Redbooks

ibm.com/redbooks

IBM z16 Technical Introduction

Ewerson Palacio

Bill White

Octavian Lascu

IBM Z

IBM

Redbooks

IBM Redbooks

IBM z16 Technical Introduction

April 2023

SG24-8950-01

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

Second Edition (April 2023)

This edition applies to IBM z16 A01, IBM z16 A02, and IBM z16 AGZ.

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 customer 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

Preface

This IBM® Redbooks® publication introduces the latest member of the IBM Z® family, IBM z16™, which is built with the IBM Telum™ processor and available in three different configuration options:

- IBM z16 A01
- IBM z16 A02
- IBM z16 AGZ

The IBM Z platform is recognized for its security, resiliency, performance, and scale. It is relied on for mission-critical workloads and as an essential element of hybrid cloud infrastructures. IBM z16 adds capabilities and value with innovative technologies that are needed to accelerate the digital transformation journey.

This book explains how IBM z16 innovations and traditional IBM Z strengths satisfy the growing demand for cloud, analytics, and a more flexible infrastructure. With IBM z16 as the base, applications can run in a trusted, reliable, and secure environment that improves operations and lessens business risk.

Authors

This book was produced by a team of specialists from around the world working for IBM Redbooks, Poughkeepsie Center.

Octavian Lascu is an IBM Redbooks Project Leader and a Senior IT Consultant for IBM Romania with over 25 years of experience. He specializes in designing, implementing, and supporting complex IT infrastructure environments (systems, storage, and networking), including high availability and disaster recovery (HADR) solutions and high-performance computing deployments. He has developed materials for and taught over 50 workshops for technical audiences around the world. He is the author of several IBM publications.

Ewerson Palacio is an IBM Redbooks Project Leader. He holds a bachelor's degree in Math and Computer Science. Ewerson worked for IBM Brazil for over 40 years and retired in 2017 as an IBM Distinguished Engineer. Ewerson co-authored many IBM Z Redbooks publications, and created and presented ITSO seminars around the globe.

Bill White is an IBM Redbooks Project Leader and Senior IT Infrastructure Specialist at IBM Poughkeepsie, New York.

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

Brent Boisvert, John Campbell, Jeffrey Chang, Anne Dames, Richard Gagnon, Les Geer III, Ronald Geiger, Darelle Gent, Kyle Giesen, Michael Goetzner, Robert Haimowitz, Gerald Hosch, David Hutton, Seth Lederer, Cedric Lichtenau, Ed Linde II, Walter Niklaus, Thomas Pohl, Yamil Rivera, Lisa Schloemer, Christine Smith, Kenneth Stine, David Surman, John Torok, Marna Walle, Kathy Walsh, Tina Wile
IBM

Special thanks to the authors of the first edition:

Makus Ertl, Jannie Houlbjerg, Hervey Kamga, Gerard Laumay, Slav Martinksi, Kazuhiro Nakajima, Martijn Raave, Paul Schouten, Anna Shugol, Andre Spahni, John Troy, Roman Vogt, Bo Xu

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 IBM Redbooks 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, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- 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

Enduring the pace of digital transformation with IBM z16

For several years, digital transformation has been changing the way companies run by affecting every part of their business, from the infrastructure to processes, workflows, and even their culture. The recent pandemic accelerated the rate and pace of the digital transformation journey, which caused most companies to find novel ways to sustain operations while unlocking new opportunities and increasing innovation.

Open industry standard tools and agile DevOps methodologies are key to any successful digital transformation strategy and to accelerate modernization. An integrated platform that supports cloud-native development and infrastructure automation is essential.

To move the business forward while continuing to maintain the necessary levels of resiliency, compliance, and sustainability, a secure infrastructure with more flexibility and agility is needed. The latest member of the IBM Z family, IBM z16, can help with these new demands.

In this chapter, you learn about IBM z16 capabilities and the business value that it can provide.

This chapter describes the following topics:

1.1 The platform for the digital transformation journey

More than any other platform, IBM z16 offers a high-value architecture that can satisfy the growing demands that are driven by the pace of digital transformation, such as:

- More compute power for increased throughput
- Special hardware co-processors and accelerators to enhance performance
- Large-scale memory to process data faster
- Industry-unique cache design to optimize performance
- Accelerated I/O bandwidth to process massive amounts of data
- Quantum-safe cryptography to protect sensitive data now and in the future
- Flexible, on-premises hardware consumption pricing to handle the impact of unpredictable high spikes on business-critical workloads in a cost-effective manner
- Instant recovery:
  - Reduce the duration that is needed to start or shut down an OS or services.
  - Recover faster from system events.
  - Tailored short-duration boosts to mitigate the impact of IBM z/OS® SAN Volume Controller (SVC) dumps, middleware restarts, and IBM HyperSwap® recovery processes.
- Ability to dynamically shift production capacity across sites
- Focused sustainability to improve energy efficiency and cut down on space and carbon dioxide footprint requirements

The IBM Z platform is recognized for its security, resiliency, performance, and scalability. It is relied on for mission-critical workloads and as an essential element of hybrid cloud infrastructures. The IBM z16 adds more capabilities and value with innovative technologies that are needed to accelerate the digital transformation journey.

The IBM z16 is the first IBM Z platform that is built with the *IBM Telum processor*.[1] It is designed to help businesses achieve the following goals:

- Create value in every interaction and optimize decision making with the on-chip Accelerator for Artificial Intelligence (AI). The Accelerator for AI can deliver the speed and scale that is required to infuse AI inferencing into workloads with no impact on service delivery.
- Act now to protect today's data against current and future threats with quantum-safe protection through quantum-safe cryptography APIs and crypto-discovery tools.
- Enhance resiliency with flexible capacity to dynamically shift system resources across locations to proactively avoid disruptions.
- Modernize and integrate applications and data in a hybrid cloud environment with consistent and flexible deployment options to innovate with speed and agility.
- Reduce costs and keep up with changing regulations through a solution that helps simplify and streamline compliance tasks.

[1] IBM Telum Processor: the next-gen microprocessor for IBM Z and IBM LinuxONE

1.1.1 Predicting and automating with accelerated AI

An approach where data gravity and transaction gravity intersect, which co-collocates data, transactional systems, and AI inferencing, can deliver insights at speed and scale to enable decision velocity. *Decision velocity* means delivering insights faster to make decisions to help identify new business opportunities, improve customer experience, and reduce operational risk.

Consider the following points:

- The on-chip Integrated Accelerator for AI is designed for high-speed, real-time inferencing at scale. It is designed to add more than six TFLOPS of processing power shared by all cores on the chip. This centralized AI design is intended to provide high performance and consistent low-latency inferencing for processing a mix of transactional and AI workloads at speed and scale.

  Now, complex neural network inferencing that uses real-time data can be run and delivers insights within high-throughput enterprise workloads in real time while still meeting stringent service-level agreements (SLAs).
- A robust ecosystem of frameworks and open-source tools, when combined with the IBM Deep Learning Compiler that generates inferencing programs that are highly optimized for the IBM Z architecture and the Integrated Accelerator for AI, help enable rapid development and deployment of deep learning and machine learning models on IBM Z to accelerate time to market.
- IBM z16 memory is designed on a new memory buffer chip that provides up to DDR4-3200 memory speed, depending on memory size, delivering 50% more memory bandwidth per drawer than IBM z15®. This design improves overall workload performance, particularly for data-intensive analytics and AI applications. The new memory interface uses transparent memory encryption technology to protect all data leaving the processor chips before it gets stored in the memory DIMMs.

1.1.2 A cyberresilient system

A cyberresilient system can help protect against risks, vulnerabilities, attacks, and failures that might happen to your business.

With the opportunity that is created by quantum computing comes the threat to today's public key cryptography. Businesses should start now to prepare for the time when a quantum computer can break today's cryptography. In fact, today's data is at risk for future exposure through "harvest now, decrypt later" attacks.

IBM z16 is the industry-first quantum-safe system, which is protected by quantum-safe technology across multiple layers of firmware. Quantum-safe secure boot technology helps protect IBM z16 firmware from quantum attacks through a built-in dual signature scheme with no configuration changes that are required for enablement.

With the new Crypto Express8S (CEX8C), IBM z16 helps deliver quantum-safe APIs that position businesses to begin the usage of quantum-safe cryptography along with classical cryptography as they begin modernizing applications and building applications.

Discovering where and what kind of cryptography is being used is a key first step along the journey to quantum safety. IBM z16 provides instrumentation that can be used to track cryptographic instruction execution in the CP Assist for Cryptographic Functions (CPACF).

Additionally, IBM Application Discovery and Delivery Intelligence (ADDI) was enhanced with new crypto discovery capabilities.

For more information about the quantum-safe technologies that are used in IBM z16, see *Transitioning to Quantum-Safe Cryptography on IBM Z*, SG24-8525.

To move enterprises forward in a world that is constantly changing, businesses require an infrastructure that is flexible, secure, and resilient. The risk and potential disruption from extreme weather events, cyberattacks, and more continues to increase. An ever-changing and complex regulatory environment is driving up the cost to maintain and keep up with regulations.

Consider the following points:

- IBM z16 enhancements in resiliency include a new capability that is called *IBM Z® Flexible Capacity for Cyber Resiliency*. With Flexible Capacity for Cyber Resiliency, you can remotely shift capacity and production workloads between IBM z16 platforms at different sites on demand and stay at the alternative site for up to 1 year. This capability can help demonstrate compliance with regulations that require organizations to dynamically shift production to an alternative site and remain there for an extended period. This capability is designed to help you proactively avoid disruptions from unplanned events and planned scenarios, such as site facility maintenance.
- IBM System Recovery Boost (SRB) enhancements are available with IBM z16. The enhancements can provide boosted processor capacity and parallelism for specific events. Client-selected middleware starts and restarts can be boosted to expedite recovery for middleware regions and restore steady-state operations as soon as possible. z/OS SVC dump processing and HyperSwap configuration load and reload can be boosted to minimize the impact on running workloads.
- IBM Parallel Sysplex® enhancements include improved Integrated Coupling Adapter Short Reach (ICA SR) performance and Coupling Facility (CF) image scalability, technology, and protocol upgrades for coupling links, simplified Dynamic CF Dispatching (DYNDISP) support, and resiliency enhancements for CF cache and lock structures. CF Control Code (CFCC) boot and the update processes are protected with quantum-safe digital signatures with IBM z16.
- IBM Z Security and Compliance Center[2] is designed to help simplify and streamline compliance tasks. This solution provides a centralized, interactive dashboard for a consolidated view of system-generated evidence in near real time. You can check the security and compliance postures of your systems on demand and more easily identify drift so that it can be remedied quickly.

1.1.3 Modernizing for hybrid cloud

IBM z16 delivers technological innovation in AI, security, and resiliency on a flexible infrastructure that is designed for mission-critical workloads in a hybrid cloud environment. IBM z16 continues to deliver new and improved cloud capabilities.

IBM z16 provides the foundation for application modernization and hybrid cloud velocity by delivering a leading hybrid cloud infrastructure to support the optimization of mission-critical applications and data.

IBM z16 and the accompanying IBM Z and cloud software, which are developed to support a cloud-native experience, deliver a broad set of open and industry-standard tools, including an agile DevOps methodology to accelerate modernization. These capabilities deliver speed to market and agility for development and operational teams as IBM z16 integrates as a critical component of hybrid cloud.

[2] For more information, see *Keeping Up With Security and Compliance on IBM Z*, SG24-8540.

Businesses can accelerate modernization and delivery of new services by using the following key software offerings along with IBM z16:

- IBM Z and Cloud Modernization Stack helps empower developers to modernize and integrate z/OS applications with services across the hybrid cloud. This solution provides a flexible and integrated platform to support a z/OS based cloud-native development, application, and data modernization and infrastructure automation.
- Red Hat OpenShift and IBM Cloud® Paks running on an IBM z16 infrastructure provide the combination of infrastructure, hybrid cloud container platform, and middleware to modernize applications and develop cloud-native applications that integrate, extend, and supply data and workloads from IBM z16 across the hybrid cloud with Red Hat OpenShift.

1.1.4 Platform sustainability

IBM z16 marks a distinct sustainability focus across the product lifecycle, from the improved energy efficiency and enhancement of manufacturing and material sourcing to the improved packaging strategies for shipment and material recycling at product end-of-life.

IBM has a long-standing commitment to building a more sustainable, equitable future. In 1971, IBM formalized its environmental programs and commitment to leadership with the issuance of its Corporate Policy on IBM Environmental Responsibilities. This policy was released a quarter century before the first International Organization for Standardization (ISO) 14001 environmental management systems standard was published. IBM activities between then and 2021, when IBM committed to reaching net zero greenhouse gas emissions by 2030 in all 175 countries in which it operates and beyond, make it an ideal partner for the increasing number of businesses that consider sustainability a strategic direction. For more information, see IBM Commits To Net Zero Greenhouse Gas Emissions By 2030.

IBM z16 is the latest in a long line of machines that are designed for system and data center energy efficiency with differentiated architectural advantages, including on-chip compression, and encryption that is designed to sustain 90% utilization along with new embedded on-chip AI acceleration to seamlessly integrate real-time AI insights into business-critical transactions.

IBM z16 builds on a more than 24-year history of improving the system performance per watt, which is a key metric for improving the data center carbon footprint. Beginning with the first CMOS mainframe processor and continuing through IBM z16, the IBM Z platform has a 27-year history of improved mainframe system capacity per watt. The largest IBM z16, which uses an intelligent power distribution unit (PDU), improves maximum system capacity per kilowatt by approximately 18% compared to the largest z15, and an approximately 54% increase versus the largest IBM z14®.

Within each single CPC drawer, IBM z16 provides 25% greater capacity than z15 for standard models and 40% greater capacity on the max configuration model, enabling efficient scaling of partitions.

A typical IBM z14 bulk power assembly (BPA) system upgraded to a similarly configured IBM z16 PDU system reduces system power consumption and provides other significant data center advantages: 31% less weight, 50% or more reduction in floor space, and direct participation in the modern data center.

The biggest opportunity for energy savings with IBM z16 comes through workload modernization and consolidation for distributed x86 systems. Many enterprises cannot easily grow their data center size, which inhibits dealing with the workload surges that are inherent in planned or reactive digital transformation on a distributed platform. The vertical scalability of the IBM Z architecture can address this problem while reducing your power usage. For example, running Linux workloads on an IBM z16 (3931-7C5) server with 125 processor units (PUs) and 30 TB of memory shows a nearly 17:1 core consolidation ratio while reducing the usage of electricity by 75% versus x86 servers running workloads under similar conditions.

IBM z16 prioritizes how IBM contributes to the circular economy. For example, the IBM z16 publication *Product Carbon Footpring* shows the attributes of the product lifecycle that have the biggest impact on energy usage. IBM z16 continues the long history of focusing on the product lifecycle, from the improved energy efficiency and enhancement of manufacturing and material sourcing, to the improved packaging strategies for shipment and material recycling at product end-of-life.

Product Carbon Footprint reports that using the Product Attributes to Impact Algorithm (PAIA) help businesses understand the lifecycle carbon sustainability of the hardware within their data centers. IBM z16 provides telemetry information that is useful for the integration into modern data center infrastructure management (DCIM) systems through a set of secure, REST-based web service APIs.

IBM z16 energy management

IBM z16 can integrate modern DCIM systems through a set of secure, REST-based web services APIs. The HMC Web Services API is the access point for any external tools to manage the IBM Z platform. This API also supports management of the lifecycle and the configuration of various platform resources, such as logical partitions, CPU, memory, virtual switches, I/O adapters, and more.

With the Web Service API metrics groups, IBM z16 can report key environmental and power consumption data, including ambient temperature and humidity, heat load, exhaust temperature, system power consumption, and power on each power cord phase.

This data is alternatively available through a *Hardware Management Console* (HMC) *monitors dashboard*. The HMC also provides tools to support energy management. With the *Max Potential Power* task, operators can reduce the power allocation for a system because the task knows the maximum power that the system can draw even during faults or a hot data center environment. This task is valuable for reducing stranded power, and it looks like power capping to higher-level management tools.

For more information about the HMC monitors dashboard, see *IBM z16 (3931) Technical Guide*, SG24-8951.

Additionally, the `zhmcclient` package, which is available on GitHub, is a client library that is written in pure Python that interacts with the Web Services API of the HMC. The goal of this package is to make the HMC Web Services API consumable for Python programmers. In addition to these tools, IBM has a sophisticated Power, Weight, and Airflow calculator, which you can find at IBM Resource Link®.[3]

[3] An IBM Resource Link user ID is required.

1.2 IBM z16 technical overview

IBM z16 aligns with the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Class A3 data center guidelines. IBM z16 is available in three different configuration options:

- IBM z16 A01 is built with a 19-inch format that scales 1 - 4 frames, depending on the configuration. IBM z16 A01 ensures continuity and upgradeability from IBM z15 T01 and IBM z14 M0x. It has five orderable features: Max39, Max82, Max125, Max168, and Max200.
- IBM z16 A02 is built with a 19-inch format single frame. IBM z16 A02 ensures continuity and upgradeability from IBM z15 T02 and IBM z14 ZR1. There are four orderable features: Max5, Max16, Max32, and Max68.
- IBM z16 AGZ indicates a rack-mount configuration that allows the core compute, I/O, and networking features to be installed into and powered by a client-designated rack with power distribution units (PDUs), respectively. IBM z16 AGZ ensures continuity and upgradeability from IBM z15 T02 and IBM z14 ZR1. The rack-mount configuration options are under a combined AGZ warranty umbrella and orderable as Max5, Max16, Max32, and Max68.

The feature names are based on the maximum number of characterizable PUs.

Figure 1-1 shows the available IBM z16 configuration options (one to four 19-inch frames for IBM z16 A01, one 19-inch frame for IBM z16 A02 and IBM z16 AGZ, and mountable components for a client-designated 19-inch rack with IBM z16 AGZ).

Figure 1-1 IBM z16 19-inch frame configurations

Table 1-1 lists some of the IBM z16 technical capabilities.

Table 1-1 Technical highlights of IBM z16 A01, IBM z16 A02, and IBM z16 AGZ

Capability	IBM z16 A01	IBM z16 A02 and IBM z16 AGZ
Greater total system capacity and more subcapacity settings for CPs. IBM z/Architecture® ensures continuity and upgradeability from previous models.	► Up to 200 characterizable PUs. ► Up to 39 subcapacity settings for CPs. ► Up to 317 total capacity levels.	► Up to 68 characterizable PUs. ► Up to 6[a] CPs with 156 subcapacity settings.
Dual-chip modules (DCMs)[b] that use the IBM Telum processor to help improve the execution of processor-intensive workloads.	5.2 GHz.	4.6 GHz.
Memory per system, which ensures high availability (HA) in the memory subsystem by using proven redundant array of independent memory (RAIM) technology.	Up to 40 TB of addressable real memory per system.[c]	Up to 16 TB of addressable real memory per system.[d]
A large fixed hardware system area (HSA) that is managed separately from ordered memory.	256 GB.	160 GB.
Processor cache structure improvements and larger cache sizes to help with more demanding production workloads. The result is 1.5x cache capacity per core compared to z15, at reduced average access latency, through a flatter system topology and overall improved system performance and scalability.	► First-level cache (L1): 128 KB. ► Second-level cache (L2): 32 MB. ► Third-level cache (L3): 256 MB. ► Fourth-level cache (L4): 2048 MB.	► First-level cache (L1): 128 KB. ► Second-level cache (L2): 32 MB. ► Third-level cache (L3): 256 MB. ► Fourth-level cache (L4): 1024 MB (Max5 & Max16). 2048 MB (Max32 & Max68).
The channel subsystem (CSS) is built for I/O resilience. The number of logical channel subsystems (LCSSs), subchannel sets, and I/O devices are consistent with its predecessor platform, as is the number of LPARs.	► Six LCSSs. ► 85 LPARs. ► Four subchannel sets. ► 64,000 I/O devices per subchannel set.	► Three LCSSs. ► 40 LPARs. ► Three subchanel sets. ► 64,000 I/O devices per subchannel set.
Proven technology (sixth-generation high frequency and fourth-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. Special coprocessors and new hardware instructions for accelerating selected workloads.		
IBM Virtual Flash Memory (VFM) can be used to handle paging workload spikes and improve availability.		
The IBM Z Sort accelerator helps reduce CPU costs, speed up the sorting process, and improve database functions.		
Improved cryptographic functions and performance, which is achieved by having one dedicated cryptographic coprocessor per PU.		
Enhanced ICA SR coupling link protocol provides up to 10% improvement for read and lock requests, and up to 25% for write requests and duplexed write requests compared to CF service times on z15 servers.		
Improved CF processor scalability for CF images. The relative scaling of a CF image beyond a 9-way is improved, meaning that the effective capacity of IBM z16 CF images continues to increase all the way up to the max of 16 processors in a CF image.		
Dedicated on-chip AI accelerator is designed for high-speed, real-time inferencing at scale.		
Improved data compression operations are achieved by having one dedicated compression coprocessor per PU, and new hardware instructions.		

a. Up to five CPs on a Max5.
b. IBM z16 DCMs use 7-nm technology.
c. IBM z16 A01 supports up to 10 TB of memory per drawer.
d. IBM z16 A02 and IBM z16 AGZ support up to 8 TB of memory per drawer.

For more information about the IBM z16 configuration options, see Chapter 2, “IBM z16 A01 hardware overview” on page 17 and Chapter 3, “IBM z16 A02 and IBM z16 AGZ hardware overview” on page 33.

1.2.1 Storage connectivity

Storage connectivity is provided on IBM z16 by IBM FICON® Express and the IBM zHyperLink Express features.

FICON Express

FICON Express 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 the following protocols:

- FICON

  This enhanced protocol 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, IBM z/VM®, IBM z/VSE®[4], 21st Century Software VSE^n V6.3, z/TPF, and Linux on IBM Z environments.

- Fibre Channel Protocol (FCP)

  This standard protocol is used for communicating with disk and tape devices through FC switches and directors. The FCP channel can connect to FCP SAN fabrics and access FCP and SCSI devices. FCP is used by z/VM, Kernel-based Virtual Machine (KVM), z/VSE[4], 21st Century Software VSE^n V6.3, and Linux on IBM Z environments.

FICON Express32S features are implemented by using Peripheral Component Interconnect Express (PCIe) cards, and offer better port granularity and improved capabilities over the previous FICON Express features. FICON Express32S is the preferred technology for new systems. The features support a link data rate up to 32 Gbps and can auto-negotiate to 8, 16, or 32 Gbps.

zHyperLink Express

zHyperLink was created to provide fast access to data through low-latency connections between the IBM Z platform and storage.

With the zHyperLink Express1.1 feature, you can make synchronous requests for data that is in the storage cache of the IBM DS8000®. This process is done by directly connecting the zHyperLink Express1.1 port in IBM z16 to an I/O bay port of the IBM DS8000. This short-distance (up to 150 m (492 feet)), direct connection is designed for low-latency read/write, such as with IBM Db2® for z/OS synchronous I/O reads and log writes.

Working with the FICON SAN Infrastructure, zHyperLink can improve application response time, which reduces I/O-sensitive workload response time by half without requiring application changes[5].

[4] z/VSE 6.2 will not support the IBM z16 A02 (machine type 3932).

Note: The zHyperLink channels complement FICON channels, but they do *not* replace FICON channels. FICON remains the main data driver and is mandatory for zHyperLink usage.

For more information about the supported FICON Express and zHyperLink Express features, see 4.2, “Storage connectivity” on page 51.

1.2.2 Network connectivity

IBM z16 is a fully virtualized platform that can support many system images at once. Therefore, network connectivity covers the connections between the platform and external networks with Open Systems Adapter-Express (OSA-Express) and RoCE Express features, and it supports specialized internal connections for intra-system communication through IBM HiperSockets and Internal Shared Memory (ISM).

HiperSockets

HiperSockets is an integrated function of the IBM Z platform that supplies attachments to up to 32 high-speed virtual LANs, with minimal system and network impact.

HiperSockets is a function of the Licensed Internal Code (LIC). It provides LAN connectivity across multiple system images on the same IBM Z platform by performing memory-to-memory data transfers in a secure way.

The HiperSockets function eliminates the usage of I/O subsystem operations. It also eliminates the requirement to traverse an external network connection to communicate between LPARs in the same IBM Z platform. In this way, HiperSockets can help with server consolidation by connecting virtual servers and simplifying the enterprise network.

OSA-Express

The OSA-Express features provide local area network (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 PU, which provides performance benefits by offloading processing from the operating system.

OSA-Express7S 1.2 features continue to support copper and fiber optic (single-mode and multimode) environments.

RoCE Express

The RoCE Express3 features can provide local area network (LAN) connectivity for Linux on IBM Z and comply with IEEE standards. In addition, RoCE Express features assume several functions of the TCP/IP stack that normally are performed by the PU, which provides performance benefits by offloading processing from the operating system.

The 25 GbE and 10 GbE RoCE Express3 features[6] use Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) to provide fast memory-to-memory communications between two IBM Z servers.

These features help reduce consumption of CPU resources for applications that use the TCP/IP stack (such as IBM WebSphere®, which accesses an IBM Db2 database). These

[5] The performance results can vary depending on the workload. For zHyperLink planning, use the zBNA tool, found at https://www.ibm.com/support/pages/ibm-z-batch-network-analyzer-zbna-tool-0.

[6] RoCE Express features can be used as general-purpose IP interfaces with Linux on IBM Z.

features can help reduce network latency with memory-to-memory transfers by using Shared Memory Communications over RDMA (SMC-R).

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

Internal Shared Memory

ISM is a virtual Peripheral Component Express (PCI) network adapter that enables direct access to shared virtual memory. It provides a highly optimized network interconnect for IBM Z platform intra-communications. Shared Memory Communications-Direct Memory Access (SMC-D) uses ISM.

SMC-D optimizes operating systems communications in a way that is transparent to socket applications. It also reduces the CPU cost of TCP/IP processing in the data path, which enables highly efficient and application-transparent communications.

SMC-D requires no extra physical resources (such as RoCE Express features, PCIe bandwidth, ports, I/O slots, network resources, or Ethernet switches). Instead, SMC-D uses LPAR-to-LPAR communication through HiperSockets or an OSA-Express feature for establishing the initial connection.

z/OS and Linux on IBM Z support SMC-R and SMC-D. Now, data can be shared through memory-to-memory transfer between z/OS and Linux on IBM Z.

For more information about the available network connectivity features, see 4.3, “Network connectivity” on page 53.

1.2.3 Cryptography

IBM z16 provides two main cryptographic functions: CP Assist for Cryptographic Functions (CPACF), and Crypto-Express8S.

CPACF

CPACF is a high-performance, low-latency coprocessor that performs symmetric key encryption operations and calculates message digests (hashes) in hardware. The following algorithms are supported:

- Advanced Encryption Standard (AES)
- Data Encryption Standard (DES) and Triple Data Encryption Standard (TDES)
- Secure Hash Algorithm (SHA)-1
- SHA-2
- SHA-3

CPACF supports Elliptic Curve Cryptography (ECC) clear key, improving the performance of Elliptic Curve algorithms. The following algorithms are supported:

- EdDSA (Ed448 and Ed25519)
- Elliptic Curve Digital Signature Algorithm (ECDSA) (P-256, P-384, and P-521)
- Elliptic-curve Diffie-Hellman (ECDH) (P-256, P-384, P521, X25519, and X448)
- Support for protected key signature creation

Crypto-Express8S

The tamper-sensing and tamper-responding Crypto-Express8S features provide acceleration for high-performance cryptographic operations and support up to 85 domains with

IBM z16 A01 and 40 domains with IBM z16 A02 and IBM z16 AGZ. This specialized hardware performs AES, DES and TDES, RSA, Elliptic Curve (ECC), SHA-1, and SHA-2, and other cryptographic operations.

It also supports specialized high-level cryptographic APIs and functions, including those functions that are required with quantum-safe cryptography and in the banking industry. Crypto-Express8S features are designed to meet the Federal Information Processing Standards (FIPS) 140-2 Level 4 and Payment Card Industry PTS HSM (PCI-HSM) security requirements for hardware security modules (HSMs).

For more information about cryptographic features and functions, see 4.6, “Cryptographic features” on page 59 and 5.7, “Quantum-safe technology” on page 97.

1.2.4 Clustering connectivity

Parallel Sysplex is an IBM Z clustering technology that is used to make applications that are running on logical and physical IBM Z servers highly reliable and available. IBM Z servers in a Parallel Sysplex are interconnected by way of coupling links.

Coupling connectivity on IBM z16 uses Coupling Express2 Long Reach (CE2 LR) and ICA SR features. The ICA SR feature supports distances up to 150 meters (492 feet); the CE2 LR feature supports unrepeated distances of up to 10 km (6.21 miles) between IBM Z servers.[7]

For more information about coupling and clustering features, see 4.4, “Clustering connectivity” on page 56.

1.2.5 Special-purpose features and functions

IBM takes a *total systems* view regarding the design and development the IBM Z platform. The IBM Z stack is built around digital services, agile application development, connectivity, and system management. This design approach creates an integrated, diverse platform with specialized hardware and dedicated computing capabilities.

IBM z16 delivers a range of features and functions so that PUs can concentrate on computational tasks while distinct, specialized features take care of the rest. For more information about these features and other IBM z16 features, see Chapter 5, “IBM z16 system design strengths” on page 65.

1.2.6 Capacity on Demand and performance

IBM z16 enables just-in-time deployment of processor resources. The Capacity on Demand (CoD) function can dynamically change available system capacity. This function can help respond to new business requirements with flexibility and precise granularity.

The IBM Tailored Fit Pricing for IBM Z options are designed to deliver unmatched simplicity and predictability of hardware capacity and software pricing, even in the constantly evolving era of hybrid cloud.

IBM Z servers help to make embracing hybrid cloud easier with Tailored Fit Pricing for IBM Z. The pricing option delivers simplicity, flexibility, and predictability of pricing across the stack, even with constantly increasing unpredictability in business demand.

[7] The distance between systems can be increased to 100 km by using dense wavelength-division multiplexing (DWDM).

Also contributing to the extra capacity on IBM z16 are numerous improvements in processor chip design, including new instructions, multithreading, and redesigned and larger caches.

In its maximum configuration, IBM z16 A01 Max200 can deliver up to 17%[8] more capacity than IBM z15 T01 Max190. An IBM z16 A01 1-way server has approximately 11% more capacity than a IBM z15 T01 1-way server.

Within each single drawer, IBM z16 A01 provides 25% greater capacity than IBM z15 T01 for standard models and 40% greater capacity on the max config model, enabling efficient scaling of partitions. IBM z16 A02 and IBM z16 AGZ provide around 13%[9] capacity growth over IBM z15 T02 on the max config model, and around 29% greater than IBM z14 ZR1.

For more information, see 5.3, “Capacity and performance” on page 78.

1.2.7 Reliability, availability, and serviceability

IBM z16 offers the same high quality of service and reliability, availability, and serviceability (RAS) that is traditional in the IBM Z platform. The RAS strategy uses a building-block approach that meets the stringent requirements for achieving continuous, reliable operation. The following RAS building blocks are available:

- 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 functions to eliminate or minimize all sources of outages.

SRB with the optional System Recovery Boost Upgrade (temporary capacity upgrade) is a function to accelerate operating system and services start and shutdown times. IBM z16 also provides boosted processor capacity and parallelism for specific events, such as middleware starts and restarts, SVC dump processing, and HyperSwap configuration load and reload to minimize the impact on running workloads.

For more information about RAS, see 5.4, “Reliability, availability, and serviceability” on page 86.

1.3 IBM z16 software support

IBM z16 supports a wide range of IBM and ISV software solutions. This range includes traditional batch and online transaction processing (OLTP) environments, such as IBM Customer Information Control System (IBM CICS®), IBM Information Management System (IMS), and IBM Db2. It also includes the following web services (among others):

- Java platform
- Linux and open standards applications
- WebSphere

[8] Variations on all the observed increased performances depend on the configuration and workload type.

[9] The IBM z16 A02 and IBM z16 AGZ performance varies between capacity setting levels. The 13% value applies to the largest (z06) models.

- IBM z/OS Connect Enterprise Edition

The following operating systems are supported on IBM z16:

- z/OS Version 2 Release 5 with program temporary fixes (PTFs)
- z/OS Version 2 Release 4 with PTFs
- z/OS Version 2 Release 3 with PTFs
- z/OS Version 2 Release 2 with PTFs (toleration support only)
- z/VM Version 7 Release 3
- z/VM Version 7 Release 2 with PTFs
- z/VM Version 7 Release 1 with PTFs
- z/VSE[10] Version 6 Release 2 with PTFs
- 21st Century Software VSE^n V6.3
- z/TPF Version 1 Release 1 (compatibility support)

IBM supports 21st Century Software VSE^n V6.3 on IBM z16. For more information, see this web page.

IBM plans to support the following Linux on IBM Z distributions[11] on IBM z16:

- SUSE Linux Enterprise Server 15 SP3 and SUSE Linux Enterprise Server 12 SP5
- Red Hat RHEL 8.4 and Red Hat RHEL 7.9
- Ubuntu 22.04 LTS and Ubuntu 20.04.1 LTS

The support statements for IBM z16 also cover the KVM hypervisor on distribution levels that have KVM support.

For more information about the IBM z16 software support, see Chapter 6, “Operating system support” on page 99.

1.3.1 IBM compilers

Compilers are built with specific knowledge of the system architecture, which is used during code generation. Therefore, the usage of the latest compilers is essential to extract the maximum benefit of a platform’s capabilities. IBM compilers use the latest architecture enhancements and new instruction sets to deliver more 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 IT investments while smoothly incorporating new, web-based applications into the infrastructure.

z/OS, XL C/C++, and XL C/C++ for Linux on IBM Z 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 uses the z/Architecture. This transformation is possible because of hardware-tailored optimizations, built-in functions, performance-tuned libraries, and language constructs that simplify system programming and boost application runtime performance.

Compilers, such as COBOL, PL/I, and z/OS XL C/C++, are inherently optimized on IBM z16 because they use floating point registers rather than memory or fast mathematical computations. The usage of compilers that leverage hardware enhancements is key to improving application performance, reducing CPU usage, and lowering operating costs.

For more information, see 6.1.2, “Application development and languages” on page 101 and 6.1.3, “Supported IBM compilers” on page 103.

[10] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.
[11] All require extra service. For more information, see 6.2.6, “Linux on IBM Z” on page 108.

2

IBM z16 A01 hardware overview

This chapter expands on the descriptions of the key hardware elements of IBM z16 A01 that were introduced in 1.2, “IBM z16 technical overview” on page 7.

This chapter describes the following topics:

2.1 IBM z16 A01 upgrade paths

IBM z16 A01 is built by using the IBM Telum processor design. Each processor chip consists of eight cores. Two processor chips are packaged in the dual-chip module (DCM). Each DCM can have 9 - 11 or 10 - 15 active processor unit (PU) cores (depending on the configuration). With IBM z16 A01, a maximum number of characterizable processors are represented by the feature names Max39, Max82, Max125, Max168, and Max200. Spare PUs, System Assist Processors (SAPs), and two Integrated Firmware Processors (IFPs) are included in the IBM z16 A01 configuration.

The number of characterizable PUs, SAPs, and spare PUs for the various features is listed in Table 2-1. For more information about PU characterization types, see “PU characterization” on page 25.

Table 2-1 IBM z16 A01 processor unit configurations

Feature name	Number of CPC drawers	Feature Code	Characterizable processor units	Standard SAPs	Spares
Max39	1	0667	0 - 39	5	2
Max82	2	0668	0 - 82	10	2
Max125	3	0669	0 - 125	15	2
Max168	4	0670	0 - 168	20	2
Max200	4	0671	0 - 200	24	2

IBM z16 A01 ensures continuity and upgrades from IBM z15 T01and IBM z14 M0x. The supported upgrade paths are shown in Figure 2-1.

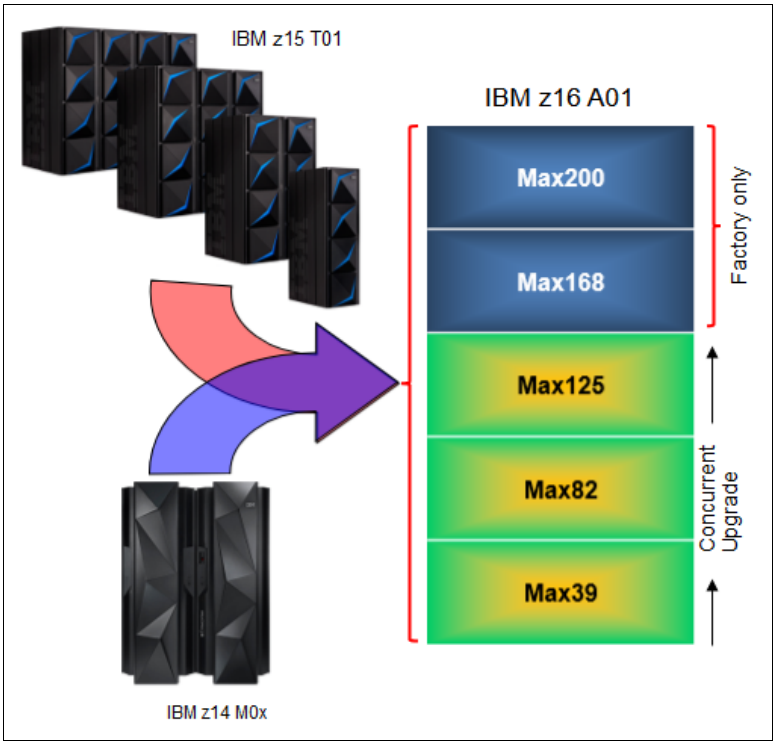

Figure 2-1 IBM z16 A01 upgrade paths

If an upgrade request cannot be fulfilled by using the existing configuration, a hardware upgrade is required in which one or more CPC drawers are added to accommodate the needed capacity. With IBM z16 A01, more CPC drawers can be installed concurrently from a Max39 to a Max82, and to a Max125.

Note: No field upgrade is available to a Max 168 or a Max 200; these two features are factory-shipped only.

With IBM z16 A01, concurrent upgrades are available for central processors (CPs), Integrated Facilities for Linux (IFLs), Integrated Coupling Facilities (ICFs), IBM Z Integrated Information Processors (zIIPs), and SAPs. However, concurrent PU upgrades require that more PUs are physically installed, but not activated previously.

In the rare event of a PU failure, one of the spare PUs is immediately and transparently activated and assigned the characteristics of the failing PU. Two spare PUs are always available on an IBM z16 A01.

IBM z16 A01 offers 317 capacity levels. In all, 200 capacity levels are based on the number of physically used CPs, plus up to 117 subcapacity models for the first 39 CPs.

For more information, see 5.3.1, “Capacity settings” on page 78.

2.2 Frames and cabling

IBM z16 A01 uses 19-inch frames and industry-standardized power and hardware. It can be configured as a one-, two-, three-, or four-frame system. Each frame takes up only two standard 24-inch floor tiles of space, which aligns with modern data center layouts.

The IBM z16 A01 configuration options as compared to previous IBM Z servers are listed in Table 2-2.

Table 2-2 IBM z16 A01 configuration options compared to IBM z15 T01 and IBM z14 M0x configurations

System	Number of frames	Number of CPC drawers	Number of I/O drawers	I/O and power connections	Power options[a]	Cooling options
IBM z16 A01	1 - 4	1 - 4	0 - 12[b]	Rear only	PDU or BPA	Radiator (air) only
IBM z15 T01	1 - 4	1 - 5	0 - 12[c]	Rear only	PDU or BPA	Radiator (air) or water-cooling unit (WCU)
IBM z14 M0x	2	1 - 4	0 - 5	Front and rear	BPA	Radiator (air) or WCU

a. The power distribution unit (PDU) option supports the air-cooling (radiator) option. The Bulk Power Assembly (BPA) option supports both air-cooling and water-cooling options.
b. Maximum of 12 if ordered with a PDU or maximum of 10 if ordered with a BPA.
c. Maximum of 12 if ordered with a PDU or maximum of 11 if ordered with a BPA.

The number of Peripheral Component Interconnect Express+ (PCIe+) I/O drawers can vary based on the number of I/O features, power options (PDU or BPA), and number of CPC drawers that is installed. For a PDU system, a maximum configuration of up to 12 PCIe+ I/O drawers can be installed. PCIe+ I/O drawers can be added concurrently.

In addition, IBM z16 A01 supports top-exit options for the fiber-optic and copper cables that are used for I/O and power. These options give you more flexibility in planning where the system is installed, eliminate the need for cables to be run under a raised floor, and increase air flow over the system.

IBM z16 A01 supports installation on raised floor and non-raised floor environments.

Figure 2-2 shows the front view of a fully configured IBM z16 A01 with radiator cooling, four CPC drawers, and 12 PCIe+ I/O drawers.

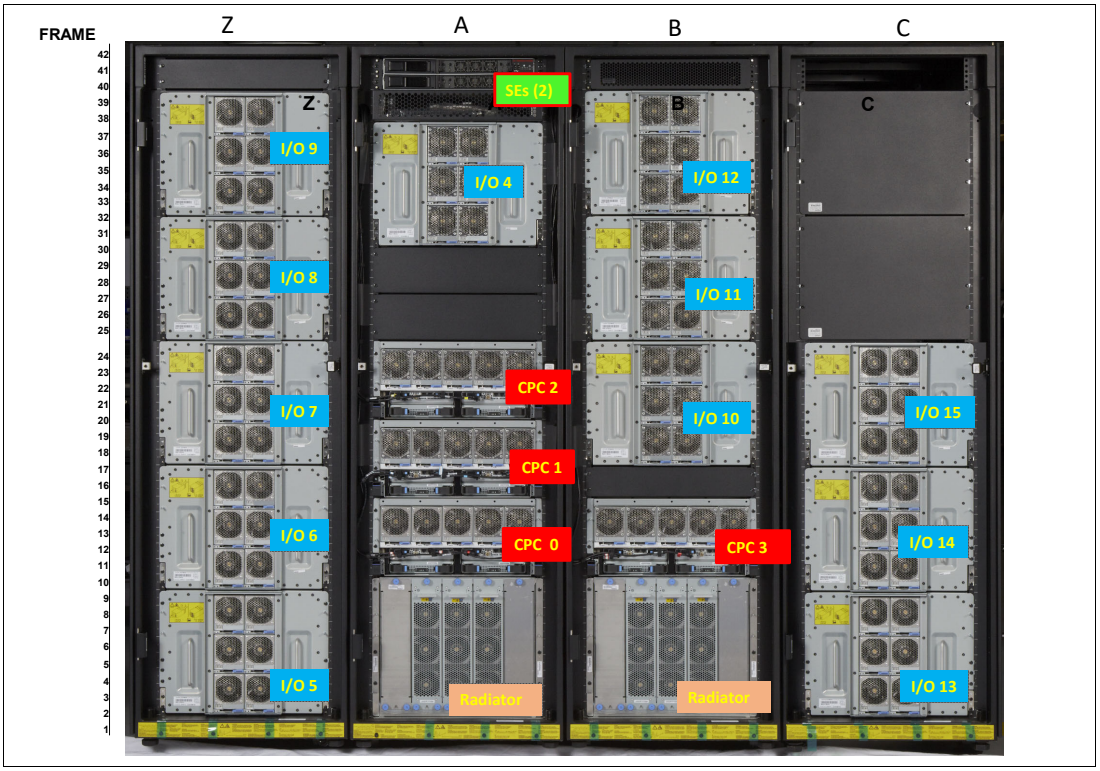

Figure 2-2 Front view of a fully configured IBM z16 A01 with radiator cooling

Figure 2-3 shows the rear view of a fully configured PDU-based IBM z16 A01 with a total of 16 drawers (I/O and CPC combined), and two radiator cooling units.

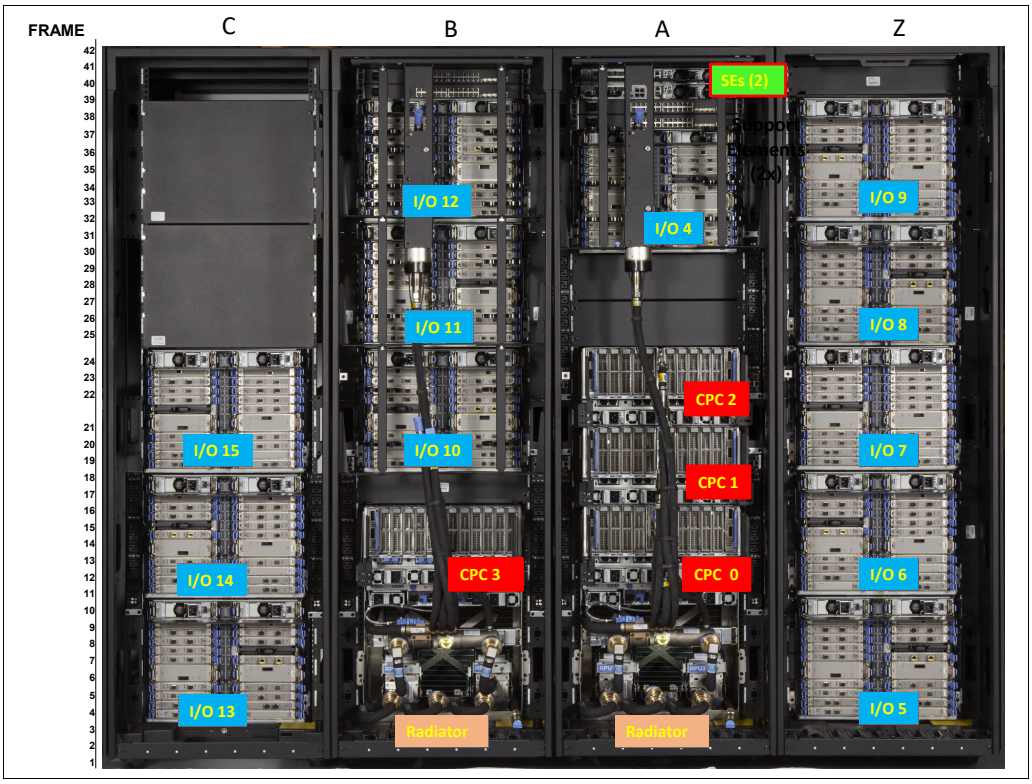

Figure 2-3 Rear view of a fully configured IBM z16 A01

The IBM configurator that is used during the order process calculates the number of frames that is required and placement of CPC and PCIe+ I/O drawers.

Factors that determine the number of frames for an IBM z16 A01 configuration include the following features:

- Number of CPC drawers
- Plan-ahead features for more CPC drawers
- Number of I/O features (determines the number of PCIe+ I/O drawers)
- PDU or BPA power

2.3 CPC drawers

IBM z16 A01 can be configured with up to four CPC drawers (three in the A Frame and one in the B Frame). Each CPC drawer contains the following elements:

- DCMs

  Four DCMs that containing eight central processor (CP) chips and 64 physical cores per drawer interconnected (each cooled by an internal water loop).

- Memory:
  - A minimum of 512 GB and a maximum of 40 TB of memory per system (excluding 256 GB for the hardware system area (HSA)) is available for use. For more information, see Table 2-3 on page 26.
  - Up to 48 dual inline memory modules (DIMMs) that are 32, 64, 128, 256, or 512 GB, which are plugged into a CPC drawer.

- Fanouts

  Each CPC drawer supports up to 12 PCIe+ fanout adapters to connect to the PCIe+ I/O drawers, and Integrated Coupling Adapter Short Reach (ICA SR) coupling links:
  - A 2-port Peripheral Component Interconnect Express (PCIe) 16 GBps I/O fanout. Each port supports one domain in the 16-slot PCIe+ I/O drawers.
  - ICA SR1.1 and ICA SR PCIe fanouts for coupling links (two links of 8 GBps each).

- Three or four Power Supply Units (PSUs), depending on the configuration (PDU or BPA), which provide power to the CPC drawer and are accessible from the rear. Loss of one PSU leaves enough power to satisfy the power requirements of the entire drawer. The PSUs can be concurrently maintained.
- Two dual-function Base Management Cards (BMCs) or Oscillator Cards (OSCs), which provide redundant interfaces to the internal management network and provide clock synchronization to the IBM Z platform.
- Two dual-function Processor Power Cards (PPCs), which control Voltage Regulation, PSUs, and Fan control. The PPCs are redundant and can be concurrently maintained.
- Five fans are installed at the front of the drawer to provide cooling airflow for the resources that are installed in the drawer (except for the PU SCMs, which are internally water-cooled).

The CPC drawer communication topology is shown in Figure 2-4 on page 23. All CPC drawers are interconnected with high-speed communications links (A-Bus) through the PU chips. Symmetric multiprocessor (SMP-9) cables are used to interconnect all the CPC drawers. The X-Bus provides connectivity between DCMs on the drawer, while the M-Bus connects the two PU chips on each DCM.

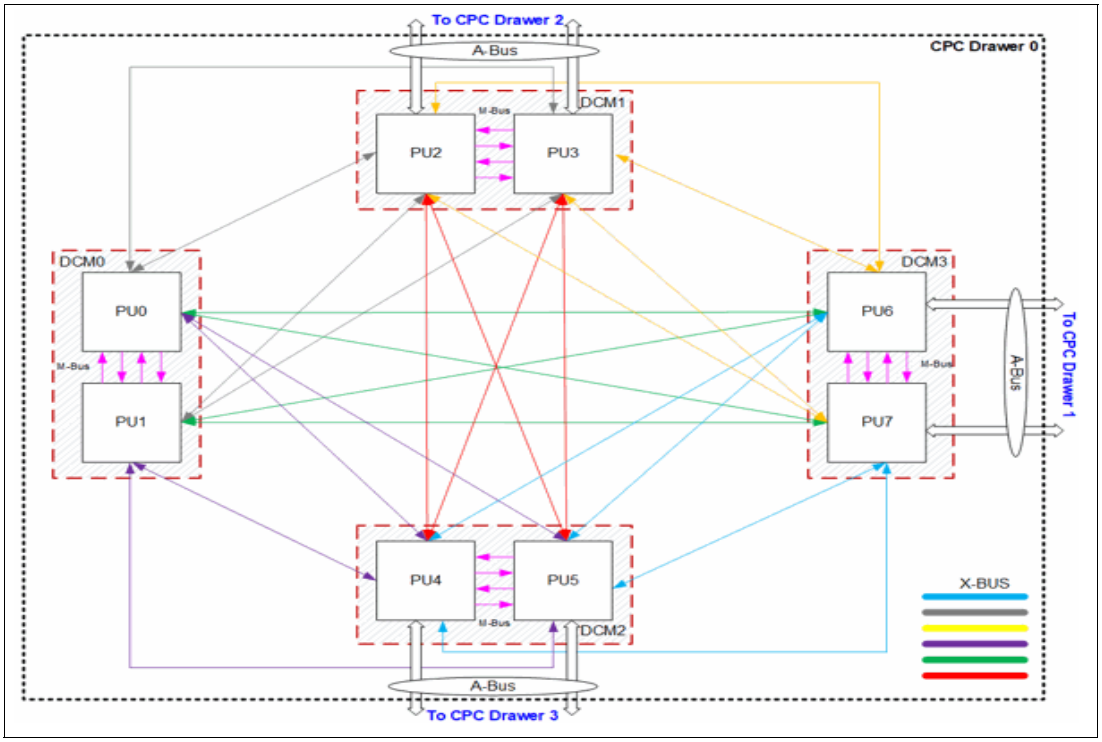

Figure 2-4 IBM z16 A01 CPC drawer communication topology

The design that is used to connect the PU 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 SMP system.

2.3.1 Dual-chip modules

The CPC drawer always includes four DCMs. Each DCM contains two PU chips. Each PU chip contains eight cores, each with a 128 KB Instruction and Data L1 cache and 32 MB semi-private L2 caches.

2.3.2 Processor unit

PU is the generic term for an IBM z/Architecture processor. Each PU is a superscalar processor with the following attributes:

- Up to six instructions can be decoded per clock cycle.
- Up to 10 instructions can run per clock cycle.
- Instructions can be issued out of order. The PU uses a high-frequency, low-latency pipeline that provides 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 a varying number of steps for different types of instructions. Several instructions can be running at any moment, and they are subject to the maximum number of decodes and completions per cycle.

Processor cache structure

The on-chip cache for the PU (core) features the following design:

- Each PU core has an L1 cache (private) that is divided into a 128 KB cache for instructions and a 128 KB cache for data.
- Each PU core has a semi-private L2 cache, which is implemented as 32 MB, near the core.

Note: L1 and L2 are physical cache and implemented in dense SRAM.

- Each PU core contains a 256 MB shared-victim virtual L3 cache. The shared-victim virtual L3 cache is a logical construction that comprises all eight semi-private L2s (8 x 32 MB = 256 MB) belonging to the other cores.
- Each CPC drawer contains a 2 GB shared-victim virtual L4, consisting of the “remote” virtual L3 caches of the DCMs in the CPC drawer.

This on-chip cache implementation optimizes system performance for high-frequency processors and includes the following features:

- Cache improvements
- New translation and TLB2 design
- Pipeline optimizations
- Better branch prediction
- New accelerators and architectures
- Secure Execution support

The IBM z16 A01 cache structure is shown Figure 2-5.

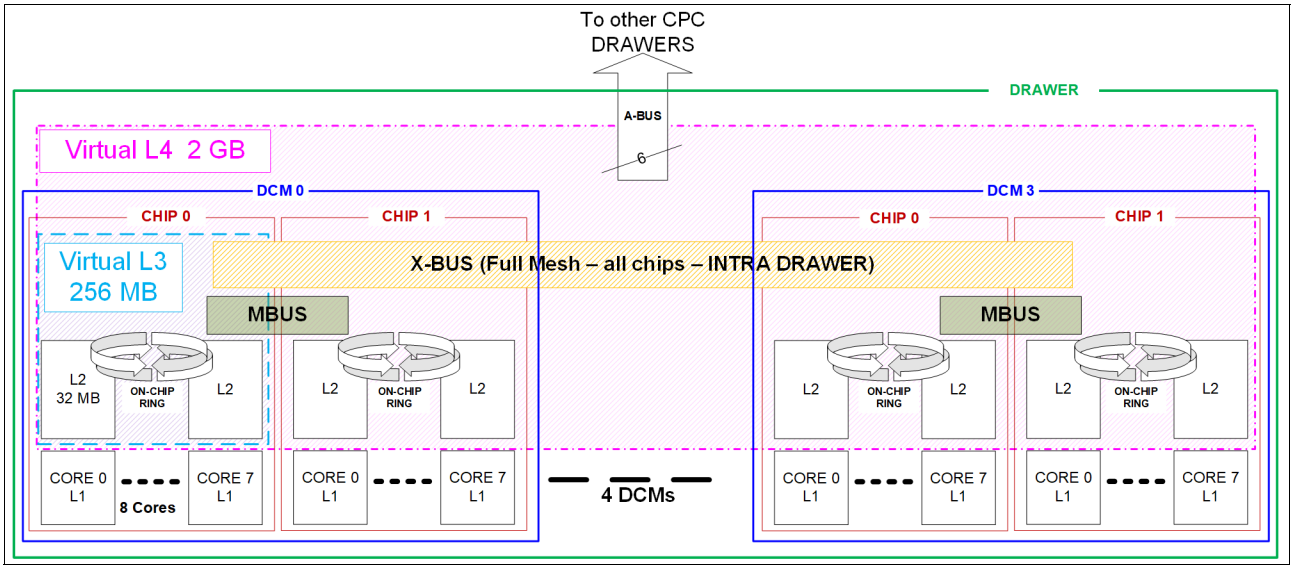

Figure 2-5 IBM z16 A01 cache structure

PU sparing

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

On-core cryptographic hardware

Dedicated on-chip cryptographic hardware for each PU core includes extended key and hash sizes for the Advanced Encryption Standard (AES) and Secure Hash Algorithm (SHA). For more information, see 4.6, “Cryptographic features” on page 59. This cryptographic hardware is available with any processor type, for example, CP, zIIP, and IFL.

On-chip functions

Consider the following points:

- IBM Integrated Accelerator for zEnterprise® Data Compression replaces the IBM zEnterprise Data Compression (zEDC) Express PCIe feature that was on previous IBM Z servers.
- The sort accelerator uses the sort instruction (**SORTL**) instruction to be used by **DFSORT** and the IBM Db2 Utilities for z/OS Suite to help reduce CPU usage and improve elapsed time for sort workloads.
- Integrated Accelerator for Artificial Intelligence Unit (AIU) is implemented on each PU chip and shared among all cores. It provides a matrix array for multiplication and convolution alongside the specialty engines for complex functions.

  The AIU provides a Neural Network Processing Assist (NNPA) instruction, which operates directly on Tensor data in user space.

Software support

The IBM z16 A01 PUs provide full compatibility with software for z/Architecture, and extend the Instruction Set Architecture (ISA) to enable enhanced functions and performance. New with IBM z16 A01 are instructions for the AIU.

PU characterization

PUs are ordered in single increments. The internal system functions are based on the configuration that is ordered. They characterize each PU into one of various types during system initialization, which is often called a power-on reset (POR) operation.

Characterizing PUs dynamically without a POR is possible by using a process that is called *Dynamic Processor Unit Reassignment*. A PU that is not characterized cannot be used. Each PU can be designated with one of the following characterizations:

- CP: These standard processors are used for general workloads.
- IFL: Designates processors to be used specifically for running the Linux application programs.
- Unassigned Integrated Facilities for Linux (UIFL): Allows you to directly purchase an IFL feature that is marked as being deactivated upon installation, which avoids software charges until the IFL is brought online for use.
- zIIP: An “Off Load Processor” that is used under z/OS for designated types of workloads. For a list of zIIP use candidates, see “Logical processors” on page 73. zIIP also is used for the IBM System Recovery Boost (SRB) feature. For more information, see the System Recovery Boost topic under 5.4, “Reliability, availability, and serviceability” on page 86.
- Unassigned zIIP: A processor that is purchased for future use as a zIIP. It is offline and cannot be used until an upgrade for the zIIP is installed. It does not affect software licenses or maintenance charges.
- Integrated Coupling Facility (ICF): Designates processors to be used specifically for coupling.

- Unassigned Coupling Facility: Allows you to directly purchase an ICF feature that is marked as being deactivated upon installation, which avoids software charges until the ICF is brought online for use.
- SAP: Designates processors to be used specifically for assisting I/O operations.
- IFP: The IFP is standard and not defined by the customer (it is used for infrastructure management).

At least one CP must be purchased before zIIPs can be purchased. The maximum of zIIPs is one less than the allowed maximum PU configuration. For example, IBM z16 A01 Max200 can have up to 199 zIIPs. These rules are also valid for unassigned zIIPs and unassigned IFLs.

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

Note: The addition of ICFs, IFLs, zIIPs, and SAPs does not change the system capacity setting or its millions of service units (MSU) rating.

2.3.3 Memory

The maximum physical memory size is directly related to the number of CPC drawers in the system. An IBM Z server includes more installed memory than was ordered because part of the installed memory is used to implement the redundant array of independent memory (RAIM) design. With IBM z16 A01, up to 10 TB of memory per CPC drawer can be ordered, and up to 40 TB for a four-CPC drawer system.

Important: z/OS requires a minimum of 8 GB of memory (2 GB of memory when running under z/VM). z/OS V2R5 can support up to 16 TB of memory in an LPAR.

The minimum and maximum memory sizes for each IBM z16 A01 feature are listed in Table 2-3.

Table 2-3 IBM z16 A01 memory per feature

Feature name	CPC drawers	Memory[a]
Max39 (Feature Code 0667	1	512 GB - 10 TB
Max82 (Feature Code 0668)	2	512 GB - 20 TB
Max125 (Feature Code 0669)	3	512 GB - 30 TB
Max168 (Feature Code 0670)	4	512 GB - 40 TB
Max200 (Feature Code 0671)	4	512 GB - 40 TB

a. There are maximum memory configurations restrictions when choosing bulk power (BPA).

The HSA on IBM z16 A01 has a fixed amount of memory (256 GB) that is managed separately from available memory. However, the maximum amount of orderable memory can vary from the theoretical number because of dependencies on the memory granularity. On IBM z16 A01, the granularity for memory is in 64, 128, 256, 512, 1024, and 2048 GB increments.

Physically, memory is organized in the following ways:

- A CPC drawer always contains a minimum of 1024 GB to a maximum of 10 TB of installed memory, of which 10 TB maximum 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 (LIC) load.
- Memory upgrades are first satisfied by using installed but unused memory capacity until it is exhausted. When no more unused memory is available from the installed cards, the cards must be upgraded to a higher capacity, or a CPC drawer with more memory must be installed.

When an LPAR is activated, PR/SM attempts to allocate PUs and the memory of an LPAR in a single CPC drawer. However, if this allocation is not possible, PR/SM uses memory resources in any CPC drawer. For example, if the allocated PUs span more than one CPC drawer, PR/SM attempts to allocate memory across that same set of CPC drawers (even if all required memory is available in only one of those CPC drawers).

No matter which CPC drawer the memory is installed in, an LPAR can access that memory after it is allocated. IBM z16 A01 is an SMP system because the PUs can access all of 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 model upgrades that involve the addition of a CPC drawer, the minimum usable memory increment (512 GB) is added to the system. During an upgrade, adding a CPC drawer and physical memory in the new drawer are concurrent operations.

Concurrent memory upgrade

When unused capacity is available on the installed memory cards, the total system memory can be upgraded concurrently through Licensed Internal Code Configuration Control (LICCC). The LICCC provides for system upgrades without hardware changes by activating extra (physically installed) unused capacity.

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 DIMMs.

The RAIM design is fully integrated in IBM z16 A01 and enhanced to include one Memory Controller Unit (MCU) per processor chip, with eight memory channels and one DIMM per channel. The MCU enables memory to be implemented as RAIM. This technology has significant reliability, availability, and serviceability (RAS) capabilities 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.

For more information about memory design and configuration options, see *IBM z16 (3931) Technical Guide*, SG24-8951.

2.3.4 Hardware system area

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

The fixed size 256 GB HSA of IBM z16 A01 is large enough to accommodate any LPAR definitions or changes, which eliminate most outage situations and the need for extensive planning.

A fixed, large HSA allows the dynamic I/O capability of IBM z16 A01 to be enabled by default. It also enables the dynamic addition and removal of the following features:

- An LPAR to a new or existing CSS
- CSSs (up to six can be defined)
- Subchannel sets (up to four can be defined)
- Devices, up to the maximum number permitted, in each subchannel set
- Logical processors by type
- Cryptographic adapters

2.4 I/O system structure

IBM z16 A01 supports the PCIe-based infrastructure for the PCIe+ I/O drawers. The PCIe I/O infrastructure consists of the dual-port PCIe fanouts in the CPC drawers that support 16 GBps connectivity to the PCIe+ I/O drawer.

Ordering of I/O features: Ordering I/O feature types determines the suitable number of PCIe+ I/O drawers.

Figure 2-6 on page 29 shows a high-level view of the I/O system structure for IBM z16 A01.

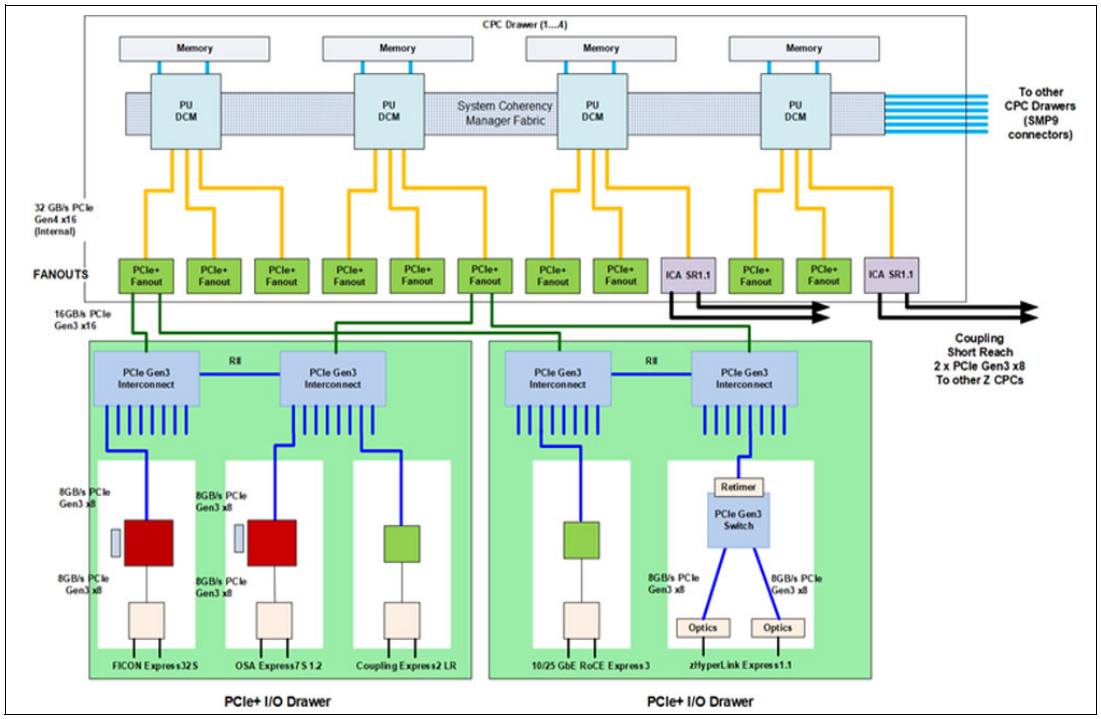

Figure 2-6 IBM z16 A01 I/O system structure

The IBM z16 A01 CPC drawer has 12 fanouts (numbered LG01 - LG12). The fanouts that are installed in these positions can be one of the following types:

- Dual-port PCIe+ fanouts for PCIe+ I/O drawer connectivity
- ICA SR fanouts for coupling
- Filler plates to assist with airflow cooling

For coupling link connectivity (IBM Parallel Sysplex and Server Time Protocol (STP) configuration), IBM z16 A01 supports the following link types:

- ICA SR 1.1 and ICA SR (installed in a CPC drawer)
- Coupling Express2 Long Reach (CE LR) (installed in a PCIe+ I/O drawer)

For systems with multiple CPC drawers, the locations of the PCIe+ fanouts are configured and plugged across all drawers for maximum availability. This configuration helps ensure that alternative paths maintain access to critical I/O devices, such as storage and networks (see Figure 2-7).

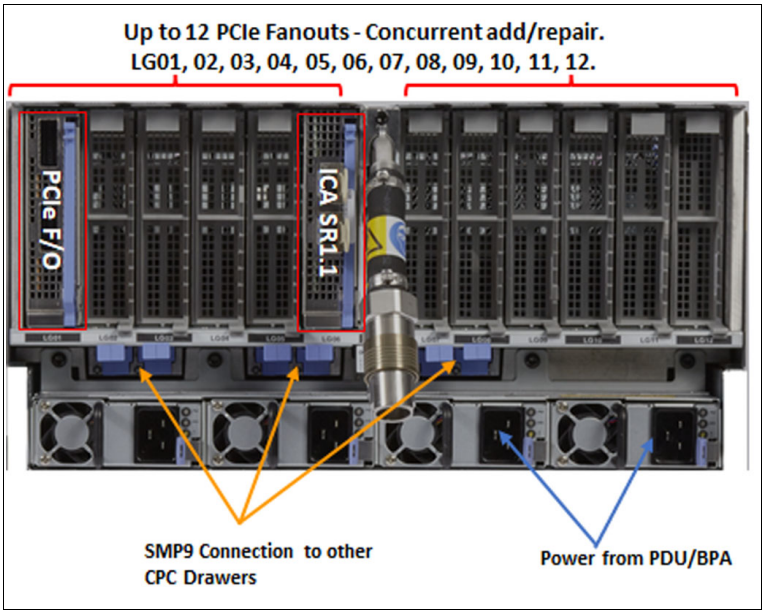

Figure 2-7 IBM z16 A01 CPC drawer: Rear view

The PCIe+ I/O drawer (see Figure 2-8), is a 19-inch single side drawer that is 8U high. I/O features are installed horizontally, with cooling air flow from front to rear. The drawer contains 16 adapter slots and 2 slots for PCIe switch cards.

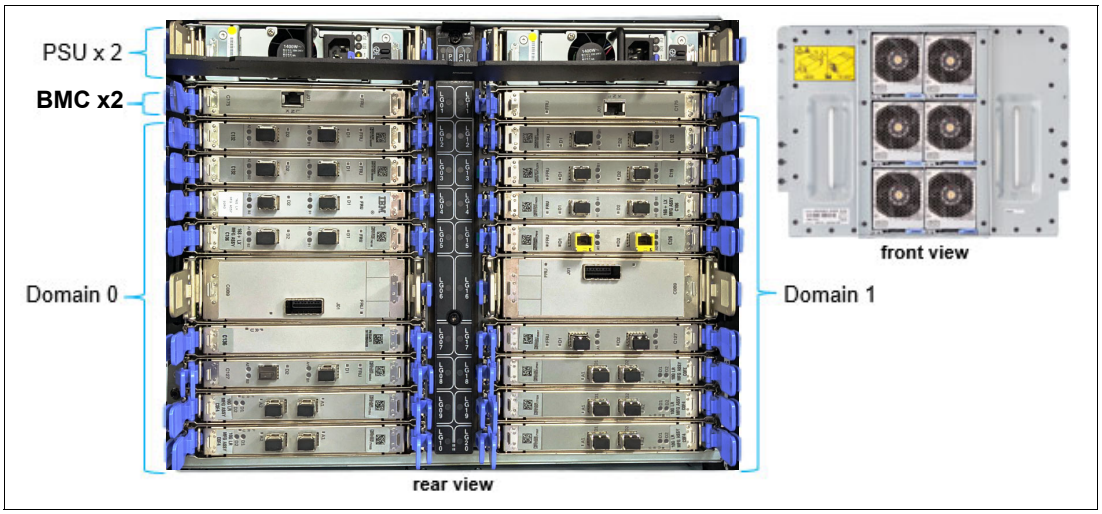

Figure 2-8 PCIe+ I/O drawer: Rear and front view

The two I/O domains per drawer each contain up to eight I/O features that support the following types:

- FICON Express32S, FICON Express16SA, or FICON Express16S+
- OSA-Express7S 1.2, OSA-Express7S, or OSA-Express6S
- Crypto-Express8S, Crypto-Express7S, or Crypto-Express6S
- RDMA over Converged Ethernet (RoCE) Express3, RoCE Express2.1, or RoCE Express2

- zHyperLink Express 1.1 and zHyperLink Express
- Coupling Express2 LR

For more information about the I/O features that are available with IBM z16 A01, see Chapter 4, “Supported features and functions” on page 49.

2.5 Power and cooling

IBM z16 A01 meets the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Class A3 specifications. ASHRAE is an organization that is devoted to the advancement of indoor-environment-control technology in the heating, ventilation, and air-conditioning industry.

2.5.1 Power options

The IBM z16 A01 19-inch frames are available with the following power options:

- PDU

  Using PDU for IBM z16 A01 can enable fewer frames, which allows for extra I/O slots and improves power efficiency to lower overall energy costs. It offers some standardization and ease of data center installation planning. PDU supports up to 12 PCIe+ I/O drawers.

- Bulk Power Assembly (BPA)

  The BPA supports up to 10 PCIe+ I/O drawers. This option is required when ordered with Balanced Power.

BPA support removal[a]: Based on the direction of the market, IBM z16 A01 is the last IBM Z platform to support BPA. Plan to migrate from BPA to PDU.

a. Statements by IBM regarding its plans, directions, and intent are subject to change or withdrawal without notice at the sole discretion of IBM. Information regarding potential future products is intended to outline general product direction and should not be relied on in making a purchasing decision.

The IBM z16 A01 operates with one or 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 that are installed. Power cords attach to a three-phase, 50/60 Hz, 200 - 480 V AC power source. The loss of one power supply per set has no effect on system operation.

The optional Balanced Power Plan Ahead feature is available for future growth, which also assures 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 BPRs and power cords.

2.5.2 Cooling options

The IBM z16 A01 cooling system is available only with the Radiator (air) cooling option. The previous water-cooling solution that was available with IBM z15 is not offered with IBM z16 A01.

DCMs are always cooled with an internal water loop. The liquid in the internal water system is cooled by using an internal radiator. The radiator, PCIe+ I/O drawers, power enclosures, and CPC drawers are cooled by chilled air with blowers.

The air-cooling system in IBM z16 A01 is redesigned for better availability and lower cooling power consumption. The radiator design is a closed-loop water-cooling pump system for the DCMs in the CPC drawers. It is designed with N+1 pumps, blowers, controls, and sensors. The radiator unit is cooled by air.

2.5.3 Power considerations

Consider the following points about power:

- A total of one to four 42U 19-inch IBM frames are used (replacing the two 24-inch frame).
- Air flow is front to rear. All blowers are mounted on the front of the frame.
- All external power cabling is at the rear of the frames (no power cabling in front).
- Top or bottom exit power is supported.
- A High-Voltage DC (HVDC) option is not available.
- No Emergency Power Off (EPO) switch is used.

Specific power requirements depend on the number of frames, CPC drawers, and type of I/O features that are installed, and the power option (PDU or BPA).

For more information about the maximum power consumption tables for the various configurations and environments, see *IBM 3931 Installation Manual for Physical Planning*, GC28-7015.

For more information about the power and weight estimation tool, see IBM Resource Link.

3

IBM z16 A02 and IBM z16 AGZ hardware overview

As the transformation of the air-cooled mainframe that began with IBM z14 ZR1continues, two new configuration options are offered:

- IBM z16 A02, which includes the core compute, I/O, and networking features installed in a traditional 19-inch IBM factory frame.
- IBM z16 AGZ, which includes the core compute, I/O, and networking features installed in and powered by a client-supplied 19-inch rack and power distribution units (PDUs), respectively.

This chapter expands on the key hardware elements of IBM z16 A02 and IBM z16 AGZ that were introduced in 1.2, “IBM z16 technical overview” on page 7.

This chapter describes the following topics:

3.1 IBM z16 A02 and IBM z16 AGZ upgrade paths

IBM z16 A02 and IBM z16 AGZ are built by using the IBM Telum processor design. Each processor chip consists of eight cores. Two processor chips are packaged in the dual-chip module (DCM). Each DCM can have 9 - 11 or 11 - 13 active processor unit (PU) cores (depending on the configuration). With IBM z16 A02 and IBM z16 AGZ, a maximum number of characterizable processors are represented by feature names Max5, Max16, Max32, and Max68.[1] Spare PUs, System Assist Processors (SAPs), and two Integrated Firmware Processors (IFPs) are included in both configurations.

The number of characterizable PUs, SAPs, and spare PUs for the various features are listed in Table 3-1. For more information about PU characterization types, see “PU characterization” on page 41.

Table 3-1 IBM z16 A02 and IBM z16 AGZ processor unit configurations

Feature name	Number of CPC drawers / PU DCMs	Feature Code	Characterizable processor units	Central processors (CPs)	Standard SAPs	Spares
Max5	1 / 2	0672	1 - 5	0 - 5	2	2
Max16	1 / 2	0673	1 - 16	0 - 6	2	2
Max32	1 / 4	0674	1 - 32	0 - 6	4	2
Max68	2 / 8	0675	0 - 68	0 - 6	8	2

IBM z16 A02 and IBM z16 AGZ ensure continuity and upgradeability from IBM z15 T02 and IBM z14 ZR1. The supported upgrade paths are shown in Figure 3-1 on page 35.

[1] The rack-mount configuration options are under a combined AGZ warranty umbrella.

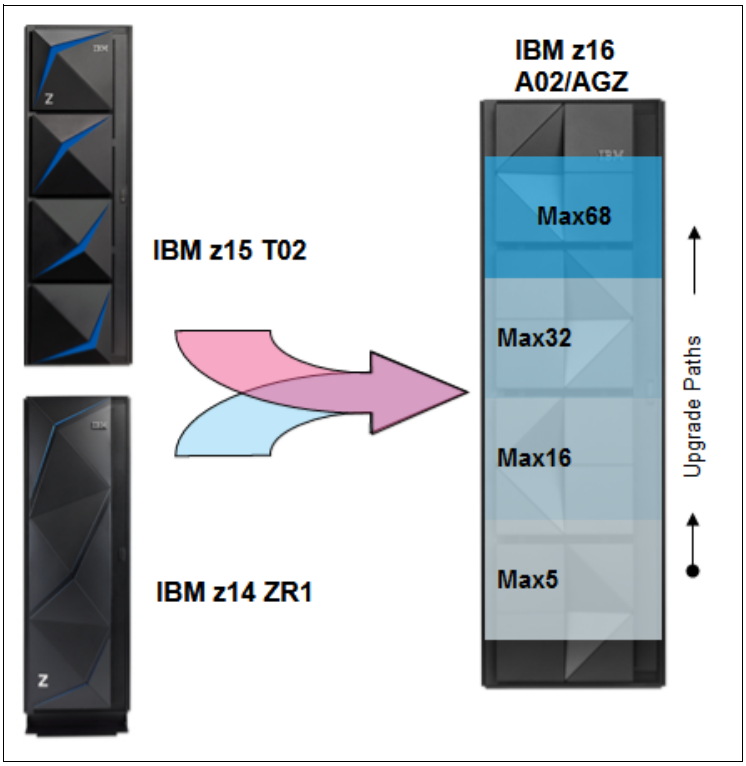

Figure 3-1 IBM z16 A02 and IBM z16 AGZ upgrade paths

If an upgrade request cannot be fulfilled by using the existing configuration, a hardware upgrade is required in which one extra CPC drawer is added to accommodate the needed capacity. With IBM z16 A02, the second CPC drawer can be installed concurrently from a Max32 to a Max68. (The second CPC drawer addition to IBM z16 AGZ requires a proper plan-ahead). Additionally, a single CPC drawer Max5 can be concurrently upgraded to a Max16.

With IBM z16 A02 and IBM z16 AGZ, concurrent upgrades are available for central processors (CPs), Integrated Facilities for Linux (IFLs), Integrated Coupling Facilities (ICFs), IBM Z Integrated Information Processors (zIIPs), and SAPs. However, concurrent PU upgrades require that more PUs are physically installed, but not activated previously.

In the rare event of a PU failure, one of the spare PUs is immediately and transparently activated and assigned the characteristics of the failing PU. Two spare PUs always are available on IBM z16 A02 and IBM z16 AGZ.

IBM z16 A02 and IBM z16 AGZ offer 156 subcapacity levels, A01 - Z06, which use 1 - 6 available CPs with features Max16, Max32, and Max68. For a Max5 feature, the subcapacity models range from A01 to Z05 (5 CPs maximum).

For more information, see 5.3.1, “Capacity settings” on page 78.

3.2 Frame and cabling

IBM z16 A02 is delivered in a 19-inch frame with industry-standardized power and hardware. It is a single frame (rack) system that takes up two standard 24-inch floor tiles of space, which aligns with modern data center layouts. IBM z16 AGZ is delivered as a bundle that is installed in a client-supplied EIA standard 19-inch rack with PDUs. Both systems share features and functions.

The IBM z16 A02 and IBM z16 AGZ configuration options, as compared to previous IBM Z servers, are listed in Table 3-2.

Table 3-2 IBM z16 A02 and IBM z16 AGZ configuration options compared to IBM z15 and IBM z14

System	Number of frames	Number of CPC drawers	Number of I/O drawers	I/O and power connections	Power options[a]	Power cord options
IBM z16 A02 or IBM z16 AGZ	1[b]	1 - 2	0 - 3[c]	Rear only	PDU	Single- or three-phase
IBM z15 T02	1	1 - 2	0 - 4[d]	Rear only	PDU	Single- or three-phase
IBM z14 ZR1	1	1	0 - 4[b]	Rear only	PDU	Single-phase only

a. The PDU option is used with different power cords for different frame configurations and countries.
b. There is no frame for IBM z16 AZG. A client-supplied rack with PDUs is required.
c. Only PCIe+ I/O drawers are available. New build only. No carry forward of any I/O drawer.
d. Maximum of four if ordered with single CPC drawer or maximum of three if ordered with two CPC drawers. Only PCIe+ I/O drawers are available.

The number of Peripheral Component Interconnect Express+ (PCIe+) I/O drawers can vary based on the number of I/O features and number of CPC drawers that are installed. For IBM z16 A02 and IBM z16 AGZ PDUs, a maximum configuration of up to three PCIe+ I/O drawers can be installed. PCIe+ I/O drawers can be added concurrently.

In addition, IBM z16 A02 and IBM z16 AGZ support top-exit options for the fiber optic and copper cables that are used for I/O and power. These options give you more flexibility in planning where the system is installed, eliminate the need for cables to be run under a raised floor, and increase air flow over the system.

IBM z16 A02 and IBM z16 AGZ support installation on raised floor and non-raised floor.

Figure 3-2 on page 37 shows a fully configured IBM z16 A02 with two CPC Drawers and three PCIe+ I/O drawers.

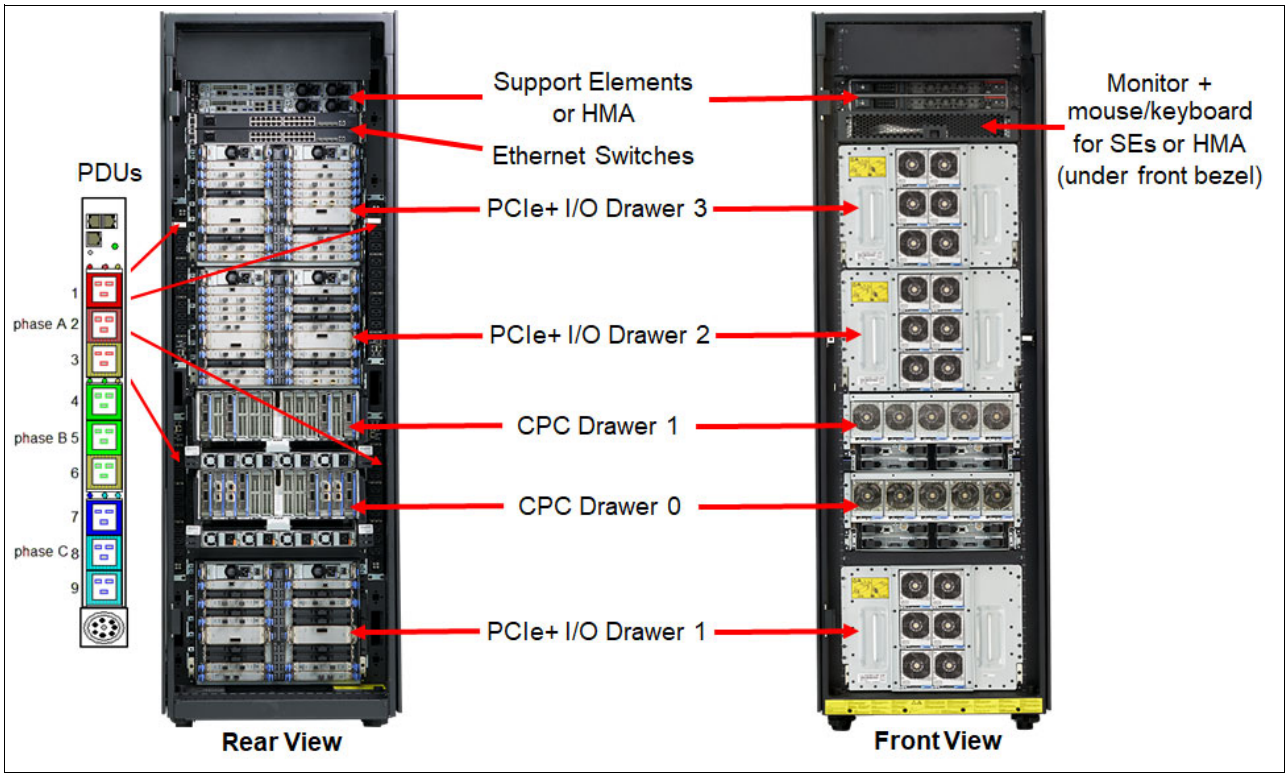

Figure 3-2 Fully configured IBM z16 A02 with two CPC drawers (Max68)

Figure 3-3 shows an IBM z16 AGZ with one CPC drawer and one PCIe+ I/O drawer.

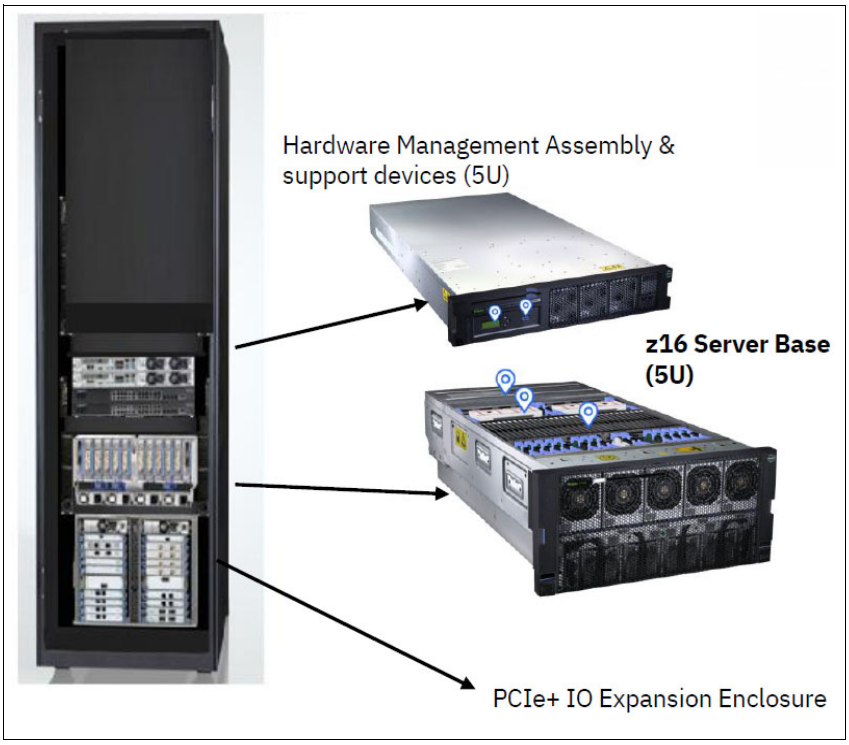

Figure 3-3 IBM z16 AGZ with one CPC drawer and one PCIe+ I/O drawer

The IBM configurator that is used during the order process calculates the placement of CPC and PCIe+ I/O drawers. Factors that determine the configuration of IBM z16 A02 and IBM z16 AGZ include the following examples:

- Feature: Max5, Max16, Max32, or Max68
- Plan-ahead features for the second CPC drawer (for single CPC drawer systems)
- Number of I/O features (Determines the number of PCIe+ I/O drawers.)
- Single- or three-phase power

3.3 CPC drawers

IBM z16 A02 and IBM z16 AGZ can be configured with one or two CPC drawers. Each CPC drawer contains the following elements:

- DCMs:
  - Two or four DCMs per drawer (model-dependent and air-cooled).
  - DCM supports two PU chips with an 8-core maximum.
  - Systems are configured by using 9 - 13 core DCMs.
  - 80 PUs maximum on a Max68 with up to 68 characterizable PUs.
- Memory:
  - A minimum of 64 GB and a maximum of 16 TB of memory per system (excluding 160 GB for the hardware system area (HSA)) is available for use. For more information, see Table 3-3 on page 42.
  - Up to 48 dual inline memory modules (DIMMs) that are 32, 64, 128, or 256 GB are plugged into a CPC drawer.
- Fanouts:

  Each CPC drawer with four DCMs supports up to 12 PCIe+ fanout adapters to connect to the PCIe+ I/O drawers, and Integrated Coupling Adapter Short Reach (ICA SR) coupling links:
  - Two-port Peripheral Component Interconnect Express (PCIe) 16 GBps I/O fanout. Each port supports one domain in the 16-slot PCIe+ I/O drawers.
  - ICA SR1.1 and ICA SR PCIe fanouts for coupling links (two links of 8 GBps each).
- Two or four Power Supply Units (PSUs), depending on the CPC drawer configuration, which provide power to the CPC drawer and are accessible from the rear. The loss of one PSU leaves enough power to satisfy the power requirements of the entire drawer. The PSUs can be concurrently maintained.
- Two dual-function Base Management Cards (BMCs) or Oscillator Cards (OSCs), which provide redundant interfaces to the internal management network and provide clock synchronization to the IBM Z platform.
- Two dual-function Processor Power Cards (PPCs), which control Voltage Regulation, and PSU and fan control. The PPCs are redundant and can be concurrently maintained.
- Five fans are installed at the front of the drawer to provide cooling airflow for the resources that are installed in the drawer, including the PU DCMs.

The CPC drawer communication topology is shown in Figure 3-4 on page 39. CPC drawers are interconnected with high-speed communications links (A-Bus) through the PU chips. Symmetric multiprocessor (SMP-9) cables are used to interconnect the CPC drawers. The X-Bus provides connectivity between DCMs within the drawer, and the M-Bus connects the two PU chips on each DCM.

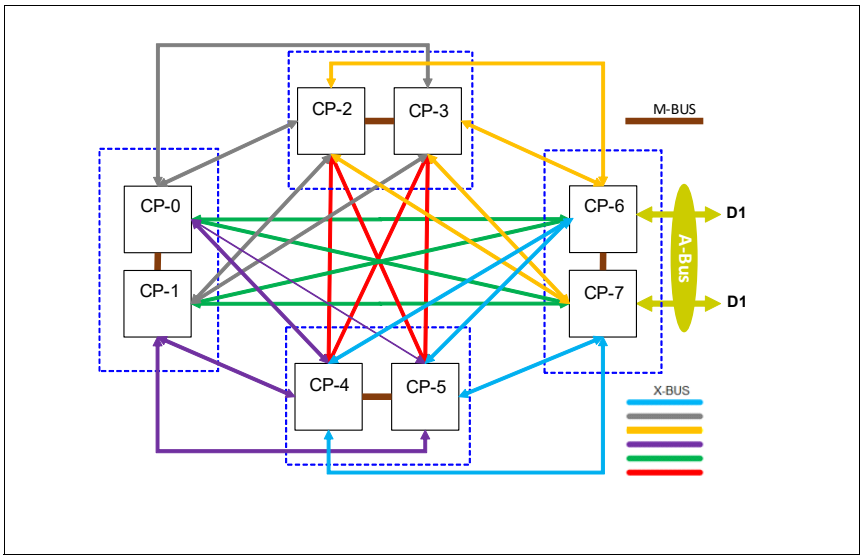

Figure 3-4 IBM z16 A02 and IBM z16 AGZ CPC with two drawers communication topology[2]

The design that is used to connect the PU 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 SMP system.

3.3.1 Dual-chip modules

The CPC drawer may include two or four DCMs. Each DCM contains two PU chips. Each PU chip contains eight cores each with 128 KB Instruction and Data L1 cache and 32 MB semi-private L2 caches.

3.3.2 Processor unit

PU is the generic term for an IBM z/Architecture processor. Each PU is a superscalar processor with the following attributes:

- Up to six instructions can be decoded per clock cycle.
- Up to 10 instructions can be running per clock cycle.
- Instructions can be issued out of order. The PU uses a high-frequency, low-latency pipeline that provides 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 running at any moment, and are subject to the maximum number of decodes and completions per cycle.

[2] IBM z16 A02 and IBM z16 AGZ Max68 is configured with two CPC drawers and four DCMs per drawer. Max5 and Max16 are configured with one CPC drawer with two DCMs per drawer. Max32 is configured with one CPC drawer with four DCMs.

Processor cache structure

The on-chip cache for the PU (core) features the following design:

- Each PU core has an L1 cache (private) that is divided into a 128 KB cache for instructions and a 128 KB cache for data.
- Each PU core has a semi-private L2 cache, which is implemented as 32 MB, near the core.

Note: L1 and L2 are physical cache and are implemented in dense SRAM.

- Each PU core contains a 256 MB shared-victim virtual L3 cache. The shared-victim virtual L3 cache is a logical construction that comprises all eight semi-private L2s (8 x 32 MB = 256 MB) belonging to the other cores.
- The IBM z16 A02 and IBM z16 AGZ Max32 and Max68 CPC drawers contain four DCMs and a maximum of 2 GB shared-victim virtual L4, consisting of the “remote” virtual L3 caches of the DCMs in the CPC drawer. Max5 and Max16 contain two DCMs per CPC drawer and a maximum of 1 GB shared-victim virtual L4.

This on-chip cache implementation optimizes system performance for high-frequency processors and includes the following features:

- Cache improvements
- New translation and TLB2 design
- Pipeline optimizations
- Better branch prediction
- New accelerators and architectures
- Secure Execution support

The IBM z16 A02 and IBM z16 AGZ cache structure is shown Figure 3-5.[3]

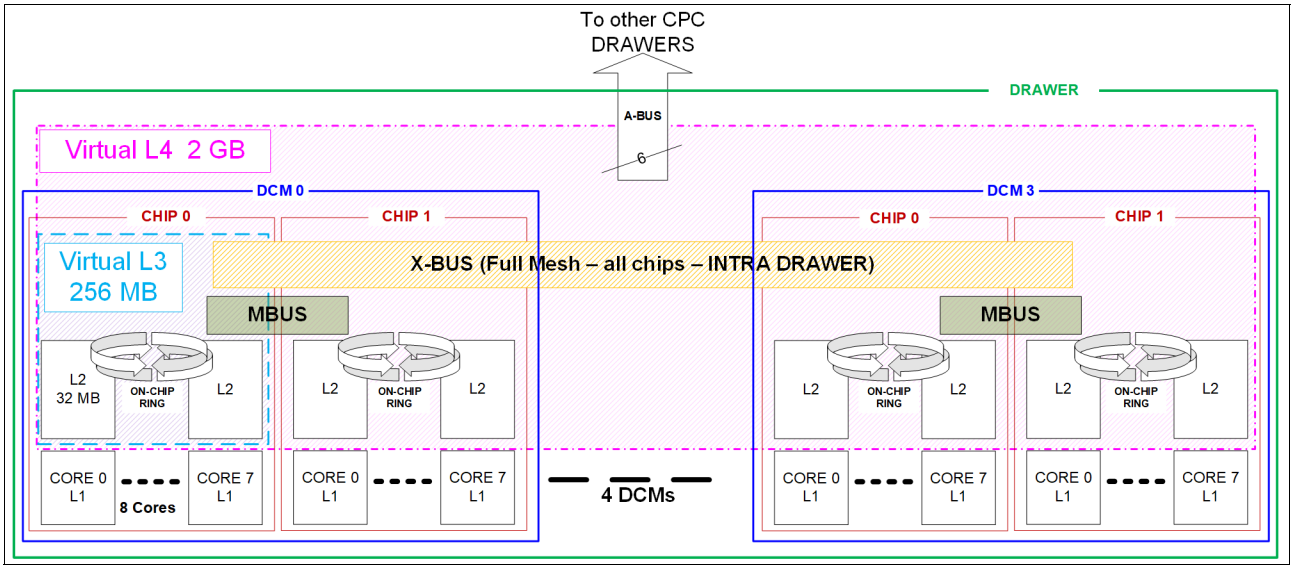

Figure 3-5 IBM z16 A02 and IBM z16 AGZ cache structure

[3] Figure 3-5 represents an IBM z16 A02 and IBM z16 AGZ Max68 with two drawers and four DCMs per CPC drawer.

PU sparing

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

On-core cryptographic hardware

Dedicated on-chip cryptographic hardware for each PU core includes extended key and hash sizes for the Advanced Encryption Standard (AES) and Secure Hash Algorithm (SHA). For more information, see 4.6, “Cryptographic features” on page 59. This cryptographic hardware is available with any processor type, for example, CP, zIIP, and IFL.

On-chip functions

Consider the following points:

- IBM Integrated Accelerator for zEnterprise Data Compression replaces the IBM zEnterprise Data Compression (zEDC) Express PCIe feature that was on previous IBM Z servers.
- The sort accelerator uses the sort instruction (**SORTL**) instruction to be used by **DFSORT** and the IBM Db2 Utilities for z/OS Suite to help reduce CPU usage and improve elapsed time for sort workloads.
- Integrated Accelerator for Artificial Intelligence Unit (AIU) is implemented on each PU chip and shared among all cores. It provides a matrix array for multiplication and convolution alongside the specialty engines for complex functions.

  The AIU provides a Neural Network Processing Assist (NNPA) instruction, which operates directly on Tensor data in user space.

Software support

The IBM z16 A02 and IBM z16 AGZ PUs provides full compatibility with software for z/Architecture, and extend the Instruction Set Architecture (ISA) to enhance functions and performance. New with IBM z16 A02 and IBM z16 AGZ are instructions for AIU.

PU characterization

PUs are ordered in single increments. The internal system functions are based on the configuration that is ordered. They characterize each PU into one of various types during system initialization, which is often called a power-on reset (POR) operation.

Characterizing PUs dynamically without a POR is possible by using a process that is called *Dynamic Processor Unit Reassignment*. A PU that is not characterized cannot be used. Each PU can be designated with one of the following characterizations:

- CP: These standard processors are used for general workloads.
- IFL: Designates processors to be used specifically for running the Linux application programs.
- Unassigned IFL (UIFL): Allows you to directly purchase an IFL feature that is marked as being deactivated upon installation, which avoids software charges until the IFL is brought online for use.
- zIIP: An “Off Load Processor” for workloads that are restricted to Db2 type applications. Also used for the IBM System Recovery Boost (SRB) feature. For more information, see 5.4, “Reliability, availability, and serviceability” on page 86.
- Unassigned zIIP: A processor that is purchased for future use as a zIIP. It is offline and cannot be used until an upgrade for the zIIP is installed. It does not affect software licenses or maintenance charges.

- Integrated Coupling Facility (ICF): Designates processors to be used specifically for coupling.
- SAP: Designates processors to be used specifically for assisting I/O operations.
- IFP: Used for infrastructure management. It is predefined and standard with the platform.

At least one CP must be purchased before zIIPs can be purchased. The maximum number of zIIPs will be one less than the allowed maximum PU configuration. For example, an IBM z16 A02 or IBM z16 AGZ Max68 can have up to 67 zIIPS. These rules also are valid for unassigned zIIPs and unassigned IFLs.

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

Note: The addition of ICFs, IFLs, zIIPs, and SAPs does not change the system capacity setting or its millions of service units (MSU) rating.

3.3.3 Memory

Maximum physical memory size is directly related to the number of CPC drawers in the system. An IBM Z server includes more installed memory than was ordered because part of the installed memory is used to implement the redundant array of independent memory (RAIM) design. With IBM z16 A02 and IBM z16 AGZ, up to 8 TB of memory per CPC drawer can be ordered and up to 16 TB for a two-CPC drawer system.

Important: z/OS requires a minimum of 8 GB of memory (2 GB of memory when running under z/VM). z/OS V2R5 can support up to 16 TB of memory in an LPAR.

The minimum and maximum memory sizes for each IBM z16 A02 and IBM z16 AGZ feature are listed in Table 3-3.

Table 3-3 IBM z16 A02 and IBM z16 AGZ memory per feature

Feature name	CPC drawers	Memory
Max5 (Feature Code 0672)	1	64 GB - 4 TB
Max16 (Feature Code 0673)	1	64 GB - 4 TB
Max32 (Feature Code 0674)	1	64 GB - 8 TB
Max68 (Feature Code 0675)	2	64 GB - 16 TB

The HSA on IBM z16 A02 and IBM z16 AGZ has a fixed amount of memory (160 GB) that is managed separately from available memory. However, the maximum amount of orderable memory can vary from the theoretical number because of dependencies on the memory granularity. On IBM z16 A02 and IBM z16 AGZ, the granularity for memory is in 64, 128, 256, and 512 GB increments.

Physically, memory is organized in the following ways:

- A CPC drawer always contains a minimum of 320 GB to a maximum of 10 TB of installed memory, of which 8 TB maximum 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 (LIC) load.
- Memory upgrades are first satisfied by using installed but unused memory capacity until it is exhausted. When no more unused memory is available from the installed cards, the cards must be upgraded to a higher capacity, or a CPC drawer with more memory must be installed.

When an LPAR is activated, PR/SM attempts to allocate PUs and the memory of the LPAR in a single CPC drawer. However, if this allocation is not possible, PR/SM uses memory resources in any CPC drawer.[4] For example, if the allocated PUs span to the additional CPC drawer, PR/SM attempts to allocate memory across that same set of CPC drawers (even if all required memory is available in a single CPC drawer).

No matter which CPC drawer the memory is installed in, an LPAR can access it after it is allocated. IBM z16 A02 and IBM z16 AGZ are SMP systems because the PUs have access to all the available memory.

Enhanced Drawer Availability

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 IBM z16 A02 and IBM z16 AGZ systems with two CPC drawers, a single CPC drawer can be concurrently removed and reinstalled for a memory upgrade or a repair with EDA.

For upgrades that involve the addition of a CPC drawer, the minimum usable memory increment (256 GB) is added to the system. During an upgrade, adding a second CPC drawer and physical memory in the new drawer are concurrent operations.

Concurrent memory upgrade

If physical memory is available, memory can be upgraded concurrently by using Licensed Internal Code Configuration Control (LICCC).

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 DIMMs.

The RAIM design is fully integrated in IBM z16 A02 and IBM z16 AGZ and enhanced to include one Memory Controller Unit (MCU) per processor chip, with eight memory channels and one DIMM per channel. The MCU enables memory to be implemented as RAIM. This technology has significant reliability, availability, and serviceability (RAS) capabilities in the area of error detection and correction. Bit, lane, DRAM, DIMM, socket, and complete memory channel failures (including many types of multiple failures) can be detected and corrected.

[4] In IBM z16 A02 or IBM z16 AGZ, the memory resources can be allocated in the second CPC drawer, if it is available.

3.3.4 Hardware system area

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

The fixed size 160 GB HSA of IBM z16 A02 and IBM z16 AGZ is large enough to accommodate any LPAR definitions or changes, which eliminates most outage situations and the need for extensive planning.

A fixed, large HSA allows the dynamic I/O capability of IBM z16 A02 and IBM z16 AGZ to be enabled by default. It also enables the dynamic addition and removal of the following features:

- LPAR to new or existing CSS
- CSS (up to three can be defined)
- Subchannel set (up to three can be defined)
- Devices, up to the maximum number permitted, in each subchannel set
- Logical processors by type
- Cryptographic adapters

3.4 I/O system structure

IBM z16 A02 and IBM z16 AGZ support the PCIe-based infrastructure for the PCIe+ I/O drawers. The PCIe I/O infrastructure consists of dual-port PCIe fanouts in the CPC drawers that support 16 GBps connectivity to the PCIe+ I/O drawer.

Ordering of I/O features: The number and type of features determine the number of I/O drawers that are needed.

Figure 3-6 on page 45 shows a high-level view of the I/O system structure for IBM z16 A02 and IBM z16 AGZ.

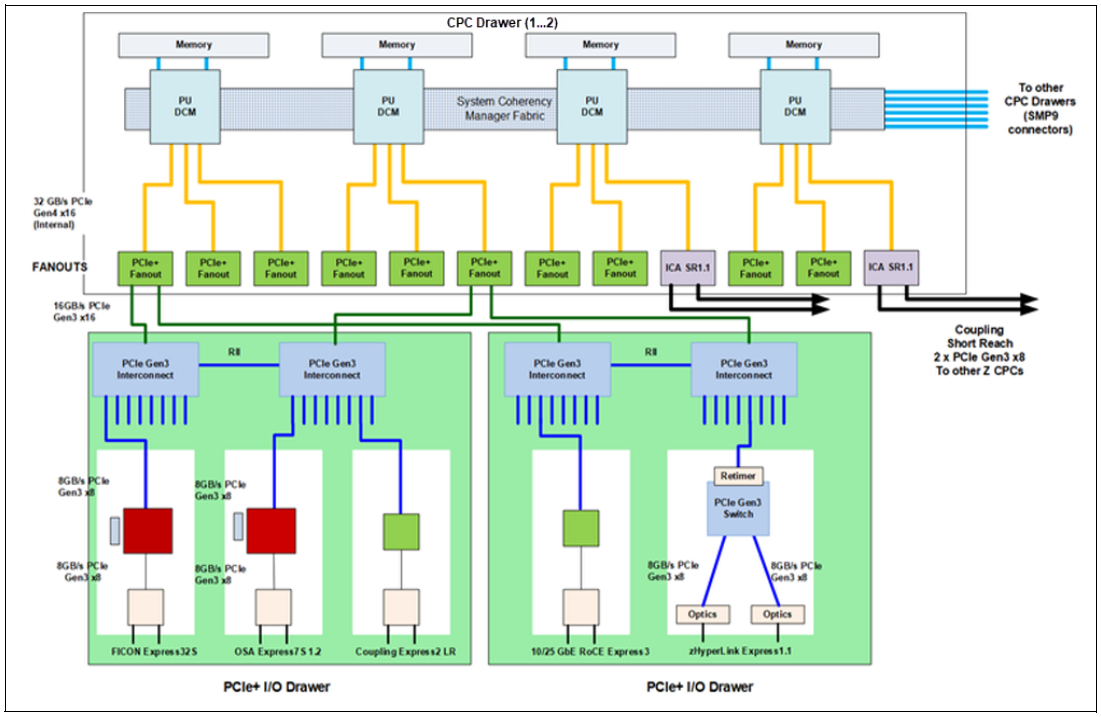

Figure 3-6 IBM z16 A02 and IBM z16 AGZ I/O system structure

The IBM z16 A02 and IBM z16 AGZ Max5 and Max16 CPC drawers have six fanouts, and the Max32 and Max68 have 12 fanouts per CPC drawer. Fanout locations are numbered LG01 - LG12 and can be one of the following types:

- Dual-port PCIe+ fanouts for PCIe+ I/O drawer connectivity
- ICA SR fanouts for coupling
- Filler plates to help with airflow cooling

For coupling link connectivity (IBM Parallel Sysplex and Server Time Protocol (STP) configuration), IBM z16 A02 and IBM z16 AGZ support the following link types:

- ICA SR1.1 and ICA SR (installed in a CPC drawer)
- Coupling Express2 Long Reach (CE LR) (installed in a PCIe+ I/O drawer)

For systems with two CPC drawers, the locations of the PCIe+ fanouts are configured and plugged across all drawers for maximum availability. This configuration helps ensure that alternative paths maintain access to critical I/O devices, such as storage and networks (see Figure 3-7).

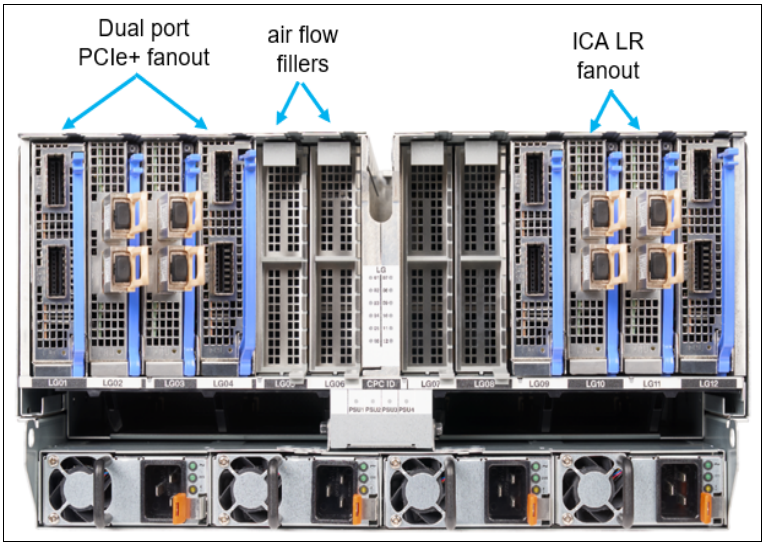

Figure 3-7 CPC drawer: Rear view

The PCIe+ I/O drawer (see Figure 3-8) is a 19-inch single side drawer that is 8U high. I/O features are installed horizontally, with cooling air flow from front to rear. The drawer contains 16 adapter slots and 2 slots for PCIe switch cards.

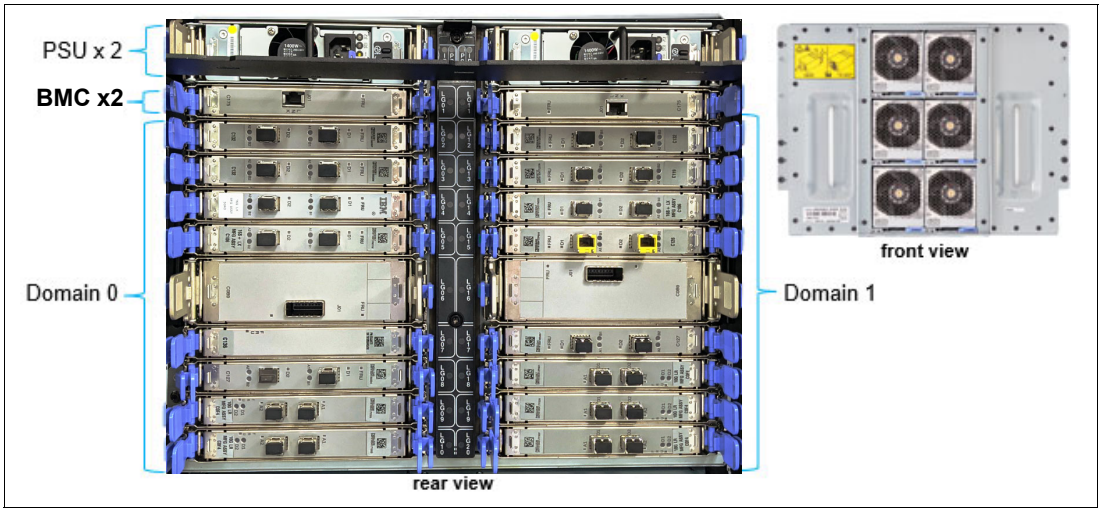

Figure 3-8 PCIe+ I/O drawer: Rear and front view

The two I/O domains per drawer each contain up to eight I/O features that support the following types:

- FICON Express32S or FICON Express16S+
- OSA-Express7S 1.2, OSA-Express7S, or OSA-Express6S
- Crypto-Express8S, Crypto-Express7S, or Crypto-Express6S
- RDMA over Converged Ethernet (RoCE) Express3, RoCE Express2.1, or RoCE Express2
- zHyperLink Express 1.1 and zHyperLink Express
- Coupling Express2 LR

For more information about the I/O features that are available with IBM z16 A02 and IBM z16 AGZ, see Chapter 4, “Supported features and functions” on page 49.

3.5 Power and cooling

IBM z16 A02 and IBM z16 AGZ meet the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Class A3 specifications. ASHRAE is an organization that is devoted to the advancement of indoor-environment-control technology in the heating, ventilation, and air-conditioning industry.

3.5.1 Power options

The IBM z16 A02 19-inch frame is available with the following power options:

- Single-phase PDU

  Single-phase power for IBM z16 A02 can be used for systems with a single CPC drawer. A system with a single CPC drawer supports up to three PCIe+ I/O drawers.

- Three-phase PDU

  Both single- and dual-CPC drawer systems support three-phase power. However, systems with two CPC drawers support only three-phase power and only up to three PCIe+ I/O drawers. IBM z16 A02 does not support plan-ahead or balanced power options.

IBM z16 A02 operates with two or four sets of redundant PDUs. Each set has its own individual power cords or pair of power cords, depending on the configuration. Power cords attach to a three-phase, 50/60 Hz, 200 - 480 V AC power source or to a single-phase 50/60 Hz, 200 - 240 V AC power source. The loss of only one power supply per set has no effect on system operation.

3.5.2 Cooling options

The IBM z16 A02 and IBM z16 AGZ cooling system is available only with one cooling option: air cooling. DCMs are cooled with forced air. PCIe+ I/O drawers and CPC drawers are cooled by chilled air with blowers.

3.5.3 Power considerations

Consider the following points about power:

- The 42U 19-inch IBM frames are used.
- Air flow is front to rear. All blowers are mounted on the front of the frame.
- All external power cabling is at the rear of the frames (no power cabling in front).
- Top or bottom exit power is supported.
- A High-Voltage DC (HVDC) option is not available.
- No Emergency Power Off (EPO) switch is used.

Specific power requirements depend on the number of CPC drawers, and the number and type of I/O features that are installed.

For more information about the maximum power consumption tables for the various configurations and environments, see *3932 Single Frame Installation Manual for Physical Planning (Models A02/LA2)*, GC28-7040-00 and *IBM z16 and LinuxONE Rockhopper 4 Rack Mount Bundle Installation Manual for Physical Planning (IMPP)*, GC28-7035-00.

For more information about the power and weight estimation tool, see IBM Resource Link.

4

Supported features and functions

IBM Z features and functions that are supported on IBM z16 are highlighted in this chapter. The information that is provided expands on the key hardware elements that are described in Chapter 1, “Enduring the pace of digital transformation with IBM z16” on page 1, Chapter 2, “IBM z16 A01 hardware overview” on page 17, and Chapter 3, “IBM z16 A02 and IBM z16 AGZ hardware overview” on page 33.

Naming: Throughout this chapter, we describe features and functions that are offered with IBM z16 A01, IBM z16 A02, and IBM z16 AGZ. The features and functions that are available across all three configurations are identified with “IBM z16”. Where features and functions differ with a given configuration, they are explicitly identified with either IBM z16 A01, IBM z16 A02, or IBM z16 AGZ.

For more information about the key capabilities and enhancements, see *IBM z16 (3931) Technical Guide*, SG24-8951. For more information about the I/O features and functions, see *IBM Z Connectivity Handbook*, SG24-5444.

This chapter describes the following topics:

4.1 IBM z16 I/O connectivity overview

IBM z16 provides a Peripheral Component Interconnect Express (PCIe)-based infrastructure for the PCIe+ I/O drawers to support the following features:

- Storage connectivity:
  - zHyperLink Express1.1 (new build and carry forward)
  - zHyperLink Express (carry forward only)
  - FICON Express32S (new build only)
  - FICON Express16SA (carry forward only)
  - FICON Express16S+ (carry forward only)
- Network connectivity:
  - OSA-Express7S 1.2 (new build only)
  - OSA-Express7S (carry forward only)
  - OSA-Express6S (carry forward only)
  - RDMA over Converged Ethernet (RoCE) Express3 (new build only)
  - RoCE Express2.1 (carry forward only)
  - RoCE Express2 (carry forward only)
- Clustering connectivity:
  - Integrated Coupling Adapter Short Reach (ICA SR) 1.1 (new build or carry forward)
  - ICA Short Reach (SR) (carry forward only)
  - Coupling Express2 Long Reach (LR) (new build only)
- Cryptographic features:
  - Crypto Express8S (CEX8C), one or two hardware security modules (HSMs)[1] (new build only)
  - Crypto Express7S, 1-port, or 2-port[2] (carry forward only)
  - Crypto Express6S (carry forward only)

Detailed specifications for these features are provided in the subsequent sections.

The following features that were supported on earlier IBM Z platforms are *not* orderable and *cannot* be carried forward to IBM z16:

- FICON Express16S
- FICON Express8S
- OSA-Express5S
- 10 GbE RoCE Express
- Crypto Express5S
- IBM zEnterprise Data Compression (zEDC)
- Coupling Express LR

[1] The CEX8C is available with one or two HSMs. The HSM is the IBM 4770 Peripheral Component Interconnect Express Cryptographic Coprocessor (PCIeCC).

[2] The Crypto Express7S comes with one (1-port) or two (2-port) HSMs. The HSM is the IBM 4769 PCIe Cryptographic Coprocessor (PCIeCC).

Note: The LC Duplex connector type is used for all fiber optic cables, except the cables that are used for zHyperLink Express, and ICA SR connections, which have multi-fiber termination push-on (MTP) connectors.

4.2 Storage connectivity

The main focus for storage connectivity is to continuously improve the latency for I/O transmission. With the introduction of zHyperLink Express, IBM ensures the optimization of the I/O infrastructure. The FICON Express32S feature offers increased speed and supports similar functions as its predecessor (FICON Express16SA).

For more information about FICON Express32S performance, see *IBM z16 FICON Express32S Performance*.

Storage connectivity options are listed in Table 4-1.

Table 4-1 Storage connectivity features

Feature	Feature Code	Bit rate per second	Cable type	Maximum unrepeated distance	Ordering information IBM z16
zHyperLink Express1.1	0451	8 GB	OM3 and OM4	See Table 4-2.	New build and carry forward
zHyperLink Express	0431				Carry forward
FICON Express32S LX	0461	8, 16, or 32 Gb	SM 9 µm	10 km [a] (6.2 miles)	New build
FICON Express32S SX	0462	8, 16, or 32 Gb	OM2, OM3, and OM4	See Table 4-2.	New build
FICON Express16SA LX[b]	0436	8 or 16 Gb	SM 9 µm	10 km (6.2 miles)	Carry forward
FICON Express16SA SX[b]	0437	8 or 16 Gb	OM2, OM3, and OM4	See Table 4-3.	Carry forward
FICON Express16S+ LX	0427	4, 8, or 16 Gb	SM 9 µm	10 km (6.2 miles)	Carry forward
FICON Express16S+ SX	0428	4, 8, or 16 Gb	OM2, OM3, and OM4	See Table 4-3.	Carry forward

a. At 32 Gbps, the distance to the first direct-connected device (other FICON adapter, SAN switch, storage device, WDM module, and so on) is limited to 5 km (3.1 miles).
b. Not supported on IBM z16 AGZ.

4.2.1 zHyperLink Express

IBM zHyperLink Express is a short-distance IBM Z I/O adapter with up to 5x lower latency than High-Performance FICON for read requests. This feature is housed in the PCIe+ I/O drawer and is a two-port adapter that is used for short distances (direct connectivity between an IBM z15 and IBM z16 and a DS8880 or newer). The zHyperLink Express is designed to support distances up to 150 meters (492 feet) at a link data rate of 8 GBps. For more information, see *Getting Started with IBM zHyperLink for z/OS*, REDP-5493.

The maximum unrepeated distances for different multimode fiber optic cable types when used with zHyperLink Express are listed in Table 4-2.

Table 4-2 Unrepeated distances for multimode fiber optic cable types for zHyperlink Express

Cable type[a] (modal bandwidth)	8 GBps
OM3 (50 µm at 2000 MHz·km)	100 meters (328 feet)
OM4 (50 µm at 4700 MHz·km)	150 meters (492 feet)

a. Fiber optic cable with 24 fibers (12 transmit plus 12 receive fibers) and MTP connectors.

A 24-fiber cable with MTP connectors is required for the ports of the zHyperLink Express feature. Internally, a single cable contains 12 fibers for transmit and 12 fibers for receive.

Note: FICON connectivity to each storage system is required. The FICON connection is used for zHyperLink initialization, I/O requests that are not eligible for zHyperLink communications, and as an alternative path if zHyperLink requests fail. For example, storage cache misses or busy storage device conditions can cause requests to fail.

4.2.2 FICON Express features

FICON Express features continue to evolve and deliver improved throughput, and reliability, availability, and serviceability (RAS). In IBM z16, these features can provide connectivity to other systems, such as Fibre Channel (FC) switches and various devices in a SAN environment.

The FICON Express features are commonly used by IBM z/OS, IBM z/VM (and guest systems), Linux on IBM Z, IBM z/VSE[3], 21st Century Software VSE^n V6.3, and IBM z/TPF.

The maximum unrepeated distances for different multimode fiber optic cable types when used with FICON SX (shortwave) features running at different bit rates are listed in Table 4-3.

Table 4-3 Unrepeated distances for multimode fiber optic cable types for FICON Express

Cable type (modal bandwidth)	2 Gbps	4 Gbps	8 Gbps	16 Gbps	32 Gbps
OM1 (62.5 µm at 200 MHz·km)	150 meters	70 meters	21 meters	N/A	N/A
	492 feet	230 feet	69 feet	N/A	N/A
OM2 (50 µm at 500 MHz·km)	300 meters	150 meters	50 meters	35 meters	20 meters
	984 feet	492 feet	164 feet	115 feet	65 feet
OM3 (50 µm at 2000 MHz·km)	500 meters	380 meters	150 meters	100 meters	70 meters
	1640 feet	1247 feet	492 feet	328 feet	229 feet
OM4 (50 µm at 4700 MHz·km)	N/A	400 meters	190 meters	125 meters	100 meters
	N/A	1312 feet	623 feet	410 feet	328 feet

[3] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

IBM Fibre Channel Endpoint Security

IBM Fibre Channel Endpoint Security was first introduced with IBM z15 for FICON Express16SA features (Feature Code 0436 and Feature Code 0437)[4]. FICON Express32S features (Feature Code 0461 and Feature Code 0462) also support IBM Fibre Channel Endpoint Security, together with Endpoint Security Enablement (Feature Code 1142). This capability adds Fibre Channel Endpoint Authentication and Encryption of Data in Flight for FICON and FC connections to IBM DS8000 storage systems.

Based tightly on the Fibre Channel–Security Protocol-2 (FC-SP-2) standard, which provides various means of authentication and essentially maps IKEv2 constructs for security association management and derivation of encryption keys to Fibre Channel Extended Link Services, the IBM Fibre Channel Endpoint Security implementation uses the IBM solution for key server infrastructure in the storage system (for data at-rest encryption).

IBM Security® Guardium® Key Lifecycle Manager provides shared secret key generation in a master-subordinate relationship between an FC initiator (the IBM Z platform) and the storage target. The solution implements authentication and key management called *IBM Secure Key Exchange* (SKE).

Data that is in-flight (to or from IBM Z servers and IBM Storage) is encrypted when it leaves either endpoint (source) and then decrypted at the destination. Encryption and decryption are done at the FC adapter level.

In endpoint security-related operations, the operating system that runs on the IBM Z platform is not involved. Tools are provided at the operating system level for displaying information about encryption status.

IBM Fibre Channel Endpoint Security is an orderable feature for IBM z16 (Feature Code 1146) and requires Central Processor Assist for Cryptographic Functions (CPACF) enablement (Feature Code 3863), specific storage (DS8900), and FICON Express32S features.

For more information and implementation details, see the *IBM Fibre Channel Endpoint Security for IBM z15 and LinuxONE III Announcement Letter*.

4.3 Network connectivity

IBM z16 offers a wide range of functions that can help consolidate or simplify the network environment. These functions are supported by HiperSockets, OSA-Express features, Shared Memory Communications (SMC), and RoCE Express features.

4.3.1 IBM HiperSockets

IBM HiperSockets are referred to as the “network in a box” because it simulates local area network (LAN) environments entirely within IBM Z servers. The data transfer is from logical partition (LPAR) memory to LPAR memory, which is mediated by IBM Z firmware.

IBM z16 supports up to 32 HiperSockets. One HiperSockets network can be shared by up to 85 LPARs (40 LPARs with IBM z16 A02 and IBM z16 AGZ). 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:

[4] FICON Express16SA Feature Codes 0436 and 0437 are not supported on IBM z16 A02 and IBM z16 AGZ.

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

Traffic can be Internet Protocol Version 4 (IPv4) or Version 6 (IPv6) or non-IP traffic. HiperSockets devices are independent of protocol and Layer 3. Each HiperSockets device has its own Layer 2 Media Access Control (MAC) address. This address 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 IBM Z environment the same way as they do for a non IBM Z environment. HiperSockets Layer 2 support is provided by Linux on IBM Z, and by z/VM for guest use.

4.3.2 OSA-Express features

OSA-Express features achieve high levels of throughput (mixed inbound/outbound) by using a data router function. The data router enables a direct host memory-to-LAN flow. This function is designed to reduce latency and increase throughput for standard Ethernet frames (1492 bytes) and jumbo frames (8992 bytes).

4.3.3 Shared Memory Communications

The SMC capabilities of IBM z16 optimize the communications between applications in cross IBM Z platform (SMC-R) or LPAR-to-LPAR (SMC-D) connectivity. With IBM z16, SMC is supported between z/OS LPARs, Linux on IBM Z LPARs, and IBM AIX® (running on IBM Power servers).

SMC is available in z/OS V2R4 (with program temporary fixes (PTFs)) and z/OS V2R5. The initial version of SMC was limited to TCP/IP connections over the same Layer 2 network; therefore, it was not routable across multiple IP subnets.

SMC Version 2 (SMCv2) supports SMC over multiple IP subnets for SMC-D and SMC-R and is referred to as SMC-Dv2 and SMC-Rv2. SMCv2 requires updates to the underlying network technology. SMC-Dv2 requires ISMv2, and SMC-Rv2 requires RoCEv2.

The SMCv2 protocol is compatible with earlier versions and allows SMCv2 hosts to continue to communicate with SMCv1 hosts.

SMC-R provides application-transparent use of the RoCE Express features that can reduce the network impact and latency of data transfers, which effectively offers the benefits of optimized network performance across processors.

The Internal Shared Memory (ISM) virtual Peripheral Component Express (PCI) function uses the capabilities of SMC-D. ISM is a virtual PCI network adapter that enables direct access to shared virtual memory, which provides a highly optimized network interconnect for IBM Z intra-system communications. Up to 32 channels for SMC-D traffic can be defined in an IBM z16 server, and each channel can be virtualized to a maximum of 255 Function IDs.[5] No other hardware is required for SMC-D.

[5] The 10 GbE RoCE features and the ISM adapters are identified by a hexadecimal Function Identifier (FID) with a range of 00 - FF.

4.3.4 RoCE Express features

The RoCE Express features help reduce the use of CPU resources for applications that use the TCP/IP stack. It also might help to reduce network latency with memory-to-memory transfers that use SMC-R in z/OS environments. It is transparent to applications, and can be used for system-to-system communication in multiple IBM Z environments.

These features are installed in the PCIe+ I/O drawer and use an SR optical transceiver. Point-to-point connections and switched connections with an Ethernet switch are supported. Ethernet switches must include enablement of the *Pause frame* as defined by the IEEE 802.3x standard.

Depending on the RoCE Express feature type, a maximum of 8 or 16 RoCE Express features can be installed in IBM z16 in any combination.

Note: The 10 GbE and 25 GbE RoCE Express3 LR are new with IBM z16. Previous RoCE generations supported SR connectivity only.

The 25 GbE RoCE Express must not be mixed with any type of 10 GbE RoCE Express in the same SMC-R link group. The 10 GbE RoCE Express adapters can be mixed in any combination in the same SMC-R link group.

The network connectivity features are listed in Table 4-4.

Table 4-4 Network connectivity features

Feature	Feature Code	Bit rate per second	Cable type	Maximum unrepeated distance[a]	Ordering information IBM z16
OSA-Express7S 1.2 25-GbE LR	0460	25 Gb	SM 9 µm	10 km (6.2 miles)	New build
OSA-Express7S 1.2 25-GbE SR	0459		MM 50 µm	70 m (2000) 100 m (4700)	
OSA-Express7S 1.2 10-GbE LR	0456	10 Gb	SM 9 µm	10 km (6.2 miles)	
OSA-Express7S 1.2 10 GbE SR	0457		MM 62.5 µm MM 50 µm	33 m (200) 82 m (500) 300 m (2000)	
OSA-Express7S 1.2 GbE LX	0454	1.25 Gb	SM 9 µm	5 km (3.1 miles)	
OSA-Express7S 1.2 GbE SX	0455		MM 62.5 µm MM 50 µm	275 m (200) 550 m (500)	
OSA-Express7S 1.2 1000BASE-T	0458	1000 Mbps	Cat 5 or 6 UTP	100 m (328 feet)	
OSA-Express7S 25 GbE SR1.1	0449	25 Gb	MM 50 µm	70 m (2000) 100 m (4700)	Carry forward
OSA-Express7S 25 GbE SR[b]	0429				
OSA-Express7S 10 GbE LR[b]	0444	10 Gb	SM 9 µm	10 km (6.2 miles)	Carry forward
OSA-Express6S 10 GbE LR	0424				

Feature	Feature Code	Bit rate per second	Cable type	Maximum unrepeated distance[a]	Ordering information IBM z16
OSA-Expess7S 10 GbE SR[b]	0445	10 Gb	MM 62.5 µm MM 50 µm	33 m (200) 82 m (500) 300 m (2000)	Carry forward
OSA-Expess6S 10 GbE SR	0425				
OSA-Express7S GbE LX[b]	0442	1.25 Gb	SM 9 µm	5 km (3.1 miles)	Carry forward
OSA-Express6S GbE LX	0422				
OSA-Express7S GbE SX[b]	0443	1.25 Gb	MM 62.5 µm MM 50 µm	275 m (200) 550 m (500)	Carry forward
OSA-Express6S GbE SX	0423				
OSA-Express7S 1000BASE-T[b]	0446	1000 Mbps	Cat 5 or 6 UTP	100 m	Carry forward
OSA-Express6S 1000BASE-T	0426	100 or 1000 Mbps			
25 GbE RoCE Express3 LR	0453	25 Gb	SM 9 µm	10 km (6.2 miles)	New build
25 GbE RoCE Express3 SR	0452		MM 50 µm	70 m (2000) 100 m (4700)	
10 GbE RoCE Express3 LR	0441	10 Gb	SM 9 µm	10 km (6.2 miles)	
10 GbE RoCE Express3 SR	0440		MM 62.5 µm MM 50 µm	33 m (200) 82 m (500) 300 m (2000)	
25 GbE RoCE Express2.1	0450	25 Gb	MM 50 µm	70 m (2000) 100 m (4700)	Carry forward
25 GbE RoCE Express2	0430				
10 GbE RoCE Express2.1	0432	10 Gb	MM 62.5 µm MM 50 µm	33 m (200) 82 m (500) 300 m (2000)	Carry forward
10 GbE RoCE Express2	0412				

a. The minimum fiber bandwidth distance in MHz-km for multi-mode fiber optic links is included in parentheses, where applicable.
b. Not supported on IBM z16 A02 and IBM z16 AGZ.

4.4 Clustering connectivity

IBM Parallel Sysplex is a clustering technology with which you can operate multiple copies of z/OS images as a single system from a user's perspective. A suitably configured Parallel Sysplex can achieve near-continuous availability. The component that enables this parallelism is the Coupling Facility (CF), which can run as a separate LPAR (internal CF) or within dedicated hardware (external CF).

What makes a group of such z/OS images into a sysplex is the inter-communication. This inter-communication is handled through coupling links. Coupling links enable all the z/OS to CF communication, CF-to-CF traffic, or Server Time Protocol (STP)[6].

For more information about options, see *Coupling Facility Configuration Options*.

[6] All external coupling links can be used to carry STP timekeeping information.

Internal coupling (IC) links are used for internal communication between LPARs on the same system that is running CFs 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.

Coupling Facility Control Code (CFCC) level 25 is available for IBM z16. Coupling link options are listed in Table 4-5.

Table 4-5 Coupling link features

Feature	Feature Code	Bit rate per second	Cable type	Maximum unrepeated distance	Ordering information
Coupling Express2 LR	0434	10 Gb	SM 9 µm	10 km (6.2 miles)	New build
ICA SR 1.1	0176	8 GB	OM3[a] and OM4[a]	150 m 100 m	New build or carry forward
ICA SR	0172				Carry forward
IC	No coupling link feature or fiber optic cable is required.				

a. A fiber optic cable with 24 fibers (12 transmit plus 12 receive fibers) and MTP connectors.

For more information about operating system-level coexistence, see the IBM Documentation.

4.4.1 Dynamic I/O configuration for stand-alone CFs, Linux on Z and z/TPF

Dynamic I/O configuration changes can be made to a stand-alone CF[7], Linux on Z and z/TPF CECs, without requiring a disruptive power on reset (POR).

This new support is applicable only when both, the driving CEC and the target CEC are z16 with the required firmware support, and when the driving system's z/OS level is 2.3 or higher with APAR OA655559 installed.

A separate LPAR with a firmware-based appliance that contains an activation service instance is used to apply I/O configuration changes to the stand-alone CF, Linux on Z and z/TPF dynamically. The firmware-based LPAR is driven by updates from an HCD instance that is running in a z/OS LPAR on a different IBM z16 server that is connected to the same Hardware Management Console (HMC).

The firmware LPAR is defined in the range of IBM reserved LPARs and does not support any attached I/O, that is, it does not take away any of your configurable resources.

4.4.2 Quantum-safe protection for CFCC

Quantum-safe protection has been expanded to cover Coupling Facility Control Code (CFCC). CFCC boot and the update processes are protected with Cryptographic Suite for Algebraic Lattices (CRYSTALS)-Dilithium, recently standardized by NIST, for quantum-safe digital signatures. For more information about quantum-safe digital signatures, see *Transitioning to Quantum-Safe Cryptography on IBM Z*, SG24-8525.

[7] A stand-alone CF does not have any running instances of z/OS or z/VM.

4.5 Server Time Protocol

STP is a message-based protocol in which timekeeping information is passed over coupling links between IBM Z servers. IBM z16 can participate in an STP Coordinated Timing Network (CTN). A CTN is a collection of IBM Z servers that are time-synchronized to a time value that is called *Coordinated Server Time* (CST).

STP is implemented in LIC as a system-wide facility of IBM z16 and other IBM Z servers. IBM z16 is enabled for STP by installing the STP feature code. Extra configuration is required for an IBM z16 to become a member of a CTN.

For high availability (HA) purposes, nondisruptive capability was implemented in IBM z16 firmware that allows two CTNs to be merged into one, or to split one CTN into two, dynamically.

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 for a multi-site Parallel Sysplex.

Note: If an IBM z16 plays a CTN role (Primary Time Server (PTS), Backup Time Server (BTS), or Arbiter), the other CTN roleplaying IBM Z server must include direct coupling connectivity to IBM z16.

4.5.1 Changes to the STP implementation in IBM z16

The following significant changes were made to the STP implementation on IBM z16 over earlier IBM Z models:

- Direct network connection to the IBM z16 CPC drawer of the External Time Source (ETS).

  ETS for Network Time Protocol (NTP), Precision Time Protocol (PTP), and Pulse per Second (PPS) network cables now connects directly to the IBM z16 central processor complex (CPC) drawers.[8] This connection provides greater accuracy to the external time reference.

- The n-mode Power STP Imminent Disruption Signal option.

  On IBM Z servers, losing a PTS results in significant consequences to the timing network and the overall workload execution environment of the IBM Z sysplex.

  Because an integrated battery facility (IBF) no longer exists, support was added to the CPC to support n-mode power conditions (wall power or power cord loss). If a power condition is detected, an automated failover occurs to the BTS. The requirement for the backup power method is to hold power for 60 seconds on the PTS to allow fail over to complete.

  HMC System Events are available for awareness through a user interface or automation.

[8] The ETS and PPS cables are connected to the Base Management Card (BMC) or Oscillator Card (OSC) adapters and from the front of the CPC drawer.

4.5.2 Network Time Protocol client support

The use of NTP servers as an ETS usually fulfills a requirement for a time source or common time reference across heterogeneous platforms. This approach also provides greater time accuracy.

NTP client support is available in the ETS or STP partition that is running on IBM z16. The code interfaces with the NTP servers. This interaction allows an NTP server to become the single-time source for IBM z16 and for other servers that have NTP clients. The NTP Ethernet cable must plug directly into the Base Management Card (BMC) or Oscillator Card (OSC) ports on the IBM z16 CPC drawer. Redundant cabling and ETS must be configured.

4.5.3 Precision Time Protocol client support

IEEE 1588 PTP is implemented on IBM z16 as an ETS for a CTN. The PTP Ethernet cable must plug directly into the BMC or OSC ports on IBM z16. Redundant cabling and ETS should be configured.

4.5.4 Pulse per Second support

Two OSCs,[9] which are included as a standard feature of IBM z16, provide a dual-path interface for the PPS signal. The cards contain a Bayonet Neill-Concelman (BNC) connector for PPS attachment at the front side of the CPC drawer. The redundant design allows continuous operation during the failure of one card and concurrent card maintenance.

STP tracks the highly stable and accurate PPS signal from the external time server. PPS maintains accuracy of 10 µs as measured at the PPS input of IBM z16.

A cable connection from the PPS port to the PPS output of an NTP server is required when IBM z16 is configured for NTP with PPS as ETS for time synchronization.

PPS is optional for PTP, but might still be required for NTP to meet financial regulations.

For more information, see *IBM z16 (3931) Technical Guide*, SG24-8951 and *IBM Z Server Time Protocol Guide*, SG24-8480.

4.6 Cryptographic features

The IBM Z platform offers cryptographic engines that provide high-speed cryptographic operations:

- CPACF

  Cryptographic functions are provided through a set of instructions, which are available in the hardware on every processor unit (PU).

- Crypto Express features

  Cryptographic functions that are provided through high-security, tamper-responding hardware security modules (HSMs).

[9] The OSCs are combined with the BMCs.

The IBM Z platform provides cryptographic functions that can be categorized in the following groups from an application program perspective:

- Symmetric cryptographic functions,[10] which are provided by the CPACF or Crypto Express features when defined as an accelerator.
- Asymmetric cryptographic functions,[11] which are provided by the Crypto Express features.

In addition, cryptographic features and functions were added to protect IBM z16 from attacks, including threats that might use quantum computers. The system includes quantum-safe technology through the many firmware layers that are loaded during the boot process. Only authentic, IBM approved firmware is accepted. This hardware-protected verification of the firmware uses a dual-signature scheme, which is a combination of quantum-safe and classical digital signatures. The protection is anchored in the IBM Z *Root of Trust*.[12]

For more information about the quantum-safe technologies that are used in IBM z16, see *Transitioning to Quantum-Safe Cryptography on IBM Z*, SG24-8525.

4.6.1 Central Processor Assist for Cryptographic Functions

The CPACF enablement feature (Feature Code 3863) offers a set of symmetric cryptographic functions for high-performance encryption and decryption with clear key operations for Secure Sockets Layer (SSL) and Transport Layer Security (TLS), VPN, and data-storing applications that do not require Federal Information Processing Standards (FIPS) 140-2 Level 4 security[13]. The CPACF is a no-charge optional feature that is integrated with the compression unit coprocessor in the IBM z16 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. A 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:

- Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys
- Improved performance of AES Galois/Counter Mode (GCM) encryption
- Data Encryption Standard (DES) data encryption and decryption with single, double, or triple length keys
- Pseudo-random number generation (PRNG)
- Deterministic Random Number Generation (DRNG)
- True-random number generator (TRNG)
- Message authentication code
- Message-Security-Assist extension 9
- Elliptic Curve Cryptography (ECC) support
- Hashing algorithms: Secure Hash Algorithm (SHA)-1, SHA-2, and SHA-3

[10] With symmetric encryption, the same cryptographic key is used for encryption and decryption of the data (the sender and receiver of the data use the same key).
[11] With asymmetric encryption, the receiver's public key is used for encryption and the receiver's private key is used for decryption.
[12] Root of Trust is a source that can always be trusted within a cryptographic system.
[13] FIPS 140-3 Security Requirements for Cryptographic Modules.

SHA-1, SHA-2, and SHA-3 support are enabled on all IBM Z servers and do not require the CPACF enablement feature. The CPACF functions are supported by z/OS, z/VM, z/VSE[14], 21st Century Software VSE^n V6.3, z/TPF, and Linux on IBM Z.

Attention: Many older cryptographic algorithms like DES or RSA, and hashing algorithms such as SHA1 are considered weak algorithms and do not provide sufficient protection against today's cyberattacks.

This risk can be mitigated by switching to stronger algorithms, such as AES-256, SHA-256, SHA-3, and CRYSTALS-Dilithium.

IBM provides several tools that can aid in the discovery process:

- IBM z/OS Integrated Cryptographic Service Facility (ICSF)
- IBM Application Discovery and Delivery Intelligence (ADDI)
- IBM Crypto Analytics Tool (CAT)
- IBM z/OS Encryption Readiness Technology (zERT)

These tools can help you identify certificates, encryption protocols, algorithms, and key lengths that are at risk in your IBM Z environment.

4.6.2 Crypto Express8S

The CEX8C represents the newest generation of the Peripheral Component Interconnect Express (PCIe) cryptographic co-processors, which are an optional feature that is available on IBM z16. These co-processors are hardware security modules (HSMs) that provide high-security cryptographic processing as required by banking and other industries.

The CEX8C provides quantum-safe APIs that enable you to use quantum-safe cryptography along with classical cryptography as existing applications are modernized and new applications are built.

This feature provides a secure programming and hardware environment where crypto-processes are performed. Each cryptographic coprocessor includes general-purpose processors, nonvolatile storage, and specialized cryptographic electronics. All these features are contained within a tamper-sensing and tamper-responsive enclosure that eliminates all keys and sensitive data on any attempt to tamper with the device. The security features of the HSM are designed to meet the requirements of FIPS 140-2 Level 4.

The CEX8C (2-port) feature (Feature Code 0908) includes two Peripheral Component Interconnect Express Cryptographic Coprocessors (PCIeCCs), and the CEX8C (1-port) feature (Feature Code 0909) includes one PCIeCC. For availability reasons, a minimum of two features is required for the one-port feature. Up to 30 CEX8C (2-port) features are supported on IBM z16 A01, and up to 20 (2-port) features are supported on IBM z16 A02 and IBM z16 AGZ. The maximum number of the 1-port features is 16. The Crypto Express8S feature occupies one I/O slot in a PCIe+ I/O drawer.

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

CEX8C provides domain support for up to 85 LPARs on IBM z16 A01 and 40 LPARs on IBM z16 A02 and IBM z16 AGZ.

[14] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

The accelerator function is designed for maximum-speed SSL and TLS acceleration, rather than for specialized financial applications for secure, long-term storage of keys or secrets. The CEX8C can be configured as one of the following configurations:

- The Secure IBM CCA coprocessor includes secure key functions with an emphasis on the specialized functions that are required for banking and payment card systems. It is optionally programmable to add custom functions and algorithms by using User Defined-Extensions (UDXs).

  A new mode, called *Payment Card Industry PTS HSM* (PCI-HSM), is available in CCA mode. PCI-HSM mode simplifies compliance with Payment Card Industry requirements for HSMs.

- The Secure IBM EP11 coprocessor implements an industry-standardized set of services that adheres to the PKCS #11 specification 2.20 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.

When the Crypto Express8S PCIe adapter is configured as a secure IBM CCA coprocessor, it still provides accelerator functions. However, up to 3x better performance for those functions can be achieved if the Crypto Express8S PCIe adapter is configured as an accelerator.

CCA enhancements include the ability to use triple-length (192-bit) Triple Data Encryption Standard (TDES) keys for operations, such as data encryption, IBM PIN® processing, and key wrapping to strengthen security. CCA also extended the support for the cryptographic requirements of the German Banking Industry Committee, Deutsche Kreditwirtschaft, and quantum-safe cryptography.

The CEX8S with CCA firmware added secure key support for quantum-safe cryptography private keys for signature and key encapsulation mechanism (KEM) use cases. The CCA interface adds support for CRYSTALS-Dilithium secure private keys of sizes (6,5) and (8,7) for both round 2 and round 3 versions of CRYSTALS-Dilithium, usable for digital signature generation and verification. The CCA interface also adds support for round2 of the CRYSTALS Kyber 1024 parameter set, a KEM, which can be used to protect 32-byte values. These pieces are also combined with Elliptic-curve Diffie-Hellman (ECDH) support to offer a complete quantum-safe cryptography hybrid key exchange scheme, which is implemented with secure CCA private keys for all involved private keys, which is helpful for scenarios where data needs enhanced authentication against future quantum computing attacks on conventional cryptography.

Several features that support the usage of the AES algorithm in banking applications also were added to CCA. These features include the addition of AES-related key management features and the AES ISO Format 4 (ISO-4) PIN blocks, as defined in the ISO 9564-1 standard. PIN block conversion is supported and used in AES PIN blocks in other CCA callable services. IBM continues to add enhancements as AES finance industry standards are released.

4.6.3 Crypto Express7S and Crypto Express6S (carry forward only)

The Crypto Express7S feature has one or two PCIeCC (HSMs) per feature (Feature Code 0898 has 2-port, and Feature Code 0899 has 1-port) and Crypto Express6S feature (Feature Code 0893) has one PCIeCC (HSM) per feature. For availability reasons, a minimum of two features is required. Up to 16 Crypto Express7S or Crypto Express6S (1-port) features and 20 (2-port) are supported on IBM z16.

Each adapter can be configured as a Secure IBM CCA coprocessor, a Secure IBM EP11 coprocessor, or as an accelerator.

Crypto Express7S and Crypto Express6S provide domain support for up to 85 LPARs on IBM z16 A01 and 40 LPARs on IBM z16 A02 and IBM z16 AGZ.

Trusted Key Entry (TKE) feature: The TKE Workstation feature is required to support the administration of the Crypto Express features when configured as an Enterprise PKCS #11 coprocessor or managing the CCA mode PCI-HSM.

Changes were made to the TKE feature to use quantum-safe cryptography when authenticating CEX8Cs, verifying replies from the CEX8S coprocessors, and protecting key parts in flight for the Common Cryptographic Architecture (CCA). Finally, the IBM Z pervasive encryption functions were updated to use quantum-safe mechanisms for key management.

4.7 IBM Virtual Flash Memory

IBM Virtual Flash Memory (VFM) is the replacement for the Flash Express features that were available on earlier IBM Z servers. With IBM z16 A01, the VFM feature (Feature Code 0644) can be ordered in 512 GB increments up to 12 VFM features for a total of 6 TB. IBM z16 A02 and IBM z16 AGZ support 4 VFM features of 512 GB for a total of 2 TB.

VFM helps to improve availability and handling of paging workload spikes. With this support, z/OS is designed to help improve system availability and responsiveness by using VFM across transitional workload events.

VFM can be used in CF images to provide extended capacity and availability for workloads that use IBM MQ shared queues. Using VFM can help availability by reducing latency from paging delays that can occur at the start of the workday or during other transitional periods. It is designed to eliminate delays that can occur when diagnostic data is collected during failures.

Therefore, VFM can help meet most demanding service-level agreements (SLAs) and compete more effectively. VFM provides rapid time to value.

4.8 Hardware Management Console and Support Element

The HMC and Support Element (SE) are appliances that provide hardware management for IBM Z servers. Hardware platform management covers a complex set of configuration, operation, monitoring, and service management tasks, and other services that are essential to the operations of the IBM Z platform.

The minimum driver level for HMC and SE for IBM z16 is Driver 51. Driver 51 is equivalent to Version 2.16.0.

Note: The HMC with Driver 51 or Version 2.16.0 can manage N-2 generations of IBM Z servers (IBM z16, IBM z15, and IBM z14).

On IBM z16, two HMCs are delivered with the Hardware Management Appliance (HMA) feature (Feature Code 0129). It is possible to order the HMA feature later. However, only new microcode is delivered without HMC hardware.

Note: The HMC code runs on the two integrated 1U rack-mounted servers on the top of the IBM z16 A frame. Stand-alone that is outside the IBM z16 HMCs (tower or rack mount) can no longer be ordered.

HMC features (Feature Code 0062, Feature Code 0063, Feature Code 0082, and Feature Code 0083) can be carried forward from previous orders, and Driver 51 or Version 2.16.0 can be installed to support IBM z16.

Also, Driver 51 or Version 2.16.0 can be installed on the two HMCs that are provided with the HMA feature (Feature Code 0100) on IBM z15. The SEs and HMCs are closed systems; therefore, no other applications can be installed on them.

With IBM z16 and HMA, the SE and HMC codes run *virtualized* on the integrated two SEs on the two integrated 1U rack-mounted servers on the top of the IBM z16 A frame.[15] One SE is the Primary SE (active) and the other is the Alternative SE (backup).

The SEs are connected to Ethernet switches for network connectivity with IBM Z servers and the HMCs. An HMC can communicate with one or more IBM Z servers.

When tasks are performed on the HMC, the commands are sent to one or more SEs, which then issue commands to their respective CPCs.

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

[15] For IBM z16 AGZ, the two SEs are installed right above the last installed component (CPC drawer or I/O drawer), in the client-supplied rack.

5

IBM z16 system design strengths

Every new generation of the IBM Z platform introduces innovative features and functions to provide more velocity, security, agility, and flexibility for building new IT solutions and services.

The IBM Z hardware, firmware, and operating systems always conform to the IBM z/Architecture[1] to ensure support of current and future workloads and services. Whenever new capabilities are implemented, the z/Architecture is extended rather than replaced. This practice helps sustain the compatibility, integrity, and longevity of the IBM Z platform. Thus, investment protection for earlier versions of workloads and solutions is ensured.

The evolution of the IBM Z platform embodies a proven architecture that is open, secure, resilient, and adaptable. From the microprocessor and memory design to the artificial intelligence (AI), sort, and cryptography capabilities, and unparalleled I/O throughput and rich virtualization, IBM z16 is built to respond with speed and versatility.

Naming: Throughout this chapter, we describe features and functions that are offered with IBM z16 A01, IBM z16 A02, and IBM z16 AGZ. The features and functions that are available across all three configurations are identified with “IBM z16”. Where features and functions differ with a given configuration, they are explicitly identified with either IBM z16 A01, IBM z16 A02, or IBM z16 AGZ.

This chapter introduces IBM z16 system design capabilities and enhancements.

This chapter describes the following topics:

[1] IBM z/Architecture is the mainframe-computational architecture notation that defines its behavior. For more information, see *IBM z/Architecture Principles of Operation*, SA22-7832.

5.1 Technology improvements

Computer systems achieve the levels of efficiency that are needed by businesses through an overall balanced design. Processor units (PUs), memory, I/O, and network communications must complement each other to achieve the required levels of performance. Hence, you can have the fastest processors available, but your workloads suffer if you cannot feed them with data.

A balanced system design also incorporates all the enhancements in software, hardware, and firmware to accelerate specific type of operations, for example, sorting, inferencing, compressing, and encrypting data.

IBM z16 provides high levels of performance, scalability, resiliency, flexibility, and security when serving as a traditional IBM Z platform, a cloud platform, or both. IBM z16 can host thousands of virtualized environments.

5.1.1 System capacity

Each generation of IBM Z servers provides more system capacity, which combines various system design enhancements. IBM z16 with the Max200 feature offers up to 25% more processing capacity than the previous IBM z15 T01Max190 feature.

IBM z16 improved the instructions per cycle (IPC) and reduced the cycles per instruction (CPI). IBM z16 has eight PUs per PU chip or 16 processors per DCM versus 12 per chip or SCM on IBM z15.

Additionally, Program Resource/System Manager (PR/SM) improved the logical partition (LPAR) placement algorithms based on IBM z15 experiences. PR/SM also allows mixed configurations (IBM z/OS and Linux) a fair share of spares between those two pools.

The family of Capacity on Demand (CoD) offerings ensures a flexible addition of capacity when it is most needed, for example, during peak workload periods, scheduled maintenance, or in disaster recovery (DR) scenarios. For more information, see 5.3.2, “Capacity on Demand offerings” on page 80.

5.1.2 Processor design highlights

IBM z16 supports 64-bit addressing mode and uses Complex Instruction Set Computer (CISC), including highly capable and complex instructions. Most of the instructions are implemented at the hardware or firmware level for the most optimal and effective execution.

PU is the generic term for the z/Architecture CPU. Each PU is a superscalar processor, which can decode up to six complex instructions per clock cycle, running instructions out-of-order. The PU uses a high-frequency, low-latency pipeline that provides robust performance across a wide range of workloads.

z/Architecture addressing modes: The z/Architecture simultaneously supports 24-bit, 31-bit, and 64-bit addressing modes. This feature delivers compatibility with earlier software versions, which provides investment protection.

Compared to its predecessors, IBM z16 features the following processor design changes, improvements, and architectural extensions:

- Redesigned processor chip that uses 7-nm silicon wafer technology
- New cache structure:
  - L1D (data) and L1I (instruction) cache: ON-core (128 KB each)
  - L2 - dense SRAM: Outside the core, semi-private to the core = 32 MB
  - Virtual L3 cache (shared victim) = 256 MB
  - Virtual L4 cache (shared victim) = 2 GB

  Note: The IBM z16 A02 and IBM z16 AGZ Max32 and Max68 CPC drawer contains four DCMs and a maximum of 2 GB shared-victim virtual L4, consisting of the “remote” virtual L3 caches of the DCMs in the CPC drawer. Max5 and Max16 contain 2 DCMs per CPC drawer and a maximum of 1 GB shared-victim virtual L4.

- New Core-Nest Interface
- A new branch prediction design that uses SRAM
- Significant architecture changes (COBOL compiler and more)
- Dedicated coprocessors and accelerators for each PU core:
  - Central Processor Assist for Cryptographic Functions (CPACF)

    This component encrypts large amounts of data in real time. With IBM z16, CPACF provides counters that track cryptographic compliance and instruction use, algorithms, bit length, and key security for a specific workload.
  - IBM Integrated Accelerator for IBM Z Sort

    This on-core accelerator for sort uses the sort instruction (**SORTL**) to accelerate the sorting of data. The IBM Integrated Accelerator for IBM Z Sort (zSort) feature helps reduce the CPU costs and improve the elapsed time for eligible workloads by speeding up the sorting process and improving database functions.

    For more information, see this IBM Documentation web page.
  - IBM Integrated Accelerator for Artificial Intelligence (AIU)

    This on-chip AIU is designed for high-speed, real-time inferencing at scale. The on-chip AI acceleration adds more than 6 TFLOPS of processing power that is shared by all cores on the chip.

    This centralized AI design provides high performance and consistent low latency inferencing for processing a mix of transactional and AI workloads at speed and scale.

    A robust system of frameworks and open-source tools, which are combined with the new IBM Deep Learning Compiler that generates inferencing programs that are highly optimized for the IBM Z architecture and the AIU, help enable rapid development and deployment of deep learning and machine learning models on IBM Z servers to accelerate time to market. For more information, see IBM Telum processor: the next-gen microprocessor for IBM Z and IBM LinuxONE
  - IBM Integrated Accelerator for z Enterprise Data Compression

    This on-chip accelerator for compression implements compression as defined by RFC1951 (DEFLATE). It performs data compression with improved performance and simplified management on a processor chip level. The on-chip compression is designed to reduce the penalty of storing, transporting, and processing data without changing applications architecture.

Neural Network Processing Assist instruction

To support the IBM Z AIU, a new memory-to-memory CISC instruction is used to operate directly on Tensor objects that are in client application program memory. AI functions and macros are abstracted by Neural Network Processing Assist (NNPA).

The integrated AI accelerator delivers more than 6 TFLOPs per chip and over 200 TFLOPs in the 32-chip system. The AI accelerator is shared by all cores on the chip. The firmware, running on the cores and accelerator, orchestrates and synchronizes the execution on the accelerator.

Guarded Storage Facility

Guarded Storage Facility (GSF) is a hardware feature that Java uses to achieve pause-less garbage collection. GSF was introduced to enable enterprise-scale Java applications to run without an extended pause for garbage collection on larger heaps. This facility improves Java performance by reducing program pauses during Java garbage collection.

Simultaneous multithreading

Simultaneous multithreading (SMT) is built into IBM z16 Integrated Facilities for Linux (IFLs), IBM Z integrated Information Processors (zIIPs), and System Assist Processors (SAPs). SMT enables more than one thread to run simultaneously in the same core and share all its resources. This function improves the usage 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 must wait for the data to be fetched before it can continue to run, cache misses affect the performance and capacity of the core to run instructions. When using SMT, 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.

Hardware decimal floating point function

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

Vector Packed Decimal Facility

Vector Packed Decimal Facility enables packed decimal operations to be performed in registers rather than in memory by using new fast mathematical computations. Compilers, such as Enterprise COBOL for z/OS V6.2, Enterprise PL/I for z/OS V5.2, and IBM XL C/C++ V2.5 for z/OS V2R5, are optimized on IBM z16.

Single-instruction, multiple-data

IBM z16 includes a set of instructions that is called *single-instruction, multiple-data* (SIMD) that can improve the performance of complex mathematical models and analytics workloads. This improvement is realized through vector processing and complex instructions that can process a large volume of data by using a single instruction.

SIMD is designed for parallel computing and can accelerate code that contains integer, string, character, and floating-point data types. This system enables better consolidation of analytics workloads and business transactions on the IBM Z platform.

Runtime Instrumentation Facility

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

Secure Execution for Linux

Secure Execution for Linux isolates and protects Kernel-based Virtual Machine (KVM) guests from hypervisor access. The hypervisor administrator can still manage and deploy workloads, but cannot view data on a guest. Multiple tenants (applications) running in an LPAR as second-level guests under KVM have fully isolated environments, which help protect intellectual property and proprietary secrets.

5.1.3 Memory

System memory is one of the core design components. Its continuous enhancements contribute to the overall system performance improvements.

Maximum physical memory size is directly related to the number of central processor complex (CPC) drawers in the system. An IBM Z server has more memory that is installed than was ordered because a portion of the installed memory is used to implement the redundant array of independent memory (RAIM) design (the technology that provides memory protection and excludes memory faults). You are *not* charged for the extra amount of memory that is needed for RAIM.

For example, with IBM z16 A01, you have up to 40 TB of memory for a four-CPC drawer configuration (25% increase per drawer compared to IBM z15 T01). With IBM z16 A02 and IBM z16 AGZ, up to 16 TB (8 TB x two CPC drawers) of configurable memory is available, which is the same as IBM z15 T02.

The IBM z16 A01 Hardware System Area (HSA) is 256 GB (same as IBM z15 T01). The IBM z16 A02 and IBM z16 AGZ HSA is 160 GB. The HSA area is a fixed size and not included in the memory that is ordered.

Important: z/OS V2R3 and later require a minimum of 8 GB of memory (2 GB of memory when running under IBM z/VM). z/OS V2R5 can support up to 16 TB of memory in an LPAR.

Each operating system can allocate the amount of main memory up to its supported limit. The amount of incremental memory is 1 GB with IBM z16.

Flexible memory

Flexible memory[2] provides the extra physical memory that is needed to support the following scenario: You need activation base memory and HSA on an IBM z16 A01 that has multiple CPC drawers and one CPC drawer that is out of service.

On IBM z16 A01, the extra resources that are required for flexible memory configurations are provided when you configure memory features and memory entitlement. Flexible memory ranges 512 GB - 30464 GB, depending on the feature ordered (Max82, Max125, Max168, or Max200).

[2] Flexible memory features are not available on IBM z16 A02 and IBM z16 AGZ.

Virtual Flash Memory

The Virtual Flash Memory (VFM) feature is offered from the main memory capacity. For IBM z16 A01, up to 12 VFM features (four VFM features for IBM z16 A02 and IBM z16 AGZ) can be ordered, each being 512 GB.

VFM can improve availability by reducing latency from paging delays that can occur during peak workload periods. It also is designed to help eliminate delays that can occur when diagnostic data is collected during failures.

VFM also can be used in Coupling Facility (CF) images to provide extended capacity and availability for workloads that use IBM MQ Shared Queues structures.

5.2 Virtualization

Virtualization is a key strength of the IBM Z platform, which is embedded in the hardware, firmware, and operating systems. All computing resources (such as CPU, memory, and I/O) are virtualized. Each set of the resources can be used independently within separate operating environments (known as *guest systems*).

The IBM Z platform is designed to concurrently run multiple virtual guest systems and provide each system with the required dynamic sharing of the resources with minimal costs and performance impact.

LPAR technology was introduced more than four decades ago on the IBM Z platform to support virtualization and provide the highest level of isolation between guest systems.

Virtualization management, which is called a *hypervisor*, operates on hardware and software levels on the IBM Z platform. The hardware hypervisor (first level or type 1) is PR/SM, which is integrated into the firmware. PR/SM runs the control code that manages the hardware resources and builds LPARs that run operating systems, middleware, and software applications.

The supported software hypervisors (type 2) are z/VM and KVM:

- z/VM supports the simultaneous execution of multiple virtual guest systems within an LPAR and nested (multi-level) virtualization. z/VM is a powerful hypervisor that can emulate various hardware devices to its guests.
- KVM provides flexibility for your hypervisor choice and the unique combination of enterprise hardware and open source.

z/VM and KVM interconnect with PR/SM and use its functions.

Multiple hypervisors can coexist and run simultaneously on the IBM Z platform so that you can create and build multiple virtualized guest systems that are running various open-source applications on the IBM Z platform with high levels of performance and integrated security.

IBM Z development is continuously improving virtualization techniques.
IBM Z development provides highly scalable dynamic platforms that can host traditional and modern solutions (such as cloud, blockchain, and containers) on the same footprint.

Figure 5-1 shows the diverse workload that is supported and virtualized on a single IBM z16, and the co-existence of multiple hypervisors: PR/SM, z/VM, and KVM.

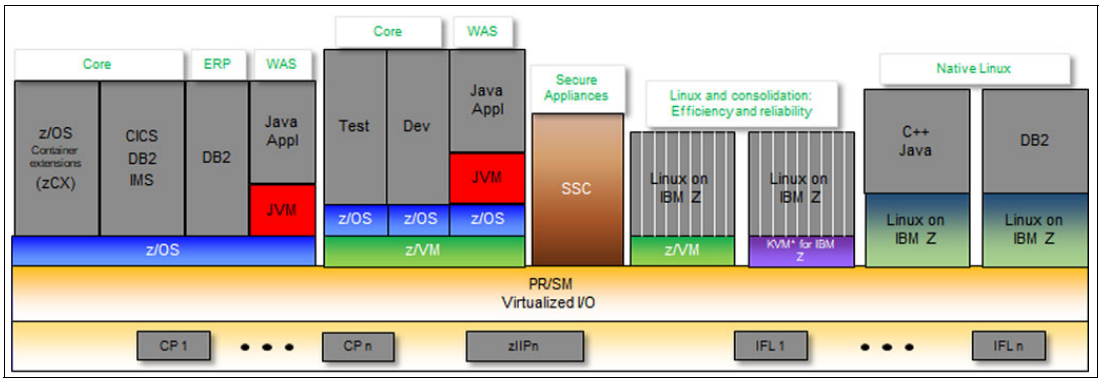

Figure 5-1 Virtualization on the IBM Z platform

Processor Resource/System Manager

PR/SM is a hardware-level hypervisor that is unique to and a vital component of the IBM Z platform.

PR/SM manages and partitions all the computing resources (CPU, memory, and I/O) among the various systems that run on the IBM Z platform. PR/SM provides each system with a required share of these resources and dynamically adjusts the share depending on the workload priority. PR/SM is integrated into the platform and runs transparently to the operating systems and applications.

Each operating system runs in its own LPAR that is managed by PR/SM, which enables a high degree of virtualization.

The goal of PR/SM is to allocate and assign (and reassign) resources to an LPAR so that a workload can achieve its best performance and throughput. Depending on the type of the workload, each LPAR can be defined as z/VM, Linux on IBM Z, z/OS, z/TPF, z/VSE[3] or 21st Century Software VSE^{n} V6.3, and runs one of these operating systems.

Initially, you might allocate a set of resources (CPU, memory, and I/O) and their quantity to the LPAR. Then, PR/SM might dynamically adjust the amount of the resources, according to the defined set of priorities. As a result, the most critical and important workload can complete within a required timeline.

PR/SM evolved over the decades on the IBM Z platform. It is a proven, secure, and fundamental IBM Z technology. Every generation of the IBM Z platform brings PR/SM improvements that are aimed to demonstrate even better system performance. With IBM z16, PR/SM is extended to support new features, such as LPAR support of IBM System Recovery Boost (SRB).

For more information about SRB features, see “SRB” on page 86.

Dynamic Partition Manager

Dynamic Partition Manager (DPM) is an infrastructure management component of the Hardware Management Console (HMC). It simplifies virtualization and configuration management. DPM is designed to be easy to use, especially for users who are new to the IBM Z platform. It does not require you to learn complex syntax or command structures.

[3] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

DPM provides partition lifecycle and integrated dynamic I/O and Peripheral Component Interconnect Express (PCIe) functions management for Linux on IBM Z that is running in an LPAR, under the z/VM or KVM. By using DPM, an environment can be created, provisioned, and modified without disrupting running workloads. It also can be monitored for troubleshooting.

Configuration note: IBM z16 can be configured in DPM mode or in PR/SM mode. It cannot be configured in both modes concurrently. DPM supports Linux on IBM Z, z/VM (only with Linux on IBM Z guests), and KVM. PR/SM mode supports a mixture of z/VM, KVM, 21st Century Software VSE^{n} V6.3, z/OS, Linux on IBM Z, z/TPF, and z/VSE[a].

a. z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

DPM provides the following capabilities through the HMC:

- Create and provision an environment (including new partitions), assign processors and memory, and configure I/O adapters to connect the system to network and storage. DPM supports the following storage and network connectivity features:
  - OSA Express
  - FICON Express (FICON channel and Fibre Channel Protocol (FCP) modes)
  - Crypto Express
  - RDMA over Converged Ethernet (RoCE) Express
  - SMC-Dv2
  - HiperSockets
  - NVMe-attached solid-state drives (SSDs)
- DPM supports the definition of FICON CTC connections between partitions running on the same CPC[4]. The main exploiter of this capability is z/VM.
  - FICON CTC support is a prerequisite for establishing z/VM Single System Image (SSI) clustering. DPM FICON CTC support allows customers to create a z/VM SSI environment within the same CPC. SSI Clustering technology is required in order for customers to exploit Live Guest Relocation (LGR) to move running Linux guests across z/VM images.
- 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.

With its intuitive user interface, DPM also exposes its capabilities through Web Services APIs, which enable the integration of the system into cloud-like management infrastructures.

DPM on IBM z16 enables the use of Shared Memory Communications (SMC)-D V2, which provides a high-performance and low-latency interconnect between the workloads that are running in LPARs within the platform.

z/VM is a native IBM Z operating system that provides virtualization services. It is a software hypervisor (type 2).

z/VM is a powerful hypervisor, and historically is the first virtual machine (VM). z/VM runs in an LPAR and manages the system hardware resources (CPU, memory, and I/O) among its guest systems efficiently.

z/VM supports z/OS, Linux on IBM Z, z/TPF, 21st Century Software VSE^{n} V6.3, and z/VSE[5] as its guest systems. It can also enable nested virtualization and host z/VM as a guest

[4] DPM FICON CTC support is available with GA 1.5.

system, enabling highly virtualized complex infrastructures for supporting containerized and cloud workloads.

z/VM can emulate and virtualize different hardware devices, such as virtual tape, and provide it to the operating systems that run under its management. z/VM is tightly coupled with PR/SM and uses its functions for the most optimized workload deployment.

z/VM is a proven, enterprise-grade hypervisor that can scale out horizontally and vertically. IBM z16 A01 supports up to 85 z/VM LPARs, while IBM z16 A02 and IBM z16 AGZ support up to 40 z/VM LPARs. Each z/VM LPAR can run thousands of guest systems.

KVM hypervisor

The KVM hypervisor is available in recent Linux on IBM Z distributions. It is a type 2 hypervisor that provides simple, cost-effective platform virtualization for Linux workloads that are running on the IBM Z platform. It enables you to share real CPUs (called IFLs), memory, and I/O resources through PR/SM.

KVM can coexist with other operating systems (Linux on IBM Z, z/OS, z/VM, z/VSE[6], 21st Century Software VSE^n V6.3, and z/TPF) that are running in different LPARs on the IBM Z platform.

The KVM hypervisor support information is provided by the Linux distribution partners. For more information, see the documentation for your distribution.

For more information about the use of KVM with IBM Z, see *Virtualization Cookbook for IBM Z Volume 5: KVM*, SG24-8463.

5.2.1 Hardware virtualization

The IBM Z platform is known for its unique virtualization capabilities. It enables you to deploy various workloads (traditional and modern) to achieve the highest performance and throughput metrics with the lowest costs and impact.

Workload separation is one of the most important parameters. PR/SM in IBM z16 is designed to meet the highest level of Common Criteria (EAL5+), similar to previous IBM Z platforms. This level of isolation ensures the integrity and security of the workloads and excludes the contamination of the running applications by other programs.

Logical processors

All physical PUs are virtualized as logical processors on the IBM Z platform and can be characterized as the following types:

- Central processors (CPs) are standard processors that support all operating systems and user workloads.
- A zIIP is used under z/OS for designated workloads. These workloads include, but are not limited to, the following examples:
  - IBM z/OS Container Extensions (zCX)
  - IBM Java virtual machine (JVM)
  - Various XML System Services
  - IPsec offload

[5] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.
[6] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

- Certain IBM Db2 for z/OS processes
- DFSMS System Data Mover for z/OS Global Mirror
- IBM HiperSockets for large messages
- IBM GBS Scalable Architecture for Financial Reporting (SAFR) enterprise business intelligence reporting
- IBM SRB

- IFL processor is used exclusively for Linux on IBM Z and for running the z/VM and KVM hypervisors in support of Linux VMs.
- An Internal Coupling Facility (ICF) processor is used for z/OS clustering and supporting the family of IBM Parallel Sysplex solutions. ICF is dedicated to this function and exclusively runs the Coupling Facility Control Code (CFCC).

The characterized PUs are aimed to streamline the specific workload. All engines architecturally and physically are the same.

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

- SAPs that run I/O operations
- Integrated Firmware Processors (IFPs) for native PCIe features and other firmware functions

PR/SM accepts requests for work by dispatching logical processors on physical processors. Physical processors can be shared across LPARs or dedicated to an LPAR. The logical processors that are assigned to an LPAR must be either all shared or all dedicated.

PR/SM ensures that the processor state is correctly saved and restored (including all registers) when you switch a physical processor from one logical processor to another one. Data isolation, integrity, and coherence inside the system are always strictly enforced.

Logical processors can be dynamically added to and removed from LPARs. Operating system support is required to use this capability. z/OS, z/VM, 21st Century Software VSE^{n} V6.3, and z/VSE[7] each can dynamically define and change the number and type of reserved PUs in an LPAR profile. No planning is required.

The newly assigned logical processors are immediately available to the operating systems and for z/VM to its guest images. Linux on IBM Z provides the Standby CPU activation and deactivation functions.

Memory

To ensure security and data integrity, memory cannot be concurrently shared by active LPARs. Strict LPAR isolation is maintained to avoid any workload contamination.

An LPAR can be defined with an initial and reserved amount of memory. At activation time, the initial amount is made available to the partition, and the reserved amount can be added later partially or totally. Those two memory zones do not have to be contiguous in real memory, but the addressing area (for initial and reserved memory) is presented as contiguous to the operating system that runs in the LPAR.

z/VM can acquire memory nondisruptively and quickly make it available to guests. z/VM virtualizes this support to its guests, which also can increase their memory nondisruptively. Releasing memory is still a disruptive operation.

[7] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

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 features 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 conversion hardware mechanism. A program's right to read/write in each page frame is validated by comparing the page key with the key of the program that is requesting access.

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.

IBM z16 supports 4 KB, 1 MB, and 2 GB pages, and an extension to the z/Architecture that is called Enhanced Dynamic Address Translation-2 (EDAT-2).

Operating systems can enable 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 VMs that use that operating system. Known as *discontiguous shared segment* (DCSS), this shared memory use for many VMs can result in significant savings of real memory and improvements in performance.

I/O virtualization

IBM z16 A01 supports six channel subsystems (CSSs), and IBM z16 A02 and IBM z16 AGZ support three CSSs. Each CSS can have up to 256 channels. In addition to the dedicated use of channels and I/O devices by an LPAR, the z/Architecture also enables sharing of 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 CSSs, 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 includes the channel configured online, PR/SM establishes one logical channel path. For availability reasons, multiple logical channel paths must be available for critical devices (for example, disks that contain vital data sets).

When more isolation is required, you can use configuration rules to restrict the access of each LPAR to specific channel paths and specific I/O devices on those channel paths.

Many installations use the parallel access volume (PAV) function, which enables access to a device through 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, which lowers the need for manually tuning the system for PAV use.

SuperPAV is an extension of the HyperPAV architecture and implements multiple logical subsystems (LSSs) within an alias management group (AMG). SuperPAV enables the following solution:

- Problem: A new I/O request occurs and no alias PAV devices are available in the alias pool for the base PAV device's LSS.
- Solution: z/OS attempts to use an alias PAV device from another LSS within the AMG subgroup.

SuperPAV can provide relief for systems that experience high I/O queue time (IOSQ) during periods of peak I/O load. When few aliases are defined in an LSS, aliases might not be available during a heavy I/O period. z/OS checks peer LSS alias pools to borrow an alias to start I/O requests. Previously, these I/O requests were left queued when aliases are not available.

In large installations, the total number of device addresses can be high. Therefore, the concept of *subchannel sets* is part of the z/Architecture.

Subchannel sets

With IBM z16 A01, up to four subchannel sets of approximately 64,000 device addresses are available. The base addresses[8] are defined to set 0 (IBM reserves 256 subchannels on set 0), and the aliases addresses are defined to set 1, set 2, and set 3.

With IBM z16 A02 and IBM z16 AGZ, up to three subchannel sets of approximately 64,000 device addresses are available. The base addresses are defined to set 0 (IBM reserves 256 subchannels on set 0) and the aliases addresses are defined to set 1 and set 2.

Subchannel sets are used by the Metro Mirror (also referred to as *synchronous Peer-to-Peer Remote Copy* (PPRC)) function by having the Metro Mirror primary devices that are defined in subchannel set 0. Secondary devices can be defined in subchannel sets 1, 2, and 3, which provide more connectivity through subchannel set 0.

To reduce the complexity of managing large I/O configurations further, the IBM Z platform introduced extended address volumes (EAVs). EAVs provide large disk volumes. In addition to z/OS, z/VM and Linux on IBM Z support EAVs.

By extending the disk volume size, potentially fewer volumes are required to hold the same amount of data, which simplifies systems and data management. EAV is supported by the IBM Storage DS8000 series.

The dynamic I/O configuration function is supported by z/OS and z/VM. This function provides the capability of concurrently changing the active I/O configuration. Changes can be made to channel paths, control units, and devices. A fixed HSA area in IBM z16 greatly eases the planning requirements and enhances the flexibility and availability of these reconfigurations.

[8] Only a z/OS base device must be in subchannel set 0. Linux on IBM Z supports base devices in the other subchannels sets.

5.2.2 Hybrid cloud environments

Virtualization is critical to the viability of cloud service offerings because it provides the elasticity that allows a platform to deal with the ebbs and flows of demands on IT resources. Because of the integration in the hardware and firmware, virtualization on IBM z16 is highly efficient. It delivers up to 100% sustained resource use, 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 IBM z16 granularity offerings, including capacity levels, Tailor Fit Pricing (for unpredictable, high spiking, and business-critical workloads), and Capacity on Demand (CoD). These characteristics and other technical leadership characteristics make the IBM Z platform the gold standard for the industry.

In addition, managing a cloud environment requires tools that can leverage 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* must enable the management of virtualized IT resources to support different types of cloud service models and cloud deployment models. OpenStack can satisfy a wide range of cloud management demands. OpenStack integrates various components to automate IT infrastructure service provisioning.

IBM z16 also can be tailored with a choice of IBM Z backed services that are delivered through IBM Cloud to help transform your infrastructure, applications, and data by exposing and connecting assets with simplified and intelligent operations across the infrastructure.

With Red Hat, the hybrid cloud capabilities on the IBM Z platform are extended. Support for running Red Hat OpenShift and IBM zCX Foundation for Red Hat OpenShift on Linux on IBM Z provides expansive cloud capabilities, including open containers, tools, and access to an extensive open community.

The new cloud-native capabilities are delivered as pre-integrated solutions that are called *IBM Cloud Paks*. This IBM certified and containerized software provides a common operating model and a common set of services.

For more information about hybrid cloud capabilities, see Hybrid cloud with IBM Z.

5.3 Capacity and performance

IBM z16 offers significant increases in capacity and performance over IBM z15. Several elements contribute to this effect, including many processors, individual processor performance, memory caches, SMT, and machine instructions, including the SIMD. Subcapacity settings continue to be offered.

Note: Capacity and performance ratios are based on measurements and projections by using standard IBM benchmarks in a controlled environment. Actual throughput can vary depending on several factors, such as the job stream, I/O and storage configurations, and workload type.

5.3.1 Capacity settings

IBM z16 expands the offer on subcapacity settings. Finer granularity in capacity levels enables the growth of installed capacity to follow more closely the enterprise growth for a smoother, pay-as-you-go investment profile. Many performance and monitoring tools are available on IBM Z environments that are coupled with the flexibility of the CoD options (see 5.3.2, “Capacity on Demand offerings” on page 80). These features help to manage growth by making capacity available when needed.

IBM z16 A01 capacity levels

Regardless of the installed model, IBM z16 A01 offers four distinct capacity levels for the first 39 CPs:

- 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[9] to a subcapacity of the full CP capacity. The clock frequency of those processors remains unchanged. The capacity adjustment is achieved by using other means.

[9] The CP is the standard processor for use with any supported operating system. It is required to run z/OS.

Full capacity CPs are identified as CP7. On IBM z16 A01, up to 200 CPs can be configured as CP7. Up to 39 CPs can have subcapacity. The three subcapacity levels are identified by CP6, CP5, and CP4, and are displayed in hardware descriptions as Feature Codes on the CPs.

If more than 39 CPs are configured to the system, all must be full capacity because all CPs must be at the same capacity level. Granular capacity adds 117 subcapacity settings to the 200 capacity settings that are available with full capacity CPs (CP7). The 317 distinct capacity settings in the system provide for a range of 1:758 in processing power.[10]

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

On IBM z16 A01, the following CP subcapacity levels are a fraction of full capacity:

- Model 7xx = 100% (2253 PCI)
- Model 6xx = 66% (1496 PCI)
- Model 5xx = 41% (937 PCI)
- Model 4xx = 12% (280 PCI)

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 200 ICFs or IFLs (which always run at full capacity) is possible.

To help size an IBM Z server to fit your requirements, IBM provides a no-cost tool that reflects the latest IBM Large Systems Performance Reference (LSPR) measurements, and it is called the IBM Z Processor Capacity Reference (zPCR). You can download the tool from Getting Started with zPCR.

For more information about LSPR measurements, see LSPR for IBM Z.

IBM z16 A02 and IBM z16 AGZ capacity levels

IBM z16 A02 and IBM z16 AGZ offer 26 distinct capacity levels for each CP in the configuration for a total of 156 (26 x 6) capacity settings. These processors deliver the scalability and granularity to meet the needs of small and medium-sized enterprises.

As in IBM z16 A01, IBM z16 A02 and IBM z16 AGZ PUs always are set to full capacity when the PUs are characterized as anything other than a CP, such as a zIIP, an IFL, or an ICF.

Figure 5-2 on page 80 shows details about IBM z16 A02 and IBM z16 AGZ Capacity Levels A01 - Z06.

[10] This value is calculated by dividing the highest capacity (200-way full speed) of 212,222 Peripheral Component Express (PCI) by the lowest subcapacity of 280 PCI (212,222 ÷ 280 = 758).

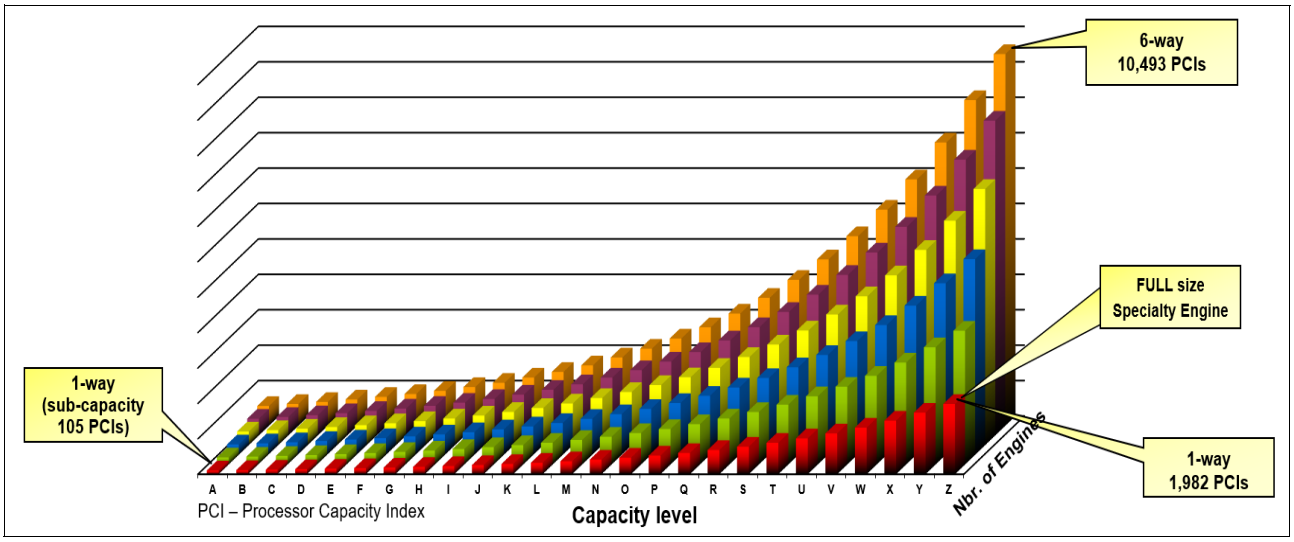

Figure 5-2 IBM z16 A02 capacity levels chart

5.3.2 Capacity on Demand offerings

IBM z16 continues to provide CoD offerings. They provide flexibility and control and ease the administrative burden in the handling of the offerings. They also give finer control over resources that are needed to meet the resource requirements in various situations.

IBM z16 can perform concurrent upgrades, which provide an increase of processor capacity with no platform outage. In most cases, with operating system support, a concurrent upgrade also can be nondisruptive to the operating system. It is important to consider that these upgrades are based on the enablement of resources that are physically present in IBM z16.

Capacity upgrades cover permanent and temporary changes to the installed capacity. The changes can be done by using the Customer Initiated Upgrade (CIU) facility without requiring the involvement of IBM service personnel. Such upgrades are started through the web by using IBM Resource Link.

Using the CIU facility requires a special contract between the customer and IBM. This contract specifies the terms and conditions for online CoD buying of upgrades, and other types of CoD upgrades are accepted. For more information, see IBM Resource Link.[11]

For more information about the CoD offerings, see *IBM z16 (3931) Technical Guide*, SG24-8951.

Permanent upgrades

Permanent upgrades of processors (CPs, IFLs, ICFs, zIIPs, and SAPs) and memory, or changes to a platform's Model-Capacity Identifier (up to the limits of the installed processor capacity on an existing IBM z16), can be performed by customers through the IBM Online Permanent Upgrade offering by using the CIU facility.

[11] IBM Resource Link ID is required.

Temporary upgrades

Temporary upgrades of an IBM z16 can be done by On/Off CoD, Capacity Backup (CBU), or Capacity for Planned Event (CPE) that is ordered from the CIU facility.

On/Off CoD function

On/Off CoD is a function that enables concurrent and temporary capacity growth of the CPC. On/Off CoD can be used for peak workload requirements for any period. On/Off CoD includes a daily hardware charge, and it can include an associated software charge. On/Off CoD offerings can be prepaid or post-paid.

When you use the On/Off CoD function, you can concurrently add processors (CPs, IFLs, ICFs, zIIPs, and SAPs), increase the CP capacity level, or both.

Prepaid On/Off CoD tokens: New prepaid On/Off CoD tokens do *not* carry forward to future systems.

Capacity Backup function

CBU enables you to perform a concurrent and temporary activation of extra CPs, ICFs, IFLs, zIIPs, and SAPs, an increase of the CP capacity level, or both. This function can be used during an unforeseen loss of IBM Z capacity within the enterprise, or to perform a test of your disaster recovery (DR) 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 as unused PUs to increase the CP capacity level on a subcapacity system, or both. A CBU contract must be in place before the Licensed Internal Code Configuration Control (LICCC) 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) and one disaster activation (up to 90 days). The record can be configured to be valid for up to 5 years. You also can order more tests for a CBU record in quantities of five tests up to a maximum of 15.

Suitable use of the CBU capability does not incur any other software charges from IBM.

Flexible Capacity for Cyber Resiliency

Flexible Capacity (Feature Code 9933 together with Feature Code 0376) dynamically shifts production capacity across sites between IBM z16 machines for flexibility and elasticity in DR testing, planned maintenance, proactive outage avoidance, and actual DR scenarios.

Flexible Capacity works with other temporary record types, for example, On/Off CoD and Tailor Fit Pricing for Hardware.

With Flexible Capacity, capacity can be transferred remotely without requiring onsite personnel (IBM or customer) after the initial set-up. Flexible Capacity provides the following benefits:

- Flexible duration of capacity transfer (up to 1 year)
- Fully automated by using solutions such as IBM GDPS®
- Simplify compliance and improve confidence DR scenarios, including test
- No need for CBU records by using this solution

For more information about IBM Z Flexible Capacity for Cyber Resiliency, see *IBM Z Flexible Capacity for Cyber Resiliency*, REDP-5702.

Capacity for Planned Event function (carry forward only)

With CPE, you can perform a concurrent and temporary activation of extra CPs, ICFs, IFLs, zIIPs, and SAPs, an increase of the CP capacity level, or both. This function can be used during a planned outage of IBM Z capacity within the enterprise (for example, data center changes or system or power maintenance). CPE cannot be used for peak workload management and can be active for a maximum of 3 days.

The CPE feature is optional and requires unused capacity to be available on CPC drawers of the backup system as unused PUs to increase the CP capacity level on a subcapacity system, or both. A CPE contract must be in place before the LICCC that enables this capability can be loaded on the system.

z/OS capacity provisioning

Capacity provisioning helps you to manage the CP and zIIP capacity of an IBM z16 that is running one or more instances of the z/OS operating system. By using the 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 IBM z16 provisioning capability gives you a flexible, automated process to control the configuration and activation of On/Off CoD offerings.

System Recovery Boost Upgrade

System Recovery Boost Upgrade can optionally enable the temporary activation of more physical zIIP processors on IBM z16 A01, but only by using a priced boost temporary capacity record. This record requires Feature Codes 6802 and 9930.

Note: SRB functions are embedded in the IBM z16 firmware. These functions can be used without ordering extra zIIP capacity (Feature Code 6802 and Feature Code 9930). For more information about more SRB options, see System Recovery Boost under 5.4, “Reliability, availability, and serviceability” on page 86.

You can temporarily increase the number of physical zIIPs to use for SRB. This feature is the new System Recovery Boost Upgrade record that you can activate from the **HMC/SE Perform Model Conversion** menu, or by using automation that calls the hardware API services.

Activation of these processors uses unused processing cores in IBM z16 to provide more zIIP processing capacity that accelerates execution of their workload (general-purpose workload and workload that is zIIP-eligible). After the boost temporary capacity record is activated for use during maintenance (for example, a planned maintenance window or for a planned site-switch activity), up to 20 other zIIP engines can become available for up to 6 hours for use on IBM z16. This extra zIIP capacity is shared across images in accordance with PR/SM management controls, which makes more zIIP capacity available to individual system images.

Images that want to use this extra zIIP capacity predefine reserved logical zIIP capacity in their PR/SM image profiles. Therefore, the operating system can bring those extra logical zIIP processors (with physical backing from the added physical zIIPs that were activated) online for use during the boost period. This configuration provides the image with increased zIIP capacity and parallelism to accelerate the workload.

IBM Tailored Fit Pricing for IBM Z

The new Tailored Fit Pricing for IBM Z Hardware Consumption Solution provides an always-on, consumption-based capacity corridor that provides hybrid cloud flexibility and control for unpredictable workload spikes throughout the day. It helps to scale IT demands and control behavior with a pay-for-use buffer and granular usage measurements.

For more information about the CoD offerings, see *IBM z16 (3931) Technical Guide*, SG24-8951 and I*BM z16 (3932) Technical Guide,* SC24-8952.

5.3.3 IBM z16 performance

The IBM Z microprocessor chip for IBM z16 has a high-frequency design that uses IBM leading microprocessor 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 IBM z16 A01 is a result of the enhancements that are described in Chapter 2, “IBM z16 A01 hardware overview” on page 17, Chapter 3, “IBM z16 A02 and IBM z16 AGZ hardware overview” on page 33, and 5.1, “Technology improvements” on page 66.

LSPR workload suite: IBM z16 changes

To help you better understand workload variations, IBM provides a no-cost tool, zPCR, which is available at the IBM Presentation and Tools website.

IBM continues to measure the 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 millions of service units (MSU) ratings for reference.Capacity performance is closely associated with how a workload uses and interacts with a specific processor hardware design. Workload capacity performance is sensitive to the following major factors:

- Instruction path length
- Instruction complexity
- Memory hierarchy

The CPU Measurement Facility (CPUMF) data offers insight into the interaction of workload with the hardware design. CPUMF 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. The workload category is determined by the Level 1 Misses per 100 instructions (L1MP) and the *relative nest intensity* (RNI).

With the IBM Z platform, the LSPR introduced the following workload capacity categories that replace all prior primitives and mixes:

- LOW (RNI): A workload category that represents light use of the memory hierarchy.
- AVERAGE (RNI): A workload category that represents average use of the memory hierarchy. This category is expected to represent most production workloads.
- HIGH (RNI): A workload category that represents heavy use of the memory hierarchy.

These categories are based on the RNI, 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. CPU MF data can be collected by z/OS SMF record type 113 or z/VM Monitor.

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 that are 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 that are used to set the single-number values, the LSPR tables contain information that pertains to Linux on IBM Z and z/VM environments.

The LSPR contains the internal throughput rate ratios (ITRRs) for IBM z16 and the previous generations of processors. The report is based on measurements and projections by using standard IBM benchmarks in a controlled environment. The actual throughput that any user might experience varies depending on several factors, such as the amount of multiprogramming in the job stream, the I/O configuration, and the workload that is processed.

Experience demonstrates that IBM Z servers can run at up to 100% usage levels sustained. However, most users prefer to leave some white space and run at 90% or slightly under. For any capacity comparison, the use of “one number” (such as the MIPS or MSU metrics) is not a valid method. Therefore, use zPCR and involve IBM Support when you are planning for capacity.

For more information about IBM z16 performance, see *IBM z16 (3931) Technical Guide*, SG24-8951.

Throughput optimization with IBM z16

The memory and cache structure implementation in the CPC drawers of IBM z16 were significantly enhanced compared to previous generations to provide sustained throughput and performance improvements. The memory is distributed throughout the CPC drawers and the CPC drawers have individual levels of cache that are private to the cores and shared by the cores. Nonetheless, all processors can access the highest level of cache and all the memory. Therefore, the system is managed as a memory-coherent symmetric multiprocessor (SMP).

Processors within the IBM z16 CPC drawer structure have different distance-to-memory attributes. CPC drawers are fully interconnected 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 redispatch work to a group of physical CPUs that share cache levels.

PR/SM manages the usage of physical processors by LPARs by dispatching the logical processors on the physical processors. However, 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 enable PR/SM and WLM to work closely together. They can cooperate to create an affinity between task and physical processor rather than between LPAR and physical processor, which is known as *HiperDispatch*.

HiperDispatch

HiperDispatch combines two functional enhancements: in the z/OS dispatcher and in PR/SM. This function is intended to improve computing efficiency in the hardware, z/OS, and 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 IBM z16, PR/SM attempts to group the logical processors into the same logical cluster or in the neighboring logical cluster in the same CPC drawer and, if possible, in the same chip. This configuration results in reduction of multi-processor effects, maximizing the usage 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, which improve 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 IBM z16. The entire IBM z16 stack (hardware, firmware, and software) tightly collaborates to obtain the full potential of the hardware. z/VM HiperDispatch provides support similar to the one in z/OS. It is possible to dynamically turn on and off HiperDispatch without requiring an initial program load (IPL).

IBM z16 includes several HiperDispatch algorithm enhancements, including the following examples:

- Improved memory affinity
- Improved LPAR placement based on IBM z15 experience
- Usage of a new chip configuration

Note: HiperDispatch is required if SMT is enabled. All IBM Z LSPR measurements are provided for z/OS environments with HiperDispatch on. A best practice is to turn on HiperDispatch for production workloads.

5.4 Reliability, availability, and serviceability

The IBM Z platform is known for its reliability, availability, and serviceability (RAS) capabilities. RAS is built in to the hardware and software stacks of the IBM z/Architecture, where mean time between failures (MTBF) is measured in decades. The RAS strategy is to manage change by learning from previous generations of IBM Z servers and investing in new RAS functions to eliminate or minimize all sources of outages.

The IBM Z family presents numerous enhancements in 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 Licensed Internal Code (LIC) driver updates that are performed to support new features and functions. Firmware updates can be performed by IBM remotely through Remote Code Load upgrades. Enhanced Driver Maintenance (EDM) can help reduce the necessity and eventual duration of a scheduled outage.

When suitably configured, IBM z16 can concurrently activate a new LIC Driver level. Concurrent activation of the select new LIC Driver level is supported at released synchronization points. However, a concurrent update or upgrade might not be possible for specific LIC updates.

IBM z16 builds on the RAS characteristics of the IBM Z family, with the following RAS improvements:

- SRB

  SRB is delivered with IBM z16 for maximize service availability by using tailored short-duration boosts to mitigate the impact of these recovery processes:

  - The z/OS SAN Volume Controller (SVC) dump boost boosts the system on which the SVC dump is taken to reduce the system impact and expedite diagnostic capture. It is possible to enable, disable, or set thresholds for this option.
  - Middleware restart, or recycle boost, boosts the system on which a middleware instance is being restarted to expedite resource recovery processing, release retained locks, and so on. This boost applies to planned restarts or restarts after failure, automated, or ARM-driven restarts. SRB does not boost any system address spaces by default and must be configured by the WLM policy specification.
  - The HyperSwap configuration load boost boosts the system in which the HyperSwap configuration and policy information are being loaded or reloaded. This boost applies to Copy Services Manager (CSM) and GDPS. HyperSwap Configuration Load boost is enabled by default. No thresholds or criteria are applied to the boost request based on the size or number of devices that is present in the HyperSwap configuration.

  Through SRB, IBM z16 continues to offer more CP capacity during particular system recovery operations to accelerate system (operating system and services) start when the system is started or shut down. SRB is operating system-dependent. No other hardware, software, or maintenance charges are required during the boost period for the base functions of SRB.

  SRB can be used during LPAR IPL or LPAR shutdown to make the running operating system and services available in a shorter period.

SRB provides the following options for the capacity increase:

- Subcapacity Boost: During the boost period, subcapacity engines are transparently activated at their full capacity (for CP engines only).
- zIIP Capacity Boost: During the boost period, all active zIIPs that are assigned to an LPAR are used to extend the CP capacity.

For more information about the optional temporary capacity upgrade, see “System Recovery Boost Upgrade” on page 82.

SRB can also be used for recovery processes, such as IBM HyperSwap[12] by using a Recovery Process Boost. Recovery Process Boosts are short-term accelerations for specific sysplex recovery events in z/OS. Sysplex recovery events often cause short-duration workload impacts or workload spikes that can affect the normal running of workloads in the sysplex until recovery processing completes.

With Recovery Process Boosts, boosted processor capacity is made available to mitigate short-term recovery impacts and restore normal steady-state sysplex operation as quickly as possible after the recovery events.

Also, boosted processor capacity is provided for a short period after the restoration of steady-state operation for workload “catch-up”.

At the time of writing, the main SRB users are z/OS, z/VM, z/VSE[13], 21st Century Software VSE^n V6.3, and z/TPF. z/VM uses SRB if it runs on subcapacity CP processors only (IFLs are always at their full clock speed). Second-level z/VM guest operating systems can inherit the boost if they are running on CPs.

z/OS that is configured as a guest system under z/VM management does not use the boost. Inheritance of the boost applies only during z/VM workload initialization and shutdowns. Starts and shutdowns of the second-level guests, in isolation from z/VM, are not boosted.

SRB support is available with GDPS 4.2[14] and newer versions through firmware enhancements that support greater parallelism and performance improvements in the hardware API services. These enhancements are used by GDPS to speed up the orchestration of shutdown and restart activities. The boost of CP capacity does not contribute to other software license charges.

For more information, see *Introducing IBM Z System Recovery Boost*, REDP-5563.

- IBM z16 Level 3 and Level 4 cache structures are virtualized, that is, the physical level of shared cache is implemented in a robust SRAM wipe-out correction and sparing design that benefits the virtualization levels. The cache allocation is fluid, dynamically changing the cache size as needed, which benefits processor cores, CBU, and On/Off CoD.
- Two PU chips are packaged in a dual-chip module (DCM), which enhances thermal conductivity and improves reliability. The PU chip uses 7-nm technology and consists of 22.5 billion transistors (compared to 9.1 billion for IBM z15). Up to eight active cores per chip that run at 5.2 GHz on IBM z16 A01 and 4.6 GHz on IBM z16 A02 and IBM z16 AGZ.
- VFM: Flash Express PCIe feature replacement with memory dual inline memory modules (DIMMs), which is a more robust solution that uses RAIM protection against memory faults.

[12] IBM HyperSwap is a high availability (HA) feature that provides dual-site, active-active access to a volume. HyperSwap functions are available on systems that can support more than one I/O group.
[13] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.
[14] SRB does not require GDPS V4.2. GDPS V4.2 provides extra enhancements that are considered part of SRB.

IBM z16 continues to support Enhanced Drawer Availability (EDA), which minimizes the effects of CPC drawer repair and upgrade actions. In a multiple CPC drawer system, a single CPC drawer can be concurrently removed and reinstalled for an upgrade[15] or repair. To ensure that the IBM z16 configuration supports removal of a CPC drawer with minimal effect to the workload, consider the flexible memory option (see “Flexible memory” on page 69).

IBM z16 implements a new RAIM[16] design that provides a method to increase memory availability where a fully redundant memory system can identify and correct memory errors without stopping. The implementation is similar to the RAID concept that has been used in storage systems for several years. For more information about RAS features, see *IBM z16 (3931) Technical Guide*, SG24-8951.

IBM z16 A01 can be configured with of a maximum of four CPC drawers that are designed as a field-replaceable unit (FRU). IBM z16 A02 and IBM z16 AGZ can be configured with a maximum of two CPC drawers. Connections among the CPC drawers are established by using SMP cables. Each CPC drawer consists of four PU DCMs (two DCMs on IBM z16 A02 and IBM z16 AGZ Max5 and Max16) and up to 48 DIMMs (protected by RAIM). In addition to the DCMs and memory, CPC drawers host the oscillators and connectors for I/O and the Server Time Protocol (STP) interface.

IBM z16 inherits I/O infrastructure reliability improvements from IBM z15, including Forward Error Correction (FEC) technology that enables better recovery of FICON channels. FICON Express32S features continue to provide a new standard for transmitting data over 32 Gbps links by using 256b/257b encoding.

The IBM z16 A01 configuration includes an improved front-to-rear radiator-cooling system. The radiator pumps, blowers, controls, and sensors are N+1 redundant. In normal operation, one active pump supports the system. A second pump is turned on and the original pump is turned off periodically, which improves the reliability of the pumps. The replacement of pumps or blowers is concurrent with no effect on performance. IBM z16 A02 and IBM z16 AGZ are air-cooled machines, so they do not use the radiator pumps.

RAS also includes the following enhancements:

- Integrated sparing
- Error detection and recovery improvements in caches and memory
- 25 GbE RoCE Express3 (optics as FRU)
- PCIe coupling links (improved diagnostics)
- Enhanced channel logging
- OSA-Express firmware changes to increase the capability of concurrent maintenance change level (MCL) updates
- System power cycle management (servicing capability)
- CFCC level 25 (various enhancements for improving CF resiliency)
- IBM RMF reporting improvements

IBM z16 continues to support concurrent addition of resources, such as processors or I/O cards, to an LPAR to achieve better serviceability. If another SAP is required on an IBM z16 server (for example, as a result of a DR situation), the SAPs can be concurrently added to the CPC configuration.

---

[15] Adding a fourth drawer to the IBM z16 configuration is not supported in the field (manufacturing only).

[16] Meaney, P.J., et al. “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.

It is possible to concurrently add CP, zIIP, IFL, and ICF processors to an LPAR. This function is supported by z/VM, 21 Century Software VSE^{n} V6.3 and (with suitable program temporary fixes (PTFs)) by z/OS and z/VSE[17]. It is possible to concurrently add memory to an LPAR. This feature is supported by z/OS and z/VM.

IBM z16 supports adding Crypto-Express features to an LPAR dynamically by changing the cryptographic information in the image profiles. Users also can dynamically delete or move Crypto-Express features. This enhancement is supported by z/OS, z/VM, and Linux on IBM Z.

5.4.1 RAS capability for the Support Element and Hardware Management Appliance

The Hardware Management Appliance (HMA) has two redundant, physical 1U rack-mounted servers with increased capacity that run virtual instances of the HMC and the Support Element (SE) on each physical device; therefore, the HMC and SE virtual instances also are redundant.

The two 1U trusted servers are inside the IBM z16 A-frame[18]: one always host the primary SE, and the other host the alternative SE. The primary SE is the active SE. The alternative acts as the backup, and SE information is mirrored once per day. The HMC's virtual instances have peer relationships that enable data replication to be set up.

The HMC's virtual instances can be accessed locally by using the physical KMM display by the service personnel or through a remote web browser. When access to the HMC is granted, the SE virtual instance can be accessed through "Single Object Operations".

Note: Because the SE is running as a guest under the HMC, any HMC Console Restart (restart) is disruptive to the SE running on that HMC.

The HMA's servers offer RAS improvements, such as ECC memory, redundant physical networks for networking requirements, redundant power modules, N+1 redundant power supplies, and better thermal characteristics.

The optional HMA (Feature Code 0129) feature provides HMC functions by using the same physical two 1U servers that run the SE code. The HMA feature provides two HMCs.

5.5 High availability with Parallel Sysplex

The Parallel Sysplex technology is an IBM Z clustering technology that enables users to build a highly resilient, highly scalable, dynamic, and robust IBM Z environment to achieve near-continuous services and application availability. Hardware, middleware, and software tightly cooperate to achieve this result.

Parallel Sysplex is an active-active cluster with up to 32 members (z/OS systems). The underlying structure of the Parallel Sysplex remains transparent to users, networks, and applications.

[17] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

[18] For IBM z16 AGZ, the two SEs are installed right above the last installed component (CPC drawer or I/O drawer) in the client-supplied rack.

The Parallel Sysplex features the following minimum components:

- CF

  The CF is the cluster management center that enables workload distribution, inter-communications, and other system tasks. It can be implemented as an LPAR of a dedicated IBM Z platform, or within the platform running alongside other LPARs. PUs that are characterized as CPs or ICFs are used to configure CF LPARs. ICFs are preferred because of software licensing reasons. Two or more CFs are recommended for HA.

- CFCC

  This IBM LIC runs inside CF (no other code runs there). The code is used to create and maintain the CF structures. These structures are used under z/OS by software components, such as z/OS, Db2 for z/OS, IBM Customer Information Control System Transaction Server (CICS TS) for z/OS, and IBM WebSphere MQ for synchronizing the access to the shared data and resources.

  z/VM can emulate CFCC as a guest VM, which enables users to build a z/OS sysplex that consists of z/OS instances (images). Such a setup is useful for testing and developing purposes, but not suitable for the production environments.

- Coupling links

  These high-speed links interconnect the several system images (each running in its own LPAR) that participate in the Parallel Sysplex. At least two connections between each physical platform and the CF must exist. Internal coupling (IC) links are used when all the system images run on the same physical platform.

  On the software side, z/OS components participate in building the Parallel Sysplex.

  z/OS and CF are highly connected. For example, they provide CF structure duplexing, which is a general-purpose, hardware-assisted mechanism for duplexing structure data that is held in CFs. This function provides a robust recovery mechanism for failures, such as the loss of a single structure on CF or the loss of connectivity to a single CF. The recovery is done through rapid failover to the other structure instance of the duplex pair.

For more information about deploying system-managed CF structure duplexing, see the technical paper *System-Managed CF Structure Duplexing*. The paper is available by clicking **Learn more** at the Parallel Sysplex website.

Normally, two or more z/OS images are clustered to create a Parallel Sysplex. Multiple clusters can span several IBM Z servers, although a specific image (LPAR) can belong to only one Parallel Sysplex.

A z/OS Parallel Sysplex implements shared-all access to data. This configuration is facilitated by IBM Z I/O virtualization capabilities, such as MIF. MIF allows several LPARs to share I/O paths in a secure way, which maximizes usage and greatly simplifies the configuration and connectivity.

A suitably configured Parallel Sysplex cluster is designed to maximize availability at the application level. Quick recovery after a failure is an essential part of sysplex recovery, and that is why SRB boosts sysplex recovery events. Sysplex also provides redundancy so that the workload can continue to run on other systems in the sysplex. Rather than a quick recovery from a failure, the Parallel Sysplex design objective is zero application downtime.

Parallel Sysplex includes the following features:

- Data sharing with integrity

  The CF is key to the implementation of share-all access to data. Every z/OS system image can access all the data. Subsystems in z/OS declare resources to the CF. The CF accepts and manages lock and unlock requests on those resources, which help to ensure the data integrity. A duplicate CF further enhances the availability. Key users of the data-sharing capability are Db2, WebSphere MQ, IBM WebSphere ESB, IBM Information Management System (IMS), and CICS.

  Because these components 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.

- Near-continuous (application) availability

  Changes, such as software upgrades and patches, can be introduced to one image at a time, while the remaining images continue to process the workload. For more information, see *Improving z/OS Application Availability by Managing Planned Outages*, SG24-8178.

- High capacity

  Parallel Sysplex scales up to 32 images. The scalability is near linear as z/OS images are added to a sysplex. This structure contrasts with other forms of clustering that use n-to-n messaging, which leads to rapidly degrading performance with a growing number of nodes.

- Dynamic workload balancing

  The incoming workload can be automatically directed to any of the Parallel Sysplex cluster operating system images where capacity is available transparently to the applications. The z/OS WLM component is a key workload distributor, ensuring that the required service-level agreement (SLA) goal is achieved.

- Systems management

  This architecture provides the infrastructure to satisfy a requirement for continuous availability, and enables techniques for achieving simplified systems management consistent with this requirement.

- Resource sharing

  The Global Resource Serialization (GRS) component of z/OS manages access to the shared resources (such as CPU, memory, network, and storage) across all members of the Parallel Sysplex.

- Single-system image (SSI)

  The collection of system images in the Parallel Sysplex is displayed as a single entity to the operator, user, database administrator, and others. An SSI ensures reduced complexity from the operational and definition perspectives. You can rapidly scale out your workload without any added infrastructure costs by adding members to the Parallel Sysplex.

- N-2 support[19]

  Three hardware generations often are supported in the same Parallel Sysplex. IBM z16 can coexist and connect with IBM z15 and IBM z14.

  Software support for multiple releases or versions is provided.

[19] The provided coupling and timing links are not InfiniBand.

- CF encryption support

  Supports encrypted data while it is being transferred to and from the CF and in the CF Structure. Consider the following points:

  - z/OS systems must have the cryptographic hardware configured and activated to perform cryptographic functions and hold Advanced Encryption Standard (AES) master keys within a secure boundary. Feature Code 3863 CPACF DES and TDES Enablement must be installed to use the Crypto-Express6 Coprocessor (CEX6C), the Crypto-Express7 Coprocessor (CEX7C), or Crypto-Express8 Coprocessor (CEX8C) features.
  - Support can be enabled only when all systems are at z/OS 2.3 or later.

- Dynamic activation of I/O configurations for stand-alone CFs

  Dynamic I/O configuration changes can be made to a stand-alone CF without requiring a disruptive power on reset. (A stand-alone CF does not have any running instances of z/OS or z/VM).

  An LPAR with a firmware-based appliance that contains an activation service is used to apply the I/O configuration changes. The LPAR on an IBM z16 is driven by an updated HCD instance that is running in a z/OS LPAR on a remote IBM z16, IBM z15, or IBM z14.

- Improved performance and resilience:[20]

  - Fair Latch Manager 2 is an enhancement to the internals of the CFCC dispatcher. It provides CF work management efficiency and processor scalability improvements, and improves the “fairness” of arbitration for internal CF resource latches across tasks, which results in CF efficiency.
  - The CF provides more information to z/OS about every message path that appears active. Namely, the current system ID (SYID) with which the message path is registered in the CF. Whenever z/OS interrogates the state of the message paths to the CF, z/OS checks this SYID information for currency and correctness. This feature improves the delivery of signals between CF and z/OS.

[20] The two enhancements that are listed were initially implemented on IBM z15.

The components of a Parallel Sysplex as implemented within the z/Architecture are shown in Figure 5-3. The configuration is one of many possible Parallel Sysplex configurations.

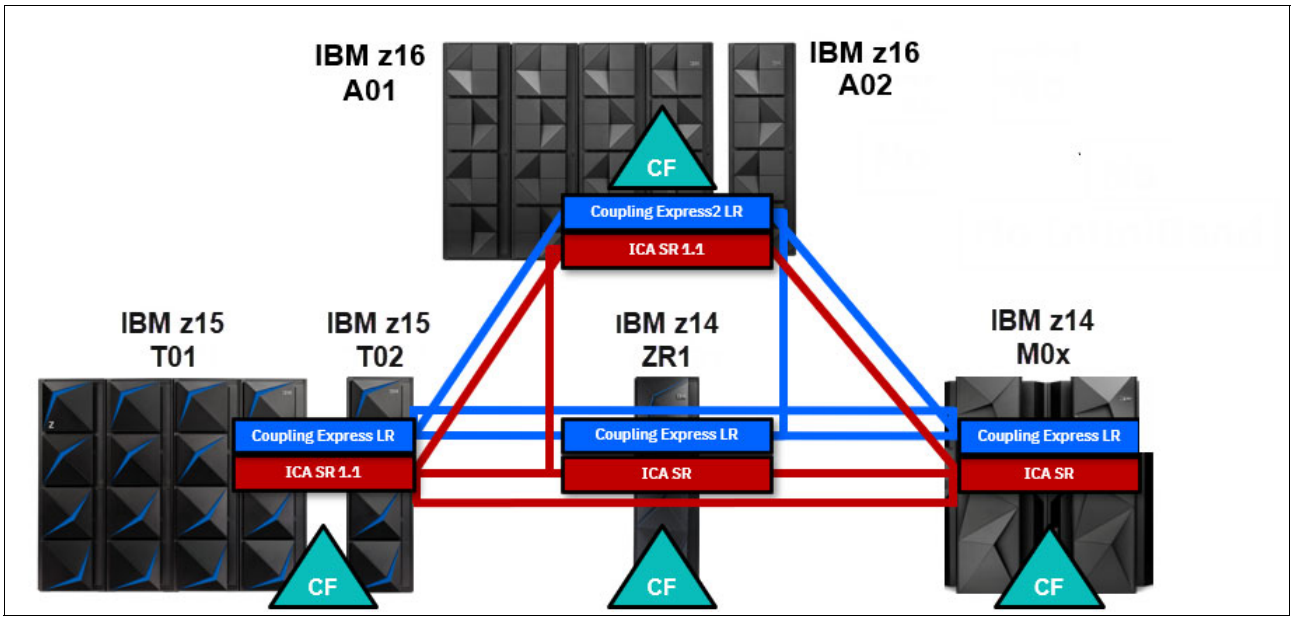

Figure 5-3 A sample Parallel Sysplex deployment

STP over coupling links provides time synchronization to all systems in the sysplex.

For more information about coupling link options, see 4.4, “Clustering connectivity” on page 56.

5.6 Pervasive encryption

Cryptography and corporate security always were a fundamental aspect of IBM Z platform development. The IBM Z platform continues to enhance and introduce new cryptography features and functions. These changes ensure the highest level of protection of your data and applications, making the platform compliant with the mandatory industry and regulatory standards.

The corporate security consists of many levels, with the strategies and policies for all components of the infrastructure: applications, databases, network, storage, and others. The traditional approach lays in defining the data perimeters and applying encryption selectively only to those perimeters. Because the data is almost never static in the system as it flows between various systems and applications, this selective approach has a significant disadvantage: After the data leaves the defined perimeter, it becomes exposed.

The selective paradigm presumes complex management of the overall infrastructure, where the failure to protect one of the components might lead to the security breach and compromise of the whole landscape.

Today, businesses demand a more comprehensive approach because applications might be subject to various cyberthreat attacks (both external and internal). Regulations around data privacy and protection also are becoming more demanding, such as the European Union (EU) General Data Protection Regulation (GDPR), Payment Card Industry Data Security Standard (PCI-DSS), and Health Insurance Portability and Accountability Act (HIPAA).

Pervasive encryption shifts this paradigm to a data-centric one: Data is becoming the new perimeter, and encryption applies to all data regardless of its origin and location. At the same time, this approach does not require costly application changes, and it is transparent to the applications and their service consumers.

Pervasive encryption provides 360 degrees of data encryption for data at-rest (stored in persistent storage) and data in-flight (transactions). This approach reduces the risks of a security breach and financial losses that are associated with it and adheres to standards and compliance.

Pervasive encryption uses hardware cryptography acceleration in IBM z16, which is proven to be more effective, performant, and stable compared to software encryption.

Pervasive encryption is enabled by using tight integration between IBM Z hardware and software, and includes the following features:

- Integrated cryptographic hardware:
  - CPACF is a coprocessor on every PU that accelerates symmetric encryption operations.
  - Crypto-Express features are hardware security modules (HSMs)[21] with the following features:
    - Complies with Federal Information Processing Standards (FIPS) 140-2 Level 4 (achieving the highest level of compliance within this standard).
    - Accelerates various cryptographic algorithms (digital signature sign and verification and many others).
    - Acting as tamper-proof storage for private keys and other highly sensitive information.
  - The CPACF and Crypto-Express usage is implemented on the hardware level, and supported natively by all IBM Z operating systems, which provides the highest encryption performance.
- Data set and file (also known as volume) encryption: Linux on IBM Z volumes and z/OS data sets are protected by using policy-controlled encryption, without any need to change or modify the applications.
- Network encryption: Network data traffic is protected by using standards-based encryption from endpoint to endpoint.
- Storage encryption: Encrypting the storage subsystem disks and its file systems.
- CF encryption: This encryption secures the Parallel Sysplex infrastructure, including the CF links and data that stored in the CF, by using policy-based encryption.
- Secure Execution for Linux on IBM Z is a capability that helps protect against cyberthreats in multitenanted cloud environments. It ensures that users and system administrators cannot access sensitive data in Linux based virtual environments.

  Secure Execution for Linux protects the confidentiality and integrity of data at enterprise scale. To achieve this goal, it isolates data at the VM level, and ensures that only the people within the organization who have a “need to know” are allowed to access the data in the clear.
- Secure Boot is an enhancement that secures the booting process of an open source operating system, such as Linux on IBM Z. With the increased number of public open source repository attacks, the extra step of verifying that the operating system kernel version was introduced.

[21] An HSM is a hardware computing device that safeguards and manages digital keys for strong authentication and accelerated crypto-operations and algorithms.

Secure Boot integrity checks validate that an operating system kernel is from an official provider and not compromised. Secure Boot can be used by Linux on IBM Z running in z/VM or KVM environments.

A complete chain of trust can be established from a trusted source to a boot loader. The process enforces Common Criteria compliance, which becomes a mandatory requirement.

5.6.1 Secure Boot for z/OS

With IBM z16 and accompanying z/OS V2.5 operating system support, IBM provides basic support for performing a Validated Boot (IPL) of z/OS systems by using IPL volumes that are defined and built for Extended Count Key Data (ECKD) DASD devices. The solution uses digital signatures to provide an IPL-time check that the z/OS system, including z/OS nucleus and LPA load module executable files, is intact, untampered with, and originates from a trusted source from the time at which it was built and signed. This approach enables the detection of subsequent unauthorized changes to those software executable files, whether those changes are accidental or malicious in nature.

When a z/OS 2.5 server is built and digitally signed as part of the client's secure build process, the target system can undergo IPL by using List-Directed IPL (LD-IPL) with digital signature validation in either Enforce or Audit mode, or undergo IPL without digital signature validation by using channel command word IPL (CCW-IPL).

- In *Enforce mode*, an IPL stops if there are validation failures for any of the load modules that are protected by Validated Boot or if the necessary configuration requirements are not met.
- in *Audit mode*, the IPL continues, but audit records are produced to describe the problems that are encountered.

Using Validated Boot for z/OS is entirely optional and bimodal. A z/OS server can continue to be built without supporting or being signed for use with Validated Boot, and a z/OS server that has been built and signed for use with Validated Boot can undergo IPL either with or without validation.

5.6.2 IBM Fibre Channel Endpoint Security

IBM Fibre Channel Endpoint Security adds endpoint authentication and encryption to data in-flight. It can help reduce insider threats of unauthorized access to the data by using traces or switch logs, and it can help technicians who use Fibre Channel (FC) analyzers to examine the packets during problem determination.

IBM Fibre Channel Endpoint Security is designed to provide a means to help ensure the integrity and confidentiality of all data that flows on FC links between trusted entities within and across data centers. The trusted entities are IBM z16 and the IBM Storage subsystem (select IBM DS8000 storage systems). No application or middleware changes are required. Fibre Channel Endpoint Security supports all data in-flight from any operating system.

IBM Z Feature Code 1146, Endpoint Security Enablement, along with CPACF enablement (Feature Code 3863) and FICON Express32S (Feature Code 0461 and Feature Code 0462) and FICON Express16SA (Feature Code 0436 or Feature Code 0437, carry forward only) turn on the Fibre Channel Endpoint Security panels on the HMC so setup can be done.

IBM Security Guardium Key Lifecycle Manager acts as a trusted authority for key generation operations and as an authentication server. It provides shared secret key generation in a

relationship between an FC initiator (IBM Z server) and the IBM Storage target. The solution implements an authentication and key management solution that is called IBM Secure Key Exchange (SKE).

Before establishing the connection, each link must be authenticated, and if successful, then it becomes a trusted connection. A policy sets the rules, for example, enforcing trusted connections only. If the link goes down, the authentication process starts again. The secure connection can be enabled automatically if both the IBM Z and IBM Storage endpoints are encryption-capable.

Data in-flight (from and to IBM Z and IBM Storage servers) is encrypted when it leaves either endpoint (source), and then it is decrypted at the destination. Encryption and decryption are done at the FC adapter level. The operating system that is running on the host (IBM Z server) is not involved in Fibre Channel Endpoint Security related operations. Tools are provided at the operating system level for displaying information about the encryption status.

5.6.3 Secure Service Container

In a production environment, applications are subject to any number of external (cyberattacks) or internal (malicious software, system administrators who use their privileged rights for unauthorized access, and many others) security risks. Secure Service Containers provide trusted execution environments for applications by using tamper protection during installation and run time, restricted administrator access, and encryption of data in-flight and at-rest.

A Secure Service Container is an integrated IBM ZIBM Z appliance and hosts the most sensitive workloads and applications. It acts as a highly protected and secured digital vault that enforces security by encrypting the entire stack. The application that is running inside the Secure Service Container is isolated and protected from outsider and insider threats.

Secure Service Containers combine hardware, software, and middleware, and it is unique to the IBM Z platform. Although it is called a container, it should not be confused with purely open-source containers (such as Docker).

An LPAR is defined as a Secure Service Container through the HMC.

A Secure Service Container features the following key advantages:

- Existing applications require zero changes to use Secure Service Container. Software developers do not need to write any Secure Service Container specific programming code.
- End-to-end encryption of data-in-flight and data at-rest:
  - Automatic Network Encryption (Transport Layer Security (TLS) and IPsec): Data in-flight.
  - Automatic volume encryption (Linux Unified Key Setup (LUKS)): Data at-rest. LUKS is the standard way in Linux to provide volume encryption. A Secure Service Container encrypts all data with a key that is stored within the appliance.
  - Protected memory: Up to 16 TB can be defined per Secure Service Container LPAR.
- Encrypted Diagnostic Data.

  All diagnostic information (debug dump data and logs) are encrypted and do not contain any user or application data.

- No operating system access.

  After the Secure Service Container appliance is built, the Secure Shell (SSH) and command-line interface (CLI) are unavailable. This configuration ensures that even system administrators cannot access the contents of a Secure Service Container and do not know what application is running there.
- Applications that run inside a Secure Service Container are accessed externally by REST APIs only.
- Tamper-proof Secure Boot for a Secure Service Container.

  Eligible applications are booted into a Secure Service Container by using a verified start sequence, in which only software code that is trusted and digitally signed and verified by IBM is uploaded into the Secure Service Container.
- Vertical workload isolation is provided by PR/SM. PR/SM in IBM z16 is designed to meet the highest level of Common Criteria (EAL5+), similar to previous IBM Z servers.
- Horizontal workload isolation: Separation from the rest of the host environment.

A Secure Service Container is a powerful IBM technology for providing an extra protection of security for the most sensitive workloads.

IBM Hyper Protect Crypto-Services offerings use the Secure Service Container technology as a core layer to provide hyper-protected services in IBM Cloud and on-premises. For more information, see IBM Cloud Hyper Protect Crypto Services.

For more information about IBM Secure Service Container offerings, see IBM Hyper Protect Virtual Servers.

5.7 Quantum-safe technology

IBM recognizes that with any new technology that new threats exist, and as such, suitable counter measures must be taken.

Quantum technology can be used for incredible good, but in the hands of an adversary, it has the potential to weaken or break core cryptographic primitives that were used to secure systems and communications. This potential leaves the foundation for digital security at risk.

The National Institute of Standards and Technology (NIST) started a process to solicit, evaluate, and standardize quantum-safe public-key cryptographic algorithms to address these issues. Quantum-safe cryptography aims to provide protection against attacks that can be started by quantum computers.

IBM z16 uses quantum-safe technologies to help protect your business-critical infrastructure and data from potential attacks.

IBM z16 Secure Boot technology protects system firmware integrity by using quantum-safe and classical digital signatures to perform a hardware-protected verification of the Initial Machine Load (IML) firmware components. This protection is anchored in a hardware-based root of trust to help ensure that the system starts safely and securely by keeping unauthorized firmware (or malware) from taking over the system during startup.

IBM z16 enables the following quantum-safe capabilities:

- Key generation
- Encryption
- Key encapsulation mechanisms (KEMs)
- Hybrid key exchange schemes
- Dual digital signature schemes

In addition to the quantum-safe cryptographic capabilities, tools such as IBM Application Discovery and Delivery Intelligence (ADDI), IBM z/OS Integrated Cryptographic Service Facility (ICSF), and IBM Crypto Analytics Tool (CAT) can help you discover where and what cryptography is used in applications. These tools can aid in developing a cryptographic inventory for migration and modernization planning.

For more information, see *Transitioning to Quantum-Safe Cryptography on IBM Z*, SG24-8525.

6

Operating system support

This chapter describes the operating system requirements and support considerations for IBM z16 A01, IBM z16 A02, and IBM z16 AGZ, and their features.

Naming: Throughout this chapter, we describe features and functions that are offered with IBM z16 A01, IBM z16 A02, and IBM z16 AGZ. The features and functions that are available across all three configurations are identified with “IBM z16”. Where features and functions differ with a given configuration, they are explicitly identified with either IBM z16 A01, IBM z16 A02, or IBM z16 AGZ.

Support and usage of hardware functions depend on the operating system version and release. The information in this chapter is subject to change. Therefore, for the most current information, see *Preventive Service Planning (PSP) bucket for 3931DEVICE (IBM z16 A01)* and *Preventive Service Planning (PSP) bucket for 3932DEVICE (IBM z16 A02 and IBM z16 AGZ)* at Preventive Service Planning buckets for mainframe operating environments.

This chapter describes the following topics:

6.1 Software support

The software portfolio for IBM z16 includes various operating systems and middleware that support the most recent and significant technologies. The following major operating systems are supported:

- IBM z/OS
- IBM z/VM
- IBM z/VSE[1]
- 21st Century Software VSE^n V6R3
- IBM z/TPF
- Linux on IBM Z and the Kernel-based Virtual Machine (KVM) hypervisor

Note: Operating system levels that are no longer in service are not covered in this publication.

6.1.1 Operating system summary

The current and minimum operating system levels that are required to support IBM z16 are listed in Table 6-1 on page 101.

Note: Program temporary fixes (PTFs) and PSP buckets: The usage of several features depends on a particular operating system. In all cases, PTFs might be necessary for the operating system level that is indicated. Review the IBM Z hardware Fix Categories (FIXCAT) before the operating system is installed. PTFs for z/OS, z/VM, and z/VSE [a]can be ordered electronically from IBM Shopz.

For more information about obtaining access to download the z/TPF and z/TPFDF APAR packages, contact TPFQA@us.ibm.com.

For Linux on IBM Z distributions and the KVM hypervisor, see the distributor's support information.

Fix packs for IBM software products that are running on Linux on IBM Z can be downloaded from IBM Fix Central.

a. z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

[1] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

Table 6-1 Operating system requirements

Operating system	End of service	Notes
z/OS V2R5	Not announced.	See the z/OS, z/VM, z/VSE, and z/TPF subsets of the 3931DEVICE (for IBM z16 A01) and 3932DEVICE (for IBM z16 A02 and IBM z16 AGZ) hardware PSP buckets and Fix Categories (FIXCAT) before installing IBM z16.
z/OS V2R4	Not announced.	
z/OS V2R3	September 2022.[a]	
z/OS V2R2[b]	September 2020.[c]	
z/VM V7R3[d]	Not announced.	
z/VM V7R2	Not announced.	
z/VM V7R1	March 2023.	
z/VSE V6R2[e]	September, 2023.	
21st Century Software VSE[n] V6R3[f]	Not announced.	
z/TPF V1R1	Not announced.	
Linux on IBM Z[g]	Support information is available for SUSE[h], Red Hat[i], and Canonical.[j]	
KVM hypervisor	For more information about minimal and recommended distribution levels, see the Linux distributors' websites.	

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. Toleration support only. Must have purchased the IBM Software Support Services offering.
c. The IBM Software Support Services offering provides the ability for customers to purchase extended defect support service for those z/OS releases that are at end of service.
d. Preview announcement at the time of writing.
e. z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.
f. 21st Century Software VSE[n] V6R3 is supported by the IBM z16 a02 and IBM z16 AGZ.
g. For more information, see the Linux on IBM Z page of the IBM Z website.
h. For more information, see the Support page of the SUSE website.
i. For more information, see the Red Hat Enterprise Linux Life Cycle page of the Red Hat website.
j. For more information, see the Ubuntu for IBM Z page of the Ubuntu website.

6.1.2 Application development and languages

Software developers can leverage the multiple programming language environments that run on the IBM Z platform. Compilers and development tool support for the IBM Z platform are continuously extended to provide developers with agile and modern development methods.

Linux on IBM Z application development support is similar to the Linux on distributed platforms, so this chapter focuses on the z/OS environment.

In addition to the traditional COBOL, PL/I, Fortran, and assembly languages, the IBM Z platform supports C, C++, Java (including Java Platform, Enterprise Edition, and batch environments), Go, Swift, Python, JavaScript, and Node.js.

IBM actively embraces open-source projects to extend z/OS functions. For example, the Zowe open-source project brings together industry experts to drive innovation for the community of next-generation mainframe developers. This framework enables an ecosystem of software solutions that are intended to provide a simple, intuitive environment for various IT professionals that are performing administrative, development, test, and operation tasks on z/OS.[2] For more information, see the Zowe community page.

The IBM z/OS Container Extensions (zCX) feature enables you to run Docker containers natively under z/OS. This feature requires z/OS V2.R4 or later and IBM z14 (or later). For more information, see the z/OS Container Extensions (zCX) web page.

An extensive set of advanced integrated development environments (IDEs) and integration tools are available for continuous development, testing, and deployment of application code.

Organizations are embarking on their journey with digital transformation and entering the API economy. Therefore, it is essential to connect business-critical applications that run on the IBM Z platform with mobile and cloud applications to better engage with customers. A key step in this evolution is to understand which assets exist in the enterprise.

The IBM DevOps offerings, such as IBM Application Delivery Foundation for z/OS and IBM Rational® Team Concert®, can be coupled with IBM Application Discovery and Delivery Intelligence (ADDI) to enable developers to understand the applications, gain cognitive insights into the process, and “evolve” those valuable older assets at speed with reduced risk to the enterprise.

Modern development practices are supported by IBM Rational Team Concert and open-source-based Git Version Control Tools for IBM z/OS (ported by Rocket Software), which are modern source code managers that run on and support z/OS.

For more information about software for the IBM Z platform, see the Products catalog web page.

Note: The usage of the most recent versions of the compilers is important. The compilers use the latest technologies on the IBM Z platform and the performance benefits that are introduced. Examples of benefits include new cache structures, new machine instructions, and instruction execution enhancements.

IBM z16 inherits the features and functions from its predecessor, IBM z15, such as on-chip compression and sort, and single-instruction multiple-data (SIMD) architectural notation, which provides efficient vector processing.

Java applications benefit from Guarded Storage Facility (GSF), which enables pause-less garbage collection.

Operating systems that run on IBM z16 can use IBM System Recovery Boost (SRB) to accelerate recovery after an outage.

The following security functions were introduced to complement the IBM Z security stack:

- Secure Execution: Provides better isolation and security for second-level guest systems that are running under the KVM hypervisor for IBM Z servers.
- Secure Boot: A feature for verifying an open-source operating systems’ kernel to ensure that it comes from the trusted provider.
- Quantum-safe cryptography algorithms (Kyber and Dilithium 8.7).

[2] https://www.openmainframeproject.org/projects/zowe#

- Fully homeomorphic encryption.
- Processor Activity Instrumentation to count cryptographic operations.
- z/OS Validated Boot (initial program load (IPL)).

6.1.3 Supported IBM compilers

Each new version of the following IBM z/OS compilers 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 delivers improved performance with IBM z16 use and advanced optimization. It supports the following items:

  - COBOL to Java Interoperability (for example, Java can call COBOL and vice versa).
  - UTF-8.
  - 64-bit JSON support.
  - JSON Boolean support for Parse and Generate.

  On IBM z16, it optimizes the following items:

  - Binary coded decimal (BCP to hex floating point (HFP) conversions.
  - Numeric editing operation.
  - Zoned decimal operations.

- IBM Automatic Binary Optimizer for z/OS

  The Automatic Binary Optimizer for z/OS improves the performance of compiled COBOL programs. It does not require source code, source code migration, or the tuning of performance options. It uses modern optimization technology to accelerate the performance of COBOL objects that are compiled by VS COBOL II V1.3 and later.

- Enterprise PL/I

  The latest version of Enterprise PL/I on IBM z16 provides the following functions:

  - More hardware use: Perform Timing Facility Function (PTFF) instruction.
  - Enhanced diagnostics for SUBSCRIPTRANGE.
  - Modernization enhancements:
    - Enhanced COMPARE for XML and JSON.
    - Support for Extended Binary Coded Decimal Interchange Code (EBCDIC) JSON.

- z/OS XL C/C++

  z/OS XL C/C++ enables developing high-performance-oriented applications through the services that are provided by IBM Language Environment® and Runtime Library extension base elements. It also works with z/OS Problem Determination Tools.

  The latest version has the following features:

  - High-performance math libraries:
    - Mathematical Acceleration Subsystem (MASS).
    - Replace Atlas with OpenBLAS.
  - Metal C for modernization of IBM High Level Assembler (HLASM) applications.
  - Neural Network Processing Assist (NNPA) Facility.

- Java

The latest version of Java supports the following features:

- On-chip zEnterprise Data Compression (zEDC) support.
- Crypto: Elliptic Curve Digital Signature Algorithm (ECDSA) and Elliptic-curve Diffie-Hellman (ECDH) acceleration.
- Zoned Decimal operations in the Data Access Accelerator (DAA) library for enhance interoperability.
- Java enablement of Deep Learning Containers (DLC) models that use IBM Z Integrated Accelerator for Artificial Intelligence (AIU).

IBM Enterprise COBOL and Enterprise PL/I support are strategic components (separately orderable products) for IBM Rational Developer for IBM Z software. These features provide a robust 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 also can use IBM Developer for z/OS to help boost productivity by editing, compiling, and debugging z/OS XL C and XL C++ applications from the workstation.

- IBM Open Enterprise SDK for Python

  Also available on z/OS is the IBM Open Enterprise SDK for Python (Current Version 3.10), which is based on the popular Python interpreter from Python Software Foundation (PSF):

  - z/OS LE-based, 64-bit only.
  - Runs on z/OS UNIX System Services environment with no prerequisites.
  - Supports ASCII, EBCDIC, and Unicode (UTF-8).
  - Includes selected Python Package Index (PyPI) packages.

For more information, see this IBM Documentation web page.

6.2 Support by operating system

This section lists the support of in-service operating systems for functions of IBM z16.

For more information about IBM z16 and its features, see *IBM z16 (3931) Technical Guide*, SG24-8951.

For more information about all of the I/O features, see *IBM Z Connectivity Handbook*, SG24-5444.

6.2.1 z/OS

z/OS is a core IBM Z operating system that supports IBM z16. IBM z16 capabilities differ depending on the z/OS release.

The z/OS release cycle was extended with IBM Software Support Services to provide the ability for customers to purchase extended defect support service for previous versions of the operating system.

The minimum required version of z/OS to run on IBM z16 is V2R2 with PTFs (IBM Extended Software Support Services offering must be purchased). There is support on z/OS for these versions:

- z/OS V2R3 + PTFs

- z/OS V2R4 + PTFs
- z/OS V2R5 + PTFs

z/OS V2R2 supports toleration only and does not support new functions.

z/OS supports the following select (but not limited to) new functions:

- Coupling Facility Control Code (CFCC) Level 25
- HiperDispatch Enhancements
- New CPU MF Counters
- SRB enhancements
- zDNN library enablement for IBM Z AIU
- Up to 16 TB of real memory for z/OS image (not IBM z16 specific)
- Compilers and Automatic Binary Optimizer for z/OS usage
- Usage of new hardware instructions: XL C/C++ ARCH(14)
- Integrated Cryptographic Support Facility (ICSF)
- IBM Integrated Accelerator for IBM Z Sort (zSort) hardware sort instruction (**SORTL**)
- Integrated on-chip IBM zEnterprise Data Compression (zEDC) compression
- zDNN library enablement for IBM Z AIU
- Compliance-ready Central Processor Assist for Cryptographic Functions (CPACF) Counters support to track crypto compliance and instruction usage

Before the IBM z16 migration process is started, see the IBM z16 workflows that are provided with each release of z/OS. This information is available in the z/OS IBM z16 Upgrade Workflow for z/OSMF, which is provided with APAR OA62703 on V2R2 and higher. The z/OSMF workflow contains only the z/OS steps for upgrading to IBM z16 and installs into the `/usr/lpp/bcp/upgrade` directory.

z/OSMF is recommended because it offers interactive assistance and runs associated health checks.

z/OS zCX enablement on IBM z16 is provided by IBM Container Hosting Foundation for z/OS (5655-HZ1). On IBM z14 and IBM z15, this enablement is provided by the hardware Feature Code 0104. Feature Code 0104 is *not* available on IBM z16.

For more information, see this IBM Support web page.

For more information about z/OS downloads, see this z/OS Downloads web page.

6.2.2 z/VM

IBM z16 support is provided with PTFs for z/VM 7.1 and 7.2, and it is included in the z/VM 7.3 base.

Compatibility support enables guest use for several new facilities, including the following examples:

- Embedded Artificial Intelligence Acceleration

  Designed to reduce the overall time that is required to run CPU operations for neural networking processing functions, and help support real-time applications, such as fraud detection.
- Compliance-ready CPACF Counters support

  Means for guests to track crypto compliance and instruction use.

- Breaking Event Address Register (BEAR) Enhancement Facility

  Facilitates debugging wild branches.
- Vector Packed Decimal Enhancements 2

  Instructions that are intended to provide performance improvements.
- Reset DAT protection Facility

  Provides a more efficient way to disable DAT protection, such as during copy-on-write or page change tracking operations.
- RDMA over Converged Ethernet (RoCE) Express3 feature

  Allows guests to use Routable RoCE, Zero Touch RoCE, and Shared Memory Communications (SMC)-R V2 support.
- Guest enablement for the Crypto Express8S (CEX8C) feature and assorted crypto enhancements

  Including quantum-safe API guest use support that is available to dedicated guests.
- CPU/Core topology location information within z/VM monitor data

  Provides a better picture of the system for diagnostic and tuning purposes.
- Consolidated Boot Loader for guest IPL from SCSI

z/VM logical partitions (LPARs): IBM z16 central processors (CPs) and the Integrated Facilities for Linux (IFLs) feature increased capacity over the capacity of their predecessors. Therefore, as a best practice, review and adjust the capacity of z/VM LPARs and of any guests in terms of the *number* of IFLs and CPs (real or virtual) to achieve the capacity that you require.

For more information about PTF availability, see the z/VM Continuous Delivery News web page.

For more information about for IBM z16 migration, see the hardware PSP buckets for 3931DEVICE, and 3931DEVICE z/VM subset or 3932DEVICE, and the 3932DEVICE z/VM subset for IBM z16 A02 and IBM z16 AGZ.

For more information about all IBM z16 features and functions that are supported by the z/VM releases, see *IBM z16 (3931) Technical Guide*, SG24-8951.

6.2.3 z/VSE[3]

IBM z16 support is provided by z/VSE[3] V6R2. Consider the following points:

- z/VSE runs in z/Architecture mode only.
- z/VSE V6.2 supports High-Performance FICON for IBM Z (zHPF) and SIMD.
- SRB (subcapacity CP speed boost only) (IBM z16 A01 only).

For more information about all IBM z16 features and functions that are supported by the z/VSE[3] releases, see *IBM z16 (3931) Technical Guide*, SG24-8951.

6.2.4 VSE[n] V6.3 support from 21st Century Software

z/VSE is **not supported** on IBM z16 A01 and IBM z16 AGZ.

[3] z/VSE is not supported n IBM z16 A02 and IBM z16 AGZ.

The IBM z16™ A01 server (machine type 3931), announced in Hardware Announcement 122-001, dated April 5, 2022, is intended to be the last IBM Z server to be supported by z/VSE 6.2 (5686-VS6). It has been announced to be withdrawn from marketing on September 5, 2022.

IBM supports 21st Century Software VSE^n V6.3 on IBM z16. For more information, see this web page.

6.2.5 z/TPF

IBM z16 support is provided by z/TPF V1R1 with PTFs.

For more information about all IBM z16 features and functions that are supported by the z/TPF, see *IBM z16 (3931) Technical Guide*, SG24-8951.

6.2.6 Linux on IBM Z

The Red Hat, SUSE, and Ubuntu releases that are supported on IBM z16 are listed in Table 6-2.

Table 6-2 Linux on IBM Z distributions

Linux on IBM Z distribution	Version and release
SUSE Linux Enterprise Server	SUSE Linux Enterprise Server 15 SP3 with service
SUSE Linux Enterprise Server	SUSE Linux Enterprise Server 12 SP5 with service (not as a Secure Execution KVM host)
Red Hat Enterprise Server	RHEL 8.4 with service
Red Hat Enterprise Server	RHEL 7.9 with service (not as a KVM host)
Canonical	Ubuntu 22.04 LTS with service
Canonical	Ubuntu 20.04.1 LTS with service

For more information about all IBM z16 features and functions that are supported by the Linux on IBM Z distributions, see *IBM z16 (3931) Technical Guide*, SG24-8951.

6.2.7 Kernel-based Virtual Machine hypervisor

For IBM z16, the KVM is delivered and supported by the Linux distribution partners. For more information about KVM support for the IBM Z platform, see the following resources:

- The documentation for your distribution
- *Virtualization Cookbook for IBM Z Volume 5: KVM*, SG24-8463

6.3 Software licensing

The IBM Z software portfolio includes operating system software (that is, z/OS, z/VM, z/VSE[4], and z/TPF) and middleware that runs on these operating systems. The portfolio also includes middleware for Linux on IBM Z environments. For IBM z16, the following metric groups for software licensing are available from IBM (depending on the software product):

- Monthly license charge (MLC)

  MLC pricing metrics feature a recurring monthly charge. In addition to permission to use the product, the charge includes access to IBM product support during the support period. MLC pricing applies to z/OS, z/VSE[4], and z/TPF operating systems. Charges are based on processor capacity, which is measured in millions of service units (MSU) per hour.

- IBM Tailored Fit Pricing for IBM Z for software[5]

  Tailored Fit Pricing for IBM Z is a flexible software pricing model that dramatically simplifies the pricing landscape through flexible deployment options that are tailored to your IBM Z environment.

  Two new pricing solutions, Enterprise Consumption and Enterprise Capacity, offer alternatives to the rolling four-hour average (R4HA)-based pricing model for new and existing workloads.

- IBM International Program License Agreement (IPLA)

  IPLA metrics feature a single, up-front charge for an entitlement to use the product. An optional and separate annual charge, called *subscription and support*, entitles you to access IBM product support during the support period. With this option, you also can receive future releases and versions at no extra charge.

Software licensing references

For more information about software licensing, see the following resources:

- Learn about Software licensing
- Base license agreements
- IBM Z Software Pricing reference guide
- The IBM International Passport Advantage® Agreement can be downloaded from the Learn about Software licensing website
- IBM Tailored Fit Pricing for IBM Z for software

Subcapacity pricing terms for z/VM and select z/VM based programs

Subcapacity pricing is available to clients running on the z/VM 7 platform. Software pricing at less than full machine capacity can provide more flexibility and improved cost of computing as a client manages the volatility and growth of new workloads.

For more information about subcapacity pricing terms for z/VM and z/VM based programs, see Sub-capacity pricing terms for z/VM and select z/VM based programs help improve flexibility and price and performance.

For more information about software licensing options that are available for IBM z16, see *IBM z16 (3931) Technical Guide*, SG24-8951.

[4] z/VSE is not supported on the IBM z16 A02 and IBM z16 AGZ.

[5] Tailored Fit Pricing for IBM Z Hardware Consumption Solution also is available for IBM z16. For more information, see https://www.ibm.com/it-infrastructure/z/pricing.

6.4 References

For planning information, see the following operating system web pages:

- z/OS
- z/VM
- z/VSE[6]
- 21st Century Software VSE^n V6R3
- z/TPF
- Linux on IBM Z

[6] z/VSE is not supported on IBM z16 A02 and IBM z16 AGZ.

Abbreviations and acronyms

ADDI	IBM Application Discovery and Delivery Intelligence
AES	Advanced Encryption Standard
AI	artificial intelligence
AIU	IBM Integrated Accelerator for Artificial Intelligence
AMG	alias management group
ASHRAE	American Society of Heating, Refrigerating, and Air-Conditioning Engineers
BEAR	Breaking Event Address Register
BMC	Base Management Card
BNC	Bayonet Neill-Concelman
BPA	Bulk Power Assembly
BPR	Bulk Power Regulator
BTS	Backup Time Server
CAT	Crypto Analytics Tool
CBU	Capacity Backup
CCA	Common Cryptographic Architecture
CCW-IPL	channel command word IPL
CEX6C	Crypto-Express6 Coprocessor
CEX7C	Crypto-Express7 Coprocessor
CEX8C	Crypto-Express8 Coprocessor
CF	Coupling Facility
CFCC	Coupling Facility Control Code
CICS	IBM Customer Information Control System
CISC	Complex Instruction Set Computer
CIU	Customer Initiated Upgrade
CLI	command-line interface
CoD	Capacity on Demand
CP	central processor
CPACF	Central Processor Assist for Cryptographic Functions
CPC	central processor complex
CPE	Capacity for Planned Event
CPI	cycles per instruction
CPM	Capacity Provisioning Manager
CPUMF	CPU Measurement Facility
CRYSTALS	Cryptographic Suite for Algebraic Lattices
CSM	Copy Services Manager
CSS	channel subsystem
CST	Coordinated Server Time
CTC	channel-to-channel
CTN	Coordinated Timing Network
DCAP	data-centric audit and protection
DCIM	data center infrastructure management
DCM	dual-chip module
DCSS	discontiguous shared segment
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DIMM	dual inline memory module
DPM	Dynamic Partition Manager
DR	disaster recovery
DRAM	dynamic random access memory
DRNG	Deterministic Random Number Generation
DWDM	dense wavelength-division multiplexing
EAV	extended address volume
EBCDIC	Extended Binary Coded Decimal Interchange Code
ECC	Elliptic Curve Cryptography
ECDH	Elliptic-curve Diffie-Hellman
ECDSA	Elliptic Curve Digital Signature Algorithm
ECKD	Extended Count Key Data
EDA	Enhanced Drawer Availability
EDAT-2	Enhanced Dynamic Address Translation-2
EDM	Enhanced Driver Maintenance
EDO	Encrypted Data Object
EP11	Enterprise PKCS #11
EPO	Emergency Power Off
ETL	extract, transform, and load
ETS	External Time Source
EU	European Union
FC	Fibre Channel
FCP	Fibre Channel Protocol
FEC	Forward Error Correction

FID	Function Identifier
FIPS	Federal Information Processing Standards
FIXCAT	Fix Categories
FRU	field replaceable unit
GCM	Galois/Counter Mode
GDPR	General Data Protection Regulation
GRS	Global Resource Serialization
GSF	Guarded Storage Facility
HA	high availability
HADR	high availability and disaster recovery
HDFP	hardware decimal floating point
HIPAA	Health Insurance Portability and Accountability Act
HLASM	IBM High Level Assembler
HMA	Hardware Management Appliance
HMC	Hardware Management Console
HSA	hardware system area
HSM	hardware security module
HVDC	High-Voltage DC
IBF	integrated battery facility
IBM	International Business Machines Corporation
IC	internal coupling
ICA SR	Integrated Coupling Adapter Short Reach
ICF	Integrated Coupling Facility
ICFs	Integrated Coupling Facilities
ICSF	IBM z/OS Integrated Cryptographic Service Facility
IDE	integrated development environment
IFLs	Integrated Facilities for Linux
IFP	Integrated Firmware Processor
IML	Initial Machine Load
IMS	IBM Information Management System
IOS	I/O Supervisor
IOSQ	I/O queue time
IPC	instructions per cycle
IPL	initial program load
IPLA	IBM International Program License Agreement
ISA	Instruction Set Architecture
ISM	Internal Shared Memory
ITRR	internal throughput rate ratio
JVM	Java virtual machine
KEM	key encapsulation mechanism
KVM	Kernel-based Virtual Machine
LAN	local area network
LCSS	logical channel subsystem
LD-IPL	List-Directed IPL
LIC	Licensed Internal Code
LICCC	Licensed Internal Code Configuration Control
LPAR	logical partition
LR	Long Reach
LSPR	Large Systems Performance Reference
LSS	logical subsystem
MAC	Media Access Control
MASS	Mathematical Acceleration Subsystem
MCL	maintenance change level
MCU	Memory Controller Unit
MIF	Multiple Image Facility
MLC	monthly license charge
MSU	millions of service units
MTBF	mean time between failures
MTP	multi-fiber termination push-on
NNPA	Neural Network Processing Assist
NTI	IBM New Technology Introduction
NTP	Network Time Protocol
OLTP	online transaction processing
OSC	Oscillator Card
PAIA	Product Attributes to Impact Algorithm
PAV	parallel access volume
PCI	Peripheral Component Express
PCI-DSS	Payment Card Industry Data Security Standard
PCI-HSM	Payment Card Industry PTS HSM
PCIe	Peripheral Component Interconnect Express
PCIeCC	Peripheral Component Interconnect Express Cryptographic Coprocessor
PDU	power distribution unit
POR	power-on reset
PPC	Processor Power Card
PPRC	Peer-to-Peer Remote Copy

PPS	pulse per second
PR/SM	Program Resource/System Manager
PRNG	pseudo-random number generation
PSF	Python Software Foundation
PSP	Preventive Service Planning
PSU	Power Supply Unit
PTF	program temporary fix
PTP	Precision Time Protocol
PTS	Preferred Time Server or (Primary Time Server
PU	processor unit
PyPI	Python Package Index
R4HA	rolling four-hour average
RAIM	redundant array of independent memory
RAS	reliability, availability, and serviceability
RDMA	Remote Direct Memory Access
RNI	relative nest intensity
RoCE	RDMA over Converged Ethernet
RSF	Remote Support Facility
SAFR	Scalable Architecture for Financial Reporting
SAP	System Assist Processor
SE	Support Element
SHA	Secure Hash Algorithm
SIMD	single-instruction, multiple-data
SKE	Secure Key Exchange
SLA	service-level agreement
SMC	Shared Memory Communications
SME	subject matter expert
SMP	symmetric multiprocessor
SMT	simultaneous multithreading
SORTL	sort instruction
SR	Short Reach
SRB	IBM System Recovery Boost
SSH	Secure Shell
SSI	single-system image
SSL	Secure Sockets Layer
STP	Server Time Protocol
SVC	SAN Volume Controller
SYID	system ID
TDES	Triple Data Encryption Standard
TKE	Trusted Key Entry
TLS	Transport Layer Security
TRNG	true-random number generator
UDX	User Defined-Extension
UIFL	Unassigned Integrated Facilities for Linux
VFM	Virtual Flash Memory
VM	virtual machine
WLM	Workload Manager
zCX	z/OS Container Extensions
zEDC	zEnterprise Data Compression
zERT	IBM z/OS Encryption Readiness Technology
zHPF	High-Performance FICON for IBM Z
zIIP	IBM Z Integrated Information Processor
zSort	IBM Integrated Accelerator for IBM Z Sort

Redbooks

IBM z16 Technical Introduction

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8950-01

ISBN 0738461083

Printed in U.S.A.