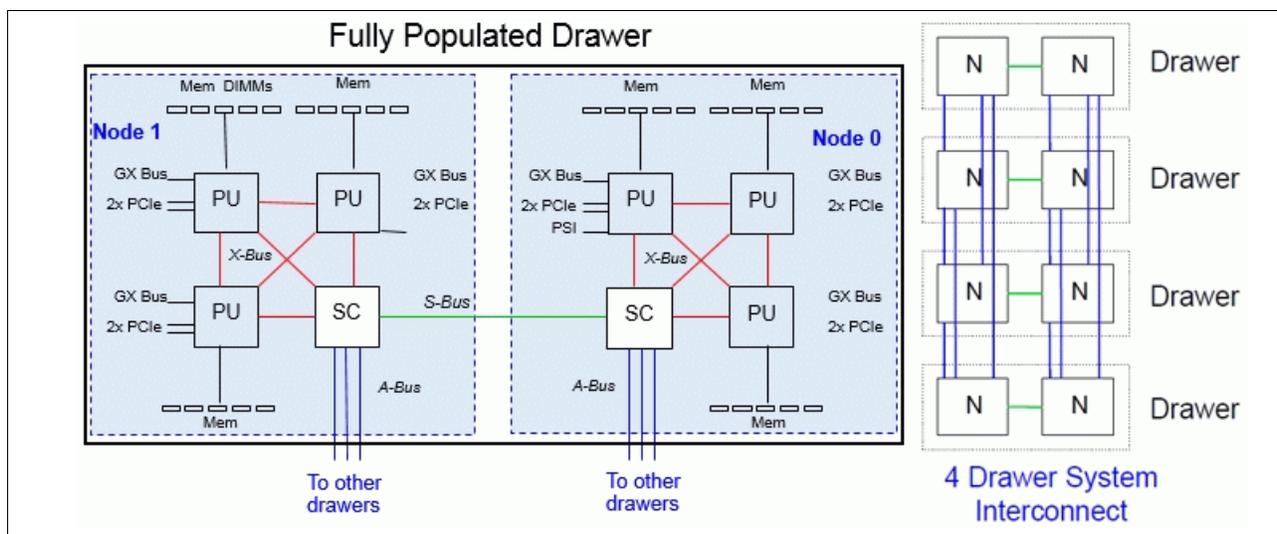


*Figure 2-4   z13 CPC drawer communication topology*

The design used to connect the Processor Unit and storage control allows the system to be operated and controlled by the IBM Processor Resource/Systems Manager™ (PR/SM™) facility as a memory-coherent symmetrical multiprocessor system (SMP).

## 2.3.1  Single chip modules

At the heart of the system are *single chip modules*, which are high-performance, glass-ceramic chips that provide the highest level of processing integration in the industry. There are eight single chip modules per CPC drawer. Six of these single chip modules hold the processor chips (Processor Unit chips), and two single chip modules hold storage control chips. Each Processor Unit chip has six, seven, or eight active cores and L1, L2, and L3 caches. The two storage control chips hold L4 caches, as shown in Figure 2-5.

For z13, two CPC drawer configurations are offered with 39 or 42 Processor Units. All the models employ CPC drawers with 39 Processor Units except for the model NE1, which has four CPC drawers with 42 active Processor Units, for a total of 168 Processor Units. Figure 2-5 also shows the CPC drawer chip and cache structure.
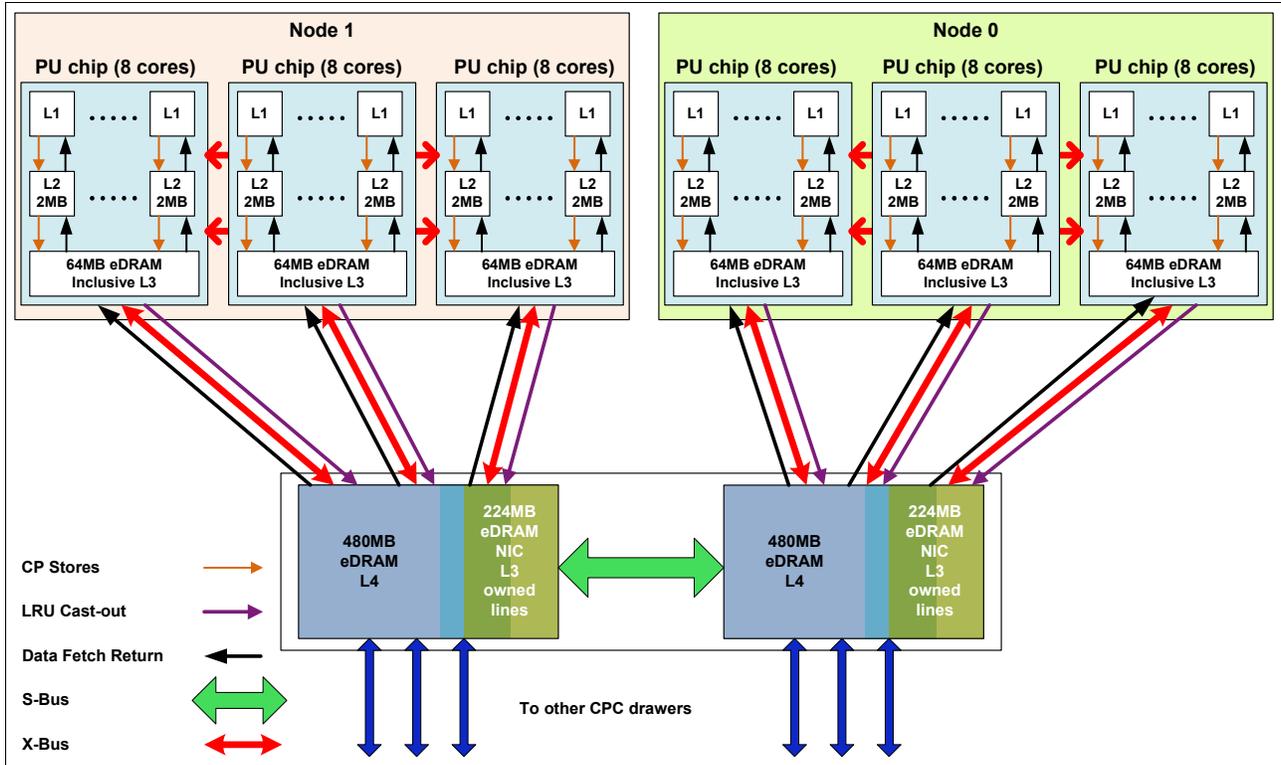


*Figure 2-5    z13 CPC drawer chips and cache structure*

Each 678.78 mm storage control chip includes 480 MB of eDRAM cache, interface logic for three Processor Unit chips each, and SMP fabric logic. The two storage control chips are configured to provide a single 960 MB L4 cache that is shared by all cores in the CPC drawer. This amount of cache provides a total of 3.8 GB of cache if all four CPC drawers are implemented, yielding outstanding SMP scalability on real-world workloads.

## 2.3.2  Processor Unit

The *Processor Unit* (PU) is the generic term for a IBM z/Architecture® processor. A Processor Unit is embedded in a z Systems chip core. Each Processor Unit is a superscalar processor with the following attributes:

► Up to six instructions can be decoded per clock cycle.

► Up to ten instructions can be in execution per clock cycle.

► Instructions can be issued out-of-order. The Processor Unit uses a high-frequency, low-latency pipeline, providing robust performance across a wide range of workloads.

► Memory accesses might not be in the same instruction order (out-of-order operand fetching).

► Most instructions flow through a pipeline with varying numbers of steps for different types of instructions. Several instructions can be in execution at any moment, subject to the maximum number of decodes and completions per cycle.

## Processor Unit chip cache

The on-chip cache for the Processor Unit (core) chip works in this way:

► Each Processor Unit has an L1 cache (private) that is divided into a 96 KB cache for instructions and a 128 KB cache for data.

► Each Processor Unit also has a private L2 cache, with 2 MB D-cache (D stands for data) and 2 MB I-cache (I stands for instruction).

► Each Processor Unit chip also contains a 64 MB L3 cache that is shared by all eight Processor Units on the chip. The shared L3 cache uses eDRAM.

This on-chip cache implementation optimizes system performance for high-frequency processors.

Each L1 cache has an associated *translation lookaside buffer* (TLB) of 512 entries. 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.

## CPU sparing

Hardware fault detection is embedded throughout the design and is combined with comprehensive instruction-level retry and dynamic CPU sparing. This function provides the reliability and availability that is required for true mainframe integrity.

## On-chip cryptographic hardware

Dedicated on-chip cryptographic hardware includes extended key and hash sizes for the Advanced Encryption Standard (AES) and Secure Hash Algorithm (SHA) as well as support for UTF8 to UTF16 conversion. This cryptographic hardware is available with any processor type, for example central processor (CP), IBM System z® Integrated Information Processor (zIIP), or Integrated Facility for Linux (IFL).

## Software support

The z13 Processor Units provide full compatibility with existing software for ESA/390 and z/Architecture and extends the Instruction Set Architecture (ISA) to enable enhanced functionality and performance. Several hardware instructions that support more efficient code generation and execution are introduced in the z13:

► Hardware decimal floating point (HDFP)
► Transactional Execution Facility
► Runtime Instrumentation Facility
► Single-instruction, multiple-data (SIMD)

These features are further described in Chapter 5, "Strengths of the z13 and z13s" on page 71.

## Processor Unit characterization

Processor Units are ordered in single increments. The internal system functions, which are based on the configuration that is ordered, characterize each Processor Unit (core) into one of various types during system initialization, which is often called a *power-on reset* (POR) operation. Characterizing Processor Units dynamically without a POR is possible using a process called *Dynamic Processor Unit Reassignment*. A Processor Unit that is not characterized cannot be used.

Each Processor Unit (core) can be characterized as follows:

► Central processor (CP)
► Integrated Facility for Linux (IFL) processor

- ► z Integrated Information Processor (zIIP)
- ► Internal Coupling Facility (ICF)
- ► System assist processor (SAP)
- ► Integrated firmware processor (IFP)

At least one CP must be purchased with a zIIP or before a zIIP can be purchased. Clients can purchase up to two zIIPs for each purchased CP (assigned or unassigned) on the system. However, an LPAR definition can go beyond the 1:2 ratio. For example, on a system with two CPs, a maximum of four zIIPs can be installed. An LPAR definition for that system can contain up to two logical CPs and four logical zIIPs. Another possible configuration is one logical CP and three logical zIIPs.

Converting a Processor Unit from one type to any other type is possible by using the *Dynamic Processor Unit Reassignment* process. These conversions happen concurrently with the operation of the system.

> **zIIPs:** The addition of ICFs, IFLs, zIIPs, and SAP to the z13 does not change the system capacity setting or its MSU rating. Only CPs work that way. IBM does not impose any software charges on work that is dispatched on zIIPs.
>
> **zAAPs:** The zEC12 and zBC12 were the last z Systems servers to offer support for zAAPs. IBM supports running zAAP workloads on zIIPs (called *zAAP on zIIP*). This change is intended to help simplify capacity planning and performance management while still supporting all the currently eligible workloads. IBM has provided a PTF for APAR OA38829 on z/OS V1R12 and V1R13. This PTF removes the restriction that prevents workloads eligible for a zAAP from running on a zIIP (when a zAAP is installed).

### 2.3.3 Memory

Maximum physical memory size is directly related to the number of CPC drawers in the system. And typically, a system has more memory installed than was ordered, because part of the installed memory is used to implement the redundant array of independent memory (RAIM) design. On the z13, this configuration results in up to 2.5 TB of available memory per CPC drawer and up to 10 TB for a four-drawer system.

The hardware system area (HSA) on the z13 has a fixed amount memory (96 GB) that is managed separately from client memory. However, the maximum amount of orderable memory can vary from the theoretical number due to dependencies on the memory granularity.

Table 2-2 lists the minimum and maximum memory sizes for each z13 model.

*Table 2-2    z13 model memory ranges*

| Model | CPC drawers | Memory range |
|-------|-------------|------------------|
| N30   | 1           | 256 GB to 2560 GB |
| N63   | 2           | 512 GB to 5120 GB |
| N96   | 3           | 768 GB to 7669 GB |
| NC9   | 4           | 1024 GB to 10 TB |
| NE1   | 4           | 1024 GB to 10 TB |

On z13 systems, the granularity for memory orders varies from 32 GB to 512 GB. Table 2-3 shows the memory increments, depending on installed memory.

*Table 2-3   z13 memory increments and ranges*

| Memory increment | Memory range (GB) |
|---|---|
| 32 GB | 64 to 192 |
| 64 GB | 256 to 448 |
| 96 GB | 544 to 928 |
| 128 GB | 1056 to 1440 |
| 256 GB | 1696 to 6048 |
| 512 GB | 6560 to 10144 |

Physically, memory is organized in these ways:

► A CPC drawer always contains a minimum of 320 GB of installed memory, of which 256 GB is usable by the operating system.

► A CPC drawer can have more installed memory than is enabled. The excess memory can be enabled by a Licensed Internal Code load.

► Memory upgrades are first satisfied using already installed but unused memory capacity, until it is exhausted. When no more unused memory is available from the installed cards, either the cards must be upgraded to a higher capacity or a CPC drawer with more memory must be installed.

When activated, IBM Processor Resource/Systems Manager (PR/SM) tries to allocate the memory of an LPAR in a single CPC drawer, but if that is not possible, it can use memory resources located in any CPC drawer. No matter which CPC drawer the memory is in, an LPAR has access to that memory, if it is allocated. Despite the CPC drawer structure, the z13 is still a symmetric multiprocessor (SMP) system because the Processor Units have access to all the available memory.

A memory upgrade is considered to be concurrent when it requires no change of the physical memory cards. A memory card change is disruptive when no use is made of *Enhanced Drawer Availability* (EDA). In a multiple CPC drawer system, a single CPC drawer can be concurrently removed and reinstalled for a repair with EDA. For a description of EDA, see *IBM z13 Technical Guide*, SG24-8251.

For model upgrades involving the addition of a CPC drawer, the minimum usable memory increment (256 GB) is added to the system. During an upgrade, adding a CPC drawer is a concurrent operation. as is adding physical memory in the new drawer.

### Concurrent memory upgrade

If physical memory is available, memory can be upgraded concurrently by using *Licensed Internal Code Configuration Control (LICCC)*, as described. The *plan ahead memory function* that is available with the z13 enables nondisruptive memory upgrades by having in the system pre-plugged memory (based on a target configuration). Pre-plugged memory is enabled through an LICCC order that is placed by the client.

### Redundant array of independent memory (RAIM)

Redundant array of independent memory (RAIM) technology makes the memory subsystem, in essence, a fully fault-tolerant N+1 design. The RAIM design automatically detects and

recovers from failures of dynamic random access memory (DRAM), sockets, memory channels, or dual inline memory modules (DIMMs).

The RAIM design is fully integrated in the z13 and has been enhanced to include one Memory Controller Unit (MCU) per processor chip, with five memory channels and one DIMM per channel. A fifth channel in each MCU enables memory to be implemented as a RAIM. This technology has significant as capabilities (reliability, availability, and serviceability) in the area of error correction. Bit, lane, DRAM, DIMM, socket, and complete memory channel failures, including many types of multiple failures, can be detected and corrected.

### 2.3.4  Hardware system area

The hardware system area (HSA) is a fixed-size, reserved area of memory that is separate from the client-purchased memory. The HSA is used for several internal functions, but the bulk of it is used by channel subsystem functions.

The fixed size 96 GB HSA for z13 is large enough to accommodate any LPAR definitions or changes, thus eliminating most outage situations and the need for extensive preplanning.

A fixed, large HSA allows the dynamic I/O capability of the z13 to be enabled by default. It also enables the dynamic addition and removal of the following features:

► LPAR to new or existing channel subsystem (CSS)
► CSS (up to six can be defined in z13)
► Subchannel set (up to four can be defined in z13)
► Devices, up to the maximum number permitted, in each subchannel set
► Logical processors by type
► Cryptographic adapters

## 2.4  z13 I/O system structure

The z13 support the following types of internal I/O infrastructures:

► Generation 3 PCIe-based infrastructure for PCIe I/O drawers (PCIe Gen3)
► InfiniBand-based infrastructure and I/O drawers (carry forward on an MES only)

The PCIe I/O infrastructure consists of the following features:

► PCIe Gen3 fanouts in the CPC drawers which support 16 Gbps connectivity to the PCIe I/O drawer (zEC12 used PCIe Gen2 fanouts at 8 GBps)

► Up to five 7U PCIe I/O drawers each with 32 slots (eight slots per I/O domain) for PCIe I/O features

The InfiniBand I/O infrastructure (carry forward only) consists of the following features:

► InfiniBand fanouts in the z13 CPC drawer, which support the 6 GBps InfiniBand I/O interconnect

► InfiniBand I/O card domain multiplexers with redundant I/O interconnect in the following configuration: up to two (5U), 8-slot, 2-domain I/O drawers (carry forward only)

► FICON Express8 (carry forward only)

**Ordering of I/O features:** Ordering of I/O feature types determines the appropriate mix of PCIe I/O drawers and I/O drawers (order-dependent).

Figure 2-6 shows a high-level view of the I/O system structure for the z13.
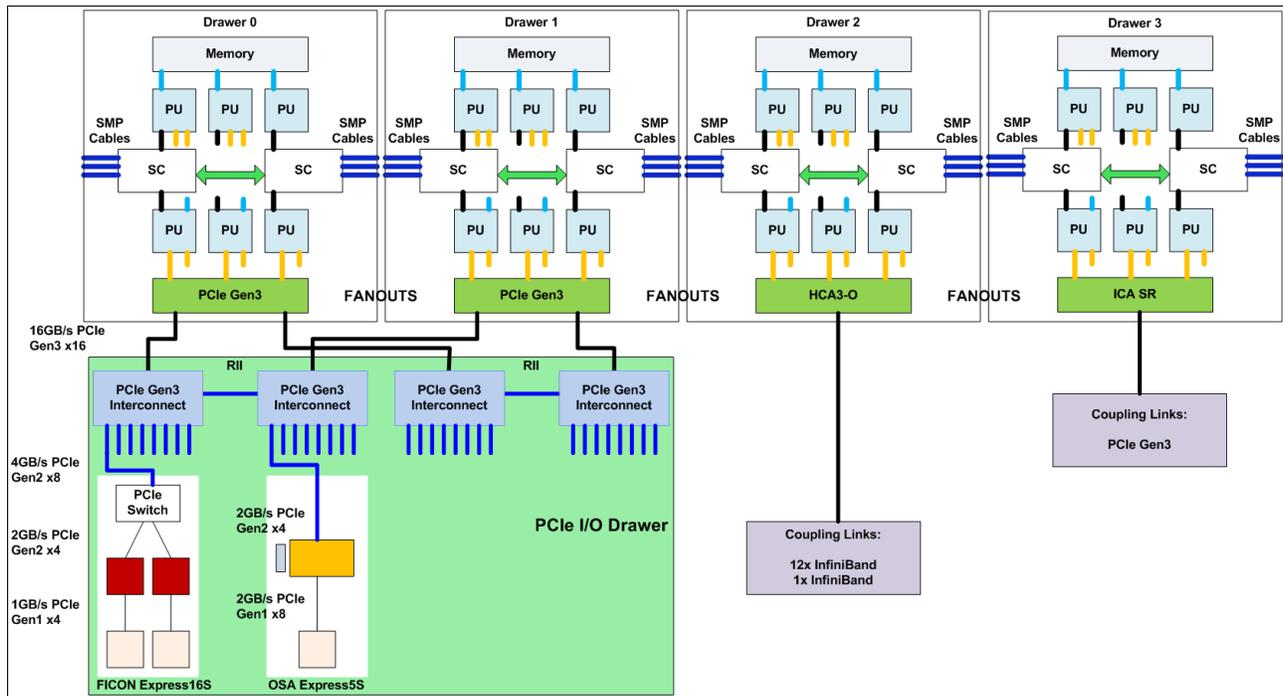


*Figure 2-6   z13 I/O system structure*

The z13 supports five fanout types (for fanout location, see Figure 2-7 on page 25), which are located at the front on the CPC drawer:

- ► ICA SR
- ► HCA3-O
- ► HCA3-O LR
- ► PCIe Gen3
- ► HCA2-C (carry forward only)

The HCA3-O LR fanout includes four ports. The PCIe Gen3 fanout has one port, and the other fanouts have two ports.

The following types of internal I/O connectivity support the PCIe I/O drawer and I/O drawer:

- ► PCIe connections to the PCIe I/O drawers from the PCIe fanouts through copper cables. The I/O features supported through the port on the fanouts are listed in 2.5, "z13 power and cooling" on page 25.

- ► InfiniBand (IFB) connections to the existing I/O drawers from the host channel adapter (HCA2-C) fanouts through copper cables in only a carry-forward MES.[1] The two ports in the fanout are dedicated to connect to an InfiniBand multiplexer (IFB-MP) card in the I/O drawer.

For coupling link connectivity (Parallel Sysplex or STP configuration), the z13 supports the following fanouts:

- ► ICA SR
- ► HCA3-O
- ► HCA3-O LR

---

[1]  MES: Miscellaneous equipment specification (upgrade or change)

The z13 CPC drawer (Figure 2-7) can have up to 10 1-port PCIe Gen3 fanouts (numbered LG02 to LG06 and LG11 to LG15) and up to four 2-port or 4-port InfiniBand fanouts for each CPC drawer, (numbered LG07 to LG10), which are used to connect to I/O drawers, PCIe I/O drawers, or for Parallel Sysplex InfiniBand and PCIe connectivity. In a system that is configured for maximum availability, alternate paths maintain access to critical I/O devices, such as disks and networks.
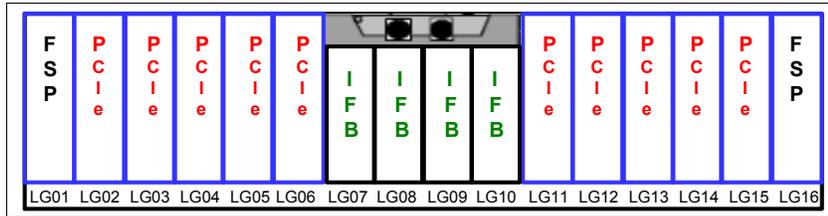


*Figure 2-7   z13 CPC drawer front view*

The PCIe I/O drawer is a two-sided drawer (I/O features are present on both sides) that is 7U high. The drawer contains 32 slots, four switch cards (two in the front and two in the rear) to support four I/O domains that each contain eight features of any type (FICON Express16S, FICON Express8S, OSA-Express5S, OSA-Express4S, Crypto Express5S, Flash Express, zEDC Express, and 10GbE RoCE Express). Two DCAs provide redundant power and two air moving devices (AMDs) provide redundant cooling to the PCIe I/O Drawer.

Each I/O drawer supports two I/O domains (A and B) for a total of eight I/O slots. Each I/O domain uses an IFB-MP card in the I/O drawer and a copper cable to connect to a host channel adapter (HCA) fanout in the CPC drawer.

All features in the I/O drawer are installed horizontally. The two DCAs distribute power to the I/O drawer.

The IFB-MP cards are installed at location 09 at the rear side of the I/O drawer. The I/O features are installed from the front and rear sides of the I/O drawer. Two I/O domains are supported. Each I/O domain has up to four FICON Express8 features. The FICON Express8 I/O features are connected to the IFB-MP card through the backplane board.

Additional details about the I/O features of the z13 are available in Chapter 4, "Supported features and functions" on page 43.

## 2.5  z13 power and cooling

The z13 meets the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Class A2 specifications. ASHRAE is an organization devoted to the advancement of indoor-environment-control technology in the heating, ventilation, and air conditioning industry.

The power and cooling system of the z13 builds on that of its predecessor, the zEC12, with several newly developed technologies. However, the underlying power service specifications of the z13 are almost identical to its predecessors. Total power consumption with the maximum system configuration has increased only by about 5% compared to those previous models.

A closed, internal water cooling loop is used to cool single chip modules in the CPC drawers of the z13. Extracting the heat from the internal water loop can be done either with a radiator (air-cooled system) or a water cooling unit (water-cooled system). Conversion from air to water cooled systems, and vice versa, is not available.

## 2.5.1 Radiator (air) cooling option

The cooling system in the z13 is redesigned for better availability and lower cooling power consumption. The radiator design is a closed-loop water cooling pump system for the single chip modules in the CPC drawers. It is designed with N+2 pumps, blowers, controls, and sensors. The radiator unit is cooled by air.

## 2.5.2 Water cooling option

The z13 continues to offer the choice of using a building's chilled water to cool the system by employing water cooling unit (WCU) technology. The single chip modules in the CPC drawer are cooled by an internal, closed, water cooling loop. In the internal closed loop, water exchanges heat with building-chilled water (provided by the client) through a cold plate.

In addition to the single chip modules, the internal water loop also circulates through two heat exchangers that are in the path of the exhaust air in the rear of the frames. These heat exchangers remove approximately 60-65% of the residual heat from the I/O drawers.

The z13 operates with two fully redundant WCUs. One water cooling unit can support the entire load and replacement of a WCU is fully concurrent. If there is a total loss of building-chilled water or if both water cooling units fail, the rear door heat exchangers cool the internal water cooling loop.

## 2.5.3 High Voltage Direct Current power feature

With the optional High Voltage Direct Current (HV DC) power feature, the z13 can directly connect to DC power input and improve data center energy efficiency by removing the need for an additional DC-to-AC inversion step. This feature can help achieve both data center UPS and power distribution energy savings.

## 2.5.4 Power considerations

The z13 operates with two sets of redundant power supplies. Each set has its own individual power cords or pair of power cords, depending on the number of Bulk Power Regulator (BPR) pairs installed. Power cords attach a 3-phase, 50/60 Hz, 200–480 V AC power source or 380–520 V DC power source. The loss of one power supply alone has no effect on system operation.

The optional Balanced Power Plan Ahead feature is available for future growth, also assuring adequate and balanced power for all possible configurations. With this feature, downtime for upgrading a system is eliminated because the initial installation includes the maximum power requirements in terms of Bulk Power Regulators (BPR) and power cords. The Balance Power Plan Ahead feature is not available with DC and 1-phase line cords.

Additional single-phase outlets (customer provided) are required for ancillary equipment such as the Hardware Management Console and its display.

Specific power requirements depend on the cooling facility that is installed, the number of CPC drawers, and the number and type of I/O units that are installed. You can find maximum power consumption tables for the various configurations and environments in *IBM z13 Installation Manual for Physical Planning*, GC28-6938.

You can also refer to the power and weight estimation tool that is available at IBM Resource Link®.

# IBM z13s hardware overview

This chapter expands on the description of the key hardware elements of the IBM z13s that was presented in Chapter 1, "Platforms matter" on page 1. It includes the following topics:

# 3.1  z13s models and upgrade paths

The z13s has an assigned machine type (MT) of 2965, which uniquely identifies the central processor complex (CPC). The z13s is offered in the following models:

- ► The *z13s N10* includes one CPC drawer with one node and a maximum of 10 customizable Processor Units.
- ► The *z13s N20* includes one or two CPC drawers with two nodes each and a maximum of 20 customizable Processor Units. The second drawer in the N20 is dedicated to I/O or memory requirements.

The last two digits of the model ID define the maximum number of Processor Units that are available for characterization. Processor Units are delivered in single-engine (core) increments. All z13s models use six or seven Processor Unit cores on two or four Processor Unit single chip modules in one CPC drawer.

Spare Processor Units, system assist processors (SAPs), and one integrated firmware processor (IFP) are integral to the system. Table 3-1 provides a summary that includes SAPs and spare Processor Units for the various models. For an explanation of Processor Unit characterization, see "Processor Unit characterization" on page 35.

*Table 3-1   z13s model summary (machine type 2965)*

| Model | Drawers/ Processor Units | Characterizable Processor Units | Standard SAPs | Spares | Integrated firmware processor |
|-------|--------------------------|----------------------------------|---------------|--------|-------------------------------|
| N10   | 1/13                     | 0–10[a]                          | 2             | 0      | 1                             |
| N20   | 1 or 2[b]/26             | 0–20[a]                          | 3             | 2      | 1                             |

a. A maximum of 6 CPs are characterizable
b. Depends on the number of I/O features and amount of memory ordered.

The z13s offers 26 capacity levels times six CPs for 156 settings. There is also one model for all *Integrated Facility for Linux* (IFL) or all *Internal Coupling Facility* (ICF) configurations. This topic is described in more detail in "Processor Unit characterization" on page 35.

Figure 3-1 summarizes the upgrade paths to the z13s.
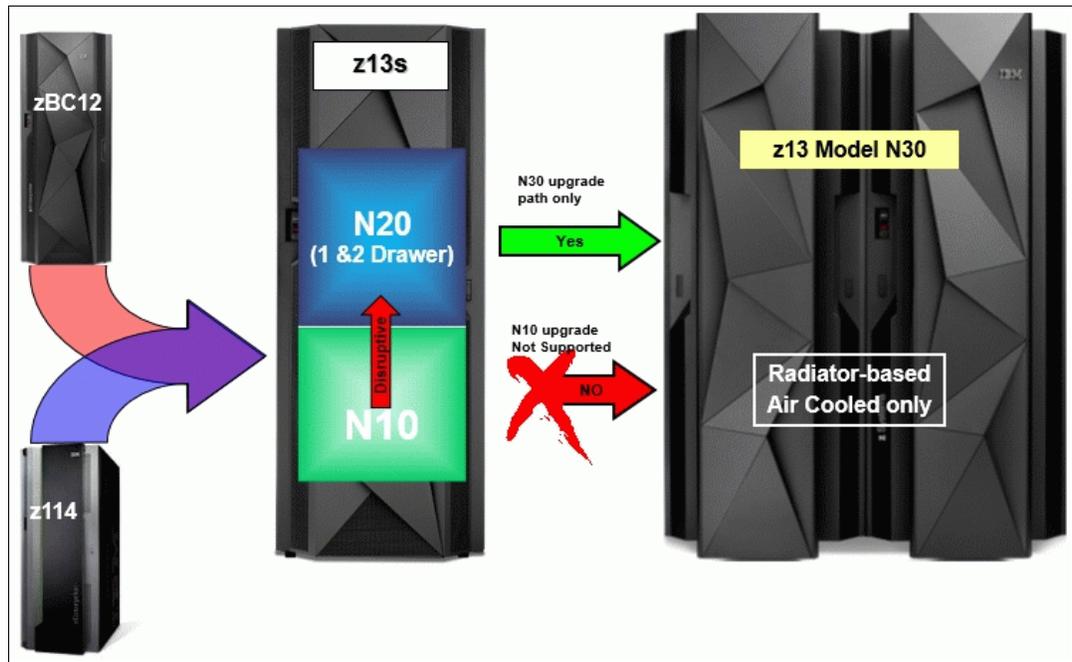


*Figure 3-1 z13s upgrade paths*

On the z13s, concurrent upgrades (explained in Chapter 5, "Strengths of the z13 and z13s" on page 71) are available for CPs, IFLs, ICFs, z Systems Integrated Information Processors (zIIPs), and SAPs. However, concurrent Processor Unit upgrades require that additional Processor Units are physically installed, but not activated, at a previous time.

If an upgrade request cannot be accomplished within the customer's existing configuration, a hardware upgrade is required in which one or more CPC drawers is added to accommodate the desired capacity. For z13s, adding an additional CPC drawer is always disruptive.

Spare Processor Units are used to replace defective Processor Units. Model N20 has two spare Processor Units, and model N10 uses unassigned Processor Units, if available, as spare Processor Units. In the rare event of a Processor Unit failure, a spare Processor Unit is immediately and transparently activated and assigned the characteristics of the failing Processor Unit.

When a z114 with a zBX Model 002 is upgraded to z13s, the zBX is converted to a Model 004. When a zBC12 with a zBX Model 003 is upgraded to z13s, the zBX is converted to a Model 004. The virtualization and configuration data are preserved, but the upgrade process requires downtime.

## 3.2  z13s frames and cabling

The z13s is always a single-frame system (labeled the *A Frame*) and is always delivered as an air-cooled system.

The frame forms the central processing complex (CPC) and contains one or two CPC drawers. The number and type of drawers dedicated to I/O features can vary and depends on the number of I/O features, the amount of memory, or both factors. For a new build system, a combination of up to two PCIe I/O drawers can be installed. A miscellaneous equipment

specification (MES) can carry forward one legacy I/O drawer (8 slot I/O drawer). This configuration provides for a maximum number of eight FICON Express8 feature cards that can be carried forward.

In addition, the z13s (both new builds and MES orders) offers top-exit options for the fiber optic and copper cables used for I/O and power. These options (*Top Exit Power* and *Top Exit I/O Cabling*) give you more flexibility in planning where the system will reside, potentially freeing you from running cables under a raised floor and increasing air flow over the system.

Figure 3-2 shows internal front and rear views of the air-cooled z13s system with one CPC drawer, two PCIe drawers, and one I/O drawer.
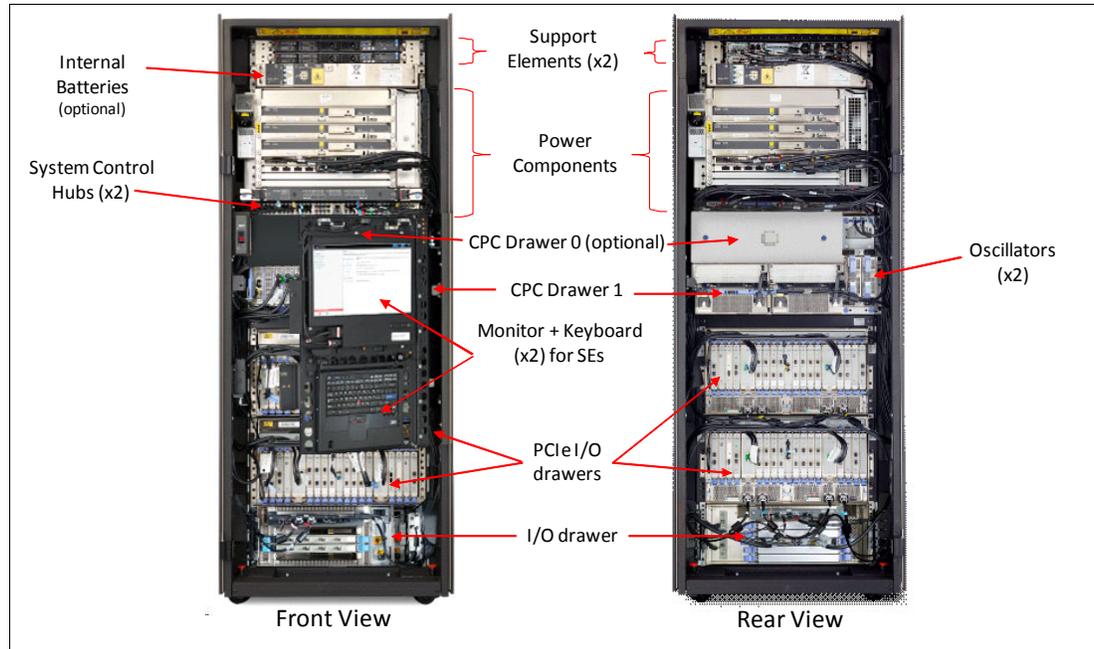


*Figure 3-2   z13s internal front and rear view: Air-cooled platform with two PCIe drawers and one I/O drawer*

## 3.3  z13s CPC drawers

The z13s can have up to two CPC drawers. Each CPC drawer contains these elements:

► Single chip modules

  – Two-Processor Unit (model N10) or four-Processor Unit (model N20) single chip modules, each containing eight Processor Unit cores (air-cooled)

  – One (model N10) or two (model N20) system controller single chip modules, with 480 MB L4 cache per single chip module

  Single chip modules are described in 3.3.1, "Single chip modules" on page 33. See Table 3-1 on page 30 for the model summary and the relation between the number of CPC drawers and number of available Processor Units.

► Memory

  – A minimum of 128 GB and a maximum of 2048 GB of memory (excluding 40 GB HSA) is available for client use. See Table 3-2 on page 36 for details.

  – Either 10 or 20 memory DIMMs are plugged in a CPC drawer.

► Fanouts

   Fanouts provide the connectivity from the CPC drawer to the I/O features via the I/O drawers. A combination of up to four InfiniBand host channel adapter fanouts (HCA3-Optical, HCA2-Copper) and eight PCIe third generation fanouts (PCIe Gen3).

   Each fanout has one, two, or four ports, so up to 40 connections are supported:

   – One-port PCIe 16 GBps I/O fanout, each supporting one domain in 32-slot PCIe I/O drawers)

   – ICA SR two-port fanout for coupling links (two links, 8 GBps each),

   – HCA3-O 12x InfiniBand fanout for coupling links (two ports, 6 GBps each)

   – HCA3-O LR 1x InfiniBand fanout for coupling links (four ports, 5 Gbps each).

   – HCA2-Copper fanouts (two ports per fanout, 6 GBps each) (supported only for I/O drawers that are carried forward, for a maximum of two features).

► Two Distributed Converter Assemblies (DCAs) that provide power to the CPC drawer

   Loss of one of the DCAs leaves enough power to satisfy the power requirements of the entire drawer. The DCAs can be concurrently maintained.

► Two Flexible Support Processors (FSPs) that provide redundant interfaces to the internal management network.

Figure 3-3 shows the inter-CPC drawer communication structure of z13s Model N20. The point-to-point connection topology shown in the figure allows direct communication between all CPC drawers. Model N10 is a single drawer, single node system (Node 0 only).
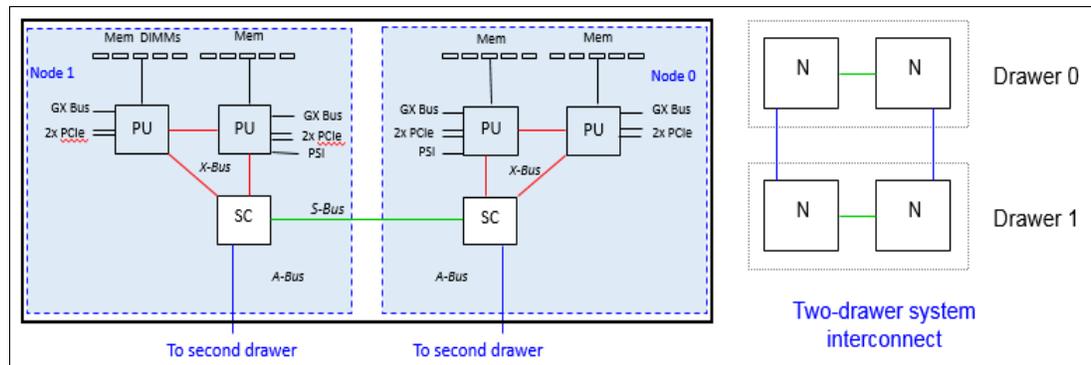


*Figure 3-3   z13s CPC drawer interconnect topology (two CPC drawers)*

The design that connects the Processor Unit and storage control allows the system to be operated and controlled by the IBM Processor Resource/Systems Manager (PR/SM) facility as a memory-coherent symmetrical multiprocessor system (SMP).

## 3.3.1  Single chip modules

At the heart of the system are *single chip modules*, which are high-performance, glass-ceramic chips that provide the highest level of processing integration in the industry.

There are three single chip modules per CPC drawer for model N10 and six single chip modules per CPC drawer for model N20. Two of these single chip modules hold the processor

chips (Processor Unit chips) and one single chip module holds the storage control chips. The CPC drawer can have one or two nodes, as explained here:

► Model N10 CPC drawer contains one node (Node 0) and provides about half the resources that the N20 CPC drawer provides.

► Model N20 CPC drawers contain two nodes each (Node 0 and 1).

Each Processor Unit chip has six or seven active cores and L1, L2, and L3 caches. The storage control chip holds L4 cache. Refer to Figure 3-4.
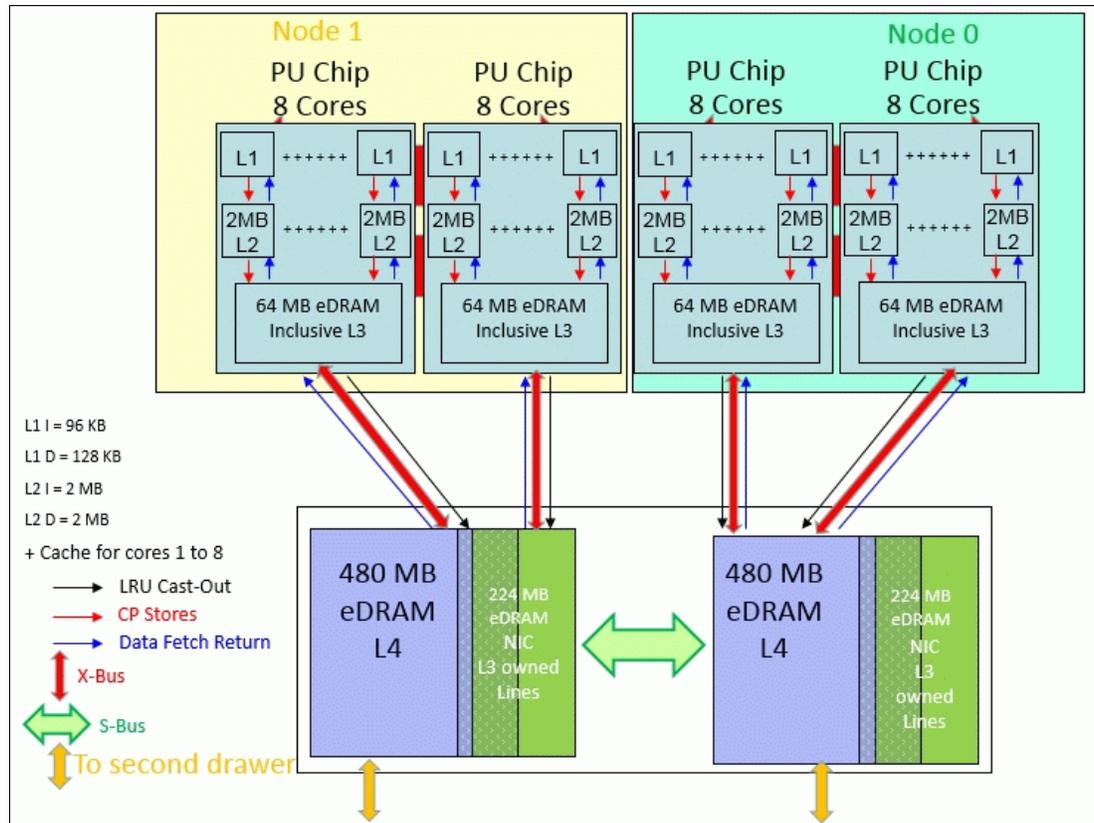


*Figure 3-4   z13s CPC drawer chips and cache structure*

## 3.3.2  Processor Unit

The *Processor Unit* (PU) is the generic term for a z/Architecture processor. A Processor Unit is embedded in a z Systems chip core. Each Processor Unit is a superscalar processor with the following attributes:

► Up to six instructions can be decoded per clock cycle.

► Up to 10 instructions can be in execution per clock cycle.

► Instructions can be issued out-of-order. The P uses a high-frequency, low-latency pipeline, providing robust performance across a wide range of workloads.

► Memory accesses might not be in the same instruction order (out-of-order operand fetching).

► Most instructions flow through a pipeline with varying numbers of steps for various types of instructions. Several instructions can be in progress at any moment, subject to the maximum number of decodes and completions per cycle.

## Processor Unit chip cache

The on-chip cache for the Processor Unit (core) chip works this way:

▶ Each Processor Unit has an L1 cache (private) that is divided into a 96 KB cache for instructions and a 128 KB cache for data.

▶ Each Processor Unit also has a private L2 cache, with 2 MB D-cache (D stands for data) and 2 MB I-cache (I stands for instruction).

▶ Each Processor Unit chip also contains a 64 MB L3 cache that is shared by all eight Processor Units on the chip. The shared L3 cache uses eDRAM.

This on-chip cache implementation optimizes performance of the system for high-frequency processors.

Each L1 cache has an associated *translation lookaside buffer* (TLB) of 512 entries. 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.

## CPU sparing

Hardware fault detection is embedded throughout the design and combined with comprehensive instruction-level retry and dynamic CPU sparing. This provides the reliability and availability that is required for true mainframe integrity.

## On-chip cryptographic hardware

Dedicated on-chip cryptographic hardware includes extended key and hash sizes for the Advanced Encryption Standard (AES) and Secure Hash (SHA) algorithms and support for UTF8 to UTF16 conversion. This cryptographic hardware is available with any processor type, for example central processor (CP), System z Integrated Information Processor (zIIP), or Integrated Facility for Linux (IFL).

## Software support

The z13s Processor Units provide full compatibility with existing software for ESA/390 and z/Architecture, while extending the Instruction Set Architecture (ISA) to enable enhanced functionality and performance. Several hardware instructions that support more efficient code generation and execution are introduced in the z13s:

▶ Hardware decimal floating point (HDFP)
▶ Transactional Execution Facility
▶ Runtime Instrumentation Facility
▶ Single-instruction, multiple-data (SIMD)

These features are further described in Chapter 5, "Strengths of the z13 and z13s" on page 71.

## Processor Unit characterization

Processor Units are ordered in single increments. The internal system functions, which are based on the configuration that is ordered, characterize each Processor Unit (core) into one of various types during system initialization, which is often called a *power-on reset* (*POR*) operation. Characterizing Processor Units dynamically without a POR is possible using a process called *Dynamic Processor Unit Reassignment*. A Processor Unit that is not characterized cannot be used.

Each Processor Unit (core) can be characterized as follows:

▶ Central processor (CP)
▶ Integrated Facility for Linux (IFL) processor

- ► z Integrated Information Processor (zIIP)
- ► Internal Coupling Facility (ICF)
- ► System assist processor (SAP)
- ► Integrated firmware processor (IFP)

At least one CP must be purchased with a zIIP or before a zIIP can be purchased. Clients can purchase up to two zIIPs for each purchased CP (assigned or unassigned) on the system. However, a logical partition definition can go beyond the 1:2 ratio. For example, on a system with two CPs, a maximum of four zIIPs can be installed. An LPAR definition for that system can contain up to two logical CPs and four logical zIIPs. Another possible configuration is one logical CP and three logical zIIPs.

Converting a Processor Unit from one type to any other type is possible by using the *Dynamic Processor Unit Reassignment* process. These conversions happen concurrently with the operation of the system.

> **zIIPs:** The addition of ICFs, IFLs, zIIPs, and SAP to the z13s does not change the system capacity setting or its MSU rating, Only CPs work that way. IBM does not impose any software charges on work that is dispatched on zIIPs.
>
> **zAAPs:** The zEC12 and zBC12 were the last z Systems servers to offer support for zAAPs. IBM supports running zAAP workloads on zIIPs (called *zAAP on zIIP*). This change can help simplify capacity planning and performance management while still supporting all the currently eligible workloads. IBM provides a PTF for APAR OA38829 on z/OS V1R12 and V1R13. This PTF removes the restriction that prevents workloads eligible for zAAP from running on zIIPs (when a zAAP is installed).

### 3.3.3  Memory

Maximum physical memory size is directly related to the number of CPC drawers in the system. And typically, a system has more memory installed than was ordered, because part of the installed memory is used to implement the redundant array of independent memory (RAIM) design. On the z13s, this configuration results in up to 2 TB of available memory per CPC drawer and up to 4 TB for a two-drawer system.

The hardware system area (HSA) on the z13s has a fixed amount of memory (40 GB) that is managed separately from client memory. However, the maximum amount of orderable memory can vary from the theoretical number due to dependencies on the memory granularity.

Table 3-2 lists the maximum and minimum memory sizes for each z13s model.

*Table 3-2   z13s model memory ranges*

| Model | Number of CPC drawers | Memory range |
|-------|-----------------------|--------------|
| N10 | 1 | 88 GB to 984 GB |
| N20 | 1 | 88 GB to 2008 GB |
| N20 | 2 | 2264 GB to 4056 GB |

With a z13s, the granularity for memory orders varies from 32 GB to 256 GB. Table 3-3 shows the memory increments, depending on the installed memory.

*Table 3-3   z13s memory increments and ranges*

| Model | Memory increment (GB) | Memory range (GB) |
|---|---|---|
| N10 and N20 | 8 | 64 to 88 |
| N10 and N20 | 32 | 120 to 344 |
| N10 and N20 | 64 | 408 to 600 |
| N10 | 128 | 728 to 984 |
| N20 | 128 | 1112 to 2008 |
| N20[a] | 256 | 2264 to 4056 |

a. With two CPC drawers

Physically, memory is organized in the following ways:

► A CPC drawer always contains a minimum of 160 GB of installed memory, of which 88 GB is usable by the operating system.

► A CPC drawer can have more installed memory than enabled. The excess amount of memory can be enabled by a Licensed Internal Code load.

► Memory upgrades are first satisfied using already installed but unused memory capacity, until it is exhausted. When no more unused memory is available from the installed cards, either the cards must be upgraded to a higher capacity or a CPC drawer with more memory must be installed.

When activated, IBM Processor Resource/Systems Manager (PR/SM) tries to allocate the memory of an LPAR in a single CPC drawer, but if that is not possible, it can use memory resources located in any CPC drawer. No matter which CPC drawer the memory is in, an LPAR has access to that memory, if it is allocated. Despite the CPC drawer structure, the z13s is still a symmetric multiprocessor (SMP) system because the Processor Units have access to all the available memory

A memory upgrade is considered to be concurrent when it requires no change of the physical memory cards or if the physically installed capacity is not yet reached. A memory card change is disruptive when physical memory needs to be changed or added.

For an upgrade that results in the addition of a CPC drawer because of Coupling, HCA, or PCIe fanout cards, the minimum memory increment is added to the system.

### Concurrent memory upgrade

If physical memory is available, memory can be upgraded concurrently by using *Licensed Internal Code Configuration Control (LICCC)*, as described. The *plan ahead memory function* that is available with the z13s (up to 2 TB), enables nondisruptive memory upgrades by having in the system pre-plugged memory (based on a target configuration). Pre-plugged memory is enabled through an LICCC order that is placed by the client.

### Redundant array of independent memory

Redundant array of independent memory (RAIM) technology makes the memory subsystem, in essence, a fully fault-tolerant N+1 design. The RAIM design automatically detects and recovers from failures of dynamic random access memory (DRAM), sockets, memory channels, or dual inline memory modules (DIMMs).

The RAIM design is fully integrated in the z13s and has been enhanced to include one Memory Controller Unit (MCU) per processor chip, with five memory channels and one DIMM per channel. A fifth channel in each MCU enables memory to be implemented as a RAIM. This technology has significant as capabilities (reliability, availability, and serviceability) in the area of error correction. Bit, lane, DRAM, DIMM, socket, and complete memory channel failures, including many types of multiple failures, can be detected and corrected.

### 3.3.4  Hardware system area

The hardware system area (HSA) is a fixed-size, reserved area of memory that is separate from the client-purchased memory. The HSA is used for several internal functions, but the bulk of it is used by channel subsystem functions.

The fixed size 40 GB HSA for z13s is large enough to accommodate any LPAR definitions or changes, thus eliminating most outage situations and the need for extensive preplanning.

A fixed, large HSA allows the dynamic I/O capability of the z13s to be enabled by default. It also enables the dynamic addition and removal of the following features:

► LPAR to new or existing channel subsystem (CSS)
► CSS (up to six can be defined in z13s)
► Subchannel set (up to four can be defined in z13s)
► Devices, up to the maximum number permitted, in each subchannel set
► Logical processors by type
► Cryptographic adapters

## 3.4  z13s I/O system structure

The z13s supports two types of internal I/O infrastructures:

► Generation 3 PCIe-based infrastructure for PCIe I/O drawers (PCIe Gen3)
► InfiniBand-based infrastructure and I/O drawers (carry forward on an MES only)

The PCIe I/O infrastructure consists of the following features:

► PCIe Gen3 fanouts in the CPC drawers, which support 16 Gbps connectivity to the PCIe I/O drawer (zBC12 used PCIe Gen2 fanouts at 8 GBps)

► Up to two 7U PCIe I/O drawers, each with 32 slots (eight slots per I/O domain) for PCIe I/O features

The InfiniBand I/O infrastructure (carry forward only) consists of the following features:

► InfiniBand fanouts in the z13s CPC drawer, which support the 6 GBps InfiniBand I/O interconnect

► InfiniBand I/O card domain multiplexers with redundant I/O interconnect in the following configuration: up to two (5U), 8-slot, 2-domain I/O drawers (carry forward only)

► FICON Express8 (carry forward only)

> **Ordering of I/O features:** Ordering of I/O feature types determines the appropriate mix of PCIe I/O drawers and I/O drawers (order-dependent).

Figure 3-5 shows a high-level view of the I/O system structure for the z13s.



*Figure 3-5   z13s I/O system structure: (N20, two CPC drawers)*

The z13s supports the following fanout types, which are located at the front on the CPC drawer:

► ICA SR
► HCA3-O
► HCA3-O LR
► PCIe Gen3
► HCA2-C (carry forward only)

The HCA3-O LR fanout includes four ports. The PCIe Gen3 fanout has one port, and the other fanouts have two ports.

The following types of internal I/O connectivity support the PCIe I/O drawer and I/O drawer:

► PCIe connections to the PCIe I/O drawers from the PCIe fanouts through copper cables. The I/O features supported through the port on the fanouts are listed in 4.1, "I/O features at a glance" on page 44.

► InfiniBand (IFB) connections to the existing I/O drawers from the host channel adapter (HCA2-C) fanouts through copper cables in only a carry-forward miscellaneous equipment specification (MES) upgrade or change. The two ports in the fanout are dedicated to connect to an InfiniBand multiplexer (IFB-MP) card in the I/O drawer.

For coupling link connectivity (Parallel Sysplex or STP configuration), the z13s supports the following fanouts:

► ICA SR
► HCA3-O
► HCA3-O LR

The model N10 CPC drawer (Figure 3-6) can have up to four 1-port PCIe Gen3 fanouts (numbered LG11 to LG14) and up to two 2-port or 4-port InfiniBand fanouts (numbered LG09 to LG10), which are used to connect to I/O drawers, PCIe I/O drawers, or for Parallel Sysplex InfiniBand and PCIe connectivity.
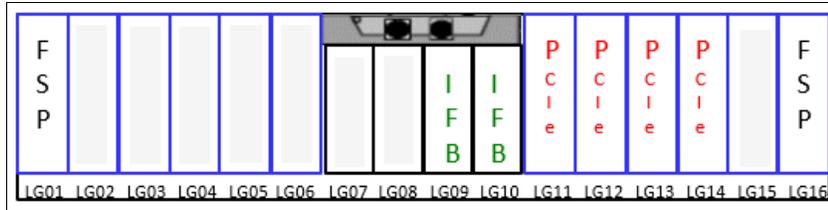


*Figure 3-6   N10 CPC drawer front view*

The model N20 CPC drawer (Figure 3-7) can have up to eight 1-port PCIe Gen3 fanouts (numbered LG03 to LG06 and LG11 to LG14) and up to four 2-port or 4-port InfiniBand fanouts for each CPC drawer (numbered LG07 to LG10), which are used to connect to I/O drawers, PCIe I/O drawers, or for Parallel Sysplex InfiniBand and PCIe connectivity.
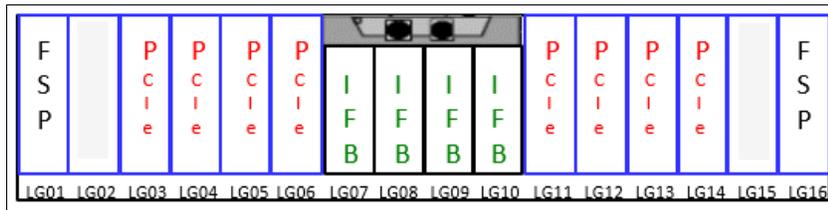


*Figure 3-7   N20 CPC drawer front view*

The PCIe I/O drawer is a two-sided drawer (I/O features are present on both sides) that is 7U high. The drawer contains 32 slots, four switch cards (two in the front and two in the rear) to support four I/O domains that each contain eight features of any type (FICON Express16S, FICON Express8S, OSA-Express5S, OSA-Express4S, Crypto Express5S, Flash Express, zEDC Express, and 10GbE RoCE Express). Two DCAs provide redundant power and two air moving devices (AMD) provide redundant cooling to the PCIe I/O Drawer.

Each I/O drawer supports two I/O domains (A and B) for a total of eight I/O slots. Each I/O domain uses an IFB-MP card in the I/O drawer and a copper cable to connect to a host channel adapter (HCA) fanout in the CPC drawer.

All features in the I/O drawer are installed horizontally. The two DCAs distribute power to the I/O drawer.

The IFB-MP cards are installed at location 09 at the rear side of the I/O drawer. The I/O features are installed from the front and rear sides of the I/O drawer. Two I/O domains are supported. Each I/O domain has up to four FICON Express8 features. The FICON Express8 I/O features are connected to the IFB-MP card through the backplane board.

Additional details about the I/O features of the z13s are available in Chapter 4, "Supported features and functions" on page 43.

## 3.5  z13s power and cooling

The z13s meets the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Class A3 specifications. ASHRAE is an organization devoted to the advancement of indoor-environment-control technology in the heating, ventilation, and air conditioning industry.

The power and cooling system of the z13s builds on that of its predecessor, the zBC12, with several newly developed technologies. However, the underlying power service specifications of the z13s are almost identical to its predecessors. Total power consumption with the maximum system configuration has increased only by about 5% compared to those previous models.

The air-cooling system in the z13s is designed with N+1 blowers, controls, and sensors for the CPC and I/O drawers.

### High Voltage Direct Current power feature

With the optional High Voltage Direct Current (HV DC) power feature, the z13s can directly connect to DC power input and improve data center energy efficiency by removing the need for an additional DC-to-AC inversion step. This feature can help achieve both data center UPS and power distribution energy savings.

### Power considerations

The z13s operates with two sets of redundant power supplies. Each set has its own individual power cords or pair of power cords, depending on the number of Bulk Power Regulator (BPR) pairs installed. Power cords attach a 3-phase, 50/60 Hz, 200–480 V AC power source or 380–520 V DC power source. The loss of one power supply alone has no effect on system operation.

The optional Balanced Power Plan Ahead feature is available for future growth, also assuring adequate and balanced power for all possible configurations. With this feature, downtime for upgrading a system is eliminated because the initial installation includes the maximum power requirements in terms of Bulk Power Regulators (BPR) and power cords. The Balance Power Plan Ahead feature is not available with DC and 1-phase line cords.

Additional single-phase outlets (customer provided) are required for ancillary equipment such as the Hardware Management Console and its display.

Specific power requirements depend on the cooling facility that is installed, the number of CPC drawers, and the number and type of I/O units that are installed. You can find maximum power consumption tables for the various configurations and environments at *IBM z13s Installation Manual for Physical Planning,* GC28-6953.

You can also refer to the power and weight estimation tool that is available at the IBM Resource Link.

**4**

# Supported features and functions

This chapter describes the I/O and other miscellaneous features and functions of the IBM z13 and IBM z13s. The information in this chapter expands upon the overview of the key hardware elements provided in Chapter 2, "IBM z13 hardware overview" on page 13 and Chapter 3, "IBM z13s hardware overview" on page 29. Only the enhanced features and functions introduced with the z13 and z13s are discussed.

Throughout the chapter, reference is made to the following helpful IBM Redbooks publications about the z13 and z13s:

► *IBM z13 Technical Guide*, SG24-8251
► *IBM z13s Technical Guide*, SG24-8294

This chapter covers these topics:

## 4.1 I/O features at a glance

The z13 and z13s support two internal I/O infrastructure types:

► A PCIe-based infrastructure for PCIe I/O drawers to support these I/O features:

- FICON Express16S
- FICON Express8S
- OSA-Express5S
- OSA-Express4S (only available when carried forward on upgrades)
- 10GbE RoCE Express
- Crypto Express5S
- Flash Express
- zEDC Express

► An InfiniBand-based infrastructure for I/O drawers to support this I/O feature:

- FICON Express8 (only available when carried forward on upgrades)

The following features that were part of earlier z Systems platforms are *not orderable* for the z13 and z13s:

► ESCON
► FICON Express4
► OSA-Express3
► ISC-3
► Crypto Express4S
► Crypto Express3

Table 4-1 lists the supported features along with required cable types, maximum unrepeated distance, and bit rates.

*Table 4-1   z13 and z13s supported features*

| Feature | Feature codes | Bit rate in Gbps (or stated) | Cable type | Maximum unrepeated distance[a] | Ordering information |
|---|---|---|---|---|---|
| **Storage connectivity features** | | | | | |
| FICON Express16S 10KM LX | 0418 | 4, 8, or 16 | SM 9 µm | 10 km (6.2 miles) | New build |
| FICON Express16S SX | 0419 | 4, 8, or 16 | OM2, OM3, OM4 | See Table 4-2 on page 46. | New build |
| FICON Express8S 10KM LX | 0409 | 2, 4, or 8 | SM 9 µm | 10 km (6.2 miles) | New build |
| FICON Express8 10KM LX | 3325 | | | | Carry forward |
| FICON Express8S SX | 0410 | 2, 4, or 8 | OM1, OM2, OM3 | See Table 4-2 on page 46. | New build |
| FICON Express8 SX | 3326 | | | | Carry forward |
| **Network connectivity features** | | | | | |
| OSA-Express5S 10 GbE LR | 0415 | 10 | SM 9 µm | 10 km (6.2 miles) | New build |
| OSA-Express4S 10 GbE LR | 0406 | | | | Carry forward |

| Feature | Feature codes | Bit rate in Gbps (or stated) | Cable type | Maximum unrepeated distance[a] | Ordering information |
|---|---|---|---|---|---|
| OSA-Expess5S 10 GbE SR | 0416 | 10 | MM 62.5 µm MM 50 µm | 33 m (200) 82 m (500) 300 m (2000) | New build |
| OSA-Express4S 10 GbE SR | 0407 | | | | Carry forward |
| OSA-Express5S GbE LX | 0413 | 1.25 | SM 9 µm | 5 km (3.1 miles) | New build |
| OSA-Express4S GbE LX | 0404 | | | | Carry forward |
| OSA-Express5S GbE SX | 0414 | 1.25 | MM 62.5 µm | 275 m (200) | New build |
| OSA-Express4S GbE SX | 0405 | | MM 50 µm | 550 m (500) | Carry forward |
| OSA-Express5S 1000BASE-T | 0417 | 100 or 1000 Mbps | Cat 5, Cat 6 unshielded twisted pair (UTP) | 100m | New build |
| OSA-Express4S 1000BASE-T | 0408 | 10, 100, or 1000 Mbps | | | Carry forward |
| 10GbE Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) Express | 0411 | 10 | OM3 | 300 m | New build |
| **Coupling and clustering** | | | | | |
| ICA SR (PCIe-O SR) | 0172 | 8 GBps | OM4 | 150 m | New build |
| | | | OM3 | 100 m | New build |
| HCA3-O (12x IFB) | 0171 | 6 GBps | OM3 | 150 m | New build |
| HCA3-O LR (1x IFB) | 0170 | 2.5 or 5 Gbps | SM 9 µm | 10 km (6.2 miles) | New build |
| IC | N/A | | N/A | N/A | N/A |
| **Special-purpose features** | | | | | |
| Crypto Express5S | 0890 | N/A | N/A | N/A | New build |
| Flash Express | 0403 | N/A | N/A | N/A | New build |
| zEDC Express | 0420 | N/A | N/A | N/A | New build |

a. Where applicable, the minimum fiber bandwidth distance in MHz-km for multi-mode fiber optic links is included in parentheses.

Note that connector type LC Duplex is used for all fiber optic cables except those for the 12x IFB and ICA SR connections, which have multifiber push-on (MPO) connectors. The MPO connector of the 12x IFB connection has one row of 12 fibers, and the MPO connector of the ICA SR connection has two rows of 12 fibers.

Table 4-2 shows the maximum unrepeated distances for different multimode fiber optic cable types when used with FICON SX (shortwave) features running at different bit rates.

*Table 4-2   Unrepeated distances for different multimode fiber optic cable types*

| Cable type (Modal bandwidth) | 2 Gbps | 4 Gbps | 8 Gbps | 16 Gbps |
|---|---|---|---|---|
| OM1 (62.5 μm at 200 MHz·km) | 150 meters | 70 meters | 21 meters | N/A |
| | 492 feet | 230 feet | 69 feet | N/A |
| OM2 (50 μm at 500 MHz·km) | 300 meters | 150 meters | 50 meters | 35 meters |
| | 984 feet | 492 feet | 164 feet | 115 feet |
| OM3 (50 μm at 2000 MHz·km) | 500 meters | 380 meters | 150 meters | 100 meters |
| | 1640 feet | 1247 feet | 492 feet | 328 feet |
| OM4 (50 μm at 4700MHz·km) | N/A | 400 meters | 190 meters | 125 meters |
| | N/A | 1312 feet | 623 feet | 410 feet |

# 4.2  Storage connectivity

IBM Fibre Connection (FICON) features continue to evolve, delivering improved throughput, reliability, availability, and serviceability. FICON features in the z13 and z13s can provide connectivity to systems, Fibre Channel switches, and various devices in a SAN environment. FICON Express16S features support a link data rate of 16 Gbps (4, 8, or 16 Gbps auto-negotiate) and are the preferred technology for new systems, providing up to a 60% increase in performance versus FICON Express8S features.

### FICON functions
The FICON protocol is fully supported on the z13 and z13s. It is commonly used with IBM z/OS, IBM z/VM (and guests), IBM z/VSE, and IBM z/TPF. The next sub-sections describe the latest FICON enhancements,

### *High Performance FICON for z Systems*
High Performance FICON for z Systems (zHPF) is implemented for protocol simplification and efficiency, which it does by reducing the number of information units (IU) that are processed. Enhancements to the z/Architecture and the FICON protocol provide optimizations for online transaction processing (OLTP) workloads. zHPF can also be used by z/OS for IBM DB2, VSAM, PDSE, and zFS.

With the z13 and z13s, zHPF has been further enhanced to allow all large write operations greater than 64 KB to be run in a single round trip to the control unit (at distances up to 100 km), thereby not elongating the I/O service for these write operations at extended distances. This is especially useful for IBM GDPS HyperSwap® configurations.

Additionally, the changes to the architecture provide end-to-end system enhancements to improve reliability, availability, and serviceability (RAS).

zHPF requires matching support by the IBM System Storage® DS8880 series or similar devices from other vendors. FICON Express16S, FICON Express8S, and FICON Express8 support both the existing FICON protocol and the zHPF protocol in the server Licensed Internal Code.

For more information about FICON channel performance, see the technical papers available at the z Systems I/O connectivity website.

### FICON Forward Error Correction (FEC)

Even with proper fiber optic cable cleaning discipline, errors can still occur on 16 Gbps links. Forward Error Correction (FEC) is a technique used for controlling errors in data transmission over lower quality communication channels. With FEC, I/O errors will be decreased, thus reducing potential impact on workload performance caused by I/O errors.

When running at 16 Gbps, FICON Express16S features can use FEC when connected to devices that support FEC, such as the IBM DS8880. FEC allows FICON Express16S channels to operate at higher speeds, over longer distances, and with reduced power and higher throughput, while retaining the same reliability and robustness for which FICON channels have traditionally been known.

### FICON Dynamic Routing

FICON Dynamic Routing (FIDR) is designed to support the dynamic routing policies supplied by FICON Director providers, such as Brocade's Exchange Based Routing (EBR) and Cisco's Open Exchange ID Routing (OxID).

With FIDR, you are no longer restricted to using static storage area network (SAN) routing policies for inter-switch links (ISLs) in a cascaded FICON Directors configuration. Performance of both FICON and FCP traffic improve due to SAN dynamic routing policies that better exploit all the available ISL bandwidth through higher utilization.

The IBM DS8880 also supports FIDR, so in a configuration with the z13 or z13s, capacity planning and management can be simplified, and provide persistent and repeatable performance and higher resiliency.

All devices in the SAN environment must support FICON Dynamic Routing to take advantage of this feature.

This z13 and z13s continue to provide the functions that were introduced on other z Systems platforms with the supported FICON features. For more information, see *IBM z Systems Connectivity Handbook*, SG24-5444.

## FCP functions

Fibre Channel Protocol (FCP) is fully supported on the z13 and z13s. It is commonly used with Linux on z Systems and supported by the z/VM and z/VSE. The next sub-sections describe the current FCP enhancements.

### 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. This is achieved by assigning a unique *worldwide port name* (WWPN) for each operating system that is 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 that are based on the assigned WWPN can be applied in the SAN environment. This function can be done by using standard mechanisms such as zoning in SAN switches and logical unit number (LUN) masking in the storage controllers.

Several recommended and allowable operating characteristic values in the FCP protocol have increased. Specifically, the recommended maximum number of NPIV hosts defined to any single physical FCP channel has increased from 32 to 64, the allowable maximum number of remote N_Ports a single physical channel can communicate with has increased from 512 to 1024, and the maximum number of LUNs addressable by a single physical channel has increased from 4096 to 8192. In support of these increases, the FCP channels have also been designed to now support 1528 concurrent I/O operations, an increase from the prior generation FCP channel limit of 764.

### Export / import physical port WWPNs for FCP channels

IBM z Systems platforms automatically assign worldwide port names (WWPNs) to the physical ports of an FCP channel, and this WWPN assignment changes when an FCP channel is moved to a different physical slot position in the I/O drawer. Now the z13 and z13s allow for the modification of these default assignments, permitting FCP channels to keep previously assigned WWPNs. This capability eliminates the need for reconfiguration of the SAN environment when a z Systems upgrade occurs or when a FICON Express feature is replaced.

### Fibre Channel Read Diagnostic Parameter

A extended link service (ELS) command called Read Diagnostic Parameter (RDP) has been added to the Fibre Channel T11 standard to allow z Systems platforms to obtain additional diagnostic data from the Small Form Factor Pluggable (SFP) optics located throughout the SAN fabric. RDP can identify a failed or failing component without unnecessarily replacing additional components in the SAN fabric (such as FICON features, optics, cables, and so on).

FICON and FCP channels will provide a means to read this additional diagnostic data for all the ports accessed in the I/O configuration and make the data available to a z Systems LPAR. For FICON channels, z/OS will display the data with a message and display command. For Linux on z Systems, z/VM, z/VSE, and KVM for IBM z this diagnostic data is made available in a panel in the SAN Explorer tool on the Hardware Management Console (HMC).

## 4.2.1 FICON Express16S

Two types of transceivers for *FICON Express16S* are supported on a new build system, one long wavelength (LX) laser version, and one short wavelength (SX) laser version:

► FICON Express16S LX feature
► FICON Express16S SX feature

Each port supports attachment to the following elements:

► FICON/FCP switches and directors that support 4 Gbps, 8 Gbps, or 16 Gbps
► Control units (storage subsystems) that support 4 Gbps, 8 Gbps, or 16 Gbps

### FICON Express16S LX feature

The *FICON Express16S LX* feature occupies one I/O slot in the PCIe I/O drawer. It has two ports, each supporting an LC duplex connector and auto-negotiated link speeds of 4 Gbps, 8 Gbps, and 16 Gbps up to an unrepeated maximum distance of 10 km.

### FICON Express16S SX feature

The *FICON Express16S SX* feature occupies one I/O slot in the PCIe I/O drawer. It has two ports, each supporting an LC duplex connector and auto-negotiated link speeds of 4 Gbps, 8 Gbps, and 16 Gbps up to an unrepeated maximum distance[1] of up to 380 meters at 4 Gbps, 150 meters at 8 Gbps, or 100 meters at 16 Gbps.

## 4.2.2  FICON Express8S

> **Statement of Direction[a]:** The IBM z13 and z13s servers will be the last z Systems server to offer ordering of FICON Express8S channel features. Enterprises that have 2 Gb device connectivity requirements must carry forward these channels.
>
> a. 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.

Two types of transceivers for *FICON Express8S* are supported on a new build[2] systems, one long wavelength (LX) laser version, and one short wavelength (SX) laser version:

► FICON Express8S 10KM LX feature
► FICON Express8S SX feature

Each port supports attachment to the following elements:

► 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

### FICON Express8S 10KM LX feature

The *FICON Express8 10KM LX* feature occupies one I/O slot in the I/O drawer. 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.

### FICON Express8S SX feature

The *FICON Express8S SX* feature occupies one I/O slot in the I/O drawer. This feature 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[1] of up to 500 meters at 2 Gbps, 380 meters at 4 Gbps, or 150 meters at 8 Gbps.

## 4.2.3  FICON Express8 (carry forward only)

> **The IBM z13 and z13s will be the last z Systems servers to support FICON Express8 channels:** IBM z13 will be the last high-end server to support FICON Express8. Enterprises should begin migrating from FICON Express8 channel features (#3325, #3326) to FICON Express16S channel features (#0418, #0419). FICON Express8 will not be supported on future high-end z Systems servers as carry forward on an upgrade.

The *FICON Express8* features are available only when carried forward on upgrades. Two types of transceivers for FICON Express8 are supported on z13 or z13s:

► FICON Express8 10KM LX feature
► FICON Express8 SX feature

---

[1]  Distances are valid for OM3 cabling. See Table 4-2 on page 46 for more options.
[2]  FICON Express8S is offered on new build to support point to point 2 Gbps attachments.

### FICON Express8 10KM LX feature

The *FICON Express8 10KM LX* feature occupies one I/O slot in the I/O drawer. 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.

### FICON Express8 SX feature

The *FICON Express8 SX* feature occupies one I/O slot in the I/O drawer. This feature 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[3] of up to 500 meters at 2 Gbps, 380 meters at 4 Gbps, or 150 meters at 8 Gbps.

# 4.3  Network connectivity

The z13 and z13s offer a wide range of functions that can help consolidate or simplify the network environment. These include OSA-Express and HiperSockets.

## 4.3.1  OSA-Express functions

Improved throughput (mixed inbound/outbound) is achieved by the data router function that was introduced in the OSA-Express3 and enhanced in OSA-Express5S, and OSA-Express4S features. 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 function avoids a hop and is designed to reduce latency and to increase throughput for standard frames (1492 bytes) and jumbo frames (8992 bytes).

The following sections describe the most current OSA-Express functions.

### Queued direct I/O (QDIO) optimized latency mode

QDIO optimized latency mode can help improve performance for applications that have a critical requirement to minimize response times for inbound and outbound data. It optimizes the interrupt processing as noted in the following configurations:

► 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-Express5S, or OSA-Express4S without requiring more program controlled interrupts (PCI).

► For outbound processing, the OSA-Express5S, or OSA-Express4S 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.

### Inbound workload queuing

Inbound workload queuing (IWQ) can help to reduce overhead and latency for inbound z/OS network data traffic and implement an efficient way for initiating parallel processing. This improvement is achieved by using an OSA-Express5S or OSA-Express4S feature in QDIO mode (CHPID types OSD and OSX) with multiple input queues and by processing network data traffic that is based on workload types. The data from a specific workload type is placed in one of four input queues (per device), and a process is created and scheduled to run on one of multiple processors, independent from the other three queues. This change can improve performance because IWQ can use the symmetric multiprocessor (SMP) architecture of the z13 and z13s.

---

[3] Distances are valid for OM3 cabling. See Table 4-2 on page 46 for more options.

## Virtual local area network support

Virtual local area network (VLAN) is a function of OSA-Express features that takes advantage of the Institute of Electrical and Electronics Engineers (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 z Systems, TCP/IP stacks can exist, potentially sharing OSA-Express features. VLAN provides a greater degree of isolation by allowing contact with a server from only the set of stations that comprise the VLAN.

## Virtual MAC support

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

This support is available for Layer 3 mode, is used by z/OS, and is supported by z/VM for guest use.

## z/VM multi-VSwitch link aggregation support

z/VM V6.3 (with PTFs) provides multi-VSwitch link aggregation support, allowing a port group of OSA-Express features to span multiple virtual switches within a single z/VM LPAR or between multiple z/VM LPARs. Sharing a link aggregation port group (LAG) with multiple virtual switches increases optimization and utilization of the OSA-Express when handling larger traffic loads. With this support, a port group is no longer required to be dedicated to a single z/VM virtual switch.

## QDIO data connection isolation for the z/VM environment

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

The QDIO data connection isolation enables the following elements:

► Adherence to security and HIPPA-security guidelines and regulations for network isolation between the instances that share physical network connectivity.

► Establishment of security zone boundaries that are 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. This feature ensures 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 port. Sharing occurs as it does today, but the ability to communicate between sharing QDIO data connections can be restricted through this support.

QDIO data connection isolation (also known as VSWITCH port isolation) applies to the z/VM environment when using the Virtual Switch (VSWITCH) function and to all of the OSA-Express5S, and OSA-Express4S features (CHPID type OSD) on z13 and z13s. z/OS supports a similar capability.

### 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 discards any packets that are 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 later and for all OSA-Express5S, and OSA-Express4S features on z13 and z13s.

### z Systems ensemble connectivity

With the IBM zEnterprise Systems, the following CHPID types support the z Systems ensemble:

- ▶ OSA-Express for Unified Resource Manager (OSM) for the intranode management network (INMN)
- ▶ OSA-Express for zBX (OSX) for the intraensemble data network (IEDN)

The INMN is one of the ensemble's two private and secure internal networks. INMN is used by the Unified Resource Manager functions in the primary HMC.

The OSM connections are through the System Control Hub (SCH) in the z13 or z13s. The INMN requires two OSA-Express5S 1000BASE-T, or OSA-Express4S 1000BASE-T ports from separate features.

The IEDN is the ensemble's other private and secure internal network. IEDN is used for communications across the virtualized images (LPARs and virtual machines). The IEDN connections use MAC addresses, not IP addresses (Layer 2 connection).

The OSX connections are from the z Systems CPC to the IEDN TOR switches in zBX. The IEDN requires two OSA-Express5S 10 GbE, or OSA-Express4S 10 GbE ports from separate features.

### OSA-ICC support for Secure Sockets Layer

The Open Systems Adapter, when configured as an integrated console controller CHPID type (OSC) on the z13 and z13s, will now support the configuration and enablement of secure connections using the Transport Layer Security (TLS) protocol versions 1.0, 1.1, and 1.2. Server-side authentication is supported using either a self-signed certificate or customer supplied certificate, which can be signed by a customer-specified Certificate Authority. The certificates used must have an RSA key length of 2048 bits, and must be signed using SHA-256. This support negotiates a cipher suite of AES-128 for the session key.

## 4.3.2  HiperSockets functions

HiperSockets feature has been referred to as the "network in a box." HiperSockets simulates LANs entirely in the hardware. The data transfer is from LPAR memory to LPAR memory, mediated by microcode. The z13 and z13s support up to 32 HiperSockets. One HiperSockets network can be shared by up to 85 LPARs on a z13 and 40 LPARs on a z13s. Up to 4096 communication paths support a total of 12,288 IP addresses across all 32 HiperSockets.

The HiperSockets internal networks can support the following transport modes:

- ▶ Layer 2 (link layer)
- ▶ Layer 3 (network or IP layer)

Traffic can be Internet Protocol Version 4 or Version 6 (IPv4, IPv6) or non-IP (such as AppleTalk, DECnet, IPX, NetBIOS, SNA, or others). HiperSockets devices are independent of protocol and Layer 3. 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 way as they do for a non-mainframe environment. HiperSockets Layer 2 support is provided by Linux on z Systems, and by z/VM for guest use.

The most current HiperSockets functions are described in the following sections.

## HiperSockets Multiple Write Facility

HiperSockets performance is enhanced to allow for the streaming of bulk data over a HiperSockets link between 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, reduces CPU use 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 are enhanced for zIIP usage. Specifically, the z/OS Communications Server allows the HiperSockets Multiple Write Facility processing for outbound large messages that originate 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 that are based on XML, HTTP, SOAP, Java, and traditional file transfer can benefit from zIIP enablement by lowering general-purpose processor use 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 unblocks the sending application.

zIIP Assisted HiperSockets for large messages is available on z13 and z13s with z/OS V1R12 and later releases.

## HiperSockets network traffic analyzer

HiperSockets network traffic analyzer (NTA) is a function that is available in the LIC of the z13 and z13s. It can simplify problem isolation and resolution by allowing Layer 2 and Layer 3 tracing of HiperSockets network traffic.

HiperSockets NTA allows Linux on z Systems to control tracing of the internal virtual LAN. It captures records into host memory and storage (file systems) that can be analyzed by system programmers and network administrators, using Linux on z Systems tools to format, edit, and process the trace records.

A customized HiperSockets NTA rule enables authorizing an LPAR to trace messages only from LPARs that are eligible to be traced by the NTA on the selected IQD channel.

## HiperSockets completion queue

The HiperSockets completion queue function allows both synchronous and asynchronous transfer of data between logical partitions. With the asynchronous support, during high

volume situations, data can be temporarily held until the receiver has buffers available in its inbound queue. This function can provide performance improvement for LPAR to LPAR communication and can be especially helpful in burst situations.

### HiperSockets integration with the intraensemble data network

The z13 and z13s servers provide the capability to integrate HiperSockets connectivity with the intraensemble data network (IEDN). Thus the reach of the HiperSockets network is extended to outside the z Systems CPC to the entire ensemble, which is displayed as a single, Layer 2 network. Because HiperSockets and IEDN are both internal z Systems networks, the combination allows z Systems virtual servers to use an optimal path for communications.

### HiperSockets virtual switch bridge support

The z/VM virtual switch is enhanced to transparently bridge a guest virtual machine network connection on a HiperSockets LAN segment. This bridge allows a single HiperSockets guest virtual machine network connection to also directly communicate with the following systems:

► Other guest virtual machines on the virtual switch
► External network hosts through the virtual switch OSA UPLINK port

A HiperSockets channel by itself is only capable of providing intra-CPC communications. The HiperSockets Bridge Port allows a virtual switch to connect z/VM guests by using real HiperSockets devices, the ability to communicate with hosts that reside externally to the CPC. The virtual switch HiperSockets Bridge Port eliminates the need to configure a separate next hop router on the HiperSockets channel to provide connectivity to destinations that are outside of a HiperSockets channel.

## 4.3.3  OSA-Express5S

This section describes the connectivity options that are offered by the *OSA-Express5S* features. The following OSA-Express5S features can be installed on z13 or z13s:

► OSA-Express5S 10 Gigabit Ethernet (GbE) Long Reach (LR)
► OSA-Express5S 10 Gigabit Ethernet (GbE) Short Reach (SR)
► OSA-Express5S Gigabit Ethernet Long Wavelength (GbE LX)
► OSA-Express5S Gigabit Ethernet Short Wavelength (GbE SX)
► OSA-Express5S 1000BASE-T Ethernet

### OSA-Express5S 10 GbE LR feature

The OSA-Express5S 10 GbE LR feature occupies one slot in a PCIe I/O drawer. It has one port that connects to a 10 Gbps Ethernet LAN through a 9 µm single mode fiber optic cable that is terminated with an LC Duplex connector. The feature supports an unrepeated maximum distance of 10 km.

### OSA-Express5S 10 GbE SR feature

The OSA-Express5S 10 GbE SR feature occupies one slot in the PCIe I/O drawer. This feature has one port that connects to a 10 Gbps Ethernet LAN through a 62.5 µm or 50 µm multimode fiber optic cable that is terminated with an LC Duplex connector. The maximum supported unrepeated distance is 33 m on a 62.5 µm multimode fiber optic cable, and 300 m on a 50 µm multimode fiber optic cable.

### OSA-Express5S GbE LX feature

The OSA-Express5S GbE LX occupies one slot in the PCIe I/O drawer. This feature has two ports, representing one channel path identifier (CHPID), that connect to a 1 Gbps Ethernet

LAN through a 9 µm single mode fiber optic cable. This cable is terminated with an LC Duplex connector, supporting an unrepeated maximum distance of 5 km. 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 multimode MCP cables reduces the supported distance of the link to a maximum of 550 meters.

### OSA-Express5S GbE SX feature
The OSA-Express5S GbE SX occupies one slot in the PCIe I/O drawer. This feature has two ports, representing one CHPID, that connect to a 1 Gbps Ethernet LAN through 50 or 62.5 µm multimode fiber optic cable. This cable is 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-Express5S 1000BASE-T feature
The OSA-Express5S 1000BASE-T occupies one slot in the PCIe I/O drawer. It has two ports, representing one CHPID, that connect to a 1000 Mbps (1 Gbps) or 100 Mbps Ethernet LAN. Each port has an RJ-45 receptacle for UTP Cat5 or Cat6 cabling, which supports a maximum distance of 100 meters.

## 4.3.4  OSA-Express4S (carry forward only)

The OSA-Express4S features offer various connectivity options. The following OSA-Express4S features can be installed on z13 or z13s:

► OSA-Express4S 10 Gigabit Ethernet (GbE) Long Reach (LR)
► OSA-Express4S 10 Gigabit Ethernet Short Reach (SR)
► OSA-Express4S Gigabit Ethernet long wavelength (GbE LX)
► OSA-Express4S Gigabit Ethernet short wavelength (GbE SX)
► OSA-Express4S 1000BASE-T Ethernet

### OSA-Express4S 10 GbE LR feature
The OSA-Express4S 10 GbE LR feature occupies one slot in a PCIe I/O drawer. It has one port that connects to a 10 Gbps Ethernet LAN through a 9 µm single mode fiber optic cable that is terminated with an LC Duplex connector. The feature supports an unrepeated maximum distance of 10 km.

### OSA-Express4S 10 GbE SR feature
The OSA-Express4S 10 GbE SR feature occupies one slot in the PCIe I/O drawer. This feature has one port that connects to a 10 Gbps Ethernet LAN through a 62.5 µm or 50 µm multimode fiber optic cable that is terminated with an LC Duplex connector. The maximum supported unrepeated distance is 33 m on a 62.5 µm multimode fiber optic cable, and 300 m on a 50 µm multimode fiber optic cable.

### OSA-Express4S GbE LX feature
The OSA-Express4S GbE LX occupies one slot in the PCIe I/O drawer. This feature has two ports, representing one channel path identifier (CHPID), that connect to a 1 Gbps Ethernet LAN through a 9 µm single mode fiber optic cable. This cable is terminated with an LC Duplex connector, supporting an unrepeated maximum distance of 5 km. 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 multimode MCP cables reduces the supported distance of the link to a maximum of 550 meters.

### OSA-Express4S GbE SX feature

The OSA-Express4S GbE SX occupies one slot in the PCIe I/O drawer. This feature has two ports, representing one CHPID, that connect to a 1 Gbps Ethernet LAN through 50 or 62.5 µm multimode fiber optic cable. This cable is 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-Express4S 1000BASE-T feature

The OSA-Express4S 1000BASE-T occupies one slot in the PCIe I/O drawer. It has two ports, representing one CHPID, 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 or Cat6 cabling, which supports a maximum distance of 100 meters.

## 4.4  Native PCIe features and integrated firmware processor

The zEC12 introduced feature types, know as native PCIe features, which require a different management design. The following native PCIe features are available on the z13 and z13s:

► Flash Express
► zEDC Express
► 10GbE RoCE Express

These features are plugged exclusively into a PCIe I/O drawer, where they coexist with the other, non-native PCIe, I/O adapters and features, but they are managed in a different way from those other I/O adapters and features. The native PCIe feature cards have a PCHID assigned according to the physical location in the PCIe I/O drawer.

For the native PCIe features supported by z13 and z13s, there are drivers included in the operating system, and the adaptation layer is not needed. The adapter management functions (such as diagnostics and firmware updates) are provided by Resource Groups partitions running on the integrated firmware processor (IFP).

The IFP is used to manage native PCIe adapters installed in a PCIe I/O drawer. The IFP is allocated from a pool of Processor Units that are available for the whole system. Because the IFP is exclusively used to manage native PCIe adapters, it is not taken from the pool of Processor Units that can be characterized for customer usage.

If a native PCIe feature is present in the system, the IFP is initialized and allocated during the system POR phase. Although the IFP is allocated to one of the physical Processor Units, it is not visible. In case of error or failover scenarios, the IFP will act like any other Processor Unit (that is, sparing is invoked).

### 4.4.1  Flash Express

Flash Express is an innovative optional feature that was introduced with the zEC12 and also available on the z13 or z13s. It can help to provide performance improvements and better availability for critical business workloads that cannot afford any hits to service levels. Flash Express is easy to configure, requires no special skills, and provides rapid time to value.

Flash Express implements storage-class memory (SCM) through an internal NAND flash solid-state drive (SSD) in a PCIe card form factor. The Flash Express feature allows each logical partition to be configured with its own SCM address space.

For availability, this feature is available in pairs of cards. Each feature offers a capacity of 1.4 TB of usable storage per pair of cards. A maximum of four pairs of cards can be installed on a z13, providing a maximum capacity of 5.6 TB of storage.

### 4.4.2 zEDC Express

The zEDC Express is an optional feature that is available to the z13 and z13s. It provides hardware-based acceleration for data compression and decompression for the enterprise, which can help to improve cross platform data exchange, reduce CPU consumption, and save disk space.

A minimum of one feature can be ordered and a maximum of eight can be installed on the system, in the PCIe I/O drawer. Up to two zEDC Express features per domain can be installed. There is one PCIe adapter/compression coprocessor per feature which implements compression as defined by RFC1951 (DEFLATE). A zEDC Express feature can be shared between up to 15 LPARs.

For resilience, there are always two independent Resource Groups (RGs) on the system, sharing the IFP. Thus, the suggestions if for a minimum of two zEDC Express features to be installed, one per RG.

### 4.4.3 The 10 Gigabit Ethernet RoCE Express

The 10 Gigabit Ethernet (10GbE) RoCE Express feature helps reduce consumption of CPU resources for applications that use the TCP/IP stack and might also help to reduce network latency with memory-to-memory transfers using Shared Memory Communications - Remote Direct Memory Access (SMC-R) in z/OS V2R1. It is transparent to applications and can be used for LPAR-to-LPAR communication on a single system or server-to-server communication in a multiple z Systems CPC environment.

This feature resides in the PCIe I/O drawer and is available to the z13 and z13s. The 10GbE RoCE Express feature has one PCIe adapter with two ports.

The 10GbE RoCE Express feature uses a short reach (SR) laser as the optical transceiver, and supports use of a multi-mode fiber optic cable terminated with an LC Duplex connector. Both point to point connection and switched connection with an enterprise-class 10 GbE switch are supported. Switch used by 10GbE RoCE Express feature must have *Pause frame* enabled as defined by the IEEE 802.3x standard.

A maximum of 16 features can be installed per system. With z13 or z13s, the 10GbE RoCE adapters can be shared by up to 31 LPARs. Also, both adapter ports are now supported by z/OS, when running on z13 or z13s.

## 4.5 Cryptographic features

The z13 and z13s provide cryptographic functions that, from an application program perspective, can be grouped as follows:

► Synchronous cryptographic functions, provided by the CP Assist for Cryptographic Function (CPACF)

► Asynchronous cryptographic functions, provided by the Crypto Express features

## 4.5.1  CP Assist for Cryptographic Function

The CP Assist for Cryptographic Function (CPACF) offers a set of symmetric cryptographic functions for high performance encryption and decryption with clear key operations for SSL/TLS, VPN, and data-storing applications that do not require FIPS[4] 140-2 level 4 security. The CPACF is integrated with the compression unit in the coprocessor (CoP) in the z13 and z13s microprocessor core.

The CPACF protected key is a function that facilitates the continued privacy of cryptographic key material while keeping the wanted high performance. CPACF ensures that key material is not visible to applications or operating systems during encryption operations. CPACF protected key provides substantial throughput improvements for large-volume data encryption and low latency for encryption of small blocks of data.

The cryptographic assist includes support for the following functions:

► Data Encryption Standard (DES) data encrypting and decrypting.

   DES supports the following key types:
   – Single-length key DES
   – Double-length key DES
   – Triple-length key DES (T-DES)

► Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys

► Pseudo random number generation (PRNG)

► Message Authentication Code (MAC)

► Hashing algorithms: SHA-1 and SHA-2 support for SHA-224, SHA-256, SHA-384, and SHA-512

SHA-1and 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, z/VSE, z/TPF, and Linux on z Systems.

## 4.5.2  Crypto Express5S

The Crypto Express5S represents the newest generation of the Peripheral Component Interconnect Express (PCIe) cryptographic coprocessors. It is an optional feature exclusive to the z13 and z13s. This feature provides a secure programming and hardware environment wherein crypto processes are performed. Each cryptographic coprocessor includes a general-purpose processor, non-volatile storage, and specialized cryptographic electronics.

The Crypto Express5S has one PCIe adapter per feature. For availability reasons, a minimum of two features is required. Up to 16 Crypto Express5S features are supported (16 PCI Express adapters per z13 and z13s). The Crypto Express5S feature occupies one I/O slot in a PCIe I/O drawer.

Each adapter can be configured as a Secure IBM CCA coprocessor, a Secure IBM Enterprise PKCS #11 (EP11) coprocessor, or as an accelerator.

Crypto Express5S is enhanced to provide domain support for up to 85 logical partitions on IBM z13 and 40 logical partitions on z13s.

---

[4] Federal Information Processing Standards (FIPS) 140-2 Security Requirements for Cryptographic Modules

The accelerator function is designed for maximum-speed Secure Sockets Layer and Transport Layer Security (SSL/TLS) acceleration, rather than for specialized financial applications for secure, long-term storage of keys or secrets. The Crypto Express5S can also be configured as one of the following configurations:

► The *Secure IBM CCA coprocessor* for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification includes secure key functions and is optionally programmable to deploy more functions and algorithms using User Defined Extension (UDX).

► The *Secure IBM Enterprise PKCS #11 (EP11) coprocessor* implements an industry-standardized set of services that adheres to the PKCS #11 specification v2.20 and more recent amendments. It was designed for extended FIPS and Common Criteria evaluations to meet industry requirements.

This cryptographic coprocessor mode introduced the PKCS #11 secure key function.

> **TKE feature:** The Trusted Key Entry (TKE) Workstation feature is required for supporting the administration of the Crypto Express5S when configured as an Enterprise PKCS #11 coprocessor.

When the Crypto Express5S PCI Express adapter is configured as a secure IBM CCA coprocessor, it still provides accelerator functions. However, up to three times better performance for those functions can be achieved if the Crypto Express5S PCI Express adapter is configured as an accelerator.

### 4.5.3 Web deliverables

For z/OS downloads, see the z/OS website.

# 4.6 Coupling and clustering

Coupling connectivity for Parallel Sysplex on z13 and z13s use Integrated Coupling Adapter (ICA SR) and InfiniBand (IFB) technology. The ICA SR is designed to support distances up to 150 m. The HCA3-O LR fanout supports longer distances between systems using the IFB technology.

ICA SR and InfiniBand technologies allow moving all of the Parallel Sysplex connectivity support to interfaces that provides high-speed interconnection at short distances and longer distance fiber optic interconnection.

See the *Coupling Facility Configuration Options* white paper for a more specific explanation regarding the coupling links technologies.

For details about all InfiniBand features, see either of these resources:

► *IBM z Systems Connectivity Handbook*, SG24-5444
► *Implementing and Managing InfiniBand Coupling Links on IBM System z*, SG24-7539

### 4.6.1 Integrated Coupling Adapter

The IBM Integrated Coupling Adapter (ICA SR), introduced on the z13 and z13s platform, is a two-port fanout used for short distance coupling connectivity and utilizes a coupling channel type: CS5. The ICA utilizes PCIe Gen3 technology, with x16 lanes that are bifurcated into x8 lanes for coupling. No performance degradation is expected compared to Coupling over InfiniBand 12X IFB3 protocol.

The ICA SR supports cable length of up to 150 m and supports a link data rate of 8 GBps. It also supports up to 4 CHPIDs per port and 7 subchannels (devices) per CHPID. The coupling links can be defined as shared between images within a CSS. They can also be spanned across multiple CSSs in a z Systems CPC. Unlike the HCA3-O 12x InfiniBand links, the ICA SR cannot define more than four CHPIDs per port.

## 4.6.2 InfiniBand coupling links

Two types of host channel adapter (HCA) fanouts are used for IFB coupling links on the z13 or z13s:

- ► HCA3-O fanout, which supports 12x InfiniBand (12x IFB)
- ► HCA3-O Long Reach (LR) fanout, which supports 1x InfiniBand (1x IFB)

HCA3s are the most recent generation of InfiniBand host channel adapters for coupling. The HCA3-O fanout for 12x InfiniBand (12x IFB) is designed for improved service times and is available for z Systems CPCs using the 12x InfiniBand3 (12x IFB3) protocol. The HCA3-O LR fanout for 1x InfiniBand (1x IFB) provides four ports and optional additional subchannels for extended-distance solutions.

> **InfiniBand coupling link data rate:** 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.

The 12x InfiniBand coupling links support double data rate (DDR) at 6 GBps for a z13 or z13s to z13 or z13s or to z Systems CPCs:

- ► InfiniBand (HCA3-O) coupling links (12x IFB), used for z/OS-to-CF communication, CF-to-CF traffic, or STP messaging at distances up to 150 meters (492 feet) by using industry standard OM3 50 µm fiber optic cables.

- ► When no more than four CHPIDs are defined per port, and an HCA3-O to HCA3-O connection is set up, the IFB3 protocol is used. When using the IFB3 protocol, synchronous service times are 40% faster than when using the IFB protocol.

- ► An HCA3-O to HCA2-O connection is supported, but the standard IFB protocol is used.

HCA3-O LR 1x InfiniBand coupling links support up to 32 subchannels (devices) per CHPID, versus the current default value of seven devices per CIB type CHPID:

- ► InfiniBand (HCA3-O LR) coupling links (1x IFB) for z/OS-to-CF communication at unrepeated distances up to 10 km (6.2 miles) using 9 µm single mode fiber optic cables and repeated distances up to 100 km (62 miles) using IBM z Systems qualified DWDM equipment. (Connectivity to HCA2-O LR is supported).

- ► The HCA3-O LR has four ports, the number of supported CHPIDs remains at 16 for the fanout card.

> **HCA2-O and HCA2-O LR features:** HCA2-O and HCA2-O LR features are not available on z13 or z13s but might be present in an existing parallel sysplex environment. These adapters were supported in a carry forward MES, migrating to the zEC12, zBC12, or to both families.

### Time source for Server Time Protocol (STP) traffic

IFB and ICA SR links can be used to carry STP timekeeping information.

For details about all InfiniBand features, see the publications mentioned earlier in this section.

## 4.6.3 Internal coupling

Internal coupling (IC) links are used for internal communication between LPARs on the same system running coupling facilities (CF) and z/OS images. The connection is emulated in Licensed Internal Code (LIC) and provides for fast and secure memory-to-memory communications between LPARs within a single system. No physical cabling is required.

## 4.6.4 Coupling Facility Control Code

Various levels of Coupling Facility Control Code (CFCC) are available.

### CFCC Level 21

CFCC Level 21 is available on z13 or z13s with driver level 27, and includes the following enhancements:

► Support for up to 141 ICF processors

The maximum number of logical processors in a Coupling Facility Partition remains at 16.

► Usability enhancement:

Enable systems management applications to collect valid CF LPAR information via z/OS BCPii:

– System type (CFCC), System Level (CFCC Level)
– Dynamic dispatch settings to indicate CF state (dedicated, shared thin interrupt)

► Availability enhancement:

Asynchronous CF duplexing for lock structures:

– Makes duplexing CF lock structures at GDPS distances practical
– Maintains robust failure recovery capability through the redundancy of duplexing
– Reduces z/OS, CF, and link utilization overhead costs associated with synchronous duplexing of lock structures

► Large memory support

– Improve availability for larger CF cache structures and data sharing performance with larger DB2 group buffer pools (GBP).

– Remove inhibitors to using large CF structures, enabling use of large memory to appropriately scale to larger DB2 local buffer pools (LBP) and group buffer pools (GBP) in data sharing environments.

– CF structure size remains at a maximum of 1 TB

► Support for IBM Integrated Coupling Adapter (ICA)

> **Statement of Direction**[a] **(coupling CHPIDs support):** IBM plans to support up to 256 coupling link CHPIDs in z13 (that is twice the 128 coupling link CHPIDs supported on zEC12). Each CF image will continue to support a maximum of 128 coupling link CHPIDs.

> a. 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.

CF structure size changes are expected to grow when going from CFCC Level 20 (or earlier) to CFCC Level 21. Review the CF LPAR size by using the following tools:

► The CFSizer tool is a web-based and is useful when changing an existing workload or introducing a new one.

► The Sizer Utility, an authorized z/OS program download, is useful when upgrading a CF.

# 4.7  IBM z BladeCenter Extension (zBX) Model 004

z13 and z13s support the IBM z BladeCenter Extension (zBX) Model 004. As a stand-alone node of an existing ensemble the zBX can house multiple environments that include IBM AIX®, Linux on System x, and Windows, supporting a "fit for purpose" application deployment.

The zBX is tested and packaged together at the IBM manufacturing site and shipped as one unit, relieving complex configuration and set up requirements. With a focus on availability, the zBX has hardware redundancy that is built in at various levels: the power infrastructure, rack-mounted network switches, power and switch units in the BladeCenter chassis, and redundant cabling for support and data connections. The IBM z BladeCenter Extension (zBX) Model 004 components are configured, managed, and serviced using a pair of internal 1U rack mounted Support Elements as a node of an ensemble, defined to the ensemble HMC as any other ensemble member.

Although the zBX processors are not z/Architecture Processor Units, the zBX is handled by z Systems firmware called *IBM z Unified Resource Manager*.

GDPS/PPRC and GDPC/GM support zBX hardware components, providing workload failover for automated multi-site recovery. These capabilities can help facilitate the management of planned and unplanned outages across IBM z13 and z13s.

## 4.7.1  IBM blades

zBX Model 004 supports IBM AIX on IBM POWER7®, Linux on System x, Microsoft Windows on System x and IBM WebSphere DataPower® Integration Appliance XI50 for zEnterprise on a blade form factor, which are connected to the z Systems CPCs through virtual LANs supported on a high-speed private network.

IBM BladeCenter PS701 Express blades virtualized by IBM PowerVM® Enterprise Edition. The virtual servers in PowerVM run the AIX operating system. PowerVM handles all the access to the hardware resources, providing a Virtual I/O Server (VIOS) function and the ability to create logical partitions. The logical partitions can be either dedicated processor LPARs, which require a minimum of one core per partition, or shared processor LPARs (micro-partitions), which in turn can be as small as 0.1 core per partition.

A select set of IBM BladeCenter HX5 (7873) blades can be used by the zBX. These blades have an integrated hypervisor, and their virtual machines run Linux on System x and Windows Server 2012.

When ordering a zBX Model 004 MES upgrade, a new entitlement record can be acquired to allow IBM System x blades or IBM POWER7 PS701 to be ordered and added to the zBX, up to the limit of available empty (not used) slots in the zBX existing blade centers.

**Unsupported:** The addition of new racks or new blade centers cannot be done and are not supported. Also the addition of the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise is not supported.

The IBM z BladeCenter Extension (zBX) Model 004 exists only as an upgrade MES from a previous existing zBX Model 003 or a zBX Model 002. The zBX Model 004 extends the z Systems qualities of service and management to integrate heterogeneous systems with high redundancy.

The zBX Model 004 (2458-004) connects to the z13 or z13s to become a *node* on its own, as part of an ensemble. The zBX stand-alone node, in turn, creates an integrated multi-platform system with advanced virtualization management (through the Unified Resource Manager) that supports diverse workloads.

The zBX is configured with the following key components:

► Model 004 incorporates its own 1U Support Elements installed in the first zBX rack (Frame B).

► One to four standard 19" 42U IBM z Systems racks with required network and power infrastructure.

► One to eight BladeCenter chassis with a combination of up to 112 blades.[5]

► Redundant infrastructure for fault tolerance and higher availability.

► Management support through the z13 or z13s HMC and zBX Model 004 SEs.

The first rack (rack B) in the zBX is the primary rack where one or two BladeCenter chassis are located. Two pairs of Top of Rack (TOR) switches are included in rack B, one pair for the intranode management network (INMN) and another pair for the intraensemble data network (IEDN) connectivity and also the new 1U SEs with their respective keyboard and display. The other three racks (C, D, and E) are expansion racks with one or two BladeCenter chassis each.

The zBX is managed from the primary ensemble HMC, connected to its internal primary SE through the existing client provided HMC/SE switch, as an ensemble node. The ensemble HMC will discover the zBX SE when the physical connectivity is established as it does for the z Systems CPCs today.

The IEDN provides private and secure 10 GbE high-speed data paths between all elements of an ensemble and the zBX node through the IEDN TOR switches. The IEDN connections use MAC addresses, not IP addresses (Layer 2 connection). The OSA-Express for zBX (OSX) CHPID type supports connectivity and access control from the z13 or z13s to the zBX node.

---

[5] Current maximum number of blades supported: 112 POWER7, 28 for DataPower XI50z, 56 for System x HX5

Figure 4-1 shows the z13 and the zBX node connections through the OSA-Express5S 10 GbE or OSA-Express4S 10 GbE features (CHPID type OSX).
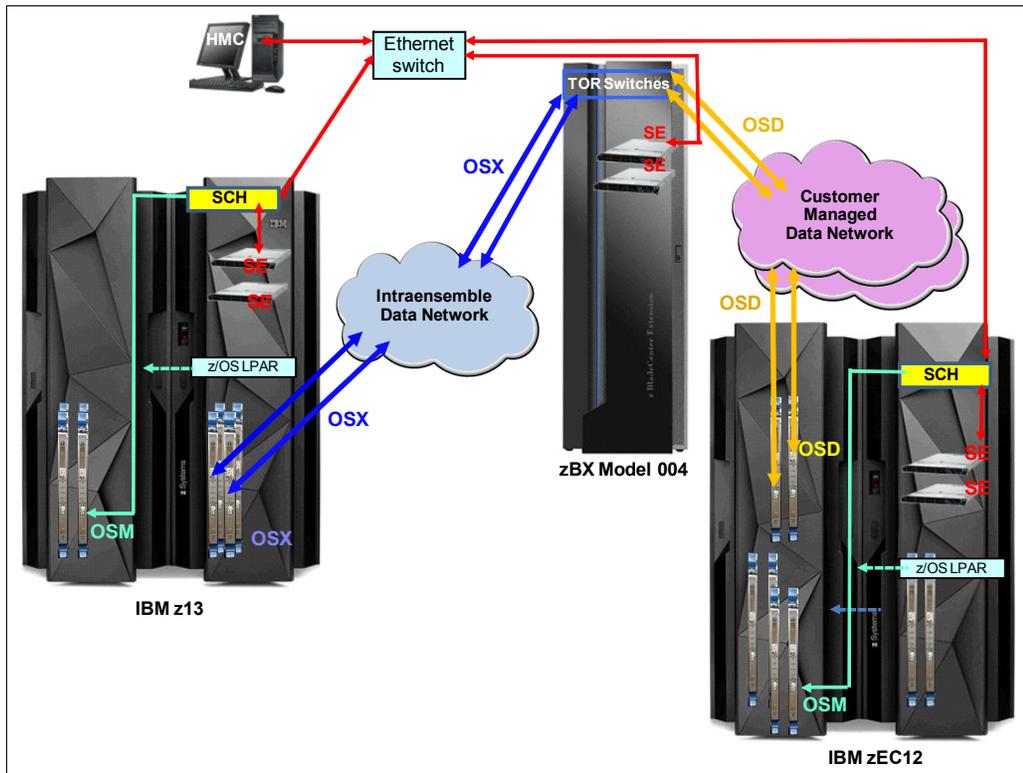


*Figure 4-1   IEDN and client-managed data networks in an ensemble*

Optionally, as part of the ensemble, any OSA-Express5S or OSA-Express4S (with CHPID type OSD) in the z13 or z13s can connect to the client-managed data network. The client-managed network can also be connected to the IEDN TOR switches in the zBX node.

In addition, each BladeCenter chassis in the zBX has two Fibre Channel (FC) switch modules that connect to FC disk storage a SAN switch. Each FC switch supports up to six external FC links to connect to SAN switches.

Figure 4-2 shows front and rear views of a zBX rack (Rack B).
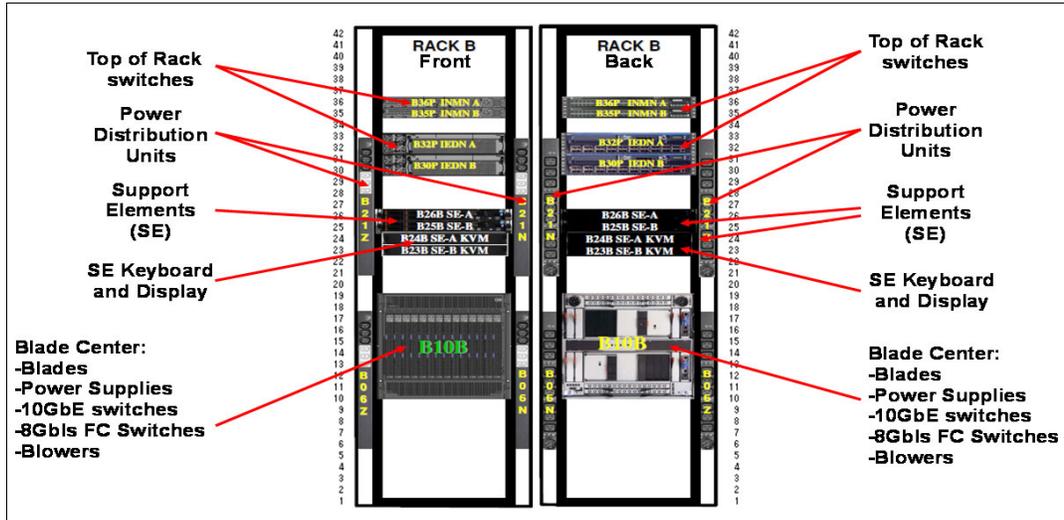


*Figure 4-2   zBX rack B front and back configuration*

The zBX racks include the following features:

► Two 1U internal Support Elements
► Two (Keyboard-Video-Mouse) SE components
► Two TOR 1000BASE-T switches (rack B only) for the INMN
► Two TOR 10 GbE switches (rack B only) for the IEDN
► Up to two BladeCenter chassis in each rack

   Each BladeCenter consists of the following features:

   – Up to 14 blade slots
   – Two Advanced Management Modules (AMM)
   – Two Ethernet Switch Modules (ESM)
   – Two 10 GbE high speed switch (HSS) modules
   – Two 8 Gbps Fibre Channel switch modules
   – Two blower modules

► Power Distribution Units (PDU)

The following blade types are supported in zBX:

► IBM POWER7 PS701 Express blades
► IBM System x blades (HX5 7873 dual-socket 16-core)
► IBM WebSphere DataPower XI50 for zEnterprise blades (double-width)

PowerVM Enterprise Edition is the hypervisor on the POWER7 blades, and the supported operating system is AIX. Linux on System x and Windows on System x are the supported operating systems for select System x blades, using the zBX integrated hypervisor for IBM System x blades (using a Kernel-based virtual machine). Both hypervisors are shipped, serviced, and deployed as z Systems LIC, booted automatically at power-on reset, and are isolated on the internal platform management network.

Client-supplied external disks are required with the zBX. For information about supported Fibre Channel Protocol (FCP) disk types and vendors with IBM blades, see the IBM System Storage Interoperation Center (SSIC).

For more information about the number of blades that are supported and configuration options, see *IBM z13 Technical Guide*, SG24-8251 and *IBM z13s Technical Guide*, SG24-8294.

# 4.8  Common time functions

Each system must have an accurate time source to maintain a time-of-day value. Logical partitions use their system's time. When system images participate in a Sysplex, coordinating the time across all system images in the sysplex is critical to its operation.

The z13 and z13s supports the Server Time Protocol (STP) and can participate in a STP-only coordinated timing network (CTN).

## 4.8.1  Server Time Protocol

Server Time Protocol (STP) 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 the supported method for maintaining time synchronization between the z13, z13s and coupling facilities (CF) in sysplex environments.

The STP design uses a concept that is called *Coordinated Timing Network (CTN)*. A CTN is a collection of z Systems CPCs that are time-synchronized to a time value called *Coordinated Server Time (CST)*. Each CPC to be configured in a CTN must be STP-enabled. STP is intended for CPCs that are configured to participate in a Parallel Sysplex or CPCs that are not in a Parallel Sysplex, but must be time-synchronized.

STP is implemented in LIC as a system-wide facility of the z13, z13s, and other z Systems CPCs. STP presents a single view of time to PR/SM and provides the capability for multiple CPCs to maintain time synchronization with each other. The z13 or z13s are enabled for STP by installing the STP feature code. Extra configuration is required for a z13 or z13s to become a member of a CTN.

STP provides the following additional value over the former used-time synchronization method by a 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 protocol allows a Parallel Sysplex to span these distances and reduces the cross-site connectivity that is required for a multi-site Parallel Sysplex.

► The STP design allows more stringent synchronization between CPCs and CFs by using communication links that are already used for the sysplex connectivity. With the z13 and z13s, STP supports coupling links over InfiniBand or Integrated Coupling Adapter links.

► STP helps eliminate infrastructure requirements, such as power and space, needed to support the Sysplex Timers and helps eliminate maintenance costs that are associated with the Sysplex Timers.

► STP can reduce the fiber optic infrastructure requirements in a multi-site configuration because it can use the coupling links that are already in use.

### 4.8.2 Network Time Protocol client support

The use of Network Time Protocol (NTP) servers as an external time source (ETS) and usually fulfills a requirement for a time source or common time reference across heterogeneous platforms and for providing a higher time accuracy.

NTP client support is available in the Support Element (SE) code of the z13 and z13s. The code interfaces with the NTP servers. This interaction allows a NTP server to become the single-time source for z13, z13s and for other servers that have NTP clients. NTP can be used only for a STP-only CTN environment.

#### Pulse per second support

Two oscillator cards (OSC), included as a standard feature of the z13 and z13s, provide a dual-path interface for the pulse per second (PPS) signal. The cards contain a BNC connector for PPS attachment at the rear side of the z Systems CPC frame A. The redundant design allows continuous operation, in case of failure of one card, and concurrent card maintenance.

STP tracks the highly stable accurate PPS signal from the NTP server. PPS maintains accuracy of 10 µs as measured at the PPS input of the z13 and z13s.

If STP uses a NTP server without PPS, a time accuracy of 100 ms to the ETS is maintained. A cable connection from the PPS port to the PPS output of a NTP server is required when the z13 or z13s is configured for using NTP with PPS as the ETS for time synchronization.

#### NTP server on HMC with security enhancements

The NTP server capability on the HMC addresses the potential security concerns that users can have for attaching NTP servers directly to the HMC/SE LAN. When using the HMC as the NTP server, the pulse per second capability is not available.

#### HMC NTP broadband authentication support for z13 and z13s

The HMC NTP authentication capability is provided by the HMC Level 2.12.0 and later. SE NTP support stays unchanged. To use this option for STP, configure the HMC as the NTP server for the SE.

The authentication support of the HMC NTP server can be set up in either of two ways:

► NTP requests are UDP socket packets and cannot pass through the proxy. If a proxy to access outside corporate data center is used, then this proxy must be configured as a NTP server to get to target servers on the web. Authentication can be set up on the client's proxy to communicate to the target time sources.

► If a firewall is used, HMC NTP requests must pass through the firewall. Clients in this configuration should use the HMC authentication to ensure untampered time stamps.

### Coupling and Server Time Protocol connectivity

Coupling connectivity in support of Parallel Sysplex environments is provided on the z13 and z13s by the following features:

- ► PCIe Gen3, Integrated Coupling Adapter (ICA SR), which allows two ports coupling links connectivity for a distance of up to 150 m (492 feet) at 8 GBps each.

- ► HCA3-O,12x InfiniBand coupling links offering up to 6 GBps of bandwidth between z13, z13s, zBC12, z196 and z114 systems, for a distance of up to 150 m (492 feet).

- ► HCA3-O LR, 1x InfiniBand up to 5 Gbps connection bandwidth between z13, z13s, zEC12, zBC12, z196 and z114 for a distance of up to 10 km (6.2 miles).

- ► Internal Coupling Channels (ICs), operating at memory speed.

All coupling link types can be used to carry Server Time Protocol (STP) messages. The z13 and z13s do not support ISC-3 connectivity. Also, HCA2-O 12x and HCA2-O LR 1x InfiniBand features are not supported in z13 and z13s.

For more details about STP, see the following books:

- ► *Server Time Protocol Planning Guide*, SG24-7280
- ► *Server Time Protocol Implementation Guide*, SG24-7281
- ► *Server Time Protocol Recovery Guide,* SG24-7380

# 4.9  Hardware Management Console functions

The Hardware Management Console (HMC) and SE are appliances that provide hardware platform management for z Systems platforms. Hardware platform management covers a complex set of setup, configuration, operation, monitoring, and service management tasks and services that are essential to the use of the z Systems hardware platform product.

When tasks are performed on the HMC, the commands are sent to one or more SEs, which issue commands to their z Systems CPCs and zBXs.

HMC/SE Version 2.13.1 is the current version available for the z13 and z13s. For more information about these HMC functions and capabilities, see *IBM z13 Technical Guide*, SG24-8251 or *IBM z13s Technical Guide*, SG24-8294.

## 4.9.1  HMC key enhancements for z13 and z13s

The HMC application has several enhancements:

- ► Tasks and panels are updated to support configuring and managing Flash Express, IBM zAware, zEDC Express, and 10GbE RoCE Express features.

- ► The Backup for HMC and SE can be saved additional to an FTP server for z13 and z13s.

- ► OSA/SF is available on the HMC for specific OSA-Express features.

- ► For STP NTP broadband security, authentication is added to the HMC's NTP communication with NTP time servers and panels to configure STP are redesigned.

- ► NTP now uses Enhanced Console Assisted Recovery (ECAR) this speeds up the process of the BTS to takeover.

- ► Modem support is removed from HMC. The Remote Support Facility (RSF) for IBM support, service, and configuration update is only possible through an Ethernet broadband connection.

- The Monitors Dashboard on the HMC and SE is enhanced with an adapter table. The Crypto Utilization percentage is displayed on the Monitors Dashboard according to the PCHID number. The adapter table also displays Flash Express. You can now display the activity for a logical partition (LPAR) by processor type and the Monitors Dashboard is enhanced with showing simultaneous multithreading (SMT) usage.

- The Environmental Efficiency Statistic Task provides historical power consumption and thermal information for z13 or z13s on the HMC. This task provides similar data along with a historical summary of processor and channel use. The initial chart display shows the 24 hours that precede the current time so that a full 24 hours of recent data is displayed. The data is presented in table form, graphical (histogram) form, and it can also be exported to a `.csv` formatted file so that it can be imported into a spreadsheet.

- The microcode update to a specific bundle is possible.

- FCP SAN Explorer is a function available through the Problem Determination panels that provides a centralized view of SAN facilities available to an FCP channel. The tool facilitates configuration setting and debugging without requiring an operating system to be running. The FCP SAN Explorer shows device numbers (hosts) assigned to a partition, the fabric zone members available to the host, the remote end port error statistics, the accessible logical unit numbers (LUNs), and basic LUN configuration information can be queried and displayed. This data can also optionally be exported in CSV format.

---

**Statements of Direction:**[a]

- **Removal of support for Classic Style User Interface on the HMC and Support Element:** IBM z13 and z13s will be the last z Systems servers to support Classic Style User Interface. In the future, user interface enhancements will be focused on the Tree Style User Interface.

- **Removal of support for the HMC Common Infrastructure Model (CIM) Management Interface:** IBM z13 and z13s will be the last z Systems servers to support the HMC CIM Management Interface. The HMC Simple Network Management Protocol (SNMP) and Web Services application programming interfaces (APIs) will continue to be supported.

a. 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.

---

For more information about the key capabilities and enhancements of the HMC, see *IBM z13 Technical Guide,* SG24-8251 or *IBM z13s Technical Guide,* SG24-8294.

## 4.9.2 Hardware Management Console and Support Element

The Hardware Management Console (HMC) and Support Element (SE) appliances together provide hardware platform management for IBM z Systems. Hardware platform management covers a complex set of configuration, operation, monitoring, service management tasks, and other services that are essential to the use of the hardware platform product.

With z13s and z13, the HMC can be a stand-alone desktop computer or an optional 1U rack-mounted computer.

The z13s or z13 is supplied with a pair of integrated 1U SEs. One, the primary SE, is always active; the other is an alternate. Power for the SEs is supplied by the z Systems CPC's power supply, and there are no additional power requirements. The SEs are connected to the external customer switches for network connectivity with the CPC and the HMCs.

The SEs and HMCs are closed systems, and no other applications can be installed on them.

The HMCs and SEs of the system are attached to a Customer LAN. An HMC communicates with one or more z Systems and with the optional zBX Model 004 own internal SEs, as shown in Figure 4-3 on page 70. When tasks are performed on the HMC, the commands are sent to one or more SEs, which then issue commands to their CPCs and optional zBXs.

The HMC Remote Support Facility (RSF) provides communication with the IBM support network for hardware problem reporting and service.

Figure 4-3 shows an example of the HMC and SE connectivity.

**RSF connection:** RSF connection through a modem is *not* supported on the z13 or z13s HMC. An Internet connection to IBM is required to have hardware problem reporting and service.
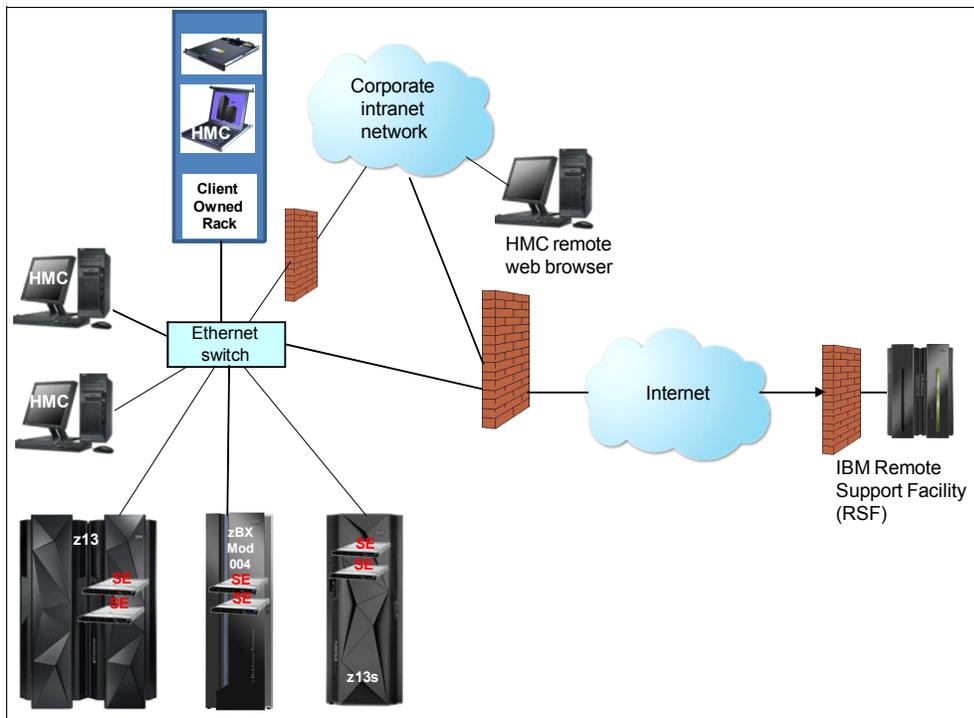


*Figure 4-3   HMC and SE connectivity*

# Strengths of the z13 and z13s

Each new z Systems platform introduces functions and features that cause it to jump ahead of its predecessor platforms. The IBM z13 and IBM z13s follow this pattern, while staying true to the underlying IBM z/Architecture.

z/Architecture is the conceptual structure of the z Systems platform that determines its basic behavior. Where appropriate, the z/Architecture itself can expand (rather than be replaced) to meet new needs, which helps sustain both the compatibility and architecture's integrity and longevity. Thus, when new capabilities are introduces in line with the architecture and at the same time existing solutions are protected, then *backward compatibility* has been achieved.

To handle new and different workloads, the scope of software and application options accommodate multiple operating systems, and the hardware and firmware components of the system must provide a viable option to integrate functionality into the architecture's capabilities. The z Systems platforms conform to the z/Architecture and its expansions ensure support of current and future workloads.

In this chapter, we highlight several z13 and z13s capabilities and strengths, and explain where they can be of value for businesses and organizations. Throughout the chapter, reference is made to the following helpful IBM Redbooks publications about the z13 and z13s:

► *IBM z13 Technical Guide*, SG24-8251
► *IBM z13s Technical Guide*, SG24-8294

This chapter includes the following topics:

# 5.1  z13 and z13s technology improvements

z13 and z13s include technology improvements that are intended to make systems integration more scalable, flexible, manageable, and secure.

The following sections provide details about the technology improvements for both the z13 and z13s.

## 5.1.1  Processor design highlights

The z/Architecture that underlies the z13 and z13s offers a rich complex instruction set computer (CISC) that supports multiple arithmetic formats. Compared to their predecessor systems, the z13 and 13s processor design includes the following improvements and architectural extensions:

►  Better performance and throughput:

  –  40% more system capacity

    Up to 141 characterizable Processor Units on the z13 (up to 20 on z13s)

  –  Faster processing (fourth generation high frequency processor)

    Larger cache (and shorter path to cache) means up to 10% faster uniprocessor performance on the z13 and up to 37% faster performance on the z13s

    Innovative core-cache design (L1 and L2), processor chip-cache design (L3), and node-cache design (L4) optimized by HiperDispatch, with focus on keeping more data closer the processor, increasing the cache sizes and decreasing the latency to access the next levels of cache

►  Re-optimized pipeline depth for power and performance:

  –  Increased instructions pipeline width per core

  –  Increased number of instructions inflight (from seven to ten)

  –  Greater integer execution bandwidth, with four fixed-point arithmetic execution units

  –  Improved fixed point and floating point divide

  –  Greater floating point execution bandwidth, with two binary and two decimal floating-point arithmetic execution units

►  Improved Instruction Fetching Unit, with new branch prediction capabilities and an instruction fetch front-end to support multithreading and to improve branch prediction throughput

►  Wider bandwidth for instruction decode, dispatch, and completion (increased to six instructions per cycle

►  Dedicated co-processor for each core:

  –  The Central Processor Assist for Cryptographic Function (CPACF) has been optimized to provide up to 2.3 times faster encryption functions than previous platforms.

  –  Hashing functions in CPACF are up to 3.9 times faster than they were in the predecessors

### Hardware decimal floating point function

The hardware decimal floating point (HDFP) function is designed to speed up calculations and provide the precision demanded by financial institutions and others. The HDFP fully implements the IEEE 754r standard.

### Simultaneous multithreading

Simultaneous multithreading (SMT) is built into the z13 and z13s IFLs and zIIPs and allows more than one thread to simultaneously execute in the same core, sharing all of its resources. This function improves utilization of the cores and increases processing capacity.

When a program accesses a memory location that is not in the cache, it is called a *cache miss*. Because the processor then must wait for the data to be fetched before it can continue to execute, cache misses affect the performance and capacity of the core to execute instructions. By using simultaneous multithreading, when one thread in the core is waiting (such as for data to be fetched from the next cache levels or from main memory), the second thread in the core can use the shared resources rather than remain idle.

Support for using SMT is provided in IBM z/OS for zIIPs and IBM z/VM for IFLs.

### Single-instruction, multiple-data

The z13 and z13s architecture includes a set of instructions called single-instruction, multiple data (SIMD) that can improve the performance of complex mathematical models and analytics workloads. This improvement is accomplished through vector processing and complex instructions that are able to process a large volume of data with a single instruction.

SIMD is designed for parallel computing and can accelerate code that contains integer, string, character, and floating point data types. enabling better consolidation of analytics workloads and business transactions on the z Systems platform.

### The Transactional Execution Facility

The Transactional Execution Facility, which is known in the industry as *hardware transactional memory*, allows instructions to be issued atomically. So either *all results* of the instructions in the group are committed or *no results* are committed, in a truly transactional manner. The execution is optimistic. The instructions are issued, but previous state values are saved in transactional memory. If the transaction succeeds, the saved values are discarded. If it fails, they are used to restore the original values. Software can test the success of execution and re-drive the code, if needed, using the same or a different path.

The Transactional Execution Facility provides several instructions, including instructions to declare the start and end of a transaction and to cancel the transaction. This capability can provide performance benefits and scalability to workloads by helping to avoid most of the locks on data. This ability is especially important for heavily threaded applications, such as Java.

### Runtime Instrumentation Facility

The Runtime Instrumentation Facility provides managed run times and just-in-time compilers with enhanced feedback about application behavior. This capability allows dynamic optimization of code generation as it is being executed.

## 5.1.2 Memory

Memory is significantly greater in the new z Systems models. The z13 can have up to 10 TB of usable memory installed, compared with the 3 TB maximum on the zEC12. The z13s can have up to 4 TB usable memory (versus 496 GB on the zBC12).

In addition, on the z13, the hardware system area (HSA) is expanded to 96 GB (was 32 GB on zEC12). On the z13s, the HSA is 40 GB (wasz13s 16 GB on zBC12). The HSA has a fixed size and is not included in the memory that the client orders.

The maximum memory size per logical partition (LPAR) has changed, too. For example, on the z13, up to 10 TB of memory can now be defined to an LPAR in the image profile. Each operating system can allocate main storage up to the maximum memory amount supported.

### Dynamic memory reassignment

LPAR memory assignment is done at the time of image activation. The defined partition memory is spread across the installed CPC drawers to maximize the use of existing memory controllers (MCs) on each CPC drawer.

On z13 and z13s, the memory allocation algorithm has evolved from previous systems. The goal of IBM Processor Resource/Systems Manager (PR/SM) memory and processors resources allocation is to assign all partition resources on a single CPC drawer, if possible. The resources, memory, and processors are assigned to the partitions when they are activated. Later, when all partitions are activated, PR/SM can dynamically move memory between CPC drawers to benefit performance, without impacting the operating system.

### Plan-ahead memory

If you anticipate someday increasing the installed memory, the initial system order can contain both starting and potential additional memory sizes. The additional memory is referred to as *plan-ahead memory*, and there is a specific memory pricing model to support it.

The starting memory size is activated when the system is installed, and the rest remains inactive. When more physical memory is required, it is fulfilled by activating the appropriate number of plan-ahead memory features. This activation is concurrent and can be nondisruptive to applications depending on the level of operating system support. z/OS and z/VM both support this function.

### Flexible memory

Flexible memory was first introduced on the IBM z9® EC as part of enhanced book availability (EBA). Flexible memory was used to temporarily replace the memory that becomes unavailable when performing maintenance on a book.

On z13, the additional resources that are required for flexible memory configurations are provided through the purchase of planned memory features and memory entitlement. Flexible memory configurations are available only on multi-CPC drawers (models N63, N96, NC9, and NE1) and range from 256 GB to 2.5 TB, depending on the model.

### Large page support

The size of pages and page frames has remained at 4 KB for a long time. Starting with the IBM System z10®, z Systems platforms are capable of having large pages of 1 MB, in

addition to supporting pages of 4 KB. This capability relates primarily to large main storage usage. Both page frame sizes can be simultaneously used.

Large pages enable 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. Large pages are better represented in the TLB and are expected to perform better.

Large pages can benefit long-running applications that are memory-access intensive and are not recommended for general use. Short-lived processes with small working sets see little to no improvement. Base the decision to use large pages on measurements of memory usage and page translation overhead for specific workloads.

### Support for 2 GB large pages

z13 and z13s use the same 2 GB page frames that were first introduced with zEC12 to increase efficiency for DB2 buffer pools, Java heaps, and other large structures. Using 2 GB pages increases TLB coverage without proportional growth in the size of the TLB:

► A 2 GB memory page is 2048 times larger than a 1 MB large page and 524,288 times larger than an ordinary 4 KB base page.

► A 2 GB page allows a single TLB entry to fulfill many more address translations than either a large page or ordinary base page.

► A 2 GB page provides users with much better TLB coverage, which improves performance:

  – Decreases the number of TLB misses that an application incurs
  – Spends less time converting virtual addresses into physical addresses
  – Uses less real storage to maintain DAT structures

## 5.2 Virtualization

The z13 and z13s are highly virtualized, with the goal of maximizing utilization of computing resources, lowering the total resources needed for defined workloads, and reducing the cost of those workloads. Virtualization is a key strength of z Systems platforms. It is embedded in the architecture and built into the hardware, firmware, and operating systems.

Virtualization requires a *hypervisor*, which is the control code that manages resources that are required for multiple independent operating system images. Hypervisors can be implemented as software or hardware, and the z13 and z13s have both. The hardware hypervisor is IBM Processor Resource/Systems Manager (PR/SM). PR/SM is implemented in firmware as part of the base system, fully virtualizes the system resources, and runs without any additional software. A software hypervisor (Type-2) is implemented with KVM for IBM z Systems and the z/VM operating system (both use PR/SM functions).

Virtualization is key to establishing flexible infrastructures capable of automated management and monitoring, such as those that underpin cloud offerings, including infrastructure as a service (IaaS) and platform as a service (PaaS).

z Systems platforms were designed with virtualization in mind. Virtualization is based on the concept of partitioning computer resources (such as CPU, memory, storage, and network resources) so that each set of features can be operated independently with its own operating environment.

IBM offers several virtualization technologies: PR/SM, z/VM, and KVM for IBM z Systems. The hypervisors are designed to enable simultaneous execution of multiple operating systems, providing operating systems with virtual resources.

Multiple hypervisors can exist on the same z Systems platform (see Figure 5-1).
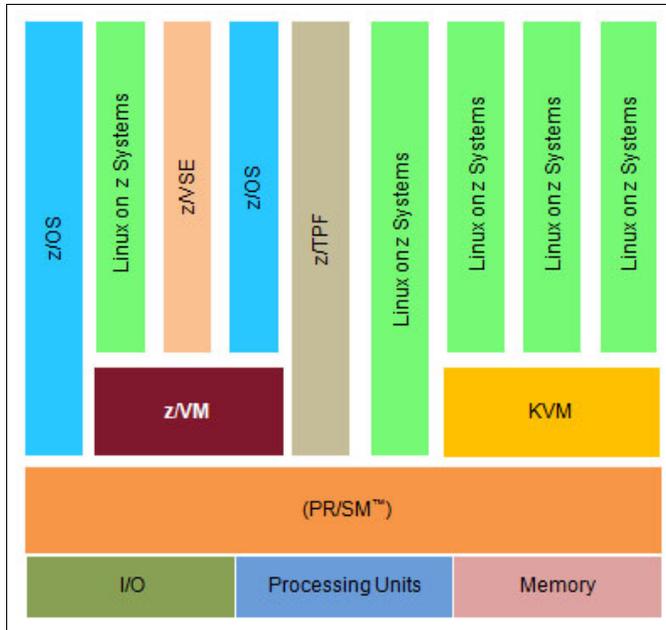


*Figure 5-1   Coexistence of different hypervisors on z Systems platforms*

The various virtualization options in z Systems platforms allow companies to migrate to new cloud-based or other virtualized environments while gaining the advantages and flexibility from open source software.

## PR/SM

Unique to z Systems platforms, PR/SM is a Type-1 hypervisor that runs directly on bare metal, allowing you to create multiple logical partitions (LPARs) on the same physical server. PR/SM is a highly stable, proven and secure, firmware-encapsulated virtualization technology that allows multiple operating systems to run on the same physical platform, with each operating system running in its own logical partition.

PR/SM logically partitions the platform across the various LPARs to share resources, such as Processor Units, memory, and I/O (for networks and storage), allowing for a high degree of virtualization.

## z/VM

z/VM is a Type-2 hypervisor that allows sharing the mainframe's physical resources, such as disk, memory, network adapters and Processor Units (CPs or IFLs). These resources are managed by the z/VM hypervisor, which typically executes on an LPAR and other virtual machines (VMs) that run on top of the hypervisor. Typically, the z/VM hypervisor is used to run Linux virtual servers. but other operating systems can also execute on z/VM. z/VM is a proven and well established virtualization platform. It supports up to 8000 virtual Linux servers in a single z13 footprint and provides industry-leading capabilities to efficiently scale both horizontally and vertically.

### KVM for IBM z Systems

KVM for IBM z Systems is an open virtualization solution that provides simple, cost-effective server virtualization for Linux workloads running on the z Systems platform. It is a Type-2 hypervisor that delivers server virtualization based on open source KVM Linux technology. KVM for IBM z Systems enables you to share real CPUs (called *IFLs*), memory, and I/O resources through server virtualization. KVM for IBM z Systems can coexist with z/VM virtualization environments, Linux on z Systems, z/OS, z/VSE, and z/TPF. KVM for IBM z Systems is optimized for scalability, performance, security, and resiliency, and provides standard Linux and KVM interfaces for simplified operational control.

## 5.2.1  z13 and z13s hardware virtualization

PR/SM was first implemented in the mainframe in the late 1980s. It allows you to define and manage LPARs. PR/SM virtualizes Processor Units, memory, and I/O features. Certain features are purely virtualized implementations.

PR/SM technology on the z13 received Common Criteria EAL5+ security certification. PR/SM is always active on the system and is enhanced to provide better performance and platform management benefits.

The LPAR definition includes a number of logical Processor Units (LPUs), memory, and I/O devices. IBM z/Architecture is designed to meet requirements with low overhead and also has achieved Common Criteria EAL5+ (the highest security certification in the industry) with a Specific Target of Evaluation (Logical Partitions). This design has been proven in many client installations over several decades.

Up to 85 LPARs can be defined on z13 (up to 40 LPARs on z13s) and hundreds or even thousands of virtual servers can be run under z/VM or KVM for IBM z. Thus, you can expect a high rate of context switching. In addition, accesses to the memory, caches, and virtual I/O devices must be kept isolated.

### Logical processors

Logical processors are defined and managed by PR/SM and are perceived by the operating systems as real processors. These processors can be characterized as follows:

► *Central processors* (CP) are standard processors for use with any supported operating system and user applications.

► *IBM System z Integrated Information Processor* (zIIP)[1] is used under z/OS for designated workloads, which include IBM Java Virtual Machine (JVM), various XML System Services, IPSec offload, certain parts of IBM DB2 DRDA®, star schema, IBM HiperSockets for large messages, and the IBM GBS Scalable Architecture for Financial Reporting

► *Integrated Facility for Linux* (IFL) is exclusively used with Linux on z Systems and for running the z/VM and KVM for z Systems hypervisors in support of Linux virtual machines (also called *guests*).

► *Internal Coupling Facility* (ICF) is used for z/OS clustering. ICF is dedicated to this function and exclusively run the *Coupling Facility Control Code* (CFCC).

In addition, the following pre-characterized processors are part of the base system configuration and are always present:

► System assist processors (SAP) that execute I/O operations
► Integrated firmware processors (IFP) for native PCIe features

---

[1] The zEC12 and zBC12 were the last z Systems servers to offer support for zAAPs. IBM supports running zAAP workloads on zIIP processors ("zAAP on zIIP").

These processors provide support for all LPARs but are never part of an LPAR configuration.

PR/SM accepts requests for work on logical processors by dispatching logical processors on physical processors. Physical processors can be shared across LPARs, but can also be dedicated to an LPAR. However, the logical processors of an LPAR must be either all shared or all dedicated.

The sum of logical processors defined in all active LPARs in a z Systems CPC might be higher than the number of physical processor units. The maximum number of LPUs that can be defined in a single LPAR cannot exceed the total number physical Processor Units that are available in the CPC. To achieve optimal ITR performance in sharing LPUs, the total number of online LPUs should be kept to a minimum, which reduces software and hardware overhead.

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

Logical processors can be dynamically added to and removed from LPARs. Operating system support is required to take advantage of this capability. z/OS, z/VM, and z/VSE each can dynamically define and change the number and type of reserved Processor Units in an LPAR profile. No pre-planning is required.

The new resources are immediately available to the operating systems and, in the case of z/VM, to its guest images. Linux on z Systems provides the Standby CPU activation/deactivation function.

## Memory

To ensure security and data integrity, memory cannot be concurrently shared by active LPARs. In fact, a strict isolation is maintained.

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 in real memory but are displayed as logically contiguous to the operating system that runs in the LPAR.

z/OS can take advantage of this support by nondisruptively acquiring and releasing memory from the reserved area. z/VM can acquire memory nondisruptively and quickly 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.

LPAR memory is said to be virtualized in the sense that, within each LPAR, memory addresses are contiguous and start at address zero. LPAR memory addresses are different from the system's absolute memory addresses, which are contiguous and have a single address of zero. Do not confuse this capability with the operating system that virtualizes its LPAR memory, which is done through the creation and management of multiple address spaces.

The z/Architecture has a robust virtual storage architecture that allows LPAR by LPAR definition of an unlimited number of address spaces and the simultaneous use by each program of up to 1023 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. A program's right to read or write in each page frame is validated by comparing the page key with the key of the program that is requesting access. This mechanism has been in use since the System/370. Memory keys were part of, and used by, the original System/360 systems. Definition and management of the address spaces is under operating system control. Three addressing modes (24-bit, 31-bit, and 64-bit) are simultaneously supported, which provides compatibility with earlier versions and investment protection.

z13 and z13s supports 2 GB pages, introduced with EC12 and zBC12, in addition to 4 KB and 1 MB pages. They also support an extension to the z/Architecture called Enhanced Dynamic Address Translation-2 (EDAT-2). With additional hardware, 1 MB pages can be pageable.[2]

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

Using the plan-ahead option, memory can be physically installed without being enabled. It can then be enabled when it is necessary. z/OS and z/VM support dynamically increasing LPAR memory size.

## I/O virtualization

The z13 supports six logical channel subsystems (LCSS) each with 256 channels, for a total of 1536 channels. The z13s supports three LCSS each with 256 channels, for a total of 768 channels. In addition to the dedicated use of channels and I/O devices by an LPAR, I/O virtualization allows concurrent sharing of channels. This architecture also allows sharing the I/O devices that are accessed through these channels, by several active LPARs. This function is known as *multiple image facility* (MIF). The shared channels can 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 that contain vital data sets).

When more 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 addresses (normally one base address and an average of three aliases). This feature increases the throughput of the device by using more device addresses. HyperPAV takes the technology a step further by allowing the I/O Supervisor (IOS) in z/OS (and the equivalent function in the Control Program of z/VM) to create PAV structures dynamically. The structures are created depending on the current I/O demand in the system, lowering the need for manually tuning the system for PAV use.

In large installations, the total number of device addresses can be high. Thus, the concept of *channel sets* as part of the z/Architecture.

### z13 subchannel sets

On the z13 up to four sets of approximately 64,000 device addresses are available. This availability allows the base addresses to be defined on set 0 (IBM reserves 256 subchannels

---

[2] z/OS support only

on set 0) and the aliases on set 1, set 2, and set 3. In total, 261,885 subchannel addresses are available per channel subsystem. Channel sets are used by the *Metro Mirror* (also referred as synchronous Peer-to-Peer Remote Copy (PPRC)) function by the ability to have the Metro Mirror primary devices defined in channel set 0. Secondary devices can be defined in channel sets 1, 2, and 3, providing more connectivity through channel set 0.

### z13s subchannel sets

On the z13s up to three sets of approximately 64,000 device addresses are available. This availability allows the base addresses to be defined on set 0 (IBM reserves 256 subchannels on set 0) and the aliases on set 1 and set 2. In total, 196350 subchannel addresses are available per channel subsystem. Channel sets are used by the *Metro Mirror* (also referred as synchronous Peer-to-Peer Remote Copy (PPRC)) function by the ability to have the Metro Mirror primary devices defined in channel set 0. Secondary devices can be defined in channel sets 1 and 2, providing more connectivity through channel set 0.

To reduce the complexity of managing large I/O configurations further, z Systems introduced *extended address volumes* (EAV). EAV provides large disk volumes. In addition to z/OS, both z/VM and Linux on z Systems support EAV.

By extending the disk volume size, potentially fewer volumes can be required to hold the same amount of data, making systems and data management less complex. EAV is supported by the IBM DS8000® series. Devices from other vendors should be checked for EAV compatibility.

The health checker function in z/OS has 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 z13 and z13s greatly eases the planning requirements and enhances the flexibility and availability of these reconfigurations.

## Dynamic Partition Manager

Dynamic Partition Manager (DPM) is a management infrastructure tool introduced with the z13 and z13s. It is intended to simplify virtualization management and is easy to use, especially for those who have less experience with z Systems. DPM provides simplified hardware and virtual infrastructure management including integrated dynamic I/O management without requiring you to learn complex syntax or command structures.

DPM provides a guided management interface to define z Systems hardware and virtual infrastructure including integrated dynamic I/O management that run KVM on IBM z Systems as a hypervisor or running bare metal Linux on z Systems. DPM provides consumable, enhanced partition life-cycle and integrated dynamic I/O management capabilities, as well.

DPM provides simplified, partition life-cycle and I/O management capabilities via the HMC to:

► Create and provision an environment, including new partitions, assignment of processors and memory, and configuration of I/O adapters

► Manage the environment, including the ability to modify system resources without disrupting workloads

► Monitor and troubleshoot the environment to identify system events that might lead to degradation

**Configuration note:** The z13 and z13s can be configured in DPM mode or in PR/SM mode but cannot be configured in both modes at the same time.

### 5.2.2  z Systems based clouds

Cloud computing is a paradigm for providing IT services. It capitalizes on the ability to rapidly and securely deliver standardized offerings, while retaining the capacity for customizing the environment. Elasticity, allowing to accompany the ebbs and flows of demand, and using just-in-time provisioning is another requirement. We make no distinction here between private and public clouds, because they are both well addressed by z Systems.

Virtualization is critical to the economic and financial viability of those offerings, because it allows minimizing the over-provisioning of resources, and reusing them at the end of the virtual server lifecycle.

Because of the extreme integration in the hardware, virtualization on z13 and z13s is highly efficient (the best in the industry) and encompasses computing and also I/O resources, including the definition of *internal virtual networks with switches*. These are all characteristics of *software defined environments* and allow supporting on a single real server, dense sets of virtual servers and server networks, with up to 100% sustained resource utilization and the highest levels of isolation and security. Therefore, the cloud solution costs, whether hardware, software, or management, are minimized.

Cloud elasticity requirements are covered by the z13 and z13s granularity offerings, including capacity levels and Capacity on Demand. These and other technologic leadership characteristics that make the z Systems platforms the server golden standard.

In addition, managing a cloud environment requires tools that can take advantage of a pool of virtualized compute, storage, and network resources, and present them to the consumer as a service in a secure way. A cloud management system should also help with these tasks:

► Offering open cloud management and application programming interfaces (APIs)
► Improving the usage of the infrastructure
► Lowering administrative overhead and improving operations productivity
► Reducing management costs and improving responsiveness to changing business needs
► Automating resource allocation
► Providing a self-service interface
► Tracking and metering resource usage

A cloud management system must also allow for the management of virtualized IT resources to support different types of cloud service models and cloud deployment models. IBM Cloud Manager with OpenStack (offered for z/VM and KVM for IBM z Systems) can satisfy a wide range of cloud management demands. It integrates various components to automate IT infrastructure service provisioning and delivers access to OpenStack APIs.

## 5.3  Capacity and performance

The z13 and z13s offers significant increases in capacity and performance over its predecessors, the zEC12 and zBC12. Many factors contribute to this effect, including the larger number of processors, individual processor performance, memory caches, simultaneous multithreading (SMT) and machine instructions, including the single-instruction, multiple-data (SIMD). Subcapacity settings continue to be offered.

### 5.3.1  Capacity settings

The z13 and z13s expand the offer on subcapacity 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. There are many performance and monitoring tools that are available on z Systems environments that are coupled with the flexibility of the capacity on-demand options (see 5.3.2, "Capacity on Demand (CoD)" on page 84). These features help to manage growth by making capacity available when needed.

#### z13 capacity levels
Regardless of the installed model, the z13 offers four distinct capacity levels for the first 30 central processors (CP):

- ► One full capacity
- ► Three subcapacities

These processors deliver the scalability and granularity to meet the needs of medium-sized enterprises, while also satisfying the requirements of large enterprises that have large-scale, mission-critical transaction and data-processing requirements.

A capacity level is a setting of each CP[3] to a subcapacity of the full CP capacity. The clock frequency of those processors remains unchanged. The capacity adjustment is achieved through other means.

Full capacity CPs are identified as CP7. On the z13 server, 141 CPs can be configured as CP7. The three subcapacity levels are identified by CP6, CP5, and CP4, respectively, and are displayed in hardware descriptions as feature codes on the CPs.

If more than 30 CPs are configured to the system, then all must be full capacity because all CPs must be on the same capacity level. Granular capacity adds 90 subcapacity settings to the 141 capacity settings that are available with full capacity CPs (CP7). The 231 distinct capacity settings in the system, provide for a range of over 1:320 in processing power.

A processor that is characterized as anything other than a CP, such as a zIIP, an IFL, or an ICF, is always set at full capacity. There is, correspondingly, a separate pricing model for non-CPs regarding purchase and maintenance prices, and various offerings for software licensing.

On z13, the CP subcapacity levels are a fraction of full capacity, as follows:

- ► Model 7xx = 100%
- ► Model 6xx = 63%
- ► Model 5xx = 44%
- ► Model 4xx = 15%

---

[3] The CP is the standard processor for use with any supported operating system, but is requested to run z/OS.

For administrative purposes, systems that have only ICFs or IFLs, are now given a capacity setting of 400. For either of these systems, having up to 141 ICFs or IFLs, which always run at full capacity, is possible.

Figure 5-2 gives more details about z13 full capacity and subcapacity offerings.
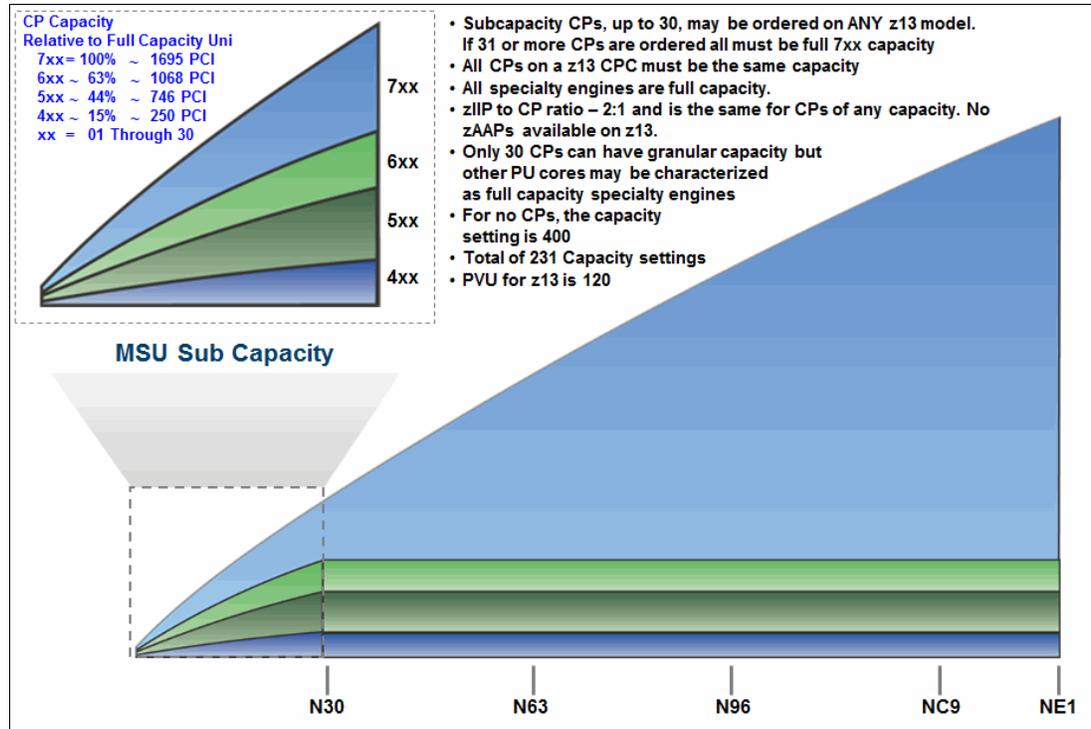


*Figure 5-2   z13 full and subcapacity CP offerings[4]*

To help size a z Systems platform to fit client requirements, IBM provides a no-cost tool that reflects the latest IBM LSPR measurements, called the *IBM Processor Capacity Reference for z Systems* (zPCR). You can download the tool here.

For more information about LSPR measurements, see 5.3.3, "z13 and z13s performance" on page 86.

## z13s capacity levels

The z13s offers 26 distinct capacity levels for each CP in the configuration for a total of 156 capacity settings (26 x 6). These processors deliver the scalability and granularity to meet the needs of small and medium-sized enterprises.

As in the z13, a z13s Processor Unit that is characterized as anything other then a CP, such as a zIIP, an IFL or an ICF, is always set to full capacity.

---

[4]  Processor Capacity Index (PCI)

Figure 5-3 gives more details about z13s capacities settings.
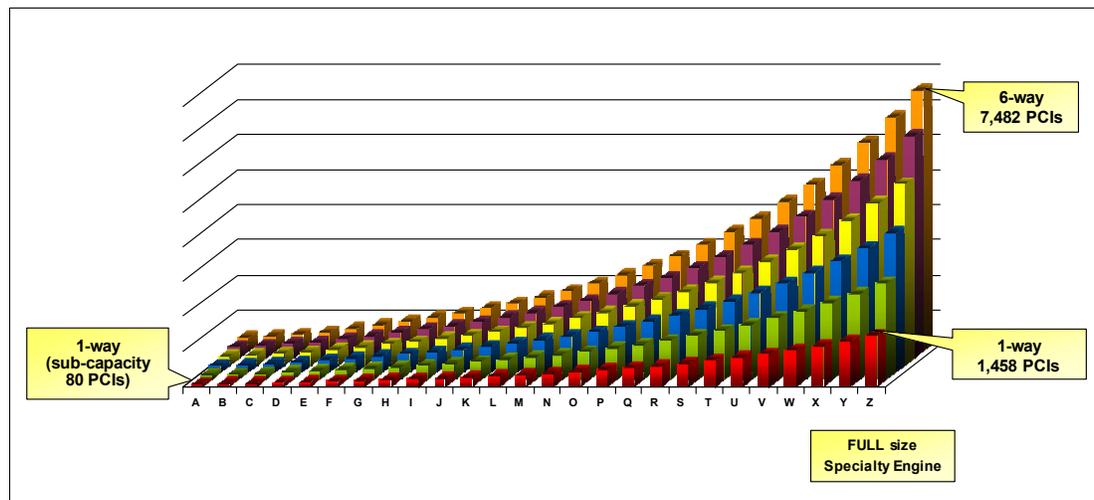


*Figure 5-3    z13s capacity settings offerings*

## 5.3.2  Capacity on Demand (CoD)

The z13 and z13s continues to provide on-demand offerings. They provide flexibility and control to the client, ease the administrative burden in the handling of the offerings, and give the client finer control over resources that are needed to meet the resource requirements in various situations.

The z13 and z13s can perform concurrent upgrades, providing an increase of processor capacity with no server outage. In most cases, with operating system support, a concurrent upgrade can also be non disruptive to the operating system. It is important to consider that these upgrades are based on the enablement of resources already physically present in the z13 or z13s.

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

For more information about the CoD offerings, see *IBM z13 Technical Guide*, SG24-8251 or *IBM z13s Technical Guide*, SG24-8294.

### Permanent upgrades
Permanent upgrades of processors (CP, IFL, ICF, zIIP, and SAP) and memory, or changes to a server's Model-Capacity Identifier, up to the limits of the installed processor capacity on an existing z13 or z13s, can be performed by the client through the IBM Online Permanent Upgrade offering by using the CIU facility.

## Temporary upgrades

Temporary upgrades of a z13 or z13s can be done by On/Off CoD, Capacity Backup (CBU), or Capacity for Planned Event (CPE) ordered from the CIU facility.

### On/Off CoD function

On/Off CoD is a function that is available on the z13 or z13s that enables concurrent and temporary capacity growth of the CPC. On/Off CoD can be used for client peak workload requirements, for any length of time, has a daily hardware charge and can have an associated software charge. On/Off CoD offerings can be pre-paid or post-paid. Capacity tokens are available on z13 and z13s. Capacity tokens are always present in prepaid offerings and can be present in post-paid if the client wants that. In both cases capacity tokens are being used to control the maximum resource and financial consumption.

When using the On/Off CoD function, the client can concurrently add processors (CP, IFL, ICF, zIIP, and SAP), increase the CP capacity level, or both.

### Capacity Backup (CBU) function

CBU allows the client to perform a concurrent and temporary activation of additional CP, ICF, IFL, zIIP, and SAP, an increase of the CP capacity level, or both. This function can be used in the event of an unforeseen loss of z Systems capacity within the client's enterprise, or to perform a test of the client'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 CPC drawers of the backup system, either as unused Processor Units or as a possibility to increase the CP capacity level on a subcapacity system, or both. A CBU contract must be in place before the LIC-CC code that enables this capability can be loaded on the system. An initial CBU record provides for one test for each CBU year (each up to 10 days in duration) and one disaster activation (up to 90 days in duration). The record can be configured to be valid for up to five years. Client can also order additional tests for a CBU record if needed, in quantities of five tests up to a maximum of 15.

Proper use of the CBU capability does not incur any additional software charges from IBM.

### Capacity for Planned Event function

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

The CPE feature is optional and requires unused capacity to be available on CPC drawers of the back-up system, either as unused Processor Units or as a possibility to increase the CP capacity level on a subcapacity system, or both. A CPE contract must be in place before the LIC-CC that enables this capability can be loaded on the system.

## z/OS capacity provisioning

Capacity provisioning helps clients manage the CP and zIIP capacity of z13 or z13s that is running one or more instances of the z/OS operating system. By using z/OS Capacity Provisioning Manager (CPM) component, On/Off CoD temporary capacity can be activated and deactivated under control of a defined policy. Combined with functions in z/OS, the z13 or z13s provisioning capability gives the client a flexible, automated process to control the configuration and activation of On/Off CoD offerings.

### 5.3.3  z13 and z13s performance

The z Systems microprocessor chip of the z13 and z13s has a high-frequency design that uses IBM leading technology and offers more cache per core than other chips. In addition, an enhanced instruction execution sequence, along with processing technologies such as SMT delivers world-class per-thread performance. z/Architecture is enhanced by providing more instructions, including SIMD, that are intended to deliver improved CPU-centric performance and analytics. For CPU-intensive workloads, more gains can be achieved by multiple compiler-level improvements. Improved performance of the z13 and z13s is a result of the enhancements that are described in Chapter 2, "IBM z13 hardware overview" on page 13 and Chapter 3, "IBM z13s hardware overview" on page 29, as well as in 5.1, "z13 and z13s technology improvements" on page 72.

The z13 Model NE1 offers up to 40% more capacity than the largest zEC12 system. Uniprocessor performance also increased significantly. A z13 Model 701 offers, based on an average workload, performance improvements of up to 10% over the zEC12 Model 701.

The z13s Model Z06 offers up to 51% more capacity than the largest zBC12 system. Uniprocessor performance also increased significantly. A z13s Model Z01 offers, based on an average workload, performance improvements of up to 34% over the zBC12 Model Z01.

However, variations on the observed performance increase depend on the workload type.

#### LSPR workload suite: z13 and z13s changes

To help you better understand workload variations, IBM provides a no-cost tool, IBM Processor Capacity Reference for z Systems (zPCR), which is available at the IBM Presentation and Tools website.

IBM continues to measure performance of the systems by using various workloads and publishes the results in the Large Systems Performance Reference (LSPR) report.

IBM also provides a list of MSU ratings for reference.

Historically, LSPR capacity tables, including pure workloads and mixes, have been identified with application names or a software characteristic. Examples are as follows:

► CICS
► IMS
► OLTP-T: Traditional online transaction processing workload (formerly known as IMS)
► CB-L: Commercial batch with long-running jobs
► LoIO-mix: Low I/O Content Mix Workload
► TI-mix: Transaction Intensive Mix Workload

However, capacity performance is more closely associated with how a workload uses and interacts with a particular processor hardware design. Workload capacity performance is sensitive to three major factors:

► Instruction path length
► Instruction complexity
► Memory hierarchy

With the availability of the CPU measurement facility (MF) data, the ability to gain insight into the interaction of workload and hardware design in production workloads has arrived. CPU MF data helps LSPR to adjust workload capacity curves that are based on the underlying hardware sensitivities, in particular the processor access to caches and memory. This is

known as *nest activity intensity*. With the IBM zEnterprise System, the LSPR introduced three workload capacity categories that replace all prior primitives and mixes:

► LOW (relative nest intensity):

A workload category that represents light use of the memory hierarchy. This category is similar to past high scaling primitives.

► AVERAGE (relative nest intensity):

A workload category that represents average use of the memory hierarchy. This category is similar to the past LoIO-mix workload and is expected to represent most of the production workloads.

► HIGH (relative nest intensity):

A workload category that represents heavy use of the memory hierarchy. This category is similar to the past TI-mix workload.

These categories are based on the relative nest intensity, which is influenced by many variables such as application type, I/O rate, application mix, CPU usage, data reference patterns, LPAR configuration, and the software configuration that is running, among others. CPU MF data can be collected by z/OS System Measurement Facility on SMF 113 records.

Guidance in converting LSPR previous categories to the new ones is provided, and built-in support is added to the IBM zPCR tool.

In addition to low, average, and high categories, the latest zPCR provides the low-average and average-high mixed categories, which allow better granularity for workload characterization.

The LSPR tables continue to rate all z/Architecture processors running in LPAR mode and 64-bit mode. The single-number values 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 values, the LSPR tables contain information that pertains to Linux and z/VM environments.

The LSPR contains the internal throughput rate ratios (ITRR) for the z13, z13s, and the previous generations of processors that are 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.

Experience demonstrates that z Systems servers can be run at up to 100% utilization levels, sustained, although most clients prefer to leave a bit of white space and run at 90% or slightly under. For any capacity comparison, using "one number" such as the MIPS or MSU metrics is not a valid method. That is why, when planning for capacity, we suggest using zPCR and involving IBM technical support. For more information about z13 or z13s performance, see *IBM z13 Technical Guide*, SG24-8251 or *IBM z13s Technical Guide*, SG24-8294.

## Throughput optimization with z13 and z13s

The z990 was the first server to use the concept of books (now called *CPC drawers*). The memory and cache structure implementation in the z13 and z13s CPC drawers are enhanced, from the z990 through successive system generations to the z13 and z13s, to provide sustained throughput and performance improvements. Although the memory is distributed throughout the CPC drawers and the CPC drawers have individual levels of caches that are private to the cores and shared by the cores, all processors have access to the highest level of caches and all of the memory. Thus, the system is managed as a memory coherent symmetric multiprocessor (SMP).

Processors within the z13 CPC drawer structure have different distance-to-memory attributes. As described earlier, CPC drawers are connected in a star configuration to minimize the distance. Other non-negligible effects result from data latency when grouping and dispatching work on a set of available logical processors. To minimize latency, the system attempts to dispatch and later re-dispatch work to a group of physical CPUs that share cache levels.

PR/SM manages the use 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 what 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 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, which is known as *HiperDispatch*.

### HiperDispatch

HiperDispatch, introduced with the z10 Enterprise Class, and evolved in z196 and zEC12, is further enhanced in z13 and z13s. It combines two functional enhancements, one in the z/OS dispatcher and one in PR/SM. This function is intended to improve efficiency in the hardware, in z/OS and in z/VM.

In general, the PR/SM dispatcher assigns work to the minimum number of logical processors that are needed for the priority (weight) of the LPAR. On z13 and z13s, PR/SM attempts to group the logical processors into the same node or in the neighbor node in the same CPC drawer and, if possible, in the same chip. This results in reducing the multi-processor effects, maximizing use of shared cache, and lowering the interference across multiple partitions.

The z/OS dispatcher is enhanced to operate with multiple dispatching queues, and tasks are distributed among these queues. Specific z/OS tasks can be dispatched to a small subset of logical processors. PR/SM ties these logical processors to the same physical processors, thus improving the hardware cache reuse and locality of reference characteristics, such as reducing the rate of cross communication.

To use the correct logical processors, the z/OS dispatcher obtains the necessary information from PR/SM through interfaces that are implemented on the z13 and z13s. The entire z13 or z13s stack (hardware, firmware, and software) now tightly collaborates to obtain the full potential of the hardware. z/VM HiperDispatch provides support similar to the z/OS one.

The HiperDispatch function is enhanced on the z13 and z13s to use the eight-core chip and improve computing efficiency. It is possible to dynamically switch on and off HiperDispatch without requiring an initial program load (IPL).

**Note:** HiperDispatch is required if SMT is enabled.

## 5.4  Reliability, availability, and serviceability

The IBM z Systems family presents numerous enhancements in reliability, availability, and serviceability (RAS). Focus was given to reducing the planning requirements, while continuing to reduce planned, scheduled, and unscheduled outages. One of the contributors to scheduled outages are License Internal Code (LIC) Driver updates that are performed in support of new features and functions. Enhanced driver maintenance (EDM) can help reduce the necessity and eventual duration of a scheduled outage. When properly configured, the z13 and z13s can concurrently activate a new LIC Driver level. Concurrent activation of the

select new LIC Driver level is supported at specifically released synchronization points. However, for certain LIC updates, a concurrent update or upgrade is not possible.

The effects of CPC drawer repair and upgrade actions are minimized on the z13 with enhanced drawer availability (EDA). In a multiple CPC drawer system, a single CPC drawer can be concurrently removed and reinstalled for an upgrade or repair. To ensure that the z13 configuration supports removal of a CPC drawer with minimal affect to the workload, consider the Flexible Memory option (see "Flexible memory" on page 74).

On a z13s concurrent repair or upgrade on the CPC drawers is *not* supported

The z13 or z13s provides a method to increase memory availability, referred to as redundant array of independent memory[5] (RAIM), where a fully redundant memory system can identify and correct memory errors without stopping. The implementation is similar to the RAID concept used in storage systems for a number of years. For a detailed description of RAS features, see *IBM z13 Technical Guide*, SG24-8251 or *IBM z13s Technical Guide*, SG24-8294.

To help prevent outages, improvements in several components of z13 and z13s are introduced. These enhancements include changes to the following components:

► Physical packaging
► Bus structures
► Processor cores
► Memory and cache hierarchy
► Power subsystem
► Thermal subsystem
► Service subsystem
► Integrated sparing

The z13 can consist of a maximum of four horizontal CPC drawers that are designed as a field replaceable unit (FRU). Connections among the CPC drawers are established using symmetric multiprocessing (SMP) cables. Each CPC drawer consists of two nodes, and each node contains three Processor Unit chips, one storage control chip, and 10 or 15 DDR3 DIMM slots.

Alternatively, the z13s central processor complex (CPC) subsystem consists of a maximum of two horizontal CPC drawers that are designed as a field replaceable unit (FRU). Connections among the CPC drawers are established using symmetric multiprocessing (SMP) cables. Each CPC drawer consists of two nodes, and each node contains two processor unit chips, one system controller chip, and 5 or 10 DDR3 DIMMs.

**Note:** The z13s model N10 CPC drawer only contains one node.

With a two-node CPC drawer structure, the z13 and z13s design supports system activation with partial-drawer resources in a degraded mode, if necessary. The Processor Unit and system controller chips are designed as single chip modules and FRUs.

A redundant pair of distributed converter assemblies (DCAs) step down the bulk power and connect to 10 point of load (POL) cards, which provide power conversion and regulation. Two redundant oscillators are connected to the drawers through an isolated backplane. Time domain reflectometry (TDR) techniques are applied to isolate failures on the SMP cables,

---

[5] Meaney, P.J.; Lastras-Montano, L.A.; Papazova, V.K.; Stephens, E.; Johnson, J.S.; Alves, L.C.; O'Connor, J.A.; Clarke, W.J., "IBM zEnterprise redundant array of independent memory subsystem," IBM Journal of Research and Development, vol.56, no.1.2, pp.4:1,4:11, Jan.-Feb. 2012, doi: 10.1147/JRD.2011.2177106

between chips (PU-PU, PU-SC, and SC-SC), and between the Processor Unit chips and DIMMs.

Additional redundancy is designed into new N+1 system control hubs (SCHs) and associated power supplies, and 1U service elements (SEs). Improvements to the z13 and z13s I/O infrastructure reliability include better recovery of FICON channels facilitated through forward error correction code (FEC) technology.

For a z13 air-cooled configuration features a fully-redundant N+2 radiator pump design that cools the Processor Unit chips through a water manifold FRU.

Further RAS enhancements include integrated sparing, error detection and recovery improvements in caches and memory, refreshes to IBM zAware, Flash Express, RoCE, and PCIe coupling, Fibre Channel Protocol support for T10-DIF, a fixed HSA with its size increased to 96 GB on the z13 and 40 GB on the z13s. OSA firmware changes to increase the capability of concurrent maintenance change level (MCL) updates, a new radiator cooling system with N+2 redundancy (z13 only), new CFCC level, and IBM RMF™ reporting.

z13 and z13s continues to support concurrent addition of resources, such as processors or I/O cards to an LPAR to achieve better serviceability. If an additional SAP is required on a z13 or z13s (for example, as a result of a disaster recovery situation), the SAPs can be concurrently added to the CPC configuration.

Concurrently adding CP, zIIP, IFL, and ICF processors to an LPAR is possible. This function is supported by z/VM, and also (with appropriate PTFs) by z/OS and z/VSE. Previously, proper planning was required to add CP, zAAP, and zIIP to a z/OS LPAR concurrently. Concurrently adding memory to an LPAR is possible. This is supported by z/OS and z/VM.

z13 and z13s supports adding Crypto Express features to an LPAR dynamically by changing the cryptographic information in the image profiles. Users can also dynamically delete or move Crypto Express features. This enhancement is supported by z/OS, z/VM, and Linux on z Systems.

## 5.4.1 RAS capability for the SE

Enhancements are made to the Support Element (SE) design for z13 and z13s. Notebooks that were used on prior generations of z Systems servers have been replaced with rack-mounted 1U System x servers in a redundant configuration on z13 and z13s. The more powerful 1U SEs offer RAS improvements such as ECC memory, redundant physical networks for SE networking requirements, redundant power modules, and better thermal characteristics.

## 5.4.2 RAS capability for the HMC

Enhancements are made to the HMC designs for z13 and z13s also. New for z13 and z13s is an option to order 1U servers for traditional and ensemble HMC configurations. This 1U HMC offers the same RAS improvements as those of the 1U SE. The 1U HMC option is a customer-supplied rack and power consolidation solution that can save space in data centers. The MiniTower design used prior to z13 will still be available.

The Unified Resource Manager is an active part of the ensemble infrastructure. Thus, the HMC has a stateful environment that needs high-availability features to ensure survival of the system in case of a HMC failure.

Each ensemble requires the following HMC workstations:

► A primary
► A backup (alternate)

The contents and activities of the primary are updated on the alternate HMC synchronously so that the alternate can take over the activities of the primary should the primary fail. Although the primary HMC can perform the classic HMC activities in addition to the Unified Resource Manager activities, the alternate HMC can be only a backup.

### 5.4.3 IBM z Advanced Workload Analysis Reporter

Introduced with the zEC12 and also available with the zBC12, the IBM z Advanced Workload Analysis Reporter (IBM zAware) feature is an integrated expert solution that uses sophisticated analytics to help identify potential problems and improve overall service levels.

IBM zAware runs analytics in a dedicated logical partition and now uses the IBM z Appliance Container Infrastructure. It intelligently examines z/OS and Linux on z Systems message logs for potential deviations, or inconsistencies, or variations from the norm, providing out-of-band monitoring and machine learning of operating system health. IBM zAware can accurately identify system anomalies in minutes. This feature analyzes massive amounts of processor data to identify problematic messages and provides information that can feed other processes or tools.

IBM zAware monitors the z/OS operations log (OPERLOG), which contains all messages that are written to the z/OS console, including application-generated messages. IBM zAware provides a graphical user interface (GUI) to help you easily drill-down into message anomalies, which can lead to faster problem resolution.

IBM zAware Version 2 was enhanced to support Linux on z Systems images running natively or as guests in KVM on IBM z or z/VM, on the z13 and z13s, identifying unusual system behavior by analyzing the syslog.

For more information about IBM zAware, see these sources:

► *IBM z13 Technical Guide*, SG24-8251
► *IBM z13s Technical Guide*, SG24-8294
► *Extending z/OS System Management Functions with IBM zAware*, SG24-8070
► *IBM z Advanced Workload Analysis Reporter (IBM zAware) Guide V2.0*, SC27-2632

## 5.5 High availability for z Systems with Parallel Sysplex

The z Systems platform is known for its reliability, availability, and serviceability capabilities, of which Parallel Sysplex is an exponent. The Parallel Sysplex technology is a clustering technology for logical and physical servers, allowing highly reliable, redundant, and robust z Systems technology to achieve near-continuous availability. Both hardware and software tightly cooperate to achieve this result.

A Parallel Sysplex has the following minimum components:

► Coupling facility (CF)

This is the cluster center. It can be implemented either as an LPAR of a stand-alone z Systems CPC or as an additional LPAR of a z Systems CPC where other loads are running. Processor units that are 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 runs in the CF. No other code runs in the CF. The code is used to create and maintain the structures, which are exploited under z/OS by software components such as z/OS itself, DB2 for z/OS, WebSphere MQ, among others.

CFCC can also run in a z/VM virtual machine (as a z/VM guest system). In fact, a complete sysplex can be set up under z/VM, allowing, for instance, testing and operations training. This setup is not recommended for production environments.

► Coupling links

These are high-speed links that connect 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 must 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 uses the hardware components to create a Parallel Sysplex. One example of z/OS and CF collaboration is the System-managed CF structure duplexing, which provides a general-purpose, hardware-assisted, easy-to-exploit mechanism for duplexing structure data hold in CFs. This function provides a robust recovery mechanism for failures (such as loss of a single structure on CF or loss of connectivity to a single CF). The recovery is done through rapid failover to the other structure instance of the duplex pair.

If you are interested in deploying system-managed CF structure duplexing, read the technical paper *System-Managed CF Structure Duplexing,* ZSW01975USEN, which you can access by clicking **Learn more** on the Parallel Sysplex website.

> **Note:** z/TPF can also use the CF hardware components. However, the term *sysplex* exclusively applies to z/OS usage of the CF.

Normally, two or more z/OS images are clustered to create a Parallel Sysplex. Multiple clusters can span several z Systems platforms, although a specific image (logical partition) can belong to only one Parallel Sysplex.

A z/OS Parallel Sysplex implements shared-all access to data. This is facilitated by z Systems I/O virtualization capabilities such as the *multiple image facility* (MIF). MIF allows several logical partitions to share I/O paths in a secure way, maximizing use and greatly simplifying the configuration and connectivity.

In short, a Parallel Sysplex comprises one or more z/OS operating system images that are coupled through one or more coupling facilities. A properly configured Parallel Sysplex cluster is designed to maximize availability *at the application level*. Rather than a quick recovery of a failure, the Parallel Sysplex design objective is *zero failure.*

The major characteristics of a Parallel Sysplex include the following features:

► 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 users of the data sharing capability are DB2, WebSphere MQ, WebSphere ESB, IMS, and CICS. Because these are major infrastructure components, applications that use them inherently benefit from sysplex characteristics. For example, many large SAP implementations have the database component on DB2 for z/OS, in a Parallel Sysplex.

► 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 more details, see *Improving z/OS Application Availability by Managing Planned Outages*, SG24-8178.

► High capacity

Parallel Sysplex scales from two to 32 images. Remember that each image can have from one to 141 processor units. CF scalability is near-linear. This structure contrasts with other forms of clustering that employ n-to-n messaging, which leads to rapidly degrading performance with a growing 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

This architecture provides the infrastructure to satisfy a client requirement for continuous availability and enables techniques for achieving simplified systems management consistent with this requirement.

► Resource sharing

A number of base z/OS components use CF shared storage. This usage enables the 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 is displayed as a single entity to the operator, user, database administrator, and so on. A single-system image ensures reduced complexity from both operational and definition perspectives.

► N-2 support

Multiple hardware generations (normally three, which are the current and the two previous ones) are supported in the same Parallel Sysplex. This configuration provides for a gradual evolution of the systems in the Sysplex, without forcing changing all simultaneously. Similarly, software support for multiple releases or versions is supported.

Figure 5-4 illustrates the components of a Parallel Sysplex as implemented within the z Systems architecture. The diagram shows one of many possible Parallel Sysplex configurations.
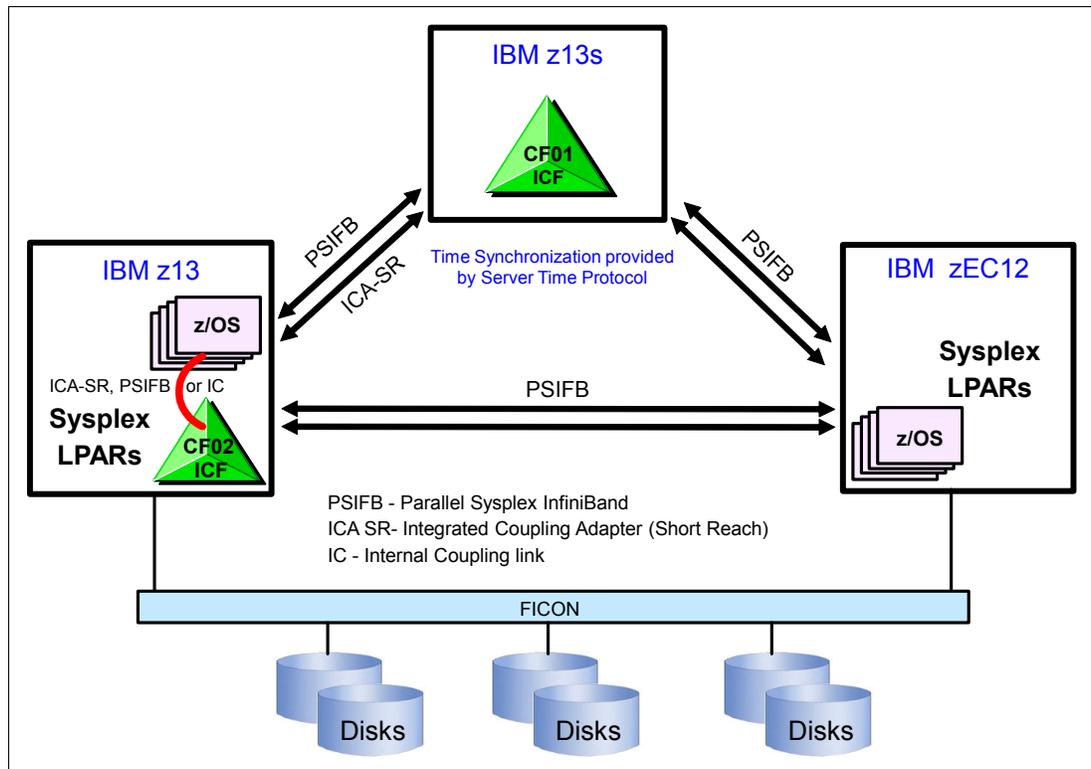


*Figure 5-4   Sysplex hardware overview*

Figure 5-4 shows a z13 system that contains multiple z/OS sysplex partitions and an internal coupling facility (CF02), a z13s server containing a stand-alone CF (CF01), and a zEC12 containing multiple z/OS sysplex partitions. STP over coupling links provides time synchronization to all servers. Appropriate CF link technology (1x IFB, 12x IFB, or ICA-SR) selection depends on server configuration and how distant they are physically located. ICA-SR links can only be used from z13 to z13s, within a short distance.

**6**

# Operating system support

This chapter describes the operating system requirements and support considerations for the IBM z13 and the IBM z13s and their features. It includes the following topics:

- ▶ 6.1, "Software support summary" on page 96
- ▶ 6.2, "Support by operating system" on page 99
- ▶ 6.3, "Software support for zBX Model 004" on page 111

Support and use of hardware functions depend on the operating system version and release. This information is subject to change. Therefore, for the most current information, see the following resources:

- ▶ For IBM z13: *Preventive Service Planning* (*PSP*) bucket for 2964DEVICE
- ▶ For IBM z13s: *Preventive Service Planning* (*PSP*) bucket for 2965DEVICE

# 6.1  Software support summary

The software portfolio for the z13 and z13s includes various operating systems and middleware that support the most recent and significant technologies. The following major operating systems are supported:

- ► z/OS
- ► z/VM
- ► z/VSE
- ► z/TPF
- ► Linux on z Systems
- ► KVM for IBM z Systems

For information about software that is supported on the zBX Model 004, see 6.3, "Software support for zBX Model 004" on page 111.

## 6.1.1  Operating systems summary

Table 6-1 lists the current and minimum operating system levels that are required to support the z13 and z13s. Operating system levels that are no longer in service are not covered in this publication. These older levels can provide support for certain features.

> **PTFs and PSP buckets:** The use of several features depends on a particular operating system. In all cases, program temporary fixes (PTF) might be necessary with the operating system level indicated.
>
> Preventive Service Planning (PSP) buckets are continuously updated and reviewed regularly when planning for installation of a new system. They contain the latest information about installation, hardware and software service levels, service recommendations, and cross-product dependencies.
>
> For Linux on z Systems distributions, consult the distributor's support information.
>
> For KVM for IBM z Systems, product can be ordered and delivered electronically using the IBM Shopz.
>
> For the current, cumulative fixes, download the latest available fix pack from IBM Fix Central.

*Table 6-1   z13 and z13s operating system requirements*

| Operating system | ESA/390 (31-bit mode) | z/Architecture (64-bit mode) | End of service | Notes |
|---|---|---|---|---|
| z/OS V2R2 | No | Yes | September 2020[a] | See the z/OS, z/VM, z/VSE, and z/TPF subsets of the 2964DEVICE and 2965DEVICE Preventive Service Planning (PSP) buckets before installing the z13 and z13s, respectively. |
| z/OS V2R1 | No | Yes | September 2018[a] | |
| z/OS V1R13 | No | Yes | September 2016[a] | |
| z/OS V1R12 | No | Yes | September 2014[b] | |
| z/VM V6R3[c] | No | Yes | December 2017[a] | |
| z/VM V6R2[c] | No | Yes | June 2017[a] | |
| z/VSE V6R1[d] | No | Yes | Not announced | |
| z/VSE V5R2[e] | No | Yes | Not announced | |
| z/VSE V5R1[e] | No | Yes | June 2016[a] | |
| z/TPF V1R1 | Yes | Yes | Not announced | |
| Linux on z Systems | No[f] | See Table 6-6 on page 108 | Support information is available for SUSE[g] and Red Hat[h] | |
| KVM for IBM z Systems | No | Yes | Not announced | |

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 Direction is at the relying party's sole risk and will not create liability or obligation for IBM.
b. z/OS V1R12 with required maintenance (compatibility support only) and extended support agreement. Service extension support is provided for up to three years beginning 1 October 2014.
c. z/VM V6R2 and V6R3 require an architectural level exclusive to System z10 and successors.
d. z/VSE V6 requires an architectural level set exclusive to System z10 and successors.
e. z/VSE V5 requires an architectural level set exclusive to IBM System z9® and successors.
f. 64-bit distributions include a 31-bit emulation layer to run 31-bit software products.
g. SUSE
h. Red Hat

## 6.1.2  Application development and languages

Several programming languages are available for the IBM z13 and IBM z13s environments. Because the Linux environment is similar to Linux on other servers, this discussion focuses on the z/OS environment.

In addition to the traditional COBOL, PL/I, Fortran, and Assembler languages, z Systems platforms support C, C++, and Java (including Java Platform, Enterprise Edition and batch environments) programming languages also.

Development can be conducted by using the latest software engineering technologies and advanced integrated development environments (IDE). 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 that are based on the Eclipse platform, are provided through the following offerings:

► IBM Rational® Application Developer for WebSphere
► IBM Rational Developer for z Systems
► IBM WebSphere Developer for z Systems
► IBM Rational Rose® product line
► IBM Rational Software Architect and Software Modeler

For more information about software for z Systems platforms, see the Products catalog website.

We cannot emphasize enough the importance of using the most recent versions of the compilers. The compilers enable the use of the latest technologies that are implemented on the system and take advantage of the performance benefits that are introduced. Examples of benefits include new cache structures, new machine instructions, and instruction execution enhancements.

For example, the z13 and z13s processors introduce the single-instruction, multiple-data (SIMD) instruction set, which uses the enhanced superscalar core to process a large number of operands (vector) through a single instruction. This function allows the development of smaller and optimized codes to improve efficiency of complex mathematical models and vector processing. This feature is fully used by the z/OS V2R1 operating system and later in conjunction with several compilers that have built-in functions for SIMD.

The z/VM guest exploitation for SIMD support is delivered with z13 and z13s general availability with PTF for APAR VM65733.

## 6.1.3 IBM compilers

Each new version of IBM z/OS compilers (Enterprise COBOL, Enterprise PL/I, XL C/C++) underscores the continuing IBM commitment to the COBOL, PL/I, and C/C++ programming languages on the z/OS platform.

► Enterprise COBOL

  The most recent version of Enterprise COBOL uses the most recent z/Architecture and performance optimization, enhanced XML parsing support, and capability of programming with Unicode, and supports Java 7 SDKs for Java interoperability.

► Enterprise PL/I

  The latest version of Enterprise PL/I provides web interoperability, which includes web services, XML parsers, and Java Platform, Enterprise Edition (Java EE). The compiler also includes the expanded support for UTF-16.

► z/OS XL C/C++

  The z/OS XL C/C++ uses the latest z/Architecture, including z13 and z13s servers. It enables developing high performance oriented applications, through the services provided by IBM Language Environment® and Runtime Library extension base elements, and works in concert with z/OS problem determination tools.

IBM Enterprise COBOL and Enterprise PL/I support are strategic components (separately orderable products) for IBM Rational Developer for IBM z Systems software. These features provide a robust, integrated development environment (IDE) for COBOL and PL/I and connecting web services, Java Platform, Enterprise Edition (Java EE) applications, and traditional business processes.

z/OS XL C/C++ programmers can also tap into Rational Developer for z Systems to help boost productivity by editing, compiling, and debugging z/OS XL C and XL C++ applications right from the workstation.

# 6.2 Support by operating system

This section lists the support by in-service operating systems of selected functions of the z13 and z13s.

For a detailed description of the z13 and its features, see *IBM z13 Technical Guide, SG24-8251*. For a detailed description of the z13s and its features, see *IBM z13s Technical Guide, SG24-8294*. For an in-depth description of all I/O features, see *IBM z Systems Connectivity Handbook*, SG24-5444.

## 6.2.1 z/OS

z/OS Version 1 Release 13 is the earliest in-service release that supports z13 and z13s. After September 2016, a fee-based service extension for defect support (for up to three years) can be obtained for z/OS V1R13. Although service support for z/OS Version 1 Release 12 ended in September of 2014, a fee-based extension for defect support (for up to three years) can be obtained by ordering IBM Software Support Services–Service Extension for z/OS 1.12.

Table 6-2 summarizes the support requirements of selected z13 and z13s functions for the currently supported z/OS releases. In the table, Y (yes) means the function is supported, N (no) means the function is not supported.

*Table 6-2   z/OS support summary*

| Function | V2R2 | V2R1 | V1R13[a] | V1R12[a] |
|---|---|---|---|---|
| Support of 141 Processor Units by a single system image[b] | Y | Y | N[c] | N[c] |
| Support of IBM zAware | Y | Y | Y | N |
| IBM *z Integrated Information Processor* (zIIP)[d] | Y[b] | Y[b] | Y[b] | Y[b] |
| 4 TB memory per LPAR exploitation | Y | Y[e] | N | N |
| Pageable 1 MB Large page support | Y | Y | Y[f] | N |
| Decimal Floating point support for packed decimal conversions | Y | Y | Y | N |
| 2 GB Large Page Support | Y | Y | Y | N |
| Support up to 85 LPARs[g] | Y | Y | Y | N |
| Support for six logical channel subsystems (CSS)[g] | Y | Y | Y | N |
| Support for four subchannel sets per CSS[g] | Y | Y | Y | N |
| Single-instruction, multiple-data (SIMD) support | Y | Y | N | N |
| Simultaneous multi threading support (SMT) | Y[h] | Y[h] | N | N |
| HiperDispatch | Y | Y | Y | Y |
| EP11 cryptography support | Y | Y | Y[i] | Y[i] |

| Function | V2R2 | V2R1 | V1R13[a] | V1R12[a] |
|---|---|---|---|---|
| Common Cryptographic architecture >16 Domain Support | Y | Y | Y[j] | N |
| CPACF | Y | Y | Y | Y |
| CPACF AES-128, AES-192, and AES-256 | Y | Y | Y | Y |
| CPACF SHA-1, SHA-224, SHA-256, SHA-384, SHA-512 | Y | Y | Y | Y |
| CPACF protected key | Y | Y | Y | Y |
| Crypto Express5S | Y[k] | Y[k] | Y[i] | Y[i] |
| Secure IBM Enterprise PKCS #11 (EP11) coprocessor mode | Y | Y | Y | Y |
| Elliptic Curve Cryptography (ECC) | Y | Y | Y | Y |
| Flash Express | Y | Y | Y | N |
| zEDC Express | Y | Y | N[l] | N[l] |
| Shared (SR-IOV) 10GbE RoCE Express | Y | Y | N | N |
| zHPF (High Performance FICON) Extended Distance II | Y | Y | Y | N |
| FICON Dynamic Routing | Y | Y | Y | N |
| FICON Express16S | Y | Y | Y | Y |
| FICON Express8S | Y | Y | Y | Y |
| FICON Express8[m] | Y | Y | Y | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR CHPID type OSD | Y | Y | Y | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR CHPID type OSX | Y | Y | Y | Y |
| OSA-Express5S Gigabit Ethernet LX and SX CHPID type OSD (using two ports per CHIPID) | Y | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSC (using two ports per CHPID) | Y | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSD (using two ports per CHPID) | Y | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSE(using one port per CHPID) | Y | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSM | Y | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSN (using one port per CHPID) | Y | Y | Y | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR CHPID type OSD | Y | Y | Y | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR CHPID type OSX | Y | Y | Y | Y |

| Function | V2R2 | V2R1 | V1R13[a] | V1R12[a] |
|---|---|---|---|---|
| OSA-Express4S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | Y | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSC (using two ports per CHPID) | Y | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | Y | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSE (using two ports per CHPID) | Y | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSM (using two ports per CHIPID) | Y | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSN (using two ports per CHPID) | Y | Y | Y | Y |
| Coupling using InfiniBand<br>CHPID type CIB | Y | Y | Y | Y |
| InfiniBand coupling links (12x IFB-SDR or 12x IFB-DDR) | Y | Y | Y | Y |
| InfiniBand coupling links (1x IFB-SDR or 1x IFB-DDR) | Y | Y | Y | Y |
| IBM Integrated Coupling Adapter support (ICA) | Y | Y | Y | Y |
| Server Time Protocol | Y | Y | Y | Y |
| CFCC Level 20 | Y | Y | Y | Y |
| Assembler instruction mnemonics | Y | Y | Y | Y |
| C/C++ exploitation of hardware instructions | Y | Y[n] | N | N |
| CPU measurement facility | Y | Y | Y | Y |

a. Service is required for support of z13 and z13s

b. Hardware limit of 6 CPs on the z13s.

c. z/OS 1.13 and z/OS 1.12 support 100 cores per LPAR

d. On an upgrade from zEC12, zBC12, z196 or z114, installed zAAPs are converted to zIIPs by default. zAAP is no longer supported on z13 and z13s.

e. Requires APAR OA47439

f. A web deliverable is required for Pageable 1M Large Page Support.

g. Support for up to 40 LPARs, three CSSs, and three subchannel sets per CSS on the z13s

h. Only for zIIP eligible workload

i. With Cryptographic support for z/OS V1R12-V1R13 web deliverable

j. With PTF and RSM enablement offering

k. Enhanced cryptographic support for z/OS V1R13-z/OS V2R2 web deliverable

l. Software decompression only.

m. Carry forward from zEC12, z196, zBC12, z114 only within an I/O drawer. This limits the LPAR memory to 1TB.

n. Exploitation Support for V2R1 XL C/C++ web deliverable

## 6.2.2 z/VM

At general availability, z/VM V6R3 and z/VM V6R2 provide compatibility support with use of some new z13 and z13s functions.

**Statements of Direction:[a]**

► *Removal of support for Expanded Storage (XSTORE):* z/VM 6.3 is the last z/VM release to support XSTORE for either host or guest usage. The IBM z13 and z13s will be the last z Systems server to support XSTORE.

► *Stabilization of z/VM V6.2 support:* The IBM z13 and IBM z13s are planned to be the last z Systems servers supported by z/VM V6.2 and the last z Systems servers that will be supported where z/VM V6.2 is running as a guest (second level). This support level is in conjunction with the statement of direction that the IBM z13 and IBM z13s will be the last z Systems servers to support ESA/390 architecture mode, which z/VM V6.2 requires. z/VM V6.2 will continue to be supported until 30 June 2017.

► *Product Delivery of z/VM on DVD/Electronic only:* z/VM 6.3 will be the last release of z/VM that will be available on tape. Subsequent releases will be available on DVD or electronically.

► *Enhanced IBM RACF® password encryption algorithm for z/VM:* In a future deliverable an enhanced RACF/VM password encryption algorithm is planned. This support will be designed to provide improved cryptographic strength using AES-based encryption in RACF/VM password algorithm processing. This planned design is intended to provide better protection for encrypted RACF password data in the event that a copy of RACF database becomes inadvertently accessible.

a. 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.

Table 6-3 summarizes the support requirements of selected functions for the currently supported z/VM releases. In the table, Y (yes) means the function is supported, N (no) means the function is not supported.

**Important:** Any new functions listed as *Y* for z/VM 6.2 and 6.3 require service (for both compatibility and exploitation).

*Table 6-3   z/VM support summary*

| Function | V6R3 | V6R2 |
|---|---|---|
| Support of up to 64 Processor Units for single thread[a] | Y[b] | N[c] |
| zIIP[d] | Y | Y |
| 1 TB real memory support | Y | N[e] |
| Large page support | N | N |
| Decimal floating point support for packed decimal numbers | Y | Y |
| Two way simultaneous multithreading (SMT) Support | Y | N |
| Single-instruction, multiple-data (SIMD) support[f] | Y | N |
| CPU measurement facility counter | Y | Y |
| HiperDispatch | Y | N |

| Function | V6R3 | V6R2 |
|---|---|---|
| CPACF | Y | Y |
| CPACF AES-128, AES-192, and AES-256 | Y | Y |
| CPACF SHA-1, SHA-224, SHA-256, SHA-384, SHA-512 | Y | Y |
| CPACF protected key | Y | Y |
| Crypto Express5S | Y | Y |
| Secure IBM Enterprise PKCS #11 (EP11) coprocessor mode | Y | Y |
| Elliptic Curve Cryptography (ECC) | Y | Y |
| Common Cryptographic Architecture (CCA) > 16 Domain Support | Y | Y |
| EP11 Support | Y | Y |
| Flash Express | N | N |
| zEDC Express[h] | Y[i] | N |
| Shared 10GbE RoCE (SR-IOV) support[j] | Y | N |
| High Performance FICON (zHPF) Extended distance[k] | N | N |
| FICON Express8S | Y | Y |
| FICON Express8 | Y | Y |
| FICON Express16S | Y | Y |
| OSA-Express QDIO data connection isolation for z/VM environments | Y | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR<br>CHPID type OSD | Y | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | N[l] | N[l] |
| OSA-Express5S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | Y | Y |
| OSA-Express5S 1000BASE-T<br>CHPID type OSC (using two ports per CHPID) | Y | Y |
| OSA-Express5S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | Y | Y |
| OSA-Express5S 1000BASE-T<br>CHPID type OSE[m] (using one or two ports per CHPID) | Y | Y |
| OSA-Express5S 1000BASE-T<br>CHPID type OSM | N[l] | N[l] |
| OSA-Express5S 1000BASE-T<br>CHPID type OSN (using two ports per CHPID) | Y | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSD | Y | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | Y | Y |

| Function | V6R3 | V6R2 |
|---|---|---|
| OSA-Express4S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSC (using two ports per CHPID) | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSE (using two ports per CHPID) | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSM | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSN (using two ports per CHPID) | Y | Y |
| Multi-Vswitch Link Aggregation | Y | N |
| Dynamic I/O support for InfiniBand CHPIDs | Y | Y |
| InfiniBand coupling links (12x IFB-SDR or 12x IFB-DDR) | Y | Y |
| InfiniBand coupling links (1x IFB-SDR or 1x IFB-DDR) | Y | Y |
| IBM Integrated Coupling Adapter support (ICA) | Y | Y |
| Dynamic I/O support for ICA CHPIDs | Y | Y |
| CFCC Level 20 | Y | Y |

a. z13s hardware limit 20 Processor Units

b. 64 cores support without multi-threading and 32 cores with multithreading-2

c. Supports only 32 Processor Unit in compatibility mode in z13 and z13s

d. zAAPs not available on z13 and z13s

e. Real memory limit is 256 GB

f. Guest exploitation only with PTF for APAR VM65733, at general availability (March 2016).

g. Service is required.

h. Minimum firmware bundle level 21

i. For z/OS guest support available through APAR

j. z/OS APAR required for exploiting SR-IOV for running z/OS as a guest in z/VM

k. zHPF(High Performance FICON) support available. But no support for extended distance.
Greater than 64 KB writes over 100 KM distance that is used by HyperSwap functions.

l. Only dynamic I/O when z/VM is the controlling LPAR.

m. A CHPID Type OSE supports SNA (LLC2) and IP connectivity over Ethernet (802.3 or DIX V2).

**z/VM logical partitions:** IBM z13 and z13s CPs and IFLs have increased capacity over that of their predecessors. Therefore, we suggest that the capacity of z/VM logical partitions and of any guests, in terms of the *number* of IFLs and CPs (real or virtual), be reviewed and adjusted to achieve the required capacity. Virtual machine might also need adjustment.

## 6.2.3  z/VSE

Table 6-4 summarizes the support requirements of selected z13 and z13s functions for the currently supported z/VSE releases. In the table, Y (yes) means the function is supported, N (no) means the function is not supported.

*Table 6-4   z/VSE support summary*

| Function | V6R1 | V5R2[a] | V5R1[a] |
|---|---|---|---|
| Support for up to ten CPs[b] [c] | Y | Y | Y |
| Large page support for data spaces (2 GB) | Y | Y | Y |
| Simultaneous multithreading (SMT) 2 support | N | N | N |
| Single-instruction, multiple-data (SIMD) support | N | N | N |
| CPACF | Y | Y | Y |
| CPACF AES-128, AES-192, and AES-256 | Y | Y | Y |
| CPACF SHA-1, SHA-224, SHA-256, SHA-384, SHA-512 | Y | Y | Y |
| CPACF protected key | N | N | N |
| Common Cryptographic Architecture (CCA) >16 Domain Support[d] | Y | Y[e] | Y[e] |
| Secure IBM Enterprise PKCS #11 (EP11) coprocessor mode | N | N | N |
| Crypto Express5S toleration | Y | Y[e] | Y[e] |
| Elliptic Curve Cryptography (ECC) | N | N | N |
| FICON Express8S | Y | Y | Y |
| FICON Express8 | Y | Y | Y |
| FICON Express16S | Y | Y | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR CHPID type OSD | Y | Y | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR CHPID type OSX | Y | Y | Y |
| OSA-Express5S Gigabit Ethernet LX and SX CHPID type OSD (using two ports) | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSC (using two ports per CHPID) | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSD (using two ports per CHPID) | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSE[f] (using two ports per CHPID) | Y | Y | Y |
| OSA-Express5S 1000BASE-T CHPID type OSM | Y | N | N |
| OSA-Express5S 1000BASE-T CHPID type OSN[g] (using two ports per CHPID) | Y | Y | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR CHPID type OSD | Y | Y | Y |

| Function | V6R1 | V5R2[a] | V5R1[a] |
|---|---|---|---|
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | Y | Y | Y |
| OSA-Express4S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSC (using two ports per CHPID) | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSE[f] (using two ports per CHPID) | Y | Y | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSM | N | N | N |
| OSA-Express4S 1000BASE-T<br>CHPID type OSN[g](using two ports per CHPID) | Y | Y | Y |
| HiperSockets Completion Queue for Linux Fast Path function in LPAR | Y | Y | Y |

a. z/VSE V5 executes in z/Architecture mode and supports 64-bit real and 64-bit virtual memory addressing. PTF support required for exploitation and compatibility

b. z/VSE Turbo Dispatcher supports up to four CPs and tolerates up to 10-way LPARs

c. z13s hardware limit of 6 CPs.

d. Coprocessor and Accelerator mode

e. Support available with PTF

f. A CHPID Type OSE supports both SNA (LLC2) and IP connectivity over Ethernet (802.3 or DIX V2).

g. One port is configured for OSN. The other port is unavailable.

## 6.2.4  z/TPF

Table 6-5 summarizes the support requirements of selected z13 and z13s functions for the currently supported z/TPF release. In the table, Y (yes) means the function is supported, N (no) means the function is not supported.

*Table 6-5   z/TPF support summary*

| Function | z/TPF V1R1[a] |
|---|---|
| Support for up to 141 Processor Units[b] | N[c] |
| 4 TB Real Storage Support | Y |
| CPACF | Y |
| CPACF AES-128, AES-192, and AES-256 | Y[d] |
| CPACF SHA-1, SHA-224, SHA-256 | Y[e] |
| CPACF protected key | N |
| Common Cryptographic Architecture (CCA) up to 85 Domain Support | Y |
| Secure IBM Enterprise PKCS #11 (EP11) coprocessor mode | N |
| Crypto Express5S | Y |
| Elliptic Curve Cryptography (ECC) | N |
| FICON Express16S | Y |

| Function | z/TPF V1R1[a] |
|---|---|
| FICON Express8S | Y |
| FICON Express8[f] | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR<br>CHPID type OSD | Y |
| OSA-Express5S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | N |
| OSA-Express5S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | Y |
| OSA-Express5S 1000BASE-T<br>CHPID type OSC (using two ports per CHPID) | N |
| OSA-Express5S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | Y |
| OSA-Express5S 1000BASE-T<br>CHPID type OSE(using two ports per CHPID) | N |
| OSA-Express5S 1000BASE-T<br>CHPID type OSM | N |
| OSA-Express5S 1000BASE-T<br>CHPID type OSN(using two ports per CHPID) | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSD | Y |
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | N |
| OSA-Express4S Gigabit Ethernet LX and SX<br>CHPID types OSD (using two ports) | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSC (using one or two ports per CHPID) | N |
| OSA-Express4S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | Y |
| OSA-Express4S 1000BASE-T<br>CHPID type OSE(using one or two ports per CHPID) | N |
| OSA-Express4S 1000BASE-T<br>CHPID type OSM | N |
| OSA-Express4S 1000BASE-T<br>CHPID type OSN (using one or two ports per CHPID) | Y |
| Coupling over InfiniBand<br>CHPID type CIB | Y[g] |
| IBM Integrated Coupling Adapter (ICA) | Y |
| CFCC Level 20 | Y |

a. PTF support required

b. z13s hardware limit of 20 Processor Units.

c. Maximum of 86 Processor Units per system image.

d. Supports only AES-128 and AES-256.

e. Supports only SHA-1 and SHA-256.

f. Carry forward from zEC12, zBC12, z196, and z114 only within an I/O drawer. This limits the LPAR memory to 1TB.

g. Compatibility is supported.

## 6.2.5  Linux on z Systems

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

None of the current versions of Linux on z Systems distributions (SUSE: SLES 11,SLES 12; Red Hat: RHEL 6, RHEL 7) require toleration support. For details about supported releases on z13 and z13s and the latest information about supported Linux distributions on z Systems, see the z Systems Tested Platforms website.

**Note:** Canonical and IBM intend to support Ubuntu 16.04 LTS on the z13 and z13s.

Table 6-6 lists the most recent service levels of the current SUSE and Red Hat supported to run on IBM z13 and IBM z13s.

*Table 6-6   Current Linux on z Systems distributions, by z/Architecture mode*

| Linux distribution | z/Architecture (64-bit mode) |
|---|---|
| SLES 11 | Yes |
| SLES 12 | Yes |
| RHEL 7 | Yes |
| RHEL 6 | Yes |

Table 6-7 lists selected z13 and z13s features, showing the minimum level of SUSE and Red Hat distributions that support each feature.

*Table 6-7   Linux on z Systems support summary*

| Function | SUSE | Red Hat |
|---|---|---|
| Maximum number of CPs or IFLs | 64/256[ab] | 64/256[ab] |
| Large page support | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| Decimal floating point for Packed decimal numbers | No | No |
| CPACF | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| CPACF AES-128, AES-192, and AES-256 | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| CPACF SHA-1, SHA-224, SHA-256, SHA-384, SHA-512 | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| CPACF protected key | Yes[c] | Yes[c] |
| Secure IBM Enterprise PKCS #11 (EP11) coprocessor mode | SLES 11,12[d] | RHEL 6,7[d] |
| Crypto Express5S[e] | SLES 12<br>SLES 11 SP3 | RHEL 7<br>RHEL 6.6 |
| Elliptic Curve Cryptography (ECC) | Yes[c] | Yes[c] |

| Function | SUSE | Red Hat |
|---|---|---|
| Common Cryptographic Architecture (CCA) > 16 domain support | SLES 12<br>SLES 11 SP3 | RHEL 7<br>RHEL 6.6 |
| HiperSockets Completion Queue | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| HiperSockets Virtual Switch Bridge | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| HiperSockets Layer 2 support | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| Shared 10GbE RoCE (SR-IOV) Express | SLES 12 SP1 | No |
| High Performance FICON (zHPF) extended distance | No | No |
| Simultaneous multithreading support (SMT) support | SLES 12 | RHEL 7 |
| Single-instruction, multiple-data (SIMD) support | SLES 12 | RHEL 7 |
| zFlash Express | SLES 12<br>SLES 11 SP3 | RHEL 7<br>RHEL 6.4 |
| zEDC Support | No | No |
| IBM zAware Support | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| FICON Express8S<br>CHPID type FC and FCP | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| FICON Express8,<br>CHPID types FC and FCP | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| FICON Express16S<br>CHPID type FC and FCP | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express5S 10 Gigabit Ethernet LR and SR<br>CHPID type OSD | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express5S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | SLES 12<br>SLES 11 SP1[f] | RHEL 7<br>RHEL 6 |
| OSA-Express5S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express5S 1000BASE-T<br>CHPID type OSC (using one or two ports per CHPID) | No | No |
| OSA-Express5S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express5S 1000BASE-T<br>CHPID type OSE[g] (using one or two ports per CHPID) | No | No |
| OSA-Express5S 1000BASE-T<br>CHPID type OSM | SLES 12<br>SLES 11 SP2 | RHEL 7<br>RHEL 6 |
| OSA-Express5S 1000BASE-T<br>CHPID type OSN (using two ports per CHPID) | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSD | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |

| Function | SUSE | Red Hat |
|---|---|---|
| OSA-Express4S 10 Gigabit Ethernet LR and SR<br>CHPID type OSX | SLES 12<br>SLES 11 SP1[f] | RHEL 7<br>RHEL 6 |
| OSA-Express4S Gigabit Ethernet LX and SX<br>CHPID type OSD (using two ports) | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express4S 1000BASE-T<br>CHPID type OSD (using two ports per CHPID) | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |
| OSA-Express4S 1000BASE-T[f]<br>CHPID type OSM | SLES 12<br>SLES 11 SP3 | RHEL 7<br>RHEL 6 |
| OSA-Express4S 1000BASE-T[f]<br>CHPID type OSN (using two ports per CHPID) | SLES 12<br>SLES 11 | RHEL 7<br>RHEL 6 |

a. SLES12 and RHEL7 supports 256 Processor Units and SLES11 and RHEL6 support 64 Processor Units. SLES and RHEL support 64 Processor Units when running as a guest under z/VM without multithreading and 32 Processor Units with multithreading.
b. z13s hardware limit of 20 Processor Units.
c. Through CCA host library
d. Requires EP11 host library
e. CEX5S toleration mode support
f. Maintenance update required.
g. A CHPID Type OSE supports both SNA (LLC2) and IP connectivity over Ethernet (802.3 or DIX V2).

IBM is working with its Linux distribution partners so that use of further IBM z13 and IBM z13s functions are provided in future Linux on z Systems distribution releases. We suggest the following updates:

► Use SLES 12 or RHEL 7 in any new projects for the z13 or z13s.

► Update any Linux distributions to their latest service level before migration to z13 or z13s.

► Adjust the capacity of any z/VM or Linux logical partitions, and of any z/VM guests, in terms of the *number* of IFLs and CPs, real or virtual, in face of the increased processor unit capacity of the z13 or z13s.

### 6.2.6 KVM for IBM z Systems

KVM is an open virtualization solution offered for z Systems servers that will provide simple, robust, cost-effective server virtualization for Linux workloads. KVM for IBM z currently requires SUSE Linux Enterprise Server (SLES) 12 SP1.

Supported functions for KVM for IBM z Systems 1.1.1 include:

► Simultaneous Multithreading (SMT) exploitation
► Guest exploitation of the Vector Facility for z/Architecture (SIMD)
► Hypervisor enhancements including support for iSCSI and NFS
► Hypervisor Crypto exploitation
► Enhanced RAS capabilities such as improved first failure data capture (FFDC)
► Improved high availability configuration
► Unattended installation of hypervisor

### 6.2.7  References

Planning information for each operating system is available at the following support websites:

- ► z/OS
- ► z/VM
- ► z/TPF
- ► z/VSE
- ► Linux on z Systems
- ► KVM for IBM z Systems

## 6.3  Software support for zBX Model 004

IBM z BladeCenter Extension (zBX) Model 004 is available as an upgrade from an existing zBX Model 002 or Model 003. The following operating systems would be supported on zBX Model 004:

- ► AIX (on POWER7 blade located in IBM BladeCenter Extension Model 004)

  AIX 5.3, AIX 6.1 and AIX 7.1 and subsequent releases and PowerVM Enterprise Edition

- ► Linux on System x (on IBM BladeCenter HX5 blade installed in zBX Model 004)

  RHEL 5.5 and later, 6.0 and later, RHEL 7.0 and later, SLES 10 (SP4) and later, SLES 11 (SP1) and later, SLES 12 and later (64-bit only)

- ► Microsoft Windows (on IBM BladeCenter HX5 blades installed in zBX Model 004)

  Microsoft Windows Server 2012, Microsoft Windows Server 2012 R2, Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008, SP2 (Datacenter Edition suggested) 64 bit only

# A

# Software licensing

Software licensing options are available for the z13 and z13s. This appendix provides information about these options and basic information about software licensing for the IBM z BladeCenter Extension (zBX) Model 004 environments.

The following topics are discussed in this appendix:

► Software licensing considerations
► Monthly License Charge pricing metrics
► zBX licensed software
► IBM z Unified Resource Manager

**113**

# Software licensing considerations

The IBM z13 and z13s 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. The portfolio also includes middleware for Linux on z Systems environments.

Use of the zBX software products are covered by the International Program License Agreement (IPLA) and additional agreements, such as the IBM International Passport Advantage® Agreement, similar to other AIX, Linux on System x, and Windows environments. PowerVM Enterprise Edition licenses must be ordered for IBM POWER7 blades.

For the z13 and z13s, two metric groups for software licensing are available from IBM, depending on the software product:

► Monthly license charge (MLC)

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 various offerings.

► International Program License Agreement (IPLA)

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 clients to access IBM product support during the support period. With this option, you can also receive future releases and versions at no additional charge.

For more details about software licensing, see the following resources:

► Learn about Software licensing
► Base license agreements
► IBM z Systems Software Pricing Reference Guide
► IBM z Systems Software Pricing

The remainder of this appendix describes the software licensing options that are available for IBM z13 and z13s.

# Monthly License Charge pricing metrics

Monthly License Charge (MLC) pricing applies to z/OS, z/VSE, or z/TPF operating systems. Any mix of z/OS, z/VM, Linux, z/VSE, and z/TPF images is allowed. Charges are based on processor capacity, which is measured in millions of service units (MSU) per hour.

### Charge models

Various Workload License Charges (WLC) pricing structures support two charge models:

► Variable charges (several pricing metrics)

Variable charges apply to products such as z/OS, z/VSE, z/TPF, DB2, IMS, CICS, and WebSphere MQ. Several pricing metrics employ the following charge types:

– Full-capacity license charges

The total number of MSUs of the central processor complex (CPC) is used for charging. Full-capacity licensing is applicable when the z Systems CPC of the client is not eligible for subcapacity.

– Subcapacity license charges

Software charges are based on the utilization of the logical partitions where the product is running.

► Flat charges

Software products that are licensed under flat charges are not eligible for subcapacity pricing.

## Subcapacity license charges

For eligible programs, subcapacity licensing allows software charges that are based on the measured utilization by logical partitions instead of the total number of MSUs of the z Systems CPC. Subcapacity licensing removes the dependency between the software charges and CPC (hardware) installed capacity.

The subcapacity licensed products are charged monthly based on the highest observed four-hour rolling average utilization of the logical partitions in which the product runs (except for products that are licensed by using the select application license charge (SALC) pricing metric). This type of charge requires measuring the utilization and reporting it to IBM.

The four-hour rolling average utilization of the logical partition can be limited by a defined capacity value on the image profile of the partition. This value activates the soft capping function of the Processor Resource/Systems Manager (PR/SM), limiting the four-hour rolling average partition utilization to the defined capacity value. Soft capping controls the maximum four-hour rolling average usage (the last four-hour average value at every five-minute interval), but does not control the maximum instantaneous partition use.

Also available is a logical partition (LPAR) group capacity limit, which sets soft capping by PR/SM for a group of logical partitions running z/OS.

Even by using the soft capping option, the use of the partition can reach up to its maximum share based on the number of logical processors and weights in the image profile. Only the four-hour rolling average utilization is tracked, allowing utilization peaks above the defined capacity value.

Some pricing metrics apply to stand-alone z Systems. Others apply to the aggregation of multiple z Systems workloads within the same Parallel Sysplex.

For more information about WLC and details of how to combine LPAR utilization, see *z/OS Planning for Workload License Charges*, SA22-7506.

## IBM z13

Metrics that are applicable to a stand-alone z13 include the following charges:

► Advanced Workload License Charges (AWLC)
► z Systems New Application License Charges (zNALC)
► Parallel Sysplex License Charges (PSLC)

Metrics that are applicable to a z13 in an actively coupled Parallel Sysplex include the following charges:

► AWLC, when all nodes are z Systems CPCs (z13, z13s, zEC12, zBC12, z196, or z114)

► Variable Workload License Charges (VWLC), allowed only under the AWLC Transition Charges for Sysplexes when not all of the nodes are z Systems CPCs

► zNALC

► PSLC

### IBM z13s

Metrics that are applicable to a stand-alone z13s include the following charges:

► Advanced Entry Workload License Charges (AEWLC)
► z Systems New Application License Charge (zNALC)
► Parallel Sysplex License Charge (PSLC)

Metrics that are applicable to a z13s in an actively coupled Parallel Sysplex include the following charges:

► AWLC, when all nodes are z Systems CPCs (z13, z13s, zEC12, zBC12, z196, or z114)
► Variable Workload License Charge (VWLC), allowed only under the AWLC Transition Charges for Sysplexes when not all of the nodes are z Systems CPC
► zNALC
► PSLC

## Advanced Workload License Charges

Advanced Workload License Charges (AWLC) type was introduced with the IBM zEnterprise 196. These charges use the measuring and reporting mechanisms, and also the existing MSU tiers, from VWLCs, although the prices for each tier were lowered.

AWLC can be implemented in full-capacity or subcapacity mode. The AWLC applies to z/OS and z/TPF and their associated middleware products such as DB2, IMS, CICS, and IBM WebSphere MQ, and IBM Lotus® IBM Domino®.

With z13 and z13s, Technology Transition Offerings are available that extend the software price and performance of the AWLC pricing metric:

► Technology Update Pricing for z13 and z13s is applicable for clients that run on a stand-alone z13/z13s or in an aggregated Parallel Sysplex consisting exclusively of z13/z13s servers.
► New Transition Charges for Sysplexes (TC3) are applicable when z13, z13s, zEC12 and zBC12 are the only servers in an actively coupled Parallel Sysplex.
► Transition Charges for Sysplexes (TC2) apply when two or more servers exist in an actively coupled Parallel Sysplex consisting of one or more z13, z13s, zEC12, zBC12, z196 or z114 servers.

For more information, see the AWLC web page.

## z Systems New Application License Charges

z Systems New Application License Charges (zNALC) offers a reduced price for the z/OS operating system on logical partitions that run a qualified new workload application. An example includes Java language business applications that run under WebSphere Application Server for z/OS or SAP.

z/OS with zNALC provides a strategic pricing model that is available on the full range of z Systems 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 more information, see the zNALC web page.

### Midrange Workload License Charges

Midrange Workload License Charges (MWLC) applies to z/VSE V5 and later when running on z13, zEC12, z196, System z10, and z9 servers. The exceptions are: the z10 BC and z9 BC servers at the capacity setting A01, to which zELC applies; and z114 and zBC12 where MWLC is not available.

Similar to workload license charges, MWLC can be implemented in full-capacity or subcapacity mode. An MWLC applies to z/VSE V5 and later, 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 V5R4 or later is required.

To report usage, the subcapacity report tool is used. One subcapacity reporting tool (SCRT) report per server is required.

For more information, see the MWLC web page.

### Parallel Sysplex License Charges

Parallel Sysplex License Charges (PSLC) applies to a large range of mainframe servers. The list is available at the Mainframe Exhibits section of the Reference web page.

Although it can be applied to stand-alone CPCs, the metric provides aggregation benefits only when applied to a group of z Systems CPCs in an actively coupled Parallel Sysplex cluster according to IBM terms and conditions.

Aggregation allows charging a product that is based on the total MSU value of the systems where the product runs (as opposed to all the systems in the cluster). In an uncoupled environment, software charges are based on the MSU capacity of the machine.

For more information, see the PSLC web page.

## z Systems International Program License Agreement

For z Systems, the following types of products are generally in the International Program License Agreement (IPLA) category:

► Data management tools
► DB2 for z/OS VUE
► CICS TS VUE V5 and CICS tools
► IMS DB VUE V12 and IMS tools
► Application development tools
► Certain WebSphere for z/OS products
► Linux middleware products
► z/VM V5 and V6

In general, three pricing metrics apply to IPLA products for z13, z13s, and z Systems:

► Value Unit (VU)

   This pricing metric 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.

- ► Engine-Based Value Unit (EBVU)

  This pricing metric enables a lower cost of incremental growth with more engine-based licenses purchased. Examples of eligible products include z/VM V5 and V6, and certain z/VM middleware, which are priced based on the number of engines.

- ► Processor Value Unit (PVU)

  This pricing metric is determined from the number of engines, under the Passport Advantage terms and conditions. Most Linux middleware is also priced based on the number of engines.

See the following web pages for more information:

- ► IPLA
- ► Engine-based Value unit prices

## zBX licensed software

The software licensing for the zBX select System x and POWER7 blades and DataPower XI50z follows the same rules as licensing for blades that are installed outside of zBX.

PowerVM Enterprise Edition *must* be licensed for POWER7 blades at the time of ordering the blades.

The hypervisor for the select System x blades for zBX is provided as part of the IBM z Unified Resource Manager (zManager).

## IBM z Unified Resource Manager

The z Unified Resource Manager is available through z13, z13s, zEC12, and zBC12 *hardware* features, either ordered with the system or later. No separate software licensing is required.

Redbooks

IBM z13 and IBM z13s Technical Introduction

Printed in U.S.A.

**Get connected**

ibm.com/redbooks