Redbooks
ibm.com/redbooks

IBM z15 Technical Introduction

Bill White
Octavian Lascu
John Troy
Jannie Houlbjerg
Frank Packheiser
Paul Schouten
Kazuhiro Nakajima
Anna Shugol
Hervey Kamga

IBM Z

IBM

Redbooks

IBM Redbooks

IBM z15 Technical Introduction

May 2021

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

Third Edition (May 2021)

This edition applies to the IBM z15 (machine type 8561 and 8562).

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

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

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

CICS®	IBM Security™	Redbooks (logo) ®
Concert®	IBM Z®	Resource Link®
Db2®	IBM z13®	WebSphere®
DB2®	IBM z13s®	z/Architecture®
DS8000®	IBM z14®	z/OS®
FICON®	Insight®	z/VM®
GDPS®	Language Environment®	z/VSE®
HyperSwap®	Parallel Sysplex®	z13®
IBM®	Passport Advantage®	z13s®
IBM Cloud®	Rational®	z15™
IBM MobileFirst®	Redbooks®	zEnterprise®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Ansible, OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Preface

This IBM® Redbooks® publication introduces the latest member of the IBM Z® platform, the IBM z15™. It includes information about the Z environment and how it helps integrate data and transactions more securely. It also provides insight for faster and more accurate business decisions.

The z15 is a state-of-the-art data and transaction system that delivers advanced capabilities, which are vital to any digital transformation. The z15 is designed for enhanced modularity, and occupies an industry-standard footprint. It is offered as a single air-cooled 19-inch frame called the z15 T02, or as a multi-frame (1 to 4 19-inch frames) called the z15 T01.

Both z15 models excel at the following tasks:

- Using hybrid cloud integration services
- Securing and protecting data with encryption everywhere
- Providing resilience with key to zero downtime
- Transforming a transactional platform into a data powerhouse
- Getting more out of the platform with operational analytics
- Accelerating digital transformation with agile service delivery
- Revolutionizing business processes
- Blending open source and IBM Z technologies

This book explains how this system uses innovations and traditional Z strengths to satisfy growing demand for cloud, analytics, and open source technologies. With the z15 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 at IBM Redbooks, Poughkeepsie Center.

Bill White is an IBM Redbooks Project Leader and Senior IT Infrastructure Specialist at 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 (DR) 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.

John Troy is an IBM Z and storage hardware National Top Gun in the northeast area of the US. He has 40 years of experience in the service field. His areas of expertise include IBM Z servers and high-end storage systems technical and customer support and services. John has also been an IBM Z hardware technical support course designer, developer, and instructor for the last eight generations of IBM high-end servers.

Jannie Houlbjerg is a Systems Programmer working at JN Data in Denmark. She has over 22 years of experience in the IBM Z field. Her areas of expertise include IBM Z hardware, Parallel Sysplex®, Performance and High Availability, and IBM GDPS®.

Frank Packheiser is a Senior zIT Specialist at the Field Technical Sales Support office in Germany. He has over 25 years of experience in IBM Z. Frank has worked for 10 years in the IBM Education Center in Germany, developing and providing professional training. He also provides professional services to IBM Z and mainframe customers. In 2008 and 2009, Frank supported customers in Middle East/North Africa (MENA) as a zIT Architect. In addition to co-authoring several IBM Redbooks publications since 1999, he has been an official ITSO presenter at ITSO workshops since 2010.

Paul Schouten is an IBM Z Client Technical Specialist based in Sydney, Australia. During his 40 years supporting mainframe systems, he has performed many roles, including Certified IT Architect, Systems Software developer, and Systems Programming. He has extensive experience developing and documenting high availability solutions for IBM's enterprise customers.

Kazuhiro Nakajima is a Senior IT Specialist in IBM Japan. He has almost 30 years of experience in IBM Japan and he has been active as an advanced Subject Matter Expert of IBM Z products for over 20 years. His areas of expertise include IBM Z hardware, performance, z/OS®, IBM Z connectivity. He is co-author of several IBM Z configuration setup IBM Redbooks publications from the IBM zEC12 to the IBM z14®.

Anna Shugol is a mainframe technical specialist at IBM UK. She has over 8 years of experience and has worked with customers in various geographies. Anna holds a Computer Science degree from Bauman Moscow State Technical University. Large-system performance is one of her key areas of expertise.

Hervey Kamga is an IBM Z Product Engineer with the EMEA I/O Connectivity Team in Montpellier, France. Before serving in his current role, he was a Support Engineer and Engineer On Site for 13 years with Sun MicroSystems and Oracle in EMEA. Hervey's areas of expertise include Oracle Solaris (Operating System and hardware products), virtualization (VMware, virtualBox), Linux (Ubuntu), and IBM Z I/O features and protocols (IBM Fibre Connection (FICON®) and OSA). He has co-authored several IBM Redbooks publications.

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

Dave Surman, Darelle Gent, David Hutton, Tom Dewkett, Frank Bosco, Patty Driever, Anthony Sofia, Brian Valentine, Purvi Patel, Les Geer III, Bill Bitner, Christine Smith, Barbara Weiler, Charlie Lawrence, Dean St Pier, Tom Morris, Ellen Carbarnes, Gregory Hutchison, Riaz Ahmad, Franco Pinto, Walter Niklaus, Martin Recktenwald, Roan Dawkins, Barbara Sannerud, David Yearack
IBM

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

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

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

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

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

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Designed to take on today's IT demands

Delivering new services efficiently and effectively with speed and at scale is crucial to any business, large or small. Managing changes and potential disruptions to the IT infrastructure while maintaining availability of services is a must. Ensuring data privacy and protection to mitigate the effects of security breaches is what every business expects from its IT infrastructure.

There is no doubt that the technology you choose determines business success and differentiates you from the competition. Technology choices must help meet the expectations of rapid dynamic development cycles and give the confidence that new and existing services are resilient and can be delivered quickly and securely. Therefore, the correct balance of technologies and the IT platform on which they run is key.

Today, many businesses are using cloud technologies to deliver cloud native services at scale and at lower cost, but often the risk of vendor lock-in and escalating costs exists. In addition, many enterprise applications and services have not moved to the cloud because of concerns about flexibility, connectivity, security, and management across multicloud and hybrid environments.

The latest members of the IBM Z family, the z15 models T01 and T02, feature a tried-and-true architecture that can satisfy today's demands. These servers can streamline your ability to integrate and manage disparate cloud environments and create a single, cohesive IT infrastructure that is open, secure, resilient, and flexible.

The z15 models can protect your data through encryption, even if it leaves the Z environment. Both models also provide the strongest workload isolation in the industry.

In this chapter, you learn about both z15 models, including their capabilities and the business value they provide. It includes the following topics:

1.1 The IBM z15: The platform for the digital economy

More than any other platform, the z15 models offer a high-value architecture that can give you the cloud that you want, with the privacy and security you need. Compared to their predecessor platforms, these models provide more of what is required to satisfy growing demands that are driven by the digital economy, including the following benefits:

- Compute power for increased throughput
- Special hardware co-processors for accelerating various workloads
- Large-scale memory to process data faster
- Industry-unique cache design to optimize performance
- Accelerated I/O bandwidth to process massive amounts of data
- Enhanced data compression to economically store and process information
- High-speed cryptographic operations to secure data with minimal extra impact
- Instant recovery:
  - to reduce the duration that is needed to start or shut down an OS or services
  - to recover faster form parallel sysplex events.
- Data privacy and protection by using end-to-end encryption for copies of eligible data
- Endpoint security to enforce the protection of data in-flight[1]
- Secure execution of kernel-based virtual machine (KVM) and Linux to isolate and protect containers and virtual servers
- Flexible, on-premises HW consumption pricing allowing customers to handle in a cost effective manner the impact of unpredictable high spikes on business critical workloads

Figure 1-1 shows the z15 configuration options that are available (one to four 19-inch frames for the z15 T01 and one 19-inch frame for the z15 T02).

[1] IBM Fibre Channel Endpoint Security is available only with the z15 T01.

Figure 1-1 IBM z15 (T01 and T02) 19-inch frame configurations

The z15 models incorporate the ability to access more processor power and larger data in memory than their predecessors. With their enhanced security and cryptography capabilities, enhanced virtualization, and industry-exclusive I/O processors for offloading data-intensive workloads, you can build highly secure hybrid cloud platforms to improve speed-to-market and competitive advantage.

The z15 models enable you to meet the needs of the digital economy with the following solutions:

- Using hybrid cloud integration services
- Securing and protecting data with encryption everywhere
- Providing resilience: The key to zero downtime
- Transforming a transactional platform into a data powerhouse
- Getting more out of the platform with operational analytics
- Accelerating digital transformation with agile service delivery
- Revolutionizing business processes
- Blending open source and IBM Z technologies

Naming: Throughout this chapter, we refer to both models (the T01 and T02) as the z15. Wherever features and functions differ across the models, they are explicitly mentioned.

1.1.1 Using hybrid cloud integration services

Cloud technology is the new norm for maximizing agility with broad data access and almost limitless collaboration possibilities. Businesses are analyzing how to best take advantage of cloud technologies for their unique use cases. In the process, they discover the challenges of managing a diverse cloud infrastructure while ensuring availability, security, accessibility, and

portability of services. The IBM Z platform and IBM software offerings work together to manage and mitigate these challenges.

In addition, businesses recognize that hybrid and private cloud deployments can allow them to take advantage of consistent access and tools across the enterprise. This consistency helps them to quickly develop, deploy, and manage services that maintain a competitive edge.

To advance cloud deployment, businesses are implementing container and orchestration technologies, and embracing agile development, microservices, and DevOps best practices. Yet, as they expand their cloud deployments, protection of their core assets and data within the hybrid cloud environment becomes even more crucial.

For application integration architectures that can be used for hybrid integration with IBM Z, including REST APIs, messaging, and event streams, see *IBM Z Integration Guide for Hybrid Cloud*, REDP-5319.

With the IBM acquisition of Red Hat, the hybrid cloud capabilities on IBM Z were extended. Support for running OpenShift[2] on Linux on Z provides expansive cloud capabilities, tools, and access to an extensive open community. For more information about hybrid cloud capabilities, see Hybrid cloud with IBM Z.

Cloud Pak for Applications extends support for Red Hat OpenShift Container Platform onto the IBM Z platform. Hybrid cloud deployments can now be extended to include Red Hat OpenShift clusters on IBM Z. In this way, you take advantage of the container orchestration platform and tools to provide a consistent experience for development of cloud-native workloads. For example, with IBM Wazi for Red Hat CodeReady Workspaces, you can write applications for z/OS by using your integrated development environment (IDE) of choice. Then, you can debug, build, and test code in your personal z/OS sandbox. For more information, see IBM Wazi for Red Hat CodeReady Workspaces.

z/OS also supports the OpenShift system through the z/OS Cloud Broker. IBM Z brings z/OS into the enterprise hybrid cloud platform with z/OS Container Extensions (zCX) and the z/OS Cloud Broker. Container Extensions allow you to deploy unmodified Linux on Z Docker images that are running inside z/OS. z/OS Cloud Broker enables integration into the cloud-native development cycle through self-service access and deployment of z/OS resources and services on IBM Cloud® for a seamless and universal cloud development experience. For more information about zCX, see z/OS Container Extensions (zCX) Content Solution.

IBM Z and Red Hat delivered support for Red Hat Ansible Automation Platform with IBM Z, specifically Ansible Collections for z/OS, which is a set of Ansible content (for example, playbooks, roles, modules, and plug-ins). Ansible content is used for development and operations on z/OS and z/OS middleware products. For more information, see Red Hat Ansible Certified Content for IBM Z Content Solution.

In addition, IBM Hyper Protect Virtual Servers can help you securely build, deploy, and manage mission-critical applications for hybrid cloud environments on IBM Z. Hyper Protect Virtual Servers protect your data and workloads from threats throughout their lifecycle. For more information, see IBM Cloud Hyper Protect Services.

Secure Execution for Linux removes the security concerns of running multiple containers in the same virtual server. Guests and containers are fully isolated and protected from Linux on Z and the hypervisor (KVM), making services immune to internal and external threats. For

[2] OpenShift is a family of containerization software that was developed by Red Hat. Its flagship product is the OpenShift Container Platform, which is an on-premises platform as a service (PaaS) built around containers that are orchestrated and managed by Kubernetes. Kubernetes is an open source container as a service (CaaS) project that originates from Google.

example, a user with system administrator authority can still manage and deploy workloads, but is unable to view data being processed in the guest virtual machine (VM) or in a container. Secure Execution for Linux provides a secure, multitenant cloud solution for both on and off premises.

1.1.2 Securing and protecting data with encryption everywhere

By placing the security controls on the data, IBM Z pervasive encryption provides the comprehensive protection that your organization and customers demand. This solution creates an envelope of protection around the data on Z. For example, pervasive encryption helps protect the at-rest and in-flight data that is on your Z infrastructure. Also, centralized, policy-based data encryption controls significantly reduce the costs that are associated with data security and regulatory compliance, including the General Data Protection Regulation (GDPR).[3]

IBM Z pervasive encryption implements this comprehensive security with your ongoing operations in mind. Therefore, it does not require you to change any application, and can be implemented by using policy-based controls, with low overhead. These capabilities can reduce costs that are associated with data security and compliance.

The z15 excels with security features that are built into the hardware, firmware, and operating systems. The built-in features range from storage protection keys and workload isolation to granular audit capabilities, and more.

The Central Processor Assist for Cryptographic Functions (CPACF), which is standard on every core, supports pervasive encryption and provides hardware acceleration for encryption operations. In addition, the hardware security module (HSM) in the Crypto-Express feature gets a performance boost on z15. Combined, these two enhancements perform encryption more efficiently on the z15 than on earlier Z platforms.

Encryption of the compressed data is proven to be more efficient. Data compression on the z15 was moved to the processor chip level. This change eliminates the need to select what data is compressed (everything is compressed) and brings cost effectiveness for data storage (see faster encryption and on-chip compression).

Building on pervasive encryption, the goal of the z15 is to protect data throughout the enterprise and into hybrid cloud environments with IBM Data Privacy Passports. Data Privacy Passports is designed to minimize the risk and impact of both data loss and privacy breaches when collecting and storing sensitive data. It manages how eligible data is shared securely through a central control of user access.

Data Privacy Passports protect eligible data wherever it goes. Security policies are kept and acknowledged whenever the data is accessed. Future data access can be revoked remotely by using Data Privacy Passports, long after data has left the system of record, and sensitive data can even be made unusable by destroying its encryption key.

IBM Enterprise Key Management Foundation - Web Edition generates and manages keys to be used for z/OS data set encryption. It provides a user interface for centralized management of multiple z/OS systems. Existing keys can be imported and managed. A data set dashboard provides an overview of data sets encryption status. EKMF Web also supports the **zkey** command on Linux on Z, as well as public cloud key management systems. See the EKMF Web Edition web page.

[3] For more information about GDPR, see the Clear the Path to the GDPR website.

Data Privacy Passports is designed to help reduce time that is spent by staff to protect data and ensure privacy throughout its lifecycle by using a central point of control. For more information, see 5.6.2, “IBM Data Privacy Passports” on page 100.

Another premier technology for IBM Z is Secure Service Container. It is an integrated IBM Z appliance and was designed to host most sensitive workloads and applications. It acts as a highly protected and secured digital vault and enforces security by encrypting the whole stack: memory, network, and data (in-flight and at-rest). Applications that are running inside the Secure Service Container are isolated and protected from outside and inside threats. For more information, see 5.6.3, “Secure Service Container” on page 100.

The z15 T01 extends the pervasive encryption approach with the introduction of Fibre Channel Endpoint Security to protect data in-flight. Data can be accessed only by trusted Z platforms and storage subsystems in and across data centers. Fibre Channel Endpoint Security encrypts all data that flows on IBM Fibre Connection (FICON) and Fibre Channel (FCP) links from IBM Z to IBM DS8900F or between Z platforms. Applications or file systems changes are not required. For more information, see 5.6.1, “IBM Fibre Channel Endpoint Security” on page 99.

For more IBM Z security information, see IBM Z Mainframe Enterprise Security.

1.1.3 Providing resilience: The key to zero downtime

For most businesses, every second of planned or unplanned downtime can mean lost revenue. It is crucial for them to keep systems that run 24 x 7, and to rapidly recover from an outage and resume business operations. For more information about the impact from downtime, such as lost revenue, regulatory penalties, and damaged reputations, see Resiliency on IBM Z.

Redundancy and the technologies that leverage it are perhaps the most prevalent mechanism to ensure resiliency, which is one of many key strengths of IBM Z. IBM designers worked through the years to remove single points of failure in the Z platform. They developed recovery routines in the firmware, operating systems, middleware, and applications to cope with slowdowns or to minimize outages and rapidly recover from them.

Also, with platform-level redundancy, the z15 is designed to handle failures while it maintains user access. Components can be repaired, maintenance is performed, and products are migrated with minimal business impacts. Some capabilities, such as capacity-on-demand, automatically turn components on and off based on current needs.

The z15 platform is highly robust. It is even more so in a Parallel Sysplex[4] environment that implements the IBM Geographically Dispersed Parallel Sysplex (IBM GDPS) family of solutions, which improve resilience in cases of unplanned outages. The GDPS solutions provide more tools to ensure Z availability, and mask or significantly reduce the effects of critical component outages or failures.

By using IBM HyperSwap® technology, you seamlessly route I/O traffic from a disk subsystem that cannot service the I/O request to a second disk subsystem that can. Also, the GDPS or Active-Active solution can route workloads from a server location that is experiencing problems to a second location that is operating well.

For more about IBM Z resiliency capabilities, see *Getting Started with IBM Z Resiliency*, SG24-8446.

[4] Parallel Sysplex is a system clustering technique for high availability. For more information, see 5.5, “High availability with Parallel Sysplex” on page 94.

IBM Instant Recovery is a capability that uses IBM Z System Recovery Boost technologies to minimize the duration and impact of downtime and short-term recovery events. System Recovery Boost offers more central processor (CP) capacity during particular recovery operations, such as accelerating OS and services startup or shutdown times, and for a short period following the restoration of steady-state operation for workload “catch-up” in sysplex environments. For more information, see *Introducing IBM Z System Recovery Boost*, REDP-5563.

1.1.4 Transforming a transactional platform into a data powerhouse

Deriving insights from that data to drive optimal business decisions can be a challenge. To maximize the value of your data, you might need to integrate more external data sources to extract hidden insights.

In the past, typically business-critical data was copied from IBM Z to other platforms, or even data lakes, to perform sophisticated analytics. This process was inefficient, expensive, time-consuming, and introduced risk on lower-security platforms and data latency.

By accessing your data in-place with minimal duplication or movement, you can reduce the cost and complexity of analytics. You can also make data accessible to analytics applications and tools by integrating transactional and analytics processing, and protect sensitive data by keeping it within the secure Z platform.

Because of the advanced infrastructure, the z15 can support the following state-of-the-art cognitive solutions with increased cache density on each chip and up to 40 TB memory for the T01, and up to 16 TB or memory for the T02:

- IBM Open Data Analytics for z/OS

  An open source, in-place analytics solution for z/OS that simplifies big data analysis. Apache Spark, Anaconda and Java, Scala, or Python gives developers and data scientists the ability to analyze business-critical z/OS data in place with no data movement. IBM Open Data Analytics for z/OS can also provide a federated view by accessing and analyzing distributed and local data.

- IBM Machine Learning for z/OS

  A comprehensive solution that manages the entire machine-learning workflow, beginning with quick ingestion and transformation of Z data where it is stored. The solution then securely creates, deploys, and manages high-quality, self-learning behavior models to help you extract hidden insights that more accurately anticipate organizational needs.

- The IBM Db2® Analytics Accelerator for z/OS

  A high-performance appliance that transforms your Z platform into a highly efficient transactional and analytics-processing environment. It supports the full lifecycle of a real-time analytics solution on a single system that integrates transactional data, historical data, and predictive analytics.

For more information about analytics on IBM Z, see Real Time Analytics and Machine Learning on IBM Z Mainframes.

1.1.5 Getting more out of the platform with operational analytics

Today, demands for 24 x 7 high-performance operations continue to rise. At the same time, allowed service windows shrink and are much less frequent. Increasing system complexity makes planning, maintaining, and troubleshooting more difficult and time-consuming. IT operations analytics represent a possible solution to this challenge.

The z15 provides the infrastructure to host real-time analytics tools so you can clearly observe your operating environment, then maximize operational efficiencies to help reduce costs.

IBM designed IBM Z Operations Insight® Suite to ensure that your Z serve operates at peak performance. To get the most out of your system, IBM Z Operations Insight Suite adds deep insights that are based on IBM industry-leading expertise into your Z operational data. For more information, see IBM Z Operations Insight Suite.

1.1.6 Accelerating digital transformation with agile service delivery

An effective DevOps solution breaks down existing development silos, unifies infrastructure platforms, and enables ongoing deliveries. z15 provides the scalable and secure infrastructure for enterprises that must rapidly create and deliver critical applications and services while meeting agreed-on levels for quality, availability, regulatory compliance, and end-use expectations.

IBM DevOps for Z solutions operate from application understanding through deployment and management. In addition, DevOps for Z solutions give you a single, cost-effective toolset to maintain and modernize valuable applications on both Z and distributed platforms.

For example, Application Discovery and Delivery Intelligence (ADDI) helps development teams understand application interdependencies, complexity, and quality across platforms, environments, and languages. This ability gives your teams an edge in identifying potential API candidates, and provides insight about maintainability and complexity. As a result, the candidate API quality rises, and the user experience improves.

In addition, the z15 provides the infrastructure to support the mission-critical workloads of cloud services. The new high-performance processors, large memory, and enhanced access to data enable the z15 to integrate business transactions, operational data, and analytics into a single workflow.

The z15 is designed as a strategic asset to power the API economy.[5] Using the API economy demands fortified clouds, which can be open, private, public, and hybrid. The z15 gives you the hardware platform necessary to support multi-clouds.

For Linux assets, Z platforms are optimized for open source software, enhanced scalability, and sharing, while focusing on business continuity to support cloud deployments. For traditional z/OS-based assets, Z offerings provide intuitive tools to help developers speed Representational State Transfer (RESTful) API development.

No matter which asset class you choose, the z15 enables cloud application developers to incorporate z/OS business-critical data and transactions into their applications without needing to understand z/OS subsystems.

Additionally, IBM Z/OS Authorized Code Scanner can provide automated system integrity testing in a development and testing environment as part of DevSecOps modernization. It scans for Program Calls and Supervisor Calls (SVCs) that are available to all address spaces in a z/OS image and generates a series of tests that dynamically scan them for integrity. The boundary between an unauthorized caller and a Program Call or SVC routine running authorized is essential to the system integrity of the z/OS solution stack. The output of this scan provides in-depth diagnostics whenever a potential vulnerability is found to facilitate remediation to further strengthen the security posture of z/OS.

[5] For more information about the API economy, see Reach new customers with the API economy.

For more information about Enterprise DevOps, see Enterprise DevOps for IBM Z Mainframes.

1.1.7 Revolutionizing business processes

Blockchain is poised to revolutionize how industries do business. It is a technology for transactional applications that establishes trust, accountability, and transparency while streamlining business processes.

In a blockchain network, members can access a distributed, shared ledger that is *cryptographically* secure, updated by consensus, and becomes an immutable, indelible record of all *transactions*. The ledger functions as a single source of "truth". Considering that blockchain is all about increasing trust in *business transactions*, it makes perfect sense to run blockchain for business on Z.

Depending on your business or regulatory policies, you can choose an on-premises installation that is supported by IBM certified images that are running on Linux on IBM Z, or you can choose the IBM Blockchain Platform for IBM Cloud. For more information about IBM Blockchain, see Blockchain transactions on IBM Z.

1.1.8 Blending open source and IBM Z technologies

The right blend and balance of open source technologies, independent software vendor (ISV) tools, and IT platform is key to enable businesses and organizations to deliver change at a much quicker pace. To this end, IBM created a system of customers, Business Partners, and ISVs who are engaging in an open source development community. This community brings the most important and most sought-after foundational open source technologies to its IT platforms. In addition, IBM is a member of many open-standard organizations and software governance consortia that help to shape the future of open source software.

The combination of a robust and securable hardware platform with the power of a Linux distribution can optimize the building, testing, and deploying of modern applications. Also, it can accommodate scale-out clusters and scalable clouds.

The z15 provides a secure, massive-capacity Linux platform that can be a stand-alone deployment, or can run side by side with z/OS or IBM z/VSE environments on a single physical platform. Therefore, you can easily integrate Linux workloads on the z15 with z/OS and IBM z/VSE® solutions that benefit from collocated data and applications with fast internal communication and improved availability.

Linux on IBM Z gives you the performance and vertical scale that you need to meet the demands of your digital enterprise while reducing the costs of server sprawl. Combined with the integration benefits, Linux on IBM Z allows you to deploy innovative new services or cognitive analytics, and consolidate x86 workloads.

In addition, deploying Linux on the z15 can benefit your bottom line. Compared to virtualized x86 alternatives and public cloud solutions, the lower costs for administration and management, software licensing, business continuity, and floor space can reduce your total cost of ownership.

Red Hat OpenShift Container Platform for IBM Z is an enterprise Kubernetes platform that enables you to develop, deploy, run, and manage your cloud-native solutions. It is the foundation for the IBM Cloud Paks, IBM z/OS Cloud Broker, and other containerized workloads.

Red Hat Ansible Certified Content for IBM Z provides the easy-to-use, industry-leading standard automation platform for automating IBM Z processes. On top of the capable and powerful Ansible Automation Platform, Ansible Certified Content for IBM Z enables users to manage IBM Z resources more simply with Ansible. This enables all IT employees to leverage a single automation platform, with support from both Red Hat and IBM. You can automate use cases today across z/OS, Linux on Z, IBM Z HMC, IBM Z System Automation, IBM Z CICS, IBM Z IMS, IBM z/OS Connect Enterprise Edition, and iBM UrbanCode Deploy.

IBM Z/VM is a powerful hypervisor that is supporting tens of thousands virtual guests on z15, leveraging pervasive encryption to guarantee security and isolation. The KVM hypervisor is also available on the z15 platform.

For more information about IBM Z virtualization, see Server virtualization.

For more information about Linux on IBM Z, see Linux on IBM Z.

1.2 z15 technical description

When compared to their predecessors, the IBM z15 models offer several improvements, such as faster, more efficient, and redesigned high-frequency chips and other granularity options. These models also deliver better availability, faster encryption, and new on-demand and system recovery options.

The footprint for the z15 is built with a 19-inch format that scales 1 - 4 frames for the z15 T01 depending on the configuration, and a single frame for the z15 T02, which means that the amount of floor space within the data center for your Z environment is reduced in most cases. More importantly, the z15 fits within your standard data center aisles. This design feature offers options for hot-aisle and cold-aisle configurations.

1.2.1 Technical highlights

The z15 is a highly scalable symmetric multiprocessor (SMP) system. The architecture ensures continuity and upgradeability from the previous z14 and IBM z13® models. The z15 T01 model has five orderable features: Max34, Max71, Max108, Max145, and Max190. The z15 T02 model has five orderable features: Max4, Max13, Max21, Max31, and Max65. The feature names are based on the maximum number of available processor units (PUs).

Table 1-1 lists the main technical enhancements in the z15 over its predecessor platforms.

Table 1-1 Technical highlights of the z15

Capability	z15 T01	z15 T02
Greater total system capacity and more subcapacity settings for CPs. The IBM z/Architecture® ensures continuity and upgradeability from previous models.	Up to 190 characterizable PUs. Up to 34 CPs support 292 subcapacity settings.	Up to 65 characterizable PUs. Up to 6 CPs with 156 subcapacity settings.
Multi-core, single-chip modules (SCMs) that are running to help improve the execution of processor-intensive workloads.	5.2 GHz (14 nm FINFET Silicon-On-Insulator [SOI]).	4.5 GHz (14 nm FINFET Silicon-On-Insulator [SOI]).
More real memory per system, which ensures high availability in the memory subsystem by using proven redundant array of independent memory (RAIM) technology.	Up to 40 TB of addressable real memory per system.	Up to 16 TB of addressable real memory per system.

Capability	z15 T01	z15 T02
A large fixed hardware system area (HSA) that is managed separately from customer-purchased memory.	256 GB.	160 GB.
Proven technology (fifth-generation high frequency and third-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 co-processors and new hardware instructions for accelerating selected workloads.		
Processor cache structure improvements and larger cache sizes to help with more of today's demanding production workloads. Both z15 models have the following levels of cache: ► First-level cache (L1 private): 128 KB for instructions, 128 KB for data ► Second-level cache (L2): 4 MB for instructions and 4 MB for data ► Third-level cache (L3): 256 MB ► Fourth-level cache (L4): 960 MB		
Improved cryptographic functions and performance over their predecessor platforms, which is achieved by having one dedicated cryptographic co-processor per PU. New cryptographic capabilities with Crypto Express7S, such as secure key enhancements for key management, digital signature creation, and verification with new keys and algorithms.		
zHyperLink provides a low-latency connection to storage subsystems for faster data retrieval. This feature is consistent with their predecessor platforms.		
Compared to their predecessor platforms, improved data compression operations are achieved by having one dedicated compression co-processor per PU, and new hardware instructions. In addition, the On-Chip compression accelerator provides a per-chip replacement for the IBM zEnterprise® Data Compression (zEDC) Express feature.		
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 logical partitions (LPARs).	► Six LCSS ► 85 LPARs ► Four subchannel sets ► 64,000 I/O devices per subchannel set	► Three LCSS ► 40 LPAR ► Three subchannel sets ► 64,000 I/O devices per subchannel set

You can compare the z15 models to the previous two IBM Z generations by using the Compare IBM Z tool.

1.2.2 IBM z15 features and functions

Common to both models (T01 and T02), z15 includes the following features and functions:

- Enhanced LPAR resource allocation algorithms for PUs and memory.
- IBM Virtual Flash Memory (VFM) can be used to handle paging workload spikes and improve availability. VFM is the replacement for the Flash Express features, which were available on the z13.
- The Crypto-Express7S feature now supports ISO-4 format PIN blocks across Common Cryptographic Architecture (CCA) financial services. Similarly, support was added for the ANSI X9.24-3 Advanced Encryption Standard (AES) Derived Unique Key Per Transaction (DUKPT) standard.
- New dedicated on-chip compression accelerator for faster compression algorithms and reduced overhead.
- Data Privacy Passports provide control and management of data privacy (for eligible data) on the z15 and as it moves to the enterprise.
- Secure Execution for Linux removes the security issue of running multiple VMs in the same server.
- Secure Service Container to build and host highly secure virtual appliances.
- The 25 GbE RDMA over Converged Ethernet (RoCE) Express2.1 provides a technology refresh for RoCE on IBM Z. This feature provides increased networking performance to take advantage of higher speed processors.
- Coupling Express Long Reach (CE LR) for coupling links that must extend up to 10 km (6.21 miles).
- Dynamic I/O for a stand-alone Coupling Facility (CF) is a means for activating a new or changed I/O configuration without disruption.
- New resiliency mechanisms for CF include Fair Latch Manager 2 and Message Path SYID resiliency enhancement.
- zHyperLink Express 1.1 is a direct-connect I/O feature that works with the FICON SAN infrastructure to reduce latency and improve workload response time.
- There is a new 19-inch format for the frame and support for American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) Class A3 data centers.

IBM z15 Model T01 (8561) specific functions

Exclusive to z15 T01, in addition with the following features and functions are available:

- 40 TB Memory (max.) with up to 16 TB per LPAR.
- Up to 12 IBM VFM features can be used to handle paging workload spikes and improve availability. VFM is the replacement for the Flash Express features, which were available on the z13.
- 85-domain support for Crypto Express features.
- System Recovery Boost with the optional System Recovery Boost Upgrade (temporary capacity upgrade) is a function to accelerate operating system and services start and shutdown times.
- Next-generation FICON Express16SA for the z15 T01 provides similar benefits for workloads. And, with the optional IBM Fibre Channel Endpoint Security feature (Feature Code 1146), it provides encryption for data in flight between z15 T01 and select IBM DS8000® models.
- OSA-Express7S 25 GbE SR1.1 meets increased networking performance demands that are driven by high-speed processors and faster network-attached devices.
- Newly designed front to rear radiator cooling system. The radiator pumps, blowers, controls, and sensors are N+2 redundant.

The z15 T01 is compared in Figure 1-2 with previous Z platforms in the following key areas:

- Memory
- System I/O bandwidth
- Single-engine (one-way) processor capacity index[6]
- Number of PUs

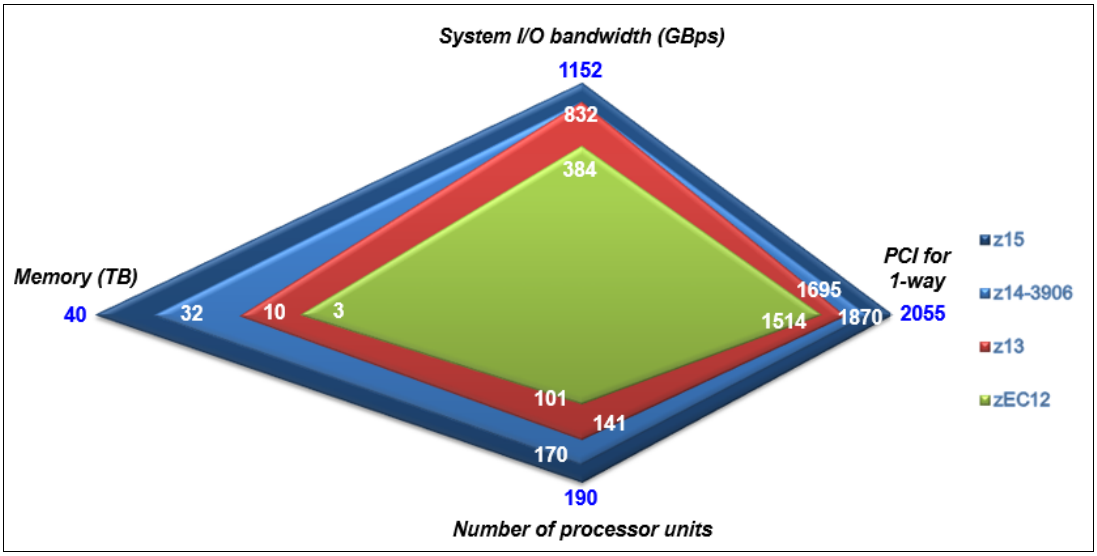

Figure 1-2 Balanced system design: z15 Model T01 versus its predecessors

[6] Based on the processor capacity index. For more information about processor capacity index values, see *Large Systems Performance Reference*, SC28-1187.

IBM z15 Model T02 (8562) specific functions

The z15 T02 includes the following features and functions:

- Up to 16 TB Memory per system with up to 8 TB per LPAR.
- Up to four IBM VFM features can be used to handle paging workload spikes and improve availability.
- 40-domain support for Crypto Express features.
- System Recovery Boost is a function to accelerate operating system start and shutdown times and speed up process recovery in a Parallel Sysplex environment.
- Three-phase power with single-phase option (depending on system configuration).
- Air-cooled system with redundant power and cooling elements.
- 16U reserved space feature increases flexibility for customer-supplied IBM DS8910F storage.

The z15 T02 is compared in Figure 1-3 on page 14 with previous Z platforms in the following key areas:

- Memory
- System I/O bandwidth
- Single-engine (one-way) PCI[6]
- CP capacity

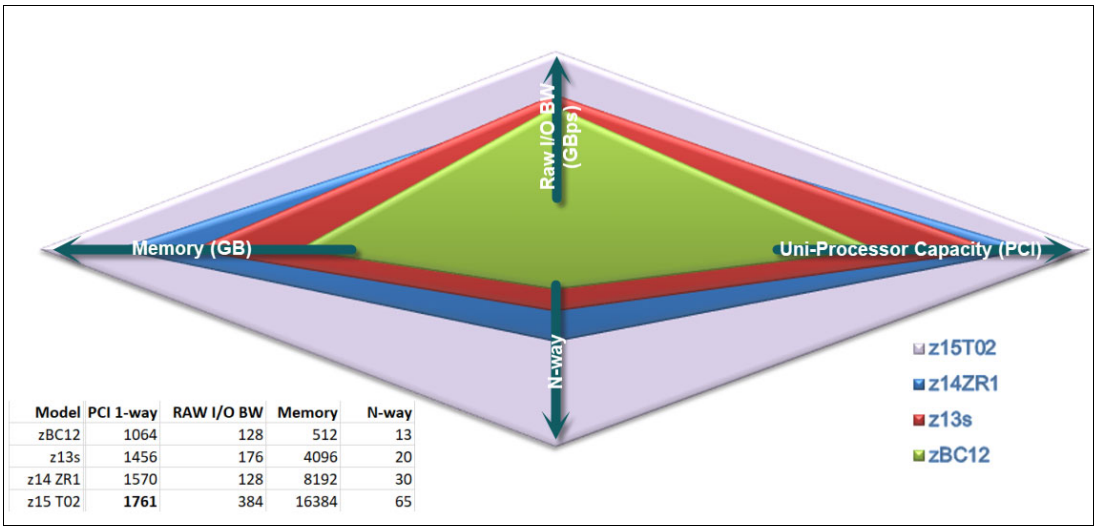

Model	PCI 1-way	RAW I/O BW	Memory	N-way
zBC12	1064	128	512	13
z13s	1456	176	4096	20
z14 ZR1	1570	128	8192	30
z15 T02	**1761**	384	16384	65

Figure 1-3 Balanced system design: z15 T02 versus its predecessors

For specific information about each z15 model, see Chapter 2, “IBM z15 Model T01 hardware overview” on page 23 and Chapter 3, “IBM z15 Model T02 hardware overview” on page 41.

For information about IBM z15 functions and features, see *IBM z15 (8561) Technical Guide*, SG24-8851 and *IBM z15 (8562) Technical Guide*, SG24-8852.

1.2.3 Storage connectivity

Storage connectivity is provided on the z15 by 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 (as compared to FC) provides for communication across channels, channel-to-channel (CTC) connectivity, and with FICON devices, such as disks, tapes, and printers. It is used in z/OS, IBM z/VM®, IBM z/VSE (Virtual Storage Extended), z/TPF (Transaction Processing Facility), 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/SCSI devices. FCP is used by z/VM, KVM, z/VSE, and Linux on IBM Z environments.

FICON Express16SA[7] features are implemented by using PCIe cards, and offer better port granularity and improved capabilities over the previous FICON Express features. FICON Express16SA supports a link data rate of 16 gigabytes per second (GBps) (8 or 16 Gbps auto-negotiate), and it is the preferred technology for new systems.

zHyperLink Express

zHyperLink was created to provide fast access to data by way of low-latency connections between the Z platform and storage.

The zHyperLink Express1.1 feature allows you to make synchronous requests for data that is in the storage cache of the IBM DS8900F. This process is done by directly connecting the zHyperLink Express1.1 port in the z15 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 reads and writes, 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 cuts I/O-sensitive workload response time in half without requiring application changes.[8]

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 available FICON Express and zHyperLink Express features, see 4.2, “Storage connectivity” on page 59.

[7] Supported by the z15 T01 for Fibre Channel Endpoint Security only. The FICON Express16S+ is supported on both z15 models, as carry forward for z15 T01 and new build or carry forward for z15 T02.

[8] The performance results can vary depending on the workload. Use zBNA tool for the zHyperLink planning.

1.2.4 Network connectivity

The z15 is a fully virtualized platform that can support many system images at once. Therefore, network connectivity covers not only the connections between the platform and external networks with Open Systems Adapter-Express (OSA-Express) and RoCE Express features, it supports specialized internal connections for intra-system communication through IBM HiperSockets and Internal Shared Memory (ISM).

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 allows significant performance benefits by offloading processing from the operating system.

OSA-Express7S features continue to support 1000BASE-T Ethernet for copper environments and 10-Gigabit Ethernet and Gigabit Ethernet fiber optic (single-mode and multimode) environments. The OSA-Express7S 25 GbE1.1[9] feature supports higher throughput by way of 25 GbE network infrastructures.

HiperSockets

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

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

RoCE Express

The 25 GbE and 10 GbE RoCE Express2.1 features[10] use Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) to provide fast memory-to-memory communications between two Z platforms.

These features are designed to help reduce consumption of CPU resources for applications that use the TCP/IP stack (such as IBM WebSphere® that accesses an IBM Db2 database). They can also 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, providing a highly optimized network interconnect for 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

[9] OSA-Express7S 25 GbE SR1.1 is supported only on IBM z15 T01 along with all OSA-Espress7S features. IBM z15 Model T02 supports OSA-Express6S features (new build and carry forward) and the OSA-Express7S 25 GbE Short Reach (SR) (new build and carry forward).

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

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 by way of memory-to-memory transfer between z/OS and Linux on Z.

For more information about the available network connectivity features, see 4.3, “Network connectivity” on page 61.

1.2.5 Cryptography

z15 provides two major groups of cryptographic functions: CPACF and Crypto-Express7S.

CPACF

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

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

With z15, CPACF supports Elliptic Curve Cryptography (ECC) clear key, improving the performance of Elliptic Curve algorithms. The following algorithms are supported:

- EdDSA (Ed448 and Ed25519)
- ECDSA (P-256, P-384, and P-521)
- ECDH (P-256, P-384, P521, X25519, and X448)
- Support for protected key signature creation

Crypto-Express7S

The tamper-sensing and tamper-responding Crypto-Express7S features provide acceleration for high-performance cryptographic operations and support up to 85 domains with the z15 T01 and 40 domains with the z15 T02. This specialized hardware performs AES, DES and TDES, RSA, Elliptic Curve (ECC), SHA-1, and SHA-2, and other cryptographic operations.

It supports specialized high-level cryptographic APIs and functions, including those functions that are required in the banking industry. Crypto-Express7S features are designed to meet the Federal Information Processing Standards (FIPS) 140-2 Level 4 and PCI HSM security requirements for hardware security modules.

The z15 offers better AES performance than the z14, a True Random Number Generator, SHA3 support, and RSA/ECC acceleration.

For more information about cryptographic features and functions, see 4.7, “Cryptographic features” on page 67.

1.2.6 Clustering connectivity

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

Coupling connectivity on the z15 use CE LR and Integrated Coupling Adapter Short Reach (ICA SR) features. The ICA SR feature supports distances up to 150 meters (492 feet), while the CE LR feature supports unrepeated distances of up to 10 km (6.21 miles) between Z platforms.

For more information about coupling and clustering features, see 4.4, “Clustering connectivity” on page 64.

1.2.7 Special-purpose features and functions

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

The z15 delivers a range of features and functions, allowing PUs to concentrate on computational tasks, while distinct, specialized features take care of the rest. The following special-purpose features and functions are offered with the z15:

- Data compression

  z15 enhanced the compression capability by taking it from the I/O device level (zEDC Express feature) and moving it to the Nest Accelerator Unit on the processor chip, adding the Deflate compliant (lossless data compression algorithm) and GZIP (GNU zip - UNIX compression utility) compression and decompression support as hardware instructions.

  This innovation performs the compression with improved performance and simplified management on a processor chip level without any of the delays that are associated with I/O requests, with minimal CPU costs.

  The enhancement preserves the compatibility between the zEDC Express feature and on-chip compression. Data that is compressed and written by zEDC is read and decompressed on the z15. This feature simplifies the migration to IBM z15 (on-chip compression replaces the zEDC Express feature).

- IBM Integrated Accelerator for Z Sort

  The Integrated Accelerator for Z Sort, which is standard on the z15, optimizes elapsed time for sort workloads, which typically occur during batch windows. By providing one sort accelerator per core, you can accelerate frequently used functions to speed up sorting, shorten batch windows, and improve select database functions, such as reorganization.

  A new SORTL instruction is used by DFSORT and the IBM Db2 Utilities for z/OS Suite to help reduce CPU usage and improve elapsed time for sort workloads.

- GDPS Virtual Appliance

  The GDPS Virtual Appliance is a fully integrated, continuous availability, and disaster recovery (DR) solution for Linux on IBM Z that can help improve availability and time-to-value.

- Guarded Storage Facility (GSF)

  GSF is an architecture that was introduced with z14 to enable enterprise-scale Java applications to run with minimal periodic pauses for garbage collection on larger heaps.

- Instruction Execution Protection Facility (IEPF)

  Instruction Execution Protection is hardware function that was introduced with z14 to enable software, such as IBM Language Environment®, to mark specific memory regions (for example, a heap or stack), as non-executable to improve the security of programs that run on Z against stack-overflow or similar attacks.

- SMT

  With SMT, you can process up to two simultaneous threads in a single core to optimize throughput. An operating system with SMT support can be configured to dispatch work to a thread on an IBM Z Integrated Information Processor (zIIP)[11] or an Integrated Facility for Linux (IFL)[12]. SAP engines also support SMT.

- SIMD

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

- Dynamic memory relocation

  This feature supports of Enhanced Drawer Availability (EDA) and Concurrent Drawer Replacement (CDR). With this technology, an algorithm is used to dynamically move memory between central processor complex (CPC) drawers to improve performance without impacting the operating system.

- Coupling and Parallel Sysplex

  Data sharing and serialization functions are offloaded to the CF by way of a special coupling link network that provides the infrastructure to run a single production workload. This workload accesses a common set of data across many z/OS system images.

For more information about these features and other z15 features, see in Chapter 5, “IBM z15 system design strengths” on page 71.

1.2.8 Capacity on Demand and performance

The z15 enables just-in-time deployment of processor resources. The Capacity on Demand (CoD) function allows users to dynamically change available system capacity. This function helps companies respond to new business requirements with flexibility and precise granularity.

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

Also contributing to the extra capacity on the z15 are numerous improvements in processor chip design, including new instructions, multithreading, and redesigned and larger caches.

In its maximum configuration, the z15 T01 Max190 can deliver up to 25 percent[13] more capacity than the largest 170-way z14 Model M05. A z15 T01 1-way system has

[11] An zIIP is used under z/OS for designated workloads, which include IBM Java virtual machine (JVM), various XML System Services, and others.

[12] An IFL is exclusively used with Linux on IBM Z and for running the z/VM or the KVM hypervisor in support of Linux.

approximately 12 percent more capacity than a z14 M0x 1-way system. In addition, the z15 T01 can be ordered with 25% more configurable memory compared to the IBM z14 M0x.

The IBM z15 T02 with the Max65 feature (65 configurable cores), is an increase of 35 over the z14 ZR1, with single processor capacity of z15 T02 being approximately 14% greater than a single processor of the z14 ZR1. In addition, the z15 T02 can be ordered with double the c

For more information, see 5.3.2, “Capacity on Demand” on page 86, and 5.3.3, “z15 performance” on page 88.

1.2.9 Reliability, availability, and serviceability

The z15 offers the same high quality of service and reliability, availability, and serviceability (RAS) that is traditional in Z platforms. 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.

For more information about RAS, see 5.4, “Reliability, availability, and serviceability” on page 91.

For IBM Z resiliency support, see *Getting Started with IBM Z Resiliency*, SG24-8446.

1.3 z15 software support

The z15 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 (IBM IMS), and IBM Db2. It also includes the following web services (among others):

- Java platform
- Linux and open standards applications
- WebSphere
- IBM MobileFirst® Platform Foundation for mobile application development
- IBM z/OS Connect Enterprise Edition

The following operating systems are supported on the z15:

- z/OS Version 2 Release 5[14]
- z/OS Version 2 Release 4
- z/OS Version 2 Release 3
- z/OS Version 2 Release 2 with program temporary fixes (PTFs)

[13] Variations on all the observed increased performance depend on the configuration workload type.
[14] Preview announcement at the time of this writing - see IBM United States Software Announcement 221-057.

- z/OS Version 2 Release 1 with PTFs[15]
- z/VM Version 7 Release 2
- z/VM Version 7 Release 1 with PTFs
- z/VM Version 6 Release 4 with PTFs[16]
- z/VSE Version 6 Release 2 with PTFs
- z/TPF Version 1 Release 1 (compatibility support)

IBM plans to support running the following Linux on IBM Z distributions on IBM z15:

- SUSE SLES 15 SP1 with service and SUSE SLES 12 SP4 with service.
- Red Hat RHEL 8.0 with service, Red Hat RHEL 7.7 with service, and Red Hat RHEL 6.10 with service.
- Ubuntu 20.04 with service, 18.04.1 LTS with service and Ubuntu 16.04.6 LTS with service.

The support statements for the z15 also cover the KVM hypervisor on distribution levels that have KVM support.

Detailed service levels are identified during toleration tests. For more information about recommended distribution levels, see IBM tested and supported Linux environments.

For more information about the z15 software support, see Chapter 6, “Operating system support” on page 103.

1.3.1 IBM compilers

Compilers are built with specific knowledge of the system architecture, which is used during code generation. Therefore, the use 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 existing IT investments, while smoothly incorporating new, web-based applications into the infrastructure.

z/OS, XL C/C++, and XL C/C++ for Linux on 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 v2.3 XL C/C++, are inherently optimized on the IBM z15 because they use floating point registers rather than memory or fast mathematical computations. The use of compilers that take advantage of hardware enhancements is key to improving application performance, reducing CPU usage, and lowering operating costs.

[15] Compatibility only. An extended support contract for IBM Software Support Services for z/OS is required with PTFs.

[16] z/VM 6.4 support has ended 3/31/2021

2

IBM z15 Model T01 hardware overview

This chapter expands on the descriptions of the key hardware elements of the z15 that were presented in 1.2, “z15 technical description” on page 10. It includes the following topics:

- 2.1, “Models and upgrade paths” on page 24
- 2.2, “Frames and cabling” on page 26
- 2.3, “CPC drawers” on page 28
- 2.4, “I/O system structure” on page 34
- 2.5, “Power and cooling” on page 37

For more information about the key capabilities and enhancements of the z15, see *IBM z15 (8561) Technical Guide*, SG24-8851.

2.1 Models and upgrade paths

The IBM z15 (machine type 8561) has one model: the T01. The maximum number of characterizable processors is represented by feature names Max34, Max71, Max108, Max145, and Max190.

Naming: Throughout this chapter, we refer to the IBM z15 Model T01 (machine type 8561) as *z15 T01*.

As with its predecessors, the z15 T01 central processor complex (CPC) is built by using a processor unit (PU) single-chip modules (SCMs). Each SCM can have 9 - 12 active PUs, or cores. Spare PUs, System Assist Processors (SAPs), and one Integrated Firmware Processor (IFP) are included in the z15 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 31.

Table 2-1 z15 T01 processor unit configurations

Feature name	Number of CPC drawers	Feature code	Characterizable processor units	Standard SAPs	Spares
Max34	1	0655	1 - 34	4	2
Max71	2	0656	1 - 71	8	2
Max108	3	0657	1 - 108	12	2
Max145	4	0658	1 - 145	16	2
Max190	5	0659	1 - 190	22	2

The supported upgrade paths for the z15 T01 are shown in Figure 2-1 on page 25.

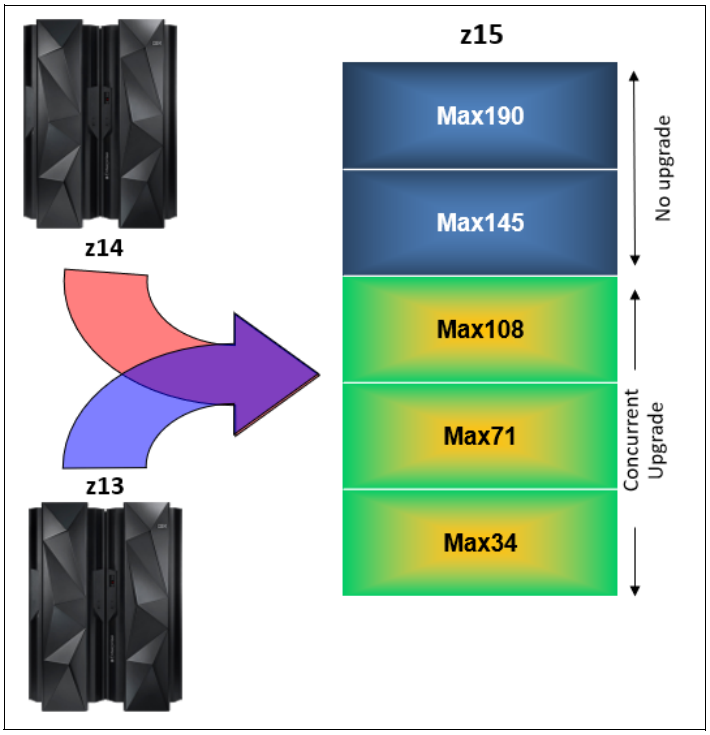

Figure 2-1 z15 T01 upgrade paths

If an upgrade request cannot be accomplished with the existing configuration, a hardware upgrade is required in which one or more CPC drawers are added to accommodate the wanted capacity. On the z15 T01, more CPC drawers can be installed concurrently from a Max34 to a Max71, and to a Max108.

Note: No field upgrade is available to a Max145 or a Max190 because these two features are factory shipped only.

On the z15 T01, 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 a z15 T01.

In addition, the z15 T01 offers 292 capacity levels. In all, 190 capacity levels, which are based on the number of physically used CPs, are available, plus up to 102 other subcapacity models for the first 34 CPs. For more information, see 5.3.1, “Capacity settings” on page 84.

2.2 Frames and cabling

The z15 T01 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 z15 T01 packaging introduces new configuration options as compared to previous Z platforms. See Table 2-2.

Table 2-2 z15 configuration options compared to z13 and z14

System	Number of frames	Number of CPC drawers	Number of I/O drawers	I/O and power connections	Power options[a]	Cooling options
z15	1 - 4	1 - 5	0 - 12[b]	Rear only	PDU or BPA	Radiator (air) or water-cooling unit (WCU)
z14	2	1 - 4	0 - 5	Front and rear	BPA	Radiator (air) or WCU
z13	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, while 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 11 if ordered with a BPA.

The number of PCIe+ I/O drawers can vary based on the number of I/O features, power options (PDU or BPA), and number of CPC drawers 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, the z15 T01 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, which potentially eliminates cables to be run under a raised floor and increases air flow over the system.

The radiator-cooled z15 T01 supports installation on raised floor and non-raised floor environments. For the water-cooled system, only the raised floor option is available.

Figure 2-2 on page 27 shows the front view of a fully configured z15 T01 with radiator cooling, five CPC drawers, and 12 PCIe+ I/O drawers.

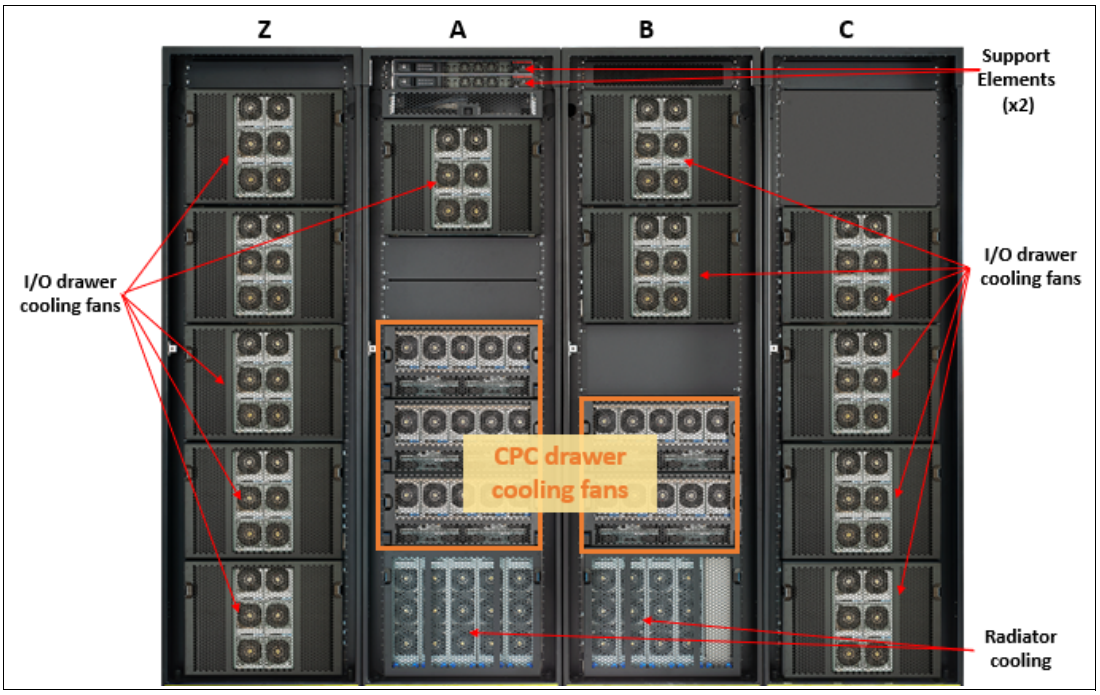

Figure 2-2 Front view of a fully configured z15 T01 with radiator cooling

Figure 2-3 shows the rear view of a fully configured z15 T01 with water cooling.

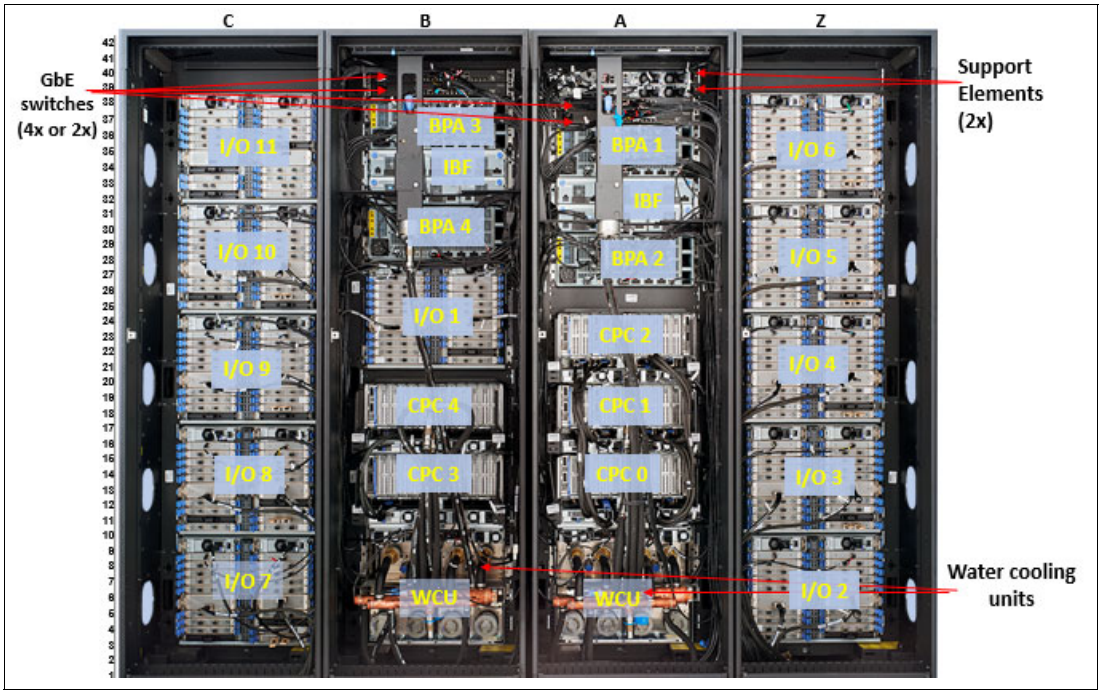

Figure 2-3 Rear view of a fully configured z15

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 z15 T01 configuration include the following examples:

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

2.3 CPC drawers

The z15 T01 can hold up to five CPC drawers (three in the A Frame and two in the B Frame). Each CPC drawer contains the following elements:

- SCMs:
  - Four PU SCMs, each containing 9 - 12 PU cores (each cooled by an internal water loop).
  - One Storage Controller (SC) SCM, with a total of 960 MB L4 cache.
- Memory:
  - A minimum of 512 GB and a maximum of 40 TB of memory (excluding 256 GB for hardware system area (HSA)) is available for use. For more information, see Table 2-3 on page 32.
  - Up to 20 dual inline memory modules (DIMMs) are plugged in a CPC drawer that are 32 GB, 64 GB, 128 GB, 256 GB, or 512 GB.
- Fanouts

  The CPC drawer provides 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 PCIe 16 gigabytes per second (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, 8 GBps each).
- Three or four Power Supply Units (PSUs), depending on the configuration (BPA or PDU), 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 Flexible Support Processor (FSP) \ oscillator cards (OSCs), which provide redundant interfaces to the internal management network and provide clock synchronization to the Z CPCs.
- 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 water-cooled.

The CPC drawer communication topology is shown in Figure 2-4 on page 29. All CPC drawers are interconnected with high-speed communications links (A-Bus) through the SC chip L4 shared caches. Symmetric multiprocessor (SMP) cables are used to interconnect all the CPC drawers. The X-Bus provides connectivity between PUs within the logical clusters and the SC on the drawer.

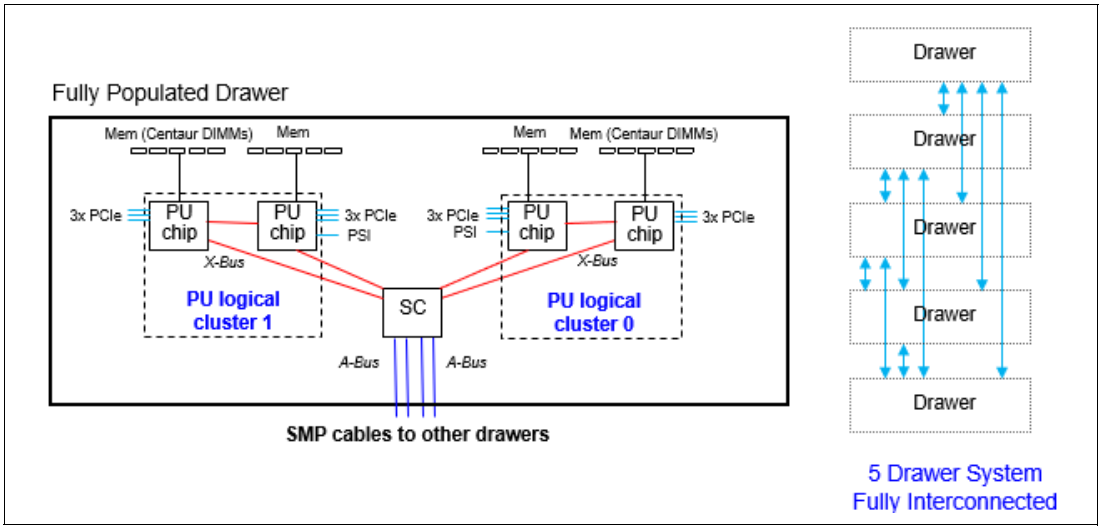

Figure 2-4 z15 T01 CPC drawer communication topology

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 Single-chip modules

The CPC drawer has four PU SCMs and one SC SCM. Each PU SCM supports up to 12 active PU cores, and L1, L2, and L3 caches.

The SC SCM includes 960 MB shared eDRAM cache, interface logic to the four PU SCMs, and SMP fabric logic. The SC SCM is configured to provide L4 cache that is shared by all PU cores in the CPC drawer.

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 be in execution 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.

PU cache

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 private L2 cache, with 4 MB D-cache (D for data) and 2 MB I-cache (I for instruction).
- Each PU SCM contains a 256 MB L3 cache that is shared by all PU cores in the SCM. The shared L3 cache uses eDRAM.

This on-chip cache implementation optimizes system performance for high-frequency processors, with cache improvements, new Translation/TLB2 design, pipeline optimizations, better branch prediction, new accelerators and architectures, and Secure Execution support.

The z15 cache structure is shown Figure 2-5.

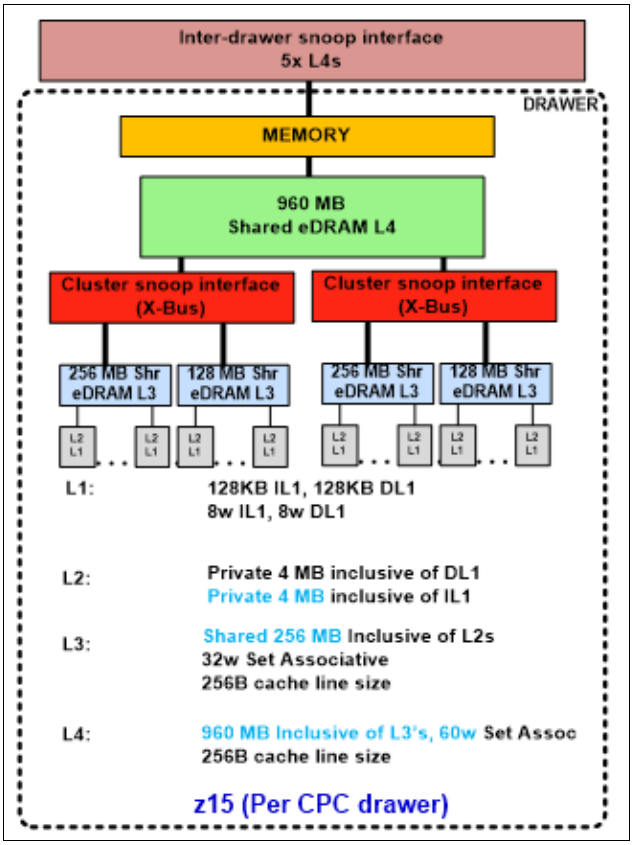

Figure 2-5 z15 T01 cache structure

Drawer cache

In addition to the on-chip (L1, L2, L3) cache, each drawer provides 960 MB L4 cache, for a total of 4,800 MB L4 cache (z15 T01 Max190).

PU sparing

Hardware fault detection is embedded throughout the system design and is 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 1.2.5, “Cryptography” on page 17. This cryptographic hardware is available with any processor type, for example, CP, zIIP, and IFL.

On-chip functions

The compression accelerator replaces the IBM zEnterprise Data Compression (zEDC) Express feature that was on previous IBM Z platforms.

The sort accelerator uses the 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.

Software support

The z15 PUs provide full compatibility with software for z/Architecture, and extend the Instruction Set Architecture (ISA) to enable enhanced functions and performance. The following hardware instructions that support more efficient code generation and execution are introduced in the z15:

- Central Processor Assist for Cryptographic Functions (CPACF)
- Compression call (CMPSC)
- CPU Measurement Facility (MF)
- Out-of-order execution
- Large page support
- IBM Virtual Flash Memory
- Instruction Execution Protection
- Guarded Storage Facility (GSF)
- Transactional Execution Facility
- Runtime Instrumentation Facility
- Single-instruction, multiple-data (SIMD)
- Secure Execution Facility for Linux

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 by using one of the following characterizations:

- CP: These standard processors are used for general workloads.
- IFLs: 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” for workloads that are restricted to Db2 type applications. Also used for System Recovery Boost, which is a z15 exclusive feature. For more information, see “Reliability, availability, and serviceability” on page 91.
- Integrated Coupling Facility (ICF): Designates processors to be used specifically for coupling.

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

Converting a 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 SAP to the z15 does not change the system capacity setting or its millions of service units (MSU) rating.

2.3.3 Memory

Maximum physical memory size is directly related to the number of CPC drawers in the system. An IBM Z platform features 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 the z15, up to 8 TB of memory per CPC drawer can be ordered and up to 40 TB for a five-CPC drawer system.

Important: z/OS V2R3 requires a minimum of 8 GB of memory (2 GB of memory when running under z/VM). z/OS can support up to 4 TB of memory in an LPAR.

The minimum and maximum memory sizes for each z15 feature are listed in Table 2-3.

Table 2-3 z15 Model T01 memory per feature

Feature name	CPC drawers	Memory
Max34 (Feature Code 0655)	1	512 GB - 8 TB
Max71 (Feature Code 0656)	2	512 GB - 16 TB
Max108 (Feature Code 0657)	3	512 GB - 24 TB
Max145 (Feature Code 0658)	4	512 GB - 32 TB
Max190 (Feature Code 0659)	5	512 GB - 40 TB

The HSA on the z15 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 z15 platforms, 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 480 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 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 has access to that memory after it is allocated. Despite the CPC drawer structure, the z15 is still 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 involving the addition of a CPC drawer, the minimum usable memory increment (256 GB) is added to the system. During an upgrade, adding a CPC drawer 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 the z15, and was enhanced to include one Memory Controller Unit (MCU) per processor chip, with five memory channels and one DIMM per channel. A fifth channel in each MCU enables memory to be implemented as 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 z15 (8561) Technical Guide*, SG24-8851.

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 z15 T01 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 the z15 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 six can be defined in z15 T01)
- Subchannel set (up to four can be defined in z15 T01)
- A total of 85 LPARs
- Devices, up to the maximum number permitted, in each subchannel set
- Logical processors by type
- Cryptographic adapters

2.4 I/O system structure

The z15 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.

The z15 CPC drawer does *not* support 12x InfiniBand and 1x InfiniBand coupling.

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

Figure 2-6 shows a high-level view of the I/O system structure for the z15 T01.

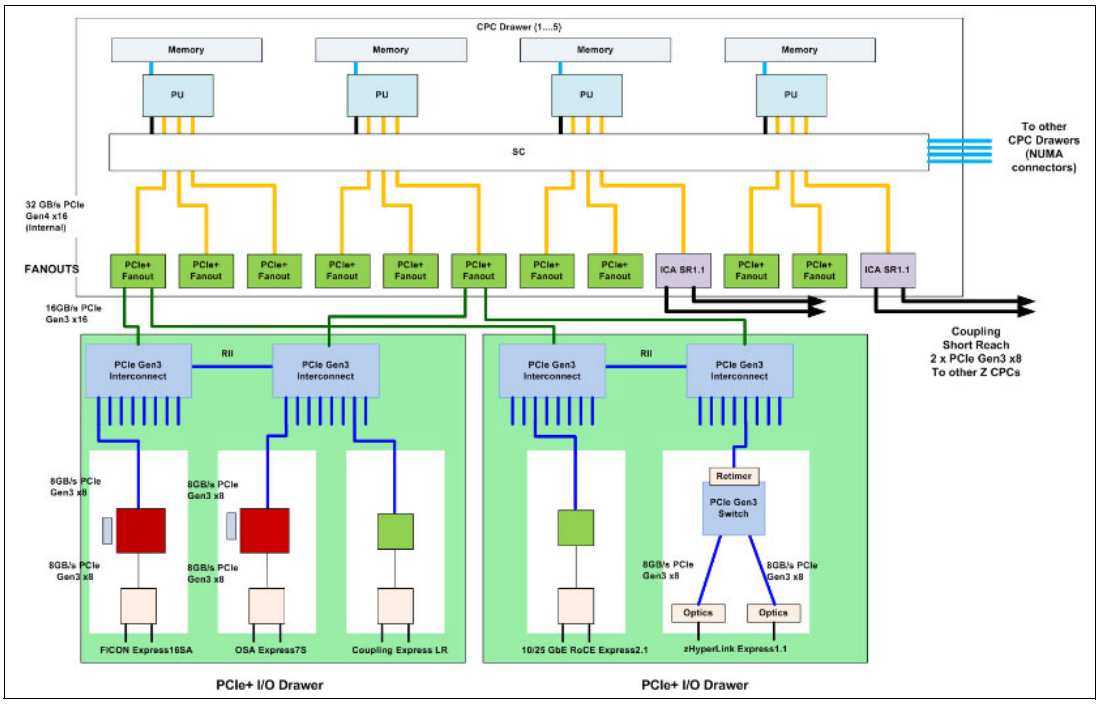

Figure 2-6 z15 T01 I/O system structure

The z15 T01 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 (Parallel Sysplex and Server Time Protocol (STP) configuration), the z15 supports the following link types:

- ICA SR1.1 and ICA SR (installed in a CPC drawer)
- Coupling Express 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).

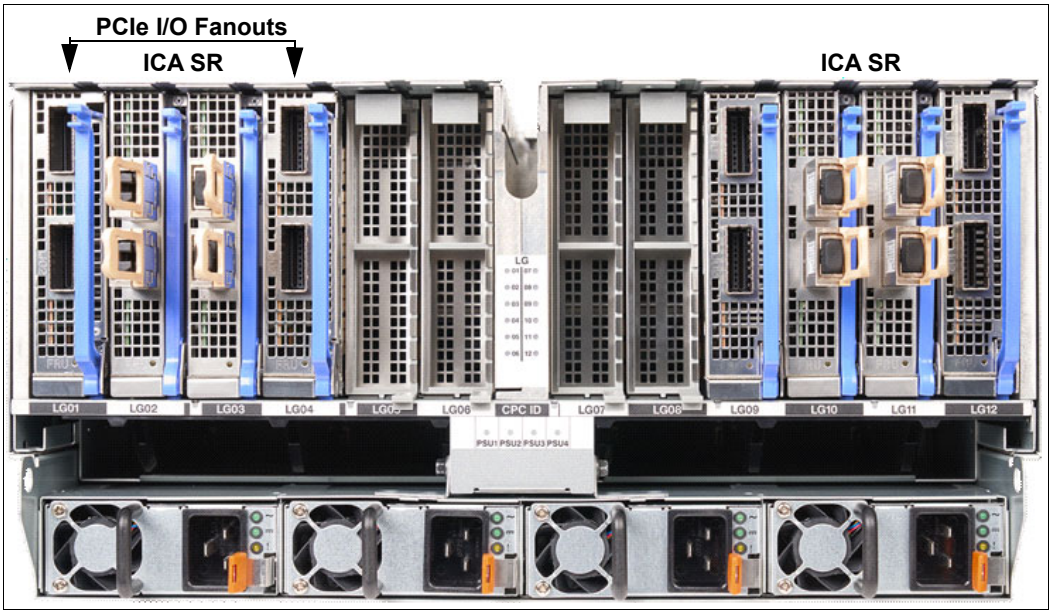

Figure 2-7 z15 T01 CPC drawer: Front view

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 two slots for PCIe switch cards.

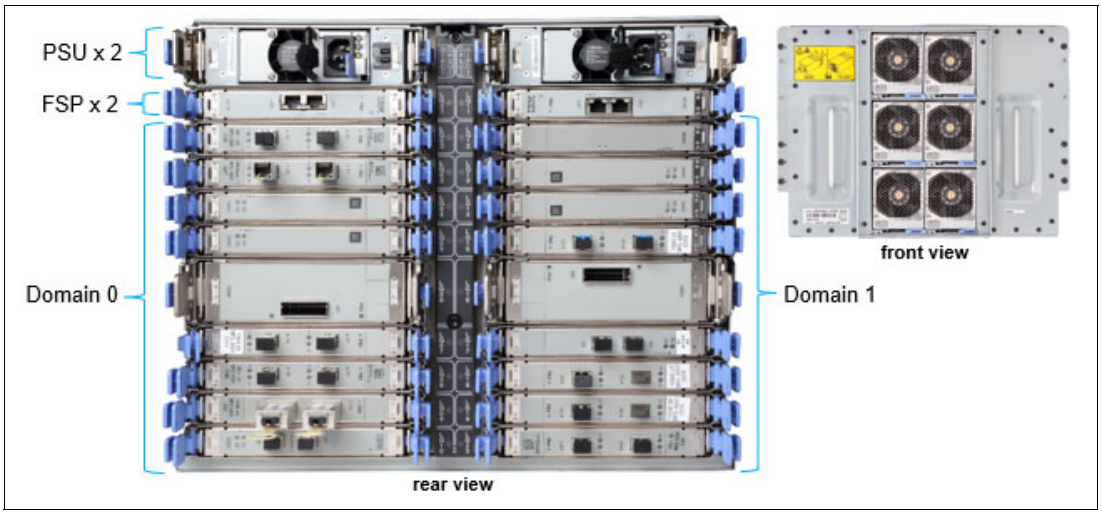

Figure 2-8 PCIe+ I/O drawer: Rear and front view

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

- FICON Express16SA, FICON Express16S+, FICON Express16S, or FICON Express8S
- OSA-Express7S, OSA-Express6S, or OSA-Express5S
- Crypto-Express7S, Crypto-Express6S, or Crypto-Express5S

- 25 GbE RDMA over Converged Ethernet (RoCE) Express2.1, 10 GbE RoCE Express2.1, or 10 GbE RoCE Express
- zHyperLink Express1.1 and zHyperLink Express
- CE LR

For more information about the I/O feature that is available with the z15, see Chapter 4, “Supported features and functions” on page 57.

2.5 Power and cooling

The z15 T01 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 z15 T01 19-inch frames are available with the following power options:

- PDU

  Use of PDU for z15 T01 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 11 PCIe+ I/O drawers. This option is required when ordered with an Internal Battery Feature (IBF), WCU, or Balanced Power.

Statement of Direction: IBM z15 is planned to be the last IBM Z server to offer an IBF. As customer data centers continue to improve power stability and uninterruptible power supply (UPS) coordination, IBM Z continues to innovate to help customers take advantage of common power efficiency and monitoring across their infrastructures. Extra support for data center power planning can be requested through your IBM Sales representative.

Bulk Power® Assembly (BPA) support removal: Based on the direction of the market, the IBM Z system *following* IBM z15™ is planned to be the last Z system to support BPA. Customers should plan to migrate from BPA to Intelligent Power Distribution Unit (iPDU).

The z15 T01 operates with one or two sets 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 is installed. Power cords attach to a three-phase, 50/60 Hz, 200 - 480 V AC power source. The loss of just 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 z15 T01 cooling system is available with two options: Radiator (air) cooling or water cooling. SCMs are always cooled with an internal water loop, no matter which cooling option that is chosen. The liquid in the internal water system can be cooled by using a radiator (for air-cooling option) or customer-supplied chilled water supply (for water-cooling option). PCIe+ I/O drawers, power enclosures, and CPC drawers are cooled by chilled air with blowers.

Conversion from air to water-cooled systems, and vice versa, is not available. The following options are available:

- Radiator (air) cooling

  The air-cooling system in the z15 T01 is redesigned for better availability and lower cooling power consumption. The radiator design is a closed-loop water-cooling pump system for the SCMs in the CPC drawers. It is designed with N+2 pumps, blowers, controls, and sensors. The radiator unit is cooled by air.

- Water cooling

Water Cooling[a]: IBM z15 is planned to be the last IBM Z server to offer customer water cooling.

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 information that is mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functions.

The z15 T01 continues to offer the choice of using a building's chilled water to cool the system by using WCU technology. The SCMs in the CPC drawer are cooled by an internal, closed, water-cooling loop. In the internal closed loop, water exchanges heat with building-chilled water (provided by the customer) through a cold plate.

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

The z15 operates with two fully redundant WCUs in Frame-A and Frame-B on separate loops (when present). One WCU in each loop can support the entire load, and replacement of a WCU is fully concurrent. During a total loss of building-chilled water or if both WCUs fail per frame, the rear door heat exchangers cool the internal water-cooling loop.

The water-cooling option is available with the BPA power option only.

2.5.3 Power considerations

Consider the following points regarding power:

- 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.
- Supports top or bottom exit power.
- 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, the number of CPC drawers, the number and type of I/O units 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 8561 Installation Manual for Physical Planning*, GC28-7002.

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

3

IBM z15 Model T02 hardware overview

This chapter expands on the descriptions of the key hardware elements of the z15 that were presented in 1.2, “z15 technical description” on page 10 with emphasis on the IBM z15 Model T02 (Machine Type 8562). It includes the following topics:

- 3.1, “Models and upgrade paths” on page 42
- 3.2, “Frames and cabling” on page 44
- 3.3, “CPC drawers” on page 46
- 3.4, “I/O system structure” on page 52
- 3.5, “Power and cooling” on page 55

For more information about the key capabilities and enhancements of the z15, see *IBM z15 (8562) Technical Guide*, SG24-8852.

Naming: Throughout this chapter, we refer to the IBM z15 Model T02 (machine type 8562) as *z15 T02*.

3.1 Models and upgrade paths

The IBM z15 Model T02 (machine type 8562) has one model: the T02. The maximum number of characterizable processors is represented by feature names Max4, Max13, Max21, Max31, and Max65.

As with its predecessors, the z15 T02 central processor complex (CPC) is built by using processor unit (PU) single-chip modules (SCMs). Each SCM can have 7 - 11 active PUs, or cores. Spare PUs, System Assist Processors (SAPs), and one Integrated Firmware Processor (IFP) are included in the z15 T02 configuration.

The number of characterizable PUs, SAPs, and spare PUs for the various features is listed in Table 3-1. For more information about PU characterization types, see “PU characterization” on page 49.

Table 3-1 z15 T02 processor unit configurations

Feature name	Number of CPC drawers/PU SCMs	Feature code	Characterizable processor units	Central processors (CPs)	Standard SAPs	Spares
Max4	1/1	0649	1 - 4	0 - 4	2	1
Max13	1/2	0650	1 - 13	0 - 6	2	1
Max21	1/3	0651	1 - 21	0 - 6	3	2
Max31	1/4	0652	1 - 31	0 - 6	4	2
Max65	2/8	0653	1 - 65	0 - 6	8	2

The supported upgrade paths from previous families for the z15 T02 are shown in Figure 3-1 on page 43.

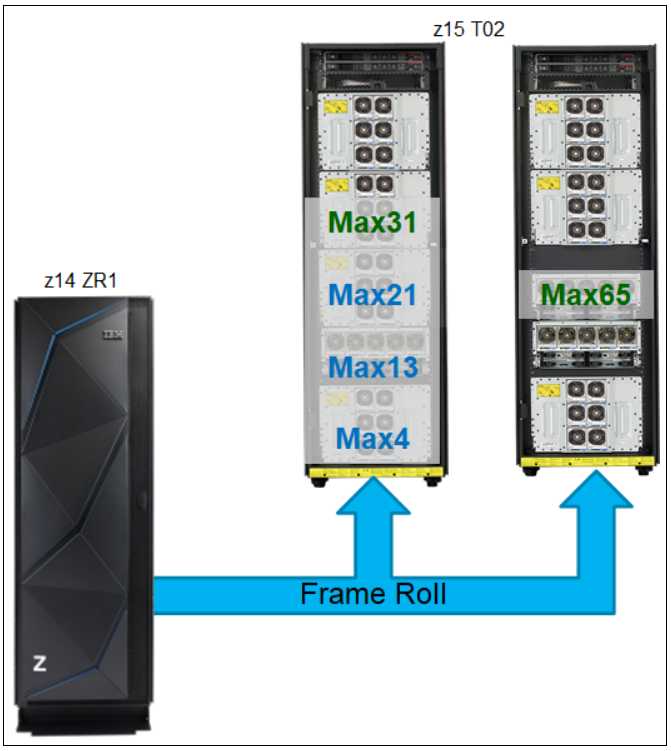

Figure 3-1 z15 T02 upgrade paths from z14 ZR1

If an upgrade request cannot be accomplished with the existing configuration, a hardware upgrade is required in which, depending on the starting point, one or more PU SCMs, or the second CPC drawer are added to accommodate the wanted capacity.

Upgrades within the z15 T02 product range

Within the z15 T02 system, upgrades are feature-based upgrades:

- Upgrades from any lower feature (Max4, Max13, and Max21) to Max31 or Max 65 are disruptive.
- Upgrades from Max31 to Max65 can be installed concurrently.

On the z15 T02, concurrent upgrades are available for 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. One or two spare PUs always are available on a z15 T02, depending on the feature.

In addition, the z15 T02 offers156 capacity levels for up to 6 CPs. For more information, see 5.3.1, “Capacity settings” on page 84.

3.2 Frames and cabling

The z15 T02 is built in a 19-inch frame and uses industry-standardized power and hardware. z15 T02 is a single frame system. The frame takes up only two standard 24-inch floor tiles of space, which aligns with modern data center layouts.

The z15 T02 packaging builds on the new format that is introduced with IBM z14 ZR1 and brings new configuration options as compared to previous Z platforms, such as IBM z13s®. See Table 3-2.

Table 3-2 z15 T02 configuration options compared to z13s and z14 ZR1

System	Number of frames	Number of CPC drawers	Number of I/O drawers	I/O and power connections	Power options	Power cord options[a]
z15 T02	One in 19-inch format	1 - 2	0 - 4[b,c]	Rear only	Power Distribution Unit (PDU)	Single- or three-phase
z14 ZR1	One in 19-inch format	1	0 - 4[b]	Rear only	PDU	Single-phase only
z13s	One in 24-inch format	1 - 2	0 - 3[d]	Front and rear	Bulk Power Assembly (BPA)	Single- or three-phase

a. Single- or three-phase is determined by the system configuration (the max. electrical power drawn).
b. Only PCIe+ I/O drawers are available. New build only. No carry forward of any I/O drawer.
c. 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.
d. Maximum of two PCIe I/O drawers.

The number of PCIe+ I/O drawers can vary based on the number of I/O features and number of CPC drawers installed. For a single CPC Drawer system, a maximum configuration of up to four PCIe+ I/O drawers can be installed. A system with two CPC Drawers supports a maximum of three PCIe+ I/O drawers. PCIe+ I/O drawers can be added concurrently.

A 16U reserved space feature (Feature Code 0151) is offered to provide increased flexibility for use with a customer-supplied IBM DS8910F system. This feature triggers single-phase power. A second CPC drawer and the third and fourth PCIe+ I/O drawers cannot be ordered with this feature.

In addition, the z15 T02 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, which potentially eliminates cables to be run under a raised floor and increases air flow over the system.

The z15 T02 supports installation on raised floor and non-raised floor environments.

Figure 3-2 on page 45 shows a fully configured z15 T02 with *a single* CPC Drawer and four PCIe+ I/O drawers.

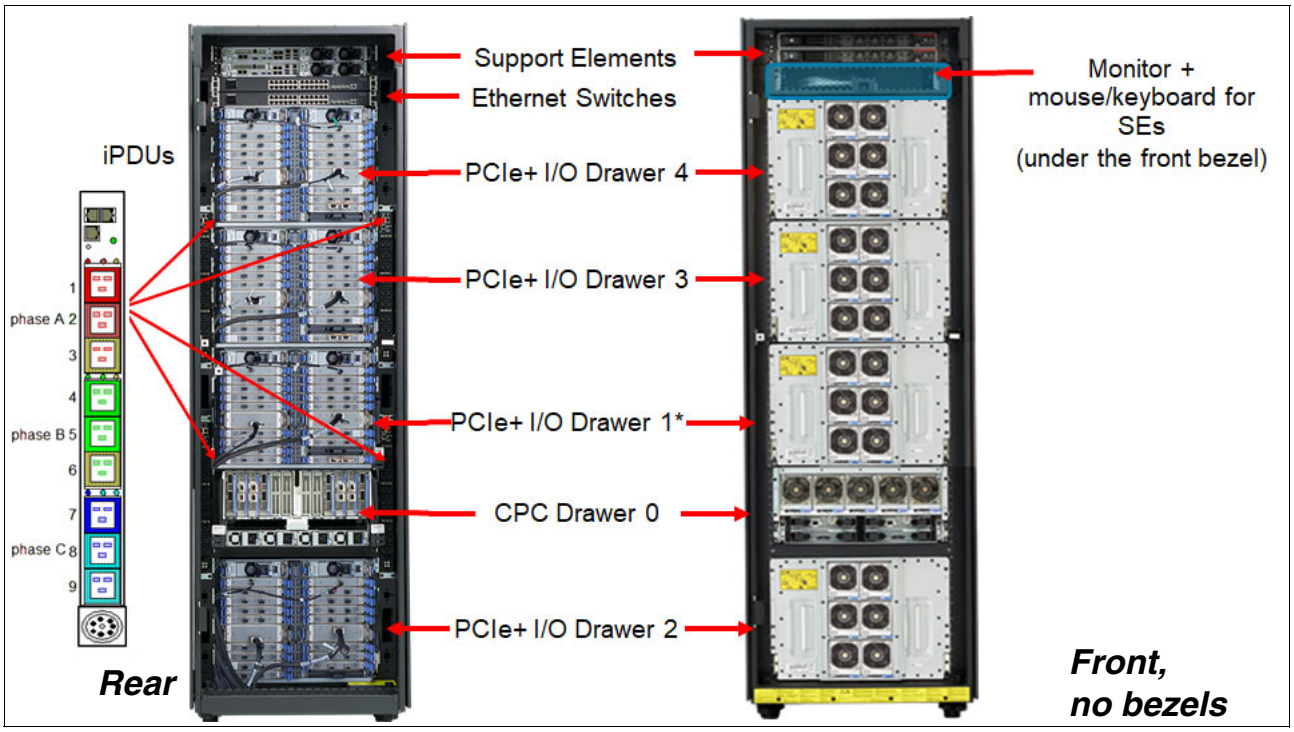

Figure 3-2 Fully configured z15 T02: Single CPC Drawer

Figure 3-3 shows a fully configured z15 T02 with *two* CPC Drawers and three PCIe+ I/O Drawers.

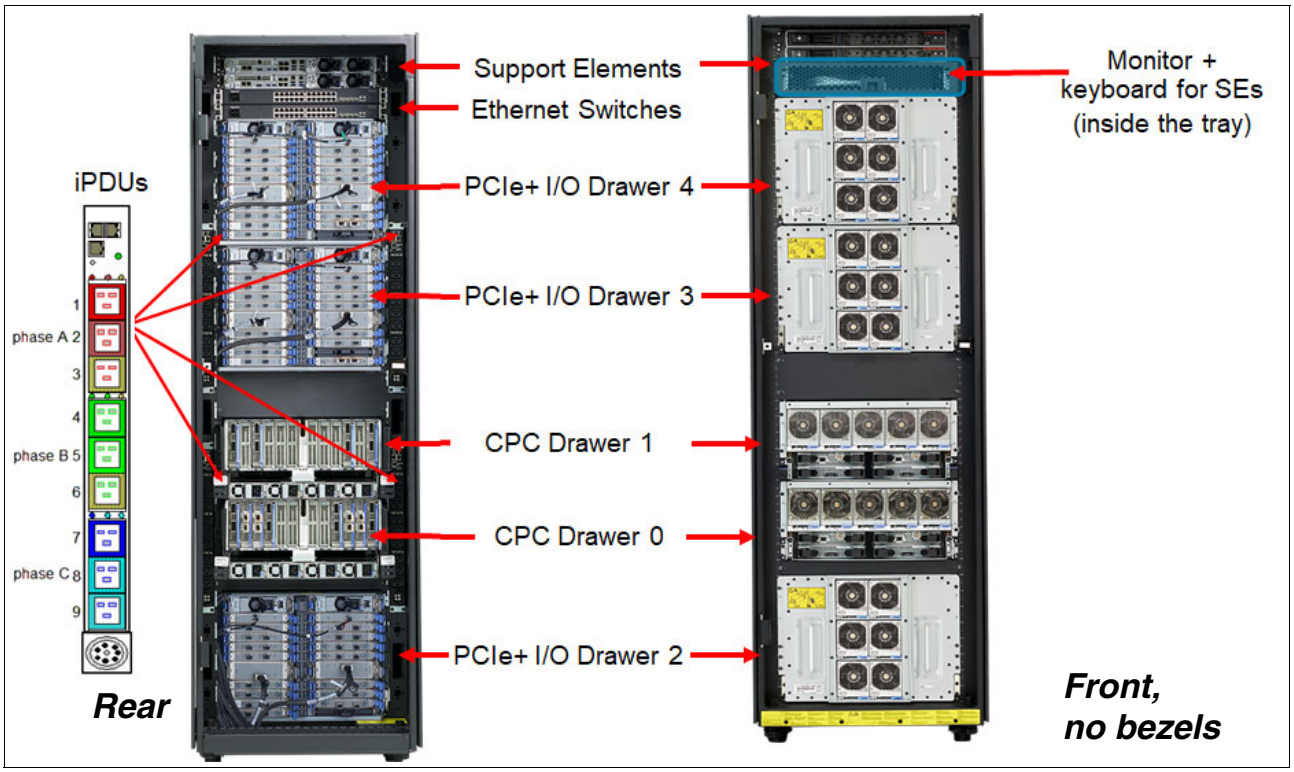

Figure 3-3 Fully configured z15 T02 with two CPC Drawers (Max65)

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 the z15 T02 include the following examples:

- Feature: Max4, Max13, Max21, Max31, and Max65
- 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

The z15 T02 can hold up to two CPC drawers. Each CPC drawer contains the following elements:

- SCMs:
  - One to four PU SCMs, each containing 7 - 11 PU cores (air-cooled).
  - One Storage Controller (SC) SCM, with a total of 960 MB L4 cache.
- Memory:
  - A minimum of 64 GB and a maximum of 16224 TB of memory (excluding 160 GB for hardware system area (HSA)) is available for use. For more information, see Table 3-4 on page 50.
  - Up to 20 dual inline memory modules (DIMMs) are plugged in a CPC drawer that are 32 GB, 64 GB, 128 GB, 256 GB, or 512 GB.
- Fanouts

  The CPC drawer provides 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 PCIe 16 gigabytes per second (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, 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.

  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 Flexible Support Processor (FSP)\ oscillator cards (OSCs), which provide redundant interfaces to the internal management network and provide clock synchronization to the Z CPCs.
- 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 water-cooled.
- z15 T02 Max31 also supports concurrent drawer add while z15 T02 Max65 supports concurrent drawer repair.

The CPC drawer communication topology is shown in Figure 3-4 on page 47. All CPC drawers are interconnected with high-speed communications links (A-Bus) through the SC chip L4 shared caches. Symmetric multiprocessor (SMP) cables are used to interconnect the CPC drawers. The X-Bus provides connectivity between PUs within the logical clusters and the SC on the drawer.

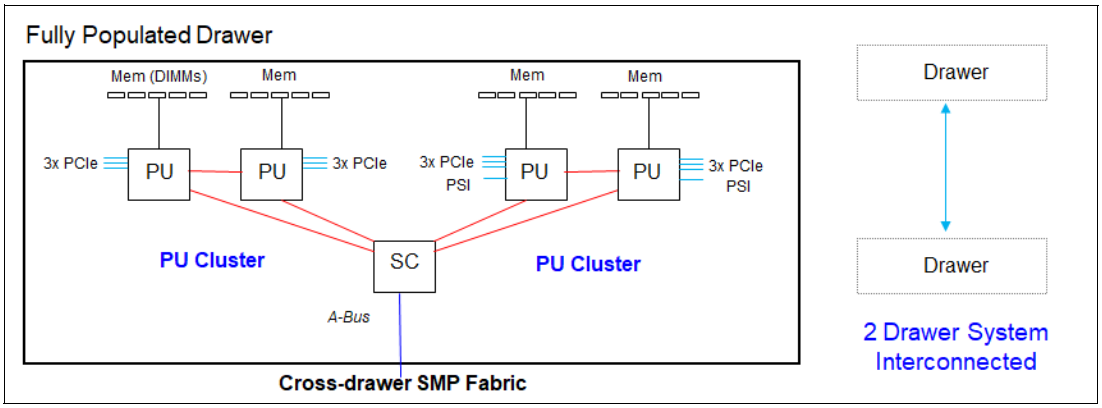

Figure 3-4 z15 T02 CPC drawer communication topology

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.

3.3.1 Single-chip modules

The CPC drawer for a single CPC drawer system has one to four PU SCMs and one SC SCM. Each PU SCM supports up to 11 active PU cores, and L1, L2, and L3 caches.

The SC SCM includes 960 MB shared eDRAM cache, interface logic to the four PU SCMs, and SMP fabric logic. The SC SCM is configured to provide L4 cache that is shared by all PU cores in the CPC drawer.

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 in execution 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 any moment, and are subject to the maximum number of decodes and completions per cycle.

PU cache

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 private L2 cache, with 4 MB D-cache (D for data) and 2 MB I-cache (I for instruction).
- Each PU SCM contains a 256 MB L3 cache that is shared by all PU cores in the SCM. The shared L3 cache uses eDRAM.

This on-chip cache implementation optimizes system performance for high-frequency processors, with cache improvements, new Translation/TLB2 design, pipeline optimizations, better branch prediction, new accelerators and architectures, and Secure Execution support.

The z15 T02 cache structure is shown Figure 3-5.

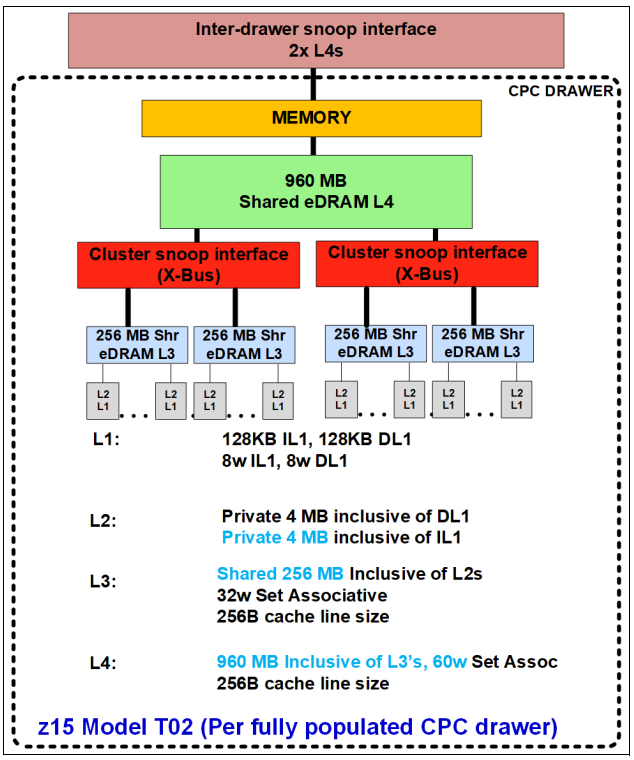

Figure 3-5 z15 T02 cache structure

Drawer cache

In addition to the on-chip (L1, L2, L3) cache, each drawer provides 960 MB L4 cache, for a total of 1920 MB L4 cache (z15 T02 Max65).

PU sparing

Hardware fault detection is embedded throughout the system design and is 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-core 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 1.2.5, “Cryptography” on page 17. This cryptographic hardware is available with any processor type, for example, CP, zIIP, and IFL.

On-chip functions

The compression accelerator replaces the IBM zEnterprise Data Compression (zEDC) Express feature that was on previous IBM Z platforms.

The sort accelerator uses the 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.

Software support

The z15 T02 PUs provide full compatibility with software for z/Architecture, and extend the Instruction Set Architecture (ISA) to enable enhanced functions and performance. The following hardware instructions support more efficient code generation and execution are introduced in the z15:

- Central Processor Assist for Cryptographic Functions (CPACF)
- Compression call (CMPSC)
- CPU Measurement Facility (MF)
- Out-of-order execution
- Large page support
- IBM Virtual Flash Memory
- Instruction Execution Protection
- Guarded Storage Facility (GSF)
- Transactional Execution Facility
- Runtime Instrumentation Facility
- Single-instruction, multiple-data (SIMD)
- Secure Execution Facility for Linux

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 by using one of the following characterizations:

- CP: These standard processors are used for general workloads.
- IFLs: 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” for workloads that are restricted to Db2-type applications. Also used for System Recovery Boost, which is a z15 exclusive feature. For more information, see “Reliability, availability, and serviceability” on page 91.
- Integrated Coupling Facility (ICF): Designates processors to be used specifically for coupling.
- System Assist Processor (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 PCIe I/O infrastructure management).

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

The IBM z15 Model T02 (machine type 8562) supports up to 6 CPs and up to 12 zIIPs. The maximum number depends on the processor feature, and is shown in Table 3-3.

Table 3-3 Customer orderable PU limits

CPC feature	CPs	zIIPs	IFLs	uIFLs	ICFs	Opt. SAPs
Max4	0 - 4	0 - 2	0 - 4	0 - 3	0 - 4	0 - 2
Max13	0 - 6	0 - 7	0 - 13	0 - 12	0 - 13	0 - 2
Max21	0 - 6	0 - 12	0 - 21	0 - 21	0 - 21	0 - 2
Max31	0 - 6	0 - 12	0 - 31	0 - 31	0 - 31	0 - 8
Max65	0 - 6	0 - 12	0 - 65	0 - 65	0 - 65	0 - 8

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 Integrated Coupling Facilities (ICFs), IFLs, zIIPs, and SAP to the z15 T02 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 platform features 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 the z15 T02, as with the z14 ZR1, up to 8 TB of memory per CPC drawer can be ordered and up to 16 TB for a two-CPC drawer system.

Important: Starting with z/OS V2R3, a minimum of 8 GB of memory is required for the LPAR (2 GB of memory when z/OS is running as a z/VM guest). z/OS can support up to 4 TB of memory in an LPAR.

The minimum and maximum memory sizes for each z15 T02 feature are listed in Table 3-4.

Table 3-4 z15 Model T02 memory per feature

Feature name	PU SCMs / CPC drawers	Customer memory
Max4 (Feature Code 0649)	1 / 1	64 GB - 2 TB
Max13 (Feature Code 0650)	2 / 1	64 GB - 4 TB
Max21 (Feature Code 0651)	3 / 1	64 GB - 4 TB
Max31 (Feature Code 0652)	4 / 1	64 GB - 8 TB
Max65 (Feature Code 0653)	8 / 2	64 GB - 16 TB

The HSA on the z15 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 z15 platforms, the granularity for memory is in 8, 32, 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 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 (if available), PR/SM attempts to allocate memory across that same set of CPC drawers (even if all required memory is available in the second CPC drawer).

No matter which CPC drawer the memory is installed in, an LPAR has access to that memory after it is allocated. Despite the CPC drawer structure, the z15 T02 is still an SMP system because the PUs can access all of the available memory.

Enhanced Drawer Availability (new for z15 T02)

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 system with two CPC drawers, a single CPC drawer can be concurrently removed and reinstalled for a repair with EDA (with resource relocation).

For model upgrades involving the addition of a CPC drawer, the minimum usable memory increment (256 GB) is added to the system. New for z15 T02, during an upgrade, adding the second CPC drawer and physical memory in the second CPC 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 the z15 T02, and was enhanced to include one Memory Controller Unit (MCU) per processor chip, with five memory channels and one DIMM per channel. A fifth channel in each MCU enables memory to be implemented as 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 z15 (8561) Technical Guide*, SG24-8851.

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 for z15 T02 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 the z15 T02 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 in z15 T02)
- Subchannel set (up to three can be defined in z15 T02)
- A total of 40 LPARs
- Devices, up to the maximum number permitted, in each subchannel set
- Logical processors by type
- Cryptographic adapters

3.4 I/O system structure

The z15 T02 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 appropriate number of PCIe+ I/O drawers.

Figure 3-6 shows a high-level view of the I/O system structure for the z15 T02.

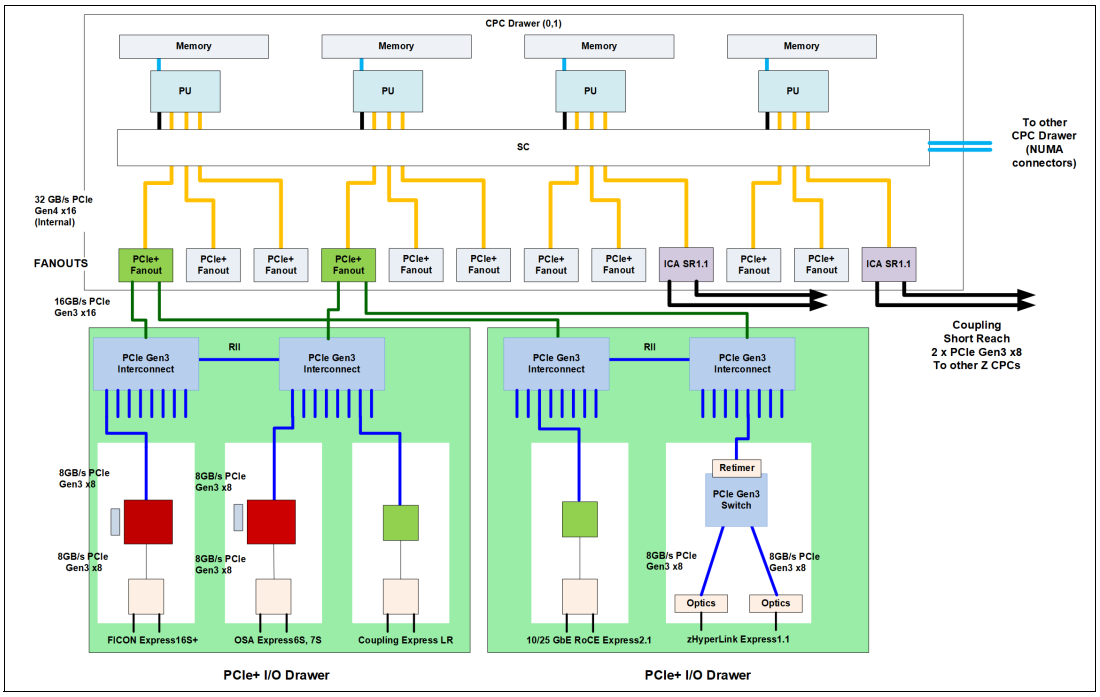

Figure 3-6 z15 T02 I/O system structure

The z15 T02 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 (Parallel Sysplex and Server Time Protocol (STP) configuration), the z15 T02 supports the following link types:

- ICA SR1.1 and ICA SR (installed in a CPC drawer)
- Coupling Express 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 3-7).

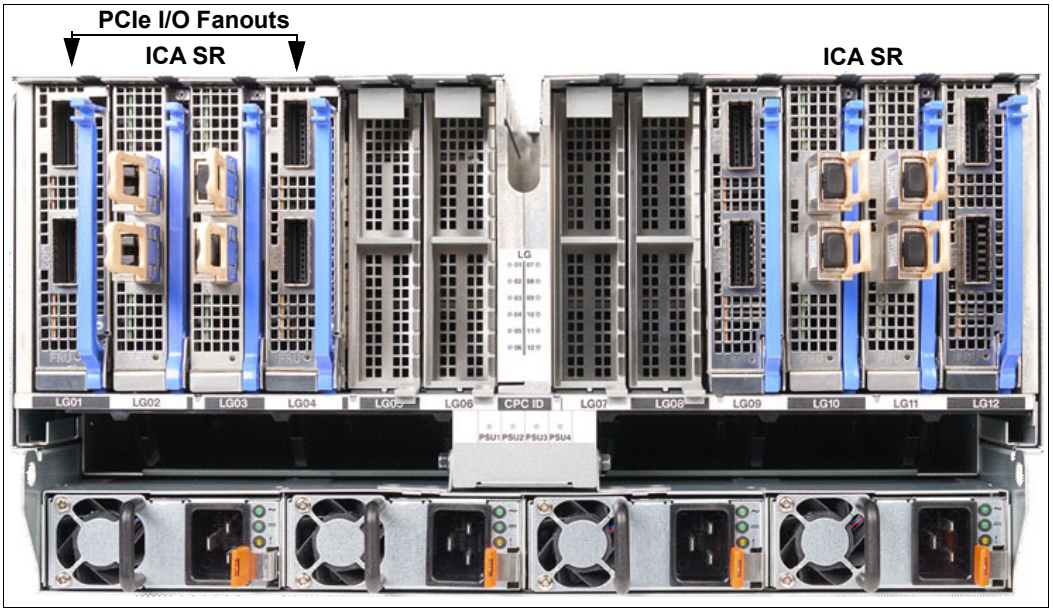

Figure 3-7 z15 T02 CPC drawer: Rear view

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 two slots for PCIe switch cards.

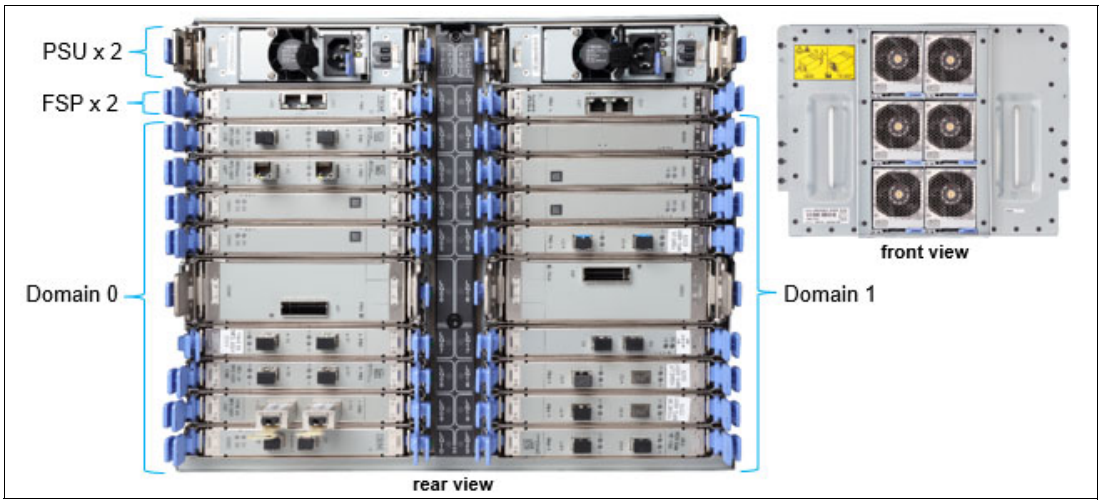

Figure 3-8 PCIe+ I/O drawer: Rear and front view

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

- IBM Fibre Connection (FICON) Express16S+, FICON Express16S, or FICON Express8S
- OSA-Express7S, OSA-Express6S, or OSA-Express5S
- Crypto Express7S, Crypto Express6S, or Crypto Express5S

- 25 GbE RDMA over Converged Ethernet (RoCE) Express2.1, 10 GbE RoCE Express2.1, 25 GbE RoCE Express2, 10 GbE RoCE Express2, or 10 GbE RoCE Express
- zHyperLink Express1.1 and zHyperLink Express
- CE LR

For more information about the I/O feature that is available with the z15, see Chapter 4, “Supported features and functions” on page 57.

3.5 Power and cooling

The z15 T02 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.

3.5.1 Power options

The z15 T02 19-inch frame is available with the following power options:

- Single-phase PDU

  Single-phase power for z15 T02 can be used for systems with a single CPC drawer (Max4 to Max31). System with single CPC drawer supports up to four 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. The z15 T02 does not support plan-ahead or balanced power options.

  The z15 T02 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 z15 T02 cooling system is available with one cooling option: air cooling. SCMs 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 regarding power:

- The 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.
- Supports top or bottom exit power.
- 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 units that are installed, and the power option (single- or three-phase power).

For more information about the maximum power consumption tables for the various configurations and environments, see *IBM 8562 Installation Manual for Physical Planning*, GC28-7011.

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

4

Supported features and functions

The I/O and other miscellaneous features and functions of the z15 are described in this chapter. The information in this chapter expands on the overview of the key hardware elements that was provided in Chapter 1, “Designed to take on today’s IT demands” on page 1, Chapter 2, “IBM z15 Model T01 hardware overview” on page 23, and Chapter 3, “IBM z15 Model T02 hardware overview” on page 41.

Naming: Throughout this chapter, we discuss features and functions available for the IBM z15 model lineup, the IBM z15 Model T01 (machine type 8561), and IBM z15 Model T02 (machine type 8562). For the common features and functions, we identify the systems as “z15”, while explicitly mentioning the model (T01 or T02) where functions and features differ.

Only the features and functions that are supported on the z15 are highlighted in this chapter. For more information about the key capabilities and enhancements of the z15, see *IBM z15 (8561) Technical Guide*, SG24-8851. For more information about the I/O features and functions, see *IBM Z Connectivity Handbook*, SG24-5444.

This chapter includes the following topics:

4.1 z15 I/O connectivity overview

The z15 provides a PCIe-based infrastructure for the PCIe+ I/O drawers to support the following features:

- zHyperLink Express1.1
- zHyperLink Express (carry forward only)
- IBM Fibre Connection (FICON) Express16SA (z15 T01 only)
- FICON Express16S+ (z15 T01 - carry forward only, z15 T02 new build and carry forward)
- FICON Express16S (carry forward only)
- FICON Express8S (carry forward only)
- OSA-Express7S GbE, 1000BASE-T and 10 GbE (z15 T01 only)
- OSA-Express7S 25 GbE1.1 (z15 T01 only)
- OSA-Express7S 25 GbE (z15 T01 carry forward only, z15 T02 - new build)
- OSA-Express6S (carry forward only)
- OSA-Express5S (carry forward only)
- 25 GbE RDMA over Converged Ethernet (RoCE) Express2.1
- 25 GbE RoCE Express2 (carry forward only)
- 10 GbE RoCE Express2.1
- 10 GbE RoCE Express2 (carry forward only)
- 10 GbE RoCE Express (carry forward only)
- Crypto Express7S (1 or 2 port[1])
- Crypto Express6S (carry forward only)
- Crypto Express5S (carry forward only)
- Integrated Coupling Adapter Short Reach1.1 (ICA SR1.1)
- ICA SR (carry forward only)
- Coupling Express Long Reach (CE LR)

Specifications for these features are provided in the subsequent sections.

The following features that were supported on earlier Z platforms are *not orderable* and *cannot be carried forward* to the z15:

- ESCON
- FICON Express8 and older
- OSA-Express4 and older
- ISC-3
- Crypto Express4S and older
- Flash Express
- Host Channel Adapter3 - Optical Long Reach (HCA3-O LR)
- Host Channel Adapter3 - Optical (HCA3-O)
- IBM zEnterprise Data Compression (zEDC)

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.

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

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 Express16SA feature offers the same functions as its predecessor.

For more information about FICON channels, see the Z I/O connectivity website. Technical papers about performance data are also available.

4.2.1 zHyperLink Express

IBM zHyperLink Express is a short-distance 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 distance (direct connectivity between a z15 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 gigabytes per second (GBps).

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 the z15, 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 fully supported on the z15. They are commonly used with IBM z/OS, IBM z/VM (and guest systems), Linux on Z, IBM z/VSE, and IBM z/TPF.

Storage connectivity options are listed in Table 4-1.

Table 4-1 Storage connectivity features

Feature	Feature codes	Bit rate in Gbps (or stated)	Cable type	Maximum unrepeated distance	Ordering information z15 T01	Ordering information z15 T02
zHyperLink Express1.1	0451	8	OM3 and OM4	See Table 4-2.	New build	New build
zHyperLink Express	0431				Carry forward	Carry forward
FICON Express16SA LX	0436	8 or 16	SM 9 µm	10 km (6.2 miles).	New build	N/A[a]
FICON Express16SA SX	0437	8 or 16	OM2, OM3, and OM4	See Table 4-3.	New build	N/A[a]
FICON Express16S+ LX	0427	4, 8, or 16	SM 9 µm	10 km (6.2 miles).	Carry forward	New build and carry forward

Feature	Feature codes	Bit rate in Gbps (or stated)	Cable type	Maximum unrepeated distance	Ordering information z15 T01	Ordering information z15 T02
FICON Express16S+ SX	0428	4, 8, or 16	OM2, OM3, and OM4	See Table 4-3.	Carry forward	New build and carry forward
FICON Express16S LX	0418	4, 8, or 16	SM 9 µm	10 km (6.2 miles).	Carry forward	Carry forward
FICON Express16S SX	0419	4, 8, or 16	OM2, OM3, and OM4	See Table 4-3.	Carry forward	Carry forward
FICON Express8S LX	0409	2, 4, or 8	SM 9 µm	10 km (6.2 miles).	Carry forward	Carry forward
FICON Express8S SX	0410	2, 4, or 8	OM2, OM3, and OM4	See Table 4-3.	Carry forward	Carry forward

a. N/A: Not available

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.

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

4.2.3 IBM Fibre Channel Endpoint Security

With IBM z15 Model T01 (machine type 8561), a new security function was added for FICON Express16SA (Feature Code 0436 and Feature Code 0437). The new capability adds Fibre Channel Endpoint Authentication and Encryption of Data in Flight.

Based tightly on the Fibre Channel–Security Protocol-2 (FC-SP-2)1 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 takes advantage of existing IBM solution for key server infrastructure in the Storage System (for data at rest encryption).

IBM Security™ Key Lifecycle Manager server provides shared secret key generation in a master-slave relationship between an FC initiator (IBM Z) and the Storage target. The solution implements authentication and key management called IBM Secure Key Exchange (SKE).

Data in flight (from/to IBM Z and IBM Storage) is encrypted when it leaves either endpoint (source) and decrypted at the destination. Encryption/decryption is done at the FC adapter level. In Endpoint Security related operations, there is no involvement of the operating system that runs on the host (IBM Z). Tools are provided at the operating system level for displaying information about encryption status.

IBM Fibre Channel Endpoint Security[2] is an orderable feature for IBM z15 T01 (Feature Code 1146) and requires also Central Processor Assist for Cryptographic Functions (CPACF) enablement (Feature Code 3863), specific storage (DS8900), and FICON Express16SA 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

The z15 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.

HiperSockets

IBM HiperSockets are referred to as the “network in a box” because it simulates local area network (LAN) environments entirely within the IBM Z platform. The data transfer is from logical partition (LPAR) memory to LPAR memory, which is mediated by IBM Z firmware. The z15 supports up to 32 HiperSockets. One HiperSockets network can be shared by up to 85 LPARs (40 for z15 T02). Up to 4096 communication paths support a total of 12,288 IP addresses across all 32 HiperSockets.

The HiperSockets internal networks can support the following transport modes:

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

[2] This feature (FC 1146) is only available for z15 T01 (M/T 8561).

Traffic can be Internet Protocol Version 4 or Version 6 (IPv4, IPv6) or non-IP. 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 Z environment the same way as they do for a non-Z environment. HiperSockets Layer 2 support is provided by Linux on IBM Z, and by z/VM for guest use.

OSA-Express features

Improved throughput (mixed inbound/outbound) is achieved by the 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).

Shared Memory Communications

The SMC capabilities of the z15 optimize the communications between applications in server-to-server (SMC-R) or LPAR-to-LPAR (SMC-D) connectivity. Before z15 this communication option was only available for z/OS-to-z/OS instances. With z15, the SMC capabilities are enhanced for support of Linux on Z environments. Therefore, SMC reduces the CPU footprint on z/OS and improves latency and throughput for network communication between Linux on Z and z/OS, or two Linux on Z instances.

SMC-R provides application transparent use of the RoCE Express features that can reduce the network overhead 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 takes advantage of 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 Z intra-system communications. Up to 32 channels for SMC-D traffic can be defined in a z15, whereby each channel can be virtualized to a maximum of 255 Function IDs[3]. No other hardware is required for SMC-D.

RoCE Express features

The RoCE Express features help reduce the use of CPU resources for applications that use the TCP/IP stack. It might also help to reduce network latency with memory-to-memory transfers that use SMC-R in z/OS V2R1 and later versions. It is transparent to applications, and can be used for system-to-system communication in a multiple Z platform environment.

These features are installed in the PCIe+ I/O drawer and use a Short Reach (SR) optical transceiver. Both 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.

A maximum of *eight* features in any combination can be installed in the z15.

Note: The 25 GbE RoCE Express should 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.

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

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

Table 4-4 Network connectivity features

Feature	Feature code	Bit rate in Gbps (or stated)	Cable type	Maximum unrepeated distance[a]	Ordering information z15 T01	Ordering information z15 T02
OSA-Express7S 25 GbE SR1.1	0449	25	MM 50 µm	70 m (2000) 100 m (4700)	New build	N/A[b]
OSA-Express7S 25 GbE SR	0429				Carry forward	New build and carry forward
OSA-Express7S 10 GbE LR	0444	10	SM 9 µm	10 km (6.2 miles)	New build	N/A[b]
OSA-Express6S 10 GbE LR	0424				Carry forward	New build and carry forward
OSA-Express5S 10 GbE LR	0415				Carry forward	Carry forward
OSA-Expess7S 10 GbE SR	0445	10	MM 62.5 µm MM 50 µm	33 m (200) 82 m (500) 300 m (2000)	New build	N/A[b]
OSA-Expess6S 10 GbE SR	0425				Carry forward	New build and carry forward
OSA-Expess5S 10 GbE SR	0416				Carry forward	Carry forward
OSA-Express7S GbE LX	0442	1.25	SM 9 µm	5 km (3.1 miles)	New build	N/A[b]
OSA-Express6S GbE LX	0422				Carry forward	New build, Carry forward
OSA-Express5S GbE LX	0413				Carry forward	Carry forward
OSA-Express7S GbE SX	0443	1.25	MM 62.5 µm MM 50 µm	275 m (200) 550 m (500)	New build	N/A[b]
OSA-Express6S GbE SX	0423				Carry forward	New build and carry forward
OSA-Express5S GbE SX	0414				Carry forward	Carry forward
OSA-Express7S 1000BASE-T	0446	1000 Mbps	Cat 5 or Cat 6 unshielded twisted pair (UTP)	100 m	New build	N/A[b]
OSA-Express6S 1000BASE-T	0426	100 or 1000 Mbps			Carry forward	New build and carry forward
OSA-Express5S 1000BASE-T	0417				Carry forward	Carry forward
25 GbE RoCE Express2.1	0450	25	MM 50 µm	70 m (2000) 100 m (4700)	New build	New build
25 GbE RoCE Express2	0430				Carry forward	Carry forward
10 GbE RoCE Express2.1	0432	10	MM 62.5 µm MM 50 µm	33 m (200) 82 m (500) 300 m (2000)	New build	New build
10 GbE RoCE Express2	0412				Carry forward	Carry forward
10 GbE RoCE Express	0411				Carry forward	Carry forward

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

b. N/A: Not available.

4.4 Clustering connectivity

A Parallel Sysplex is a clustering technology that allows you to operate multiple copies of z/OS images as a single system from a user's perspective. A properly 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 connectivity for Parallel Sysplex on z15 uses ICA SR and CE LR. ICA SR and CE LR technologies allow all of the z/OS-to-CF communication, CF-to-CF traffic, or Server Time Protocol (STP)[4].

For more information about coupling links technologies, see the *Coupling Facility Configuration Options* white paper.

Internal coupling (IC) links are used for internal communication between LPARs on the same system that is running Coupling Facilities (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 24 is available for the z15. For more information, see IBM Knowledge Center.

Coupling link options are listed in Table 4-5.

Table 4-5 Coupling and clustering features

Feature	Feature code	Bit rate	Cable type	Maximum unrepeated distance	Ordering information
CE LR	0433	10 Gbps	SM 9 µm	10 km (6.2 miles)	New build or carry forward
ICA SR 1.1	0176	8 GBps	OM3[a] and OM4[a]	150 m 100 m	New build
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.

Dynamic I/O configuration for stand-alone CFs

Dynamic I/O configuration changes can be made to a stand-alone CF[5] without requiring a disruptive power on reset. This capability is supported on the z14 and z15.

A separate z14 or z15 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 dynamically. The firmware-based LPAR is driven by updates from an HCD instance that is running in a z/OS LPAR on a different z14 or z15 that is connected to the same Hardware Management Console (HMC).

[4] All coupling links can be used to carry STP timekeeping information.
[5] A stand-alone CF does not have any running instances of z/OS or z/VM.

Although the firmware LPAR was optional on z14, it is always defined with z15. This 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.

4.5 Server Time Protocol

STP is a message-based protocol in which timekeeping information is passed over coupling links between Z platforms. The z15 can participate in an STP Coordinated Timing Network (CTN). A CTN is a collection of Z platforms 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 the z15 and other Z platforms. The z15 is enabled for STP by installing the STP feature code. Extra configuration is required for a z15 to become a member of a CTN.

For high availability purposes, nondisruptive capability was implemented in the z15 firmware that allows two CTNs to be merged into one, or to split one CTN into two, dynamically.

Note: If a z15 plays a CTN role (PIN Transaction Security (PTS), BTS, or Arbiter), the other CTN roleplaying Z platforms must include direct coupling connectivity to the z15. The z15 does not support direct coupling connectivity to zEC12, zBC12, or older Z platforms.

STP supports a multi-site timing network of up to 100 km (62 miles) over fiber optic cabling, without requiring an intermediate site. This protocol allows a Parallel Sysplex to span these distances and reduces the cross-site connectivity that is required for a multi-site Parallel Sysplex.

Network Time Protocol client support

The use of Network Time Protocol (NTP) servers as an external time source (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 Support Element (SE) code of the z15. The code interfaces with the NTP servers. This interaction allows an NTP server to become the single-time source for z15 and for other servers that have NTP clients.

Pulse per second support

Two oscillator cards[6] (OSCs), which are included as a standard feature of the z15, provide a dual-path interface for the pulse per second (PPS) signal. The cards contain a Bayonet Neill-Concelman (BNC) connector for PPS attachment at the front side of the central processor complex (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 the z15. If STP uses an NTP server without PPS, a time accuracy of 100 ms to the ETS is maintained.

A cable connection from the PPS port to the PPS output of an NTP server is required when you configure the z15 for NTP with PPS as ETS for time synchronization.

[6] The OSCs are combined with the Flexible Support Processor (FSP) cards.

Precision Time Protocol client support

The z15 introduced IEEE 1588 Precision Time Protocol (PTP) as an ETS for an IBM Z CTN. The initial implementation is for PTP connectivity by way of the IBM Z HMC/SE. Current implementation brings no change to the use of STP CTNs for time coordination, other than the potential to use a PTP-based ETS. For more information, see *IBM Z Server Time Protocol Guide*, SG24-8480.

4.6 Compression options

The following compression options are available with the z15:

- A standard internal compression coprocessor that is tightly connected to each processor core.
- An integrated on-chip compression unit on each processor chip, which is called the IBM Integrated Accelerator for zEDC.

4.6.1 On-core Compression coprocessor

The compression coprocessor (CMPSC) is a feature that works with the processor units (PUs) in the Z platform. This coprocessor works with a proprietary compression format and is used for many types of z/OS data.

4.6.2 On-chip IBM Integrated Accelerator for zEnterprise Data Compression

Introduced with z15, on-chip compression offers industry-leading compression throughput with faster single thread performance. More compression throughput than previous Z generations allows all data to be compressed with no extra CPU cost. It is designed to reduce the cost of storing, transporting, and processing data without changing applications architecture.

It provides hardware-based acceleration for data compression and decompression for the enterprise, which can help to improve cross platform data exchange, reduce CPU consumption, and save disk space.

This on-chip compression capability replaces the zEDC Express feature on the IBM z14 and earlier Z platforms, whereby all data interchanges remain compatible. The throughput per engine is improved from 1 GBps with zEDC to 16 GBps with on-chip compression. The on-chip compression can be virtualized between all LPARs and virtual machines (VMs) running on the z15. The zEDC Express feature was limited to 15 LPARs or VMs on earlier Z platforms.

The z15 on-chip compression module implements DEFLATE, gzip, and lzip algorithms and works in the following modes:

- Synchronous execution in problem state
- Asynchronous optimization for large operations under z/OS

The on-chip compression implements compression as defined by RFC1951 (DEFLATE).

4.7 Cryptographic features

The z15 provides cryptographic functions that can be categorized in the following groups from an application program perspective:

- Symmetric and asymmetric[7] cryptographic functions, which are provided by the CPACF or the Crypto Express features when defined as an accelerator.
- Asymmetric cryptographic functions, which are provided by the Crypto Express features.

4.7.1 Central Processor Assist for Cryptographic Functions

CPACF offers a set of symmetric cryptographic functions for high-performance encryption and decryption with clear key operations for SSL/TLS, VPN, and data-storing applications that do not require Federal Information Processing Standards (FIPS) 140-2 level 4 security[8]. The CPACF is a no-charge optional feature that is integrated with the compression unit in the coprocessor in the z15 microprocessor core.

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

The cryptographic assist includes support for the following functions:

- Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys
- Data Encryption Standard (DES) data encryption and decryption with single, double, or triple length keys
- Pseudo-random number generation (PRNG)
- True-random number generator (TRNG)
- Message authentication code
- Elliptic Curve Cryptography (ECC) support (z15 only)
- Hashing algorithms: Secure Hash Algorithm (SHA)-1, SHA-2, and SHA-3

SHA-1, SHA-2, and SHA-3 support are enabled on all Z platforms and do not require the CPACF enablement feature. The CPACF functions are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on IBM Z.

4.7.2 Crypto Express7S

The Crypto Express7S represents the newest generation of the Peripheral Component Interconnect Express (PCIe) cryptographic coprocessors, which are an optional feature that is available on the z15. These coprocessors are hardware security modules (HSMs) that provide high-security cryptographic processing as required by banking and other industries.

[7] Elliptic Curve Cryptography (ECC) algorithms.
[8] FIPS 140-2 Security Requirements for Cryptographic Modules.

This feature provides a secure programming and hardware environment wherein crypto-processes are performed. Each cryptographic coprocessor includes general-purpose processors, non-volatile 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, which is the highest security level defined.

The Crypto Express7S (2-port) feature (0898) includes two PCIeCCs, and the Crypto Express7S (1-port) feature (0899) includes one PCIeCC. For availability reasons, a minimum of two features is required for the one port feature. Up to 30 Crypto Express7S (2-port) features are supported on z15 T01, while z15 T02 supports 20. The maximum number of the 1-port features is 16. The Crypto Express7S 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.

Crypto Express7S provides domain support for up to 85 LPARs (40 for z15 T02).

The accelerator function is designed for maximum-speed Secure Sockets Layer and Transport Layer Security (SSL/TLS) acceleration, rather than for specialized financial applications for secure, long-term storage of keys or secrets. The Crypto Express7S can also be configured as one of the following configurations:

- The Secure IBM CCA coprocessor includes secure key functions with 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 (UDX).

  A new mode, called Payment Card Industry (PCI) PTS Hardware Security Module (HSM) (PCI-HSM), is available in CCA mode. PCI-HSM mode simplifies compliance with Payment Card Industry requirements for hardware security modules.
- The Secure IBM EP11 coprocessor implements an industry-standardized set of services that adheres to the PKCS #11 specification v2.20 and more recent amendments. It was designed for extended FIPS and Common Criteria evaluations to meet industry requirements.

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

When the Crypto Express7S PCIe adapter is configured as a secure IBM CCA co-processor, it still provides accelerator functions. However, up to 3x better performance for those functions can be achieved if the Crypto Express7S 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, 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.

Several features that support the use 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 use of AES PIN blocks in other CCA callable services. IBM continues to add enhancements as AES finance industry standards are released.

4.7.3 Crypto Express6S and Crypto Express5S (carry forward only)

The Crypto Express6S feature (0893) and Crypto Express5S feature (0890) have one PCIeCC (HSM) per feature. For availability reasons, a minimum of two features is required. Up to 16 Crypto Express6S or Crypto Express5S features are supported.

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

Crypto Express6S and Crypto Express5S provide domain support for up to 85 LPARs on IBM z15 (40 for z15 T02).

The Crypto Express6S feature and Crypto Express 5S feature support all the functions of the Crypto Express7S, except that the PCI-HSM standard is not supported by Crypto Express 5S.

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

4.8 IBM Virtual Flash Memory

IBM Virtual Flash Memory (VFM) is the replacement for the Flash Express features that were available on the z13 and IBM z13s. On z15, the VFM feature (0643) can be ordered in 512 GB increments up to 12 VFM features for to the following maximums:

- For z15 T01, 6 TB in total.
- For z15 T02, support for up to four features per CPC, for a total of 2 TB.

VFM is designed to help improve availability and handling of paging workload spikes when z/OS V2.1, V2.2, V2.3, or V2.4 is run. With this support, z/OS is designed to help improve system availability and responsiveness by using VFM across transitional workload events.

VFM can also be used in CF images to provide extended capacity and availability for workloads that use WebSphere MQ Shared Queues structures. The use of 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 also designed to eliminate delays that can occur when diagnostic data is collected during failures.

Therefore, VFM can help meet most demanding service level agreements and compete more effectively. VFM is easy to configure and provides rapid time to value.

No application changes are required to migrate from IBM Flash Express to VFM.

4.9 Hardware Management Console and Support Element

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

With z15, the HMC can be a stand-alone computer (MiniTower), or can be run a 1U rack-mounted server.

Note: The z15 HMC supports the previous two generations (z13 and z14) only.

The z15 is supplied with a pair of integrated 1U SEs. The primary SE is always active; the other SE is an alternative. Power for the SEs is supplied by the rack Power Distribution Units (PDUs) or Bulk Power Assemblies (BPAs). Each SE features dual Power Supply Units (PSUs), 1+1 redundant.

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

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

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.

Note: An RSF connection through a modem is *not* supported on the z15 HMC. An internet connection to IBM is required to enable hardware problem reporting and service.

HMC/SE Driver 41 Version 2.15.0 or later is required for the z15. The IBM Hardware Management Appliance (Feature Code 0100) was added with Driver 41.

Hardware Management Appliance

The Hardware Management appliance (Feature Code 0100) can be ordered with a new z15 system (cannot be ordered as an upgrade or Miscellaneous Equipment Specification (MES)) and it provides the management capabilities of the HMC embedded into the two SE appliances. The SE functions remain the same, and HMC functions are added. SE code runs as a guest of the HMC. The SE hardware appliance provides all resources (embedded storage, processing power, memory, and external connectivity) required to accommodate the functions of both HMC and SE.

The SE HW provides redundancy for the Hardware Management (virtual) appliance as well, while customer connectivity into the HMC is ensured by using remote browser access (as with current physical HMC appliance).

Future HMC Hardware[a]: IBM z15 T02 is planned to be the last IBM Z server to offer the ability to order stand-alone HMC hardware. For future systems, new HMC hardware can be ordered only in the form of the Hardware Management Appliance feature (Feature Code 0100) which was introduced on z15. The Hardware Management Appliance feature provides redundant HMCs and SEs that are inside the CPC frame. This feature includes the ability to eliminate stand-alone HMC hardware (tower or rack mounted) outside the CPC frame. Stand-alone HMC hardware (tower or rack mounted) can still be ordered and used with IBM z15 (T01 and T02).

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 information that is mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functions.

5

IBM z15 system design strengths

Every new generation of IBM Z introduces innovative features and functions to provide more scalable, performant, secure, resilient, and flexible capabilities for building IT solutions and services. The 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 Z platform. Thus, protection and compatibility with an earlier versions of workloads and solutions are ensured.

The evolution of the Z platform embodies a proven architecture that is open, secure, resilient, and flexible. From their microprocessor and memory design to their cryptography capabilities, unparalleled I/O throughput, and rich virtualization, the z15 models are built to respond with speed, agility, and versatility.

This chapter introduces the IBM z15 system design capabilities and enhancements and includes the following topics:

[1] IBM z/Architecture is the mainframe computational architecture notation that defines its behavior. For more information, see:
https://www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss?CTY=US&FNC=SRX&PBL=SA22-7832-09

5.1 Technology improvements

Systems achieve the levels of efficiency that are demanded 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. You can have the fastest processors in the world, but if you cannot feed them, your workloads suffer.

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

The IBM z15 models (T01 and T02) deliver the latest innovative Z technologies and features, within 19-inch frames. They provide high levels of performance, scalability, resiliency, flexibility, and security when serving as a traditional Z platform, a cloud platform, or both. The z15 models can host thousands of virtualized environments.

This section contains information about z15 technology improvements and enhancements.

Naming: Throughout this chapter, we refer to both models (T01 and T02) as the z15. Wherever features and functions differ across the models, they are explicitly mentioned.

5.1.1 System capacity

Each generation of IBM Z platforms provides more system capacity, which combines various system design enhancements. The family of Capacity-on-Demand solutions guarantees a flexible addition of capacity when it is most needed, for example, during peak workload periods or scheduled maintenance.

The IBM z15 T01 with the Max190 feature is designed to offer up to 25% more capacity and 25% more configurable memory than the previous IBM z14 Model M05. The maximum number of configurable cores also increased from 170 with the z14 to 190 with the z15 T01.

The IBM z15 T02 with the Max65 feature (65 configurable PUs) is an increase of 35 PUs over the z14 ZR1. The single-PU capacity of the z15 T02 is approximately 14% greater[2] than a single PU of the z14 ZR1. In addition, the z15 T02 can be ordered with double the configurable memory compared to the IBM z14 ZR1.

5.1.2 Processor design highlights

The z15 supports 64-bit addressing mode and leverages Complex Instruction Set Computer (CISC), including highly capable and complex instructions. Most of the instructions are implemented at the hardware or firmware level for 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 provides compatibility with earlier software versions, which provides investment protection.

[2] Some variation is possible based on workload and configuration.

Compared to their predecessors, the z15 models feature the following processor design improvements and architectural extensions:

- Better performance and throughput:
  - Fast PUs with enhanced microarchitecture
  - Larger L2, L3 caches
  - Up to 12 cores per processor chip
  - More characterizable PUs
  - Larger cache (and shorter path to cache) means faster uniprocessor performance

    Innovative core-cache design (L1 and L2 private to the processor core), processor chip-cache design (L3), and processor node design (L4), allows z15 to keep more data closer to the processor, increase cache sizes, and decrease the latency in accessing the next levels of cache.

    This on-chip cache implementation optimizes system performance for high-frequency processors, with cache improvements, new Translation/TLB2 design, pipeline optimizations, better branch prediction, new accelerators, and architecture support.
- Reoptimized design for power and performance:
  - Improved instruction delivery
  - Improved branch prediction
  - Reduced execution latency
  - Optimized third-generation SMT
  - Enhanced out-of-order execution
  - New and enhanced vector instructions
- Dedicated co-processor for each PU:
  - The Central Processor Assist for Cryptographic Function (CPACF) is well-suited for encrypting large amounts of data in real time because of its proximity to the PU.

    CPACF supports Data Encryption Standard (DES), Triple Data Encryption Standard (TDES), AES-128, AES-256, Secure Hash Algorithm (SHA)-1, SHA-2, SHA-3, and True Random Number Generator. With the z15, CPACF supports Elliptic Curve Cryptography (ECC) clear key, improving the performance of Elliptic Curve algorithms. The following algorithms are supported: EdDSA (Ed448, Ed25519), ECDSA (P-256, P-384, P-521), ECDH(P-256, P-384, P521, X25519, X448). Protected key signature creation is also supported.
  - On-chip compression: The z15 is enhancing compression as follows:
    - Taking it from the I/O device level (IBM zEnterprise Data Compression (zEDC) feature) and moving it to the Nest Accelerator Unit on the processor chip.
    - Adding the Deflate-compliant (lossless data compression algorithm) and GZIP (GNU zip UNIX compression utility) compression/decompression support as hardware instructions.

    On-chip compression provides compression services for all logical partitions (LPARs), whereas the zEDC feature can be assigned to 15 LPARs only.

This innovation allows users to perform the compression with the improved performance and simplified management (no need to manage the zEDC features), on a processor chip level, without any delays associated with I/O requests, and with minimal CPU costs. The enhancement preserves the compatibility with the earlier version with the data that is compressed by zEDC features. Data that is compressed and written by zEDC is read and decompressed on the z15. This process simplifies the migration to z15 (on-chip compression removes the need of acquiring zEDC features).

Transactional Execution Facility

The Transactional Execution Facility, which is also known as *hardware transactional memory*, allows a group of instructions to be issued automatically. Either *all results* of the instructions in the group are committed or *no results* are committed, in a truly transactional manner. The execution is optimistic.

The instructions are issued, but previous state values are saved in transactional memory. If the transaction succeeds, the saved values are discarded. If it fails, they are used to restore the original values. Software can test the success of execution and redrive the code, if needed, by using the same or a different path.

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

Guarded Storage Facility

Guarded Storage Facility (GSF) is a hardware feature that Java uses to achieve pause-less garbage collection. GSF was introduced with the z14 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 the z15 Integrated Facilities for Linux (IFLs), IBM Z Integrated Information Processors (zIIPs), and System Assist Processors (SAPs). It allows more than one thread to simultaneously run in the same core and share all of its resources. This function improves the use 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 then 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. By 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.

Third-generation SMT on z15 is adjusted with the growth in the core cache and TLB2. It improves thread balancing, supports multiple outstanding conversions, optimizes hang avoidance mechanisms, and delivers improved virtualization performance to benefit Linux on IBM Z. z15 provides economies of scale with next generation multithreading (SMT) for Linux on Z and zIIP-eligible workloads while adding support for the I/O SAP.

Hardware decimal floating point function

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

Vector Packed Decimal Facility

Vector Packed Decimal Facility allows 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, IBM XL C/C++ V2.4.1 for z/OS V2.4, the COBOL optimizer, Automatic Binary Optimizer for z/OS V1.3, and Java are optimized on z15.

Single-instruction, multiple-data

The z15 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 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 allows dynamic optimization of code generation as it is being run.

Secure Execution for Linux

Secure Execution for Linux was introduced to isolate and protect kernel-based virtual machine (KVM) guests from hypervisor access. Hypervisor admin can still manage and deploy workloads, but is unable to view data running on a guest. Multiple tenants (applications) running in LPAR (as second-level guests under KVM) have fully isolated environments, which helps 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 platform has more memory that is installed than was ordered because part 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.

- With the z15 T01, up to 40 TB of memory for a five-CPC drawer configuration (25 percent increase compared to 32 TB on the z14).
- With the z15 T02, up to 16 TB of memory for a two-CPC drawer configuration (double compared with the IBM z14 ZR1).

Hardware system area (HSA) was expanded to 256 GB for the z15 T01 (compared to 192 GB with the z14). The z15 T02 has 192 GB HSA (compared to 64 GB for the z14 ZR1). The HSA area is a fixed size and is not included into the memory ordered.

Important: z/OS V2R3 and later releases require a minimum of 8 GB of memory (2 GB of memory when running under z/VM). z/OS can support up to 4 TB of memory in an LPAR.

Each operating system can allocate the amount of main memory up to the supported limit. The amount of incremental memory was changed to 1 GB with the z15.

Flexible memory

Flexible memory provides the extra physical memory that is needed to support the following scenario: You need activation base memory and HSA on a z15 T01 that has multiple CPC drawers and one CPC drawer that is out of service.

On z15 T01, the extra resources that are required for flexible memory configurations are provided when you configure memory features and memory entitlement. Flexible memory is available only with z15 T01: Max71, Max108, Max145, and Max190 features, and range 512 GB - 32512 GB, depending on the feature.

IBM Virtual Flash Memory

The IBM Virtual Flash Memory (VFM) feature is offered from the main memory capacity. For z15 T01, up to 12 VFM features (four VFM features for the z15 T02) can be ordered, each of 512 GB. VFM replaces the Flash Express adapters that were available on the z14 and z13.

VFM provides much simpler management and better performance by eliminating the I/O of the adapters that are in the PCIe drawers. VFM does not require any application changes when moving from IBM Flash Express.

VFM can help improve availability and handling of paging workload spikes when z/OS is run. VFM can also be used in Coupling Facility (CF) images to provide extended capacity and availability for workloads that use WebSphere MQ Shared Queues structures.

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

5.2 Virtualization

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

IBM Z is designed to concurrently run multiple virtual guest systems, which provides each system with the required dynamic share of the resources with minimal costs and performance overhead.

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

Virtualization management, which is called a *hypervisor*, operates on hardware and software levels on IBM Z. The hardware hypervisor (first level or type 1), is Program Resource Sharing Manager (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 Z platform so that you can create and build multiple virtualized guest systems running various open source applications on the Z platform with high levels of performance and integrated security.

IBM Z development is continuously improving virtualization techniques, providing 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 z15, and the co-existence of multiple hypervisors: PR/SM, z/VM, and KVM.

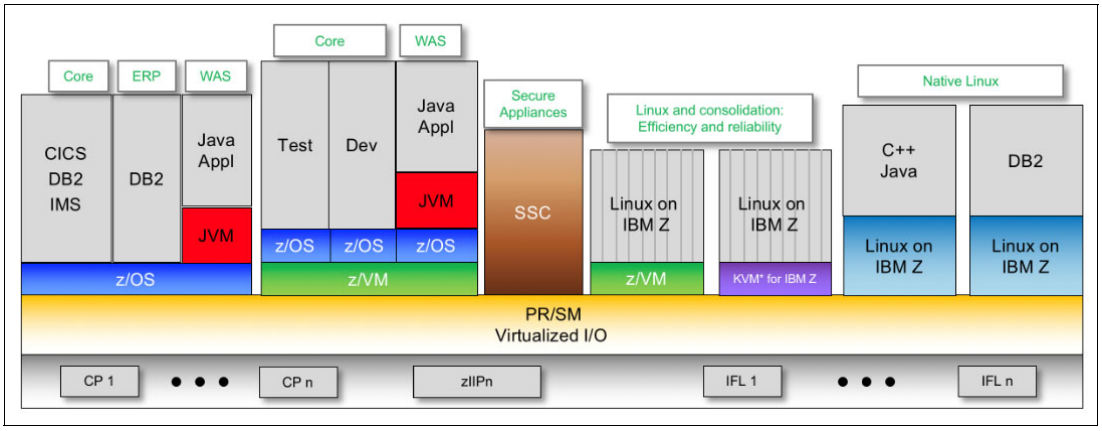

Figure 5-1 Virtualization on the IBM Z platform

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 Z platform. It provides each system with a required share of these resources and dynamically adjusts it 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 allows for a high degree of the 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 Z, z/OS, z/TPF, or z/VSE and runs one of these operating systems.

Initially, you might allocate a set of resources (CPU, memory, and I/O) and their quantity 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 is allowed to complete within a required timeline.

PR/SM evolved over the decades on Z platforms. It is a proven, secure, and fundamental IBM Z technology. Every generation of the Z platform brings PR/SM improvements that are aimed to demonstrate even better system performance. With the z15, PR/SM is extended to support new features, such as the LPAR support of the System Recovery Boost. For more information about System Recovery Boost features, see “System Recovery Boost” on page 91.

Dynamic Partition Manager

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

DPM provides partition lifecycle and integrated dynamic I/O and PCIe functions management for Linux on 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.

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.
- 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.

Enhancements to DPM on z15 simplify the installation of the Linux on Z, support more hardware features, and enable base cloud provisioning through OpenStack, including the following enhancements:

- Support for auto-configuration of devices to simplify Linux Operating System Installation, where Linux distribution installers use functions
- Secure FTP through HMC for starting and installing an Operating system by using FTP
- Support for OSA-Express, FICON Express, Crypto-Express, RDMA over Converged Ethernet (RoCE) Express features, and HiperSockets

Configuration note: The z15 can be configured in either DPM mode or in PR/SM mode. It cannot be configured in both modes concurrently. DPM supports Linux on Z, z/VM, and KVM only, and PR/SM mode supports a mixture of z/VM, KVM, z/OS, Linux on Z, z/TPF, and z/VSE.

z/VM

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, historically being the first virtual machine (VM). z/VM runs in an LPAR and manages the system hardware resources (CPU, memory, and I/O) between its guest systems in a most efficient way.

z/VM supports z/OS, Linux on Z, z/TPF, and z/VSE as its guest systems. It can also enable nested virtualization and can host z/VM as a guest system, allowing 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, having the capabilities to scale out horizontally and vertically. The z15 T01 supports up to 85 z/VM LPARs, while the z15 T02 supports 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 Z distributions. It is a type 2 hypervisor that provides simple, cost-effective platform virtualization for Linux workloads that are running on the 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 such as Linux on IBM Z, z/OS, z/VM, z/VSE, and z/TPF running in different LPARs on the 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 these resources:

- IBM Knowledge Center
- *Virtualization Cookbook for IBM Z Volume 5: KVM*, SG24-8463

5.2.1 Hardware virtualization

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

Workload separation is one of the most important parameters. PR/SM on z15 is certified at the highest level of Common Criteria (EAL5+), which confirms that PR/SM provides the isolation of the running workloads as though they were running on separate physical 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 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 not limited to, the following examples:
  - IBM z/OS Container Extensions (zCX)
  - IBM Java virtual machine (JVM)
  - Various XML System Services
  - IPSec offload
  - 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 Z System Recovery Boost
- IFL processor is used exclusively for Linux on Z and for running the z/VM and KVM hypervisors in support of Linux VMs.
- Internal Coupling Facility processor is used for z/OS clustering and supporting the family of Parallel Sysplex solutions. Integrated Coupling Facility (ICF) is dedicated to this function and exclusively run the Coupling Facility Control Code (CFCC).

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

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

- System Assist Processors (SAP) that run I/O operations
- Integrated Firmware Processors (IFPs) for native PCIe features

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 all shared or all dedicated.

PR/SM ensures that the processor state is properly saved and restored (including all registers) when you switch a physical processor from one logical processor to another. 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, and z/VSE each can dynamically define and change the number and type of reserved PUs in an LPAR profile. No pre-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 non-disruptively and quickly make it available to guests. z/VM virtualizes this support to its guests, which can also increase their memory non-disruptively. Releasing memory is still a disruptive operation.

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

The z/Architecture 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 or 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.

z15 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 allow sharing of address spaces, or parts of them, across multiple processes. For example, under z/VM, a single copy of the read-only part of a kernel can be shared by all VMs that use that operating system. Known as *discontiguous shared segment* (DCSS), this shared memory exploitation for many VMs can result in large savings of real memory and improvements in performance.

I/O virtualization

The z15 T01 supports six logical channel subsystems (LCSSs), while the z15 T02 supports three LCSSs. The z15 T01 has four subchannel sets in each LCSS with up to 256 channels, for a total of 1536 channels. The z15 T02 has three subchannel sets in each LCSS with up to 256 channels, for a total of 768 channels. In addition to the dedicated use of channels and I/O devices by an LPAR, I/O virtualization allows concurrent sharing of channels.

The z/Architecture also allows 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 channel subsystems (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 should be available for critical devices (for example, disks that contain vital data sets).

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

Many installations use the parallel access volume (PAV) function, which allows 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 were left queued when aliases are not available.

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

Subchannel sets

With the z15 T01, up to four subchannel sets of approximately 64,000 device addresses are available. The base addresses[3] 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 the z15 T02, 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 (z15 T01 only), which provides more connectivity through subchannel set 0.

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

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 DS8K series.

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

The dynamic I/O configuration function is supported by z/OS and z/VM. It provides the capability of concurrently changing the currently active I/O configuration. Changes can be made to channel paths, control units, and devices. A fixed HSA area in the z15 models greatly eases the planning requirements and enhances the flexibility and availability of these reconfigurations.

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 extreme integration in the hardware and firmware, virtualization on the z15 is highly efficient (the best in the industry). It delivers up to 100% sustained resource utilization, and the highest levels of isolation and security. Therefore, the cloud solution costs, whether hardware, software, or management, are minimized.

Cloud elasticity requirements are covered by the z15 granularity offerings, including capacity levels and Capacity on Demand (CoD). These and other technical leadership characteristics make the Z platforms the gold standard for the industry.

In addition, managing a cloud environment requires tools that can take advantage of a pool of virtualized compute, storage, and network resources and present them to the consumer as a service in a secure way.

A cloud management system must enable the management of virtualized IT resources to support different types of cloud service models and cloud deployment models. OpenStack (offered for z/VM and KVM) can satisfy a wide range of cloud management demands. It integrates various components to automate IT infrastructure service provisioning.

The z15 can also be tailored with a choice of Z-backed services that are delivered by way of IBM Cloud to help transform your infrastructure, applications, and data by exposing and connecting assets with simplified and intelligent operations across the infrastructure.

With IBM's acquisition of Red Hat, the hybrid cloud capabilities on IBM Z were extended. Support for running OpenShift on Linux on Z provides expansive cloud capabilities, including open containers, tools, and access to an extensive open community. For more information, see Announcing our direction for Red Hat OpenShift for IBM Z and LinuxONE.

The new cloud-native capabilities are delivered as pre-integrated solutions called IBM *Cloud Paks*. The 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

The z15 offers significant increases in capacity and performance over its predecessor, the z14. Several elements contribute to this effect, including the larger number of 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

The z15 expands the offer on subcapacity settings. Finer granularity in capacity levels allows the growth of installed capacity to more closely follow the enterprise growth, for a smoother, pay-as-you-go investment profile. Many performance and monitoring tools are available on Z environments that are coupled with the flexibility of the CoD options (see 5.3.2, “Capacity on Demand” on page 86). These features help to manage growth by making capacity available when needed.

z15 T01 capacity levels

Regardless of the installed model, the z15 T01 offers four distinct capacity levels for the first 34 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[4] to a subcapacity of the full CP capacity. The clock frequency of those processors remains unchanged. The capacity adjustment is achieved through other means.

Full capacity CPs are identified as CP7. On the z15 T01, up to 190 CPs can be configured as CP7. Up to 34 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 34 CPs are configured to the system, all must be full capacity because all CPs must be at the same capacity level. Granular capacity adds 102 subcapacity settings to the 190 capacity settings that are available with full capacity CPs (CP7). The 292 distinct capacity settings in the system provide for a range of over 1:716 in processing power.

A processor is always 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 z15 T01, the following CP subcapacity levels are a fraction of full capacity:

- Model 7xx = 100 percent
- Model 6xx = 56 percent
- Model 5xx = 38 percent
- Model 4xx = 13 percent

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

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

Figure 5-2 shows the z15 T01 full capacity and subcapacity offerings.

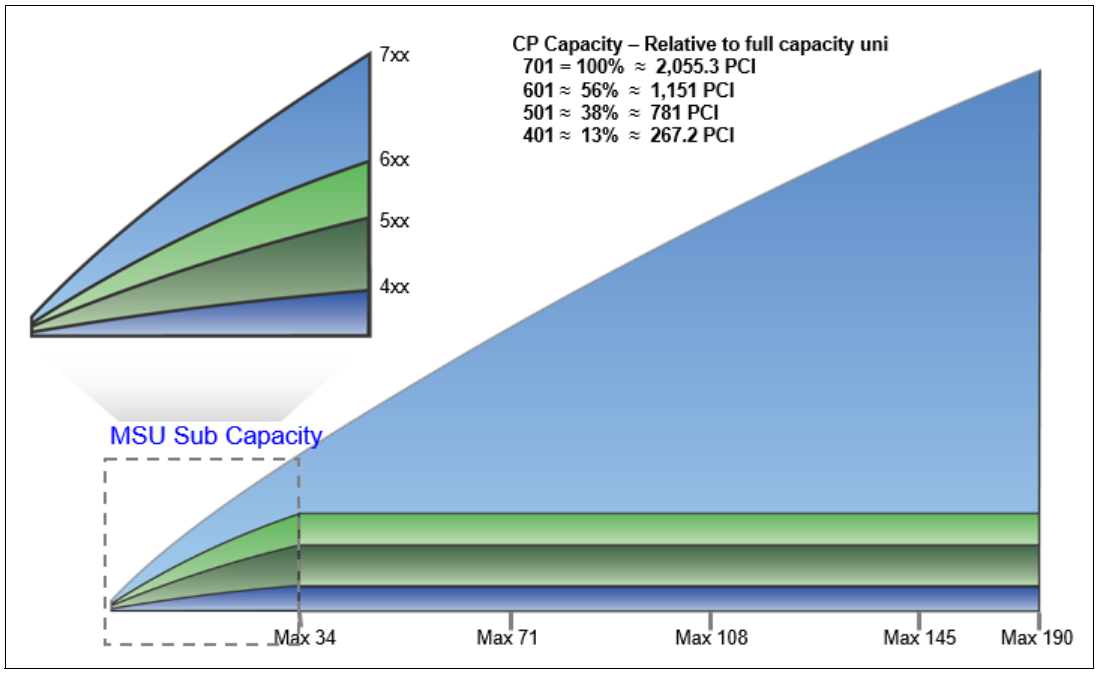

Figure 5-2 IBM z15 T01 full and subcapacity CP offerings

z15 T02 capacity levels

The z15 T02 offers 26 distinct capacity levels for each CP in the configuration for a total of 156 capacity settings (26 x 6). These processors deliver the scalability and granularity to meet the needs of small and medium-sized enterprises.

As in the z15 T01, a z15 T02 PU is always set to full capacity when it is characterized as anything other than a CP, such as a zIIP, an IFL, or an ICF.

Figure 5-3 gives more details about z15 T02 capacities settings.

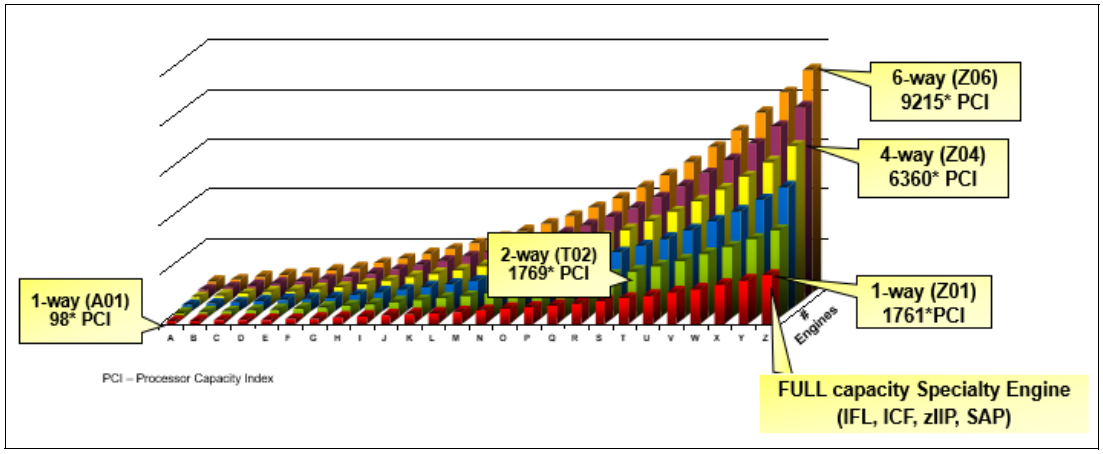

Figure 5-3 z15 T02 capacity settings offerings

To help size a Z platform to fit your requirements, IBM provides a no-cost tool that reflects the latest IBM Large Systems Performance Reference (LSPR) measurements, 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.

System Recovery Boost upgrade (optional, z15 T01 only)

System Recovery Boost upgrade can optionally enable the temporary activation of more physical zIIP processors on the z15 T01 by way of a priced Boost temporary capacity record. This record requires the use of feature codes 6802 and 9930.

Note: System Recovery Boost functionality is embedded in the z15 firmware. It can be used without ordering extra zIIP capacity (Feature Code 6802 and Feature Code 9930). System Recovery Boost upgrade option (extra temporary zIIP capacity) is only available for z15 T01.

Activation of these processors uses unused processing cores in the z15 to provide more zIIP processing capacity that accelerates execution of their workload (general-purpose workload and workload that is already 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 the z15 T01. This extra zIIP capacity is then shared across images in accordance with PR/SM management controls, which makes more zIIP capacity available to individual system images.

Images that want to take advantage of this extra zIIP capacity predefine reserved logical zIIP capacity in their PR/SM image profiles. That way, the operating system can then 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.

5.3.2 Capacity on Demand

The z15 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.

The z15 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 can also 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 the z15.

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.

Use of 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 the IBM Resource Link.

For more information about the CoD offerings, see *IBM z15 (8561) Technical Guide*, SG24-8851.

Permanent upgrades

Permanent upgrades of processors (CP, IFL, ICF, zIIP, and SAP) and memory, or changes to a platform's Model-Capacity Identifier (up to the limits of the installed processor capacity on an existing z15) can be performed by customers through the IBM Online Permanent Upgrade offering by using the CIU facility.

Temporary upgrades

Temporary upgrades of a z15 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 is available on the z15 that enables concurrent and temporary capacity growth of the CPC. On/Off CoD can be used for peak workload requirements for any length of time, includes a daily hardware charge, and 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 (CP, IFL, ICF, zIIP, and SAP), increase the CP capacity level, or both.

Prepaid OOCoD tokens: Beginning with the IBM z15, new prepaid OOCoD tokens do *not* carry forward to future systems.

Capacity Backup function

CBU allows you to perform a concurrent and temporary activation of extra CP, ICF, IFL, zIIP, and SAP, an increase of the CP capacity level, or both. This function can be used during an unforeseen loss of 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, as a possibility 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 in duration) and one disaster activation (up to 90 days in duration). The record can be configured to be valid for up to five years. You can also order more tests for a CBU record in quantities of five tests up to a maximum of 15.

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

Capacity for Planned Event function

Capacity for Planned Events (CPE)[a]: IBM z15 is planned to be the last server to offer Capacity for Planned Events.

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 information that is mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functions.

CPE allows 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 a planned outage of 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 three days.

The CPE feature is optional and requires unused capacity to be available on CPC drawers of the backup system as unused PUs, as a possibility 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 manage the CP and zIIP capacity of z15 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 z15 provisioning capability gives you a flexible, automated process to control the configuration and activation of On/Off CoD offerings.

Taylor Fit Pricing for IBM Z Hardware

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 z15 (8561) Technical Guide*, SG24-8851.

5.3.3 z15 performance

The Z microprocessor chip of the z15 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 the z15 is a result of the enhancements that are described in Chapter 2, “IBM z15 Model T01 hardware overview” on page 23, Chapter 3, “IBM z15 Model T02 hardware overview” on page 41, and in 5.1, “Technology improvements” on page 72.

LSPR workload suite: z15 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 performance of the systems by using various workloads and publishes the results in the Large Systems Performance Reference (LSPR) report.

IBM also provides a list of MSU ratings for reference.

Capacity performance is closely associated with how a workload uses and interacts with a particular processor hardware design. Workload capacity performance is sensitive to the following major factors:

- Instruction path length
- Instruction complexity
- Memory hierarchy

The CPU Measurement Facility (MF) data offers insight into the interaction of workload with the hardware design. CPU MF data helps LSPR to adjust workload capacity curves that are based on the underlying hardware sensitivities, in particular the processor access to caches and memory. This process is known as *nest activity intensity*.

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

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

These categories are based on the relative nest intensity, which is influenced by many variables, such as application type, I/O rate, application mix, CPU usage, data reference patterns, LPAR configuration, and the software configuration that is running, among others. CPU MF data can be collected by z/OS System MF on SMF 113 records 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 Z and z/VM environments.

The LSPR contains the internal throughput rate ratios (ITRRs) for the z15 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 Z platforms can run at up to 100% utilization 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 millions of service units (MSU) metrics) is not a valid method. Therefore, use zPCR and involve IBM technical support when you are planning for capacity.

For more information about z15 performance, see *IBM z15 (8561) Technical Guide*, SG24-8851.

Throughput optimization with z15

The memory and cache structure implementation in the CPC drawers of the z15 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 of the memory. Therefore, the system is managed as a memory coherent symmetric multiprocessor (SMP).

Processors within the z15 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 use 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 more 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 z15, 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 use of shared cache, and lowering the interference across multiple partitions.

The z/OS dispatcher is enhanced to operate with multiple dispatching queues, and tasks are distributed among these queues. Specific z/OS tasks can be dispatched to a small subset of logical processors. PR/SM ties these logical processors to the same physical processors, which improves the hardware cache reuse and locality of reference characteristics, such as reducing the rate of cross communication.

To use the correct logical processors, the z/OS dispatcher obtains the necessary information from PR/SM through interfaces that are implemented on the z15. The entire z15 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).

The z15 includes several HiperDispatch algorithm enhancements, including the following examples:

- Improved memory affinity
- Improved LPAR placement based on z14 experience
- Exploitation of 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. The general recommendation 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 into the hardware and software stacks of the IBM z/Architecture, where mean time between failures is measured in decades. The RAS strategy is to manage change by learning from previous generations of IBM Z 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 in support of new features and functions. Enhanced Driver Maintenance (EDM) can help reduce the necessity and eventual duration of a scheduled outage.

When properly configured, the z15 can concurrently activate a new LIC Driver level. Concurrent activation of the select new LIC Driver level is supported at released synchronization points. However, for certain LIC updates, a concurrent update or upgrade might not be possible.

z15 builds on the RAS characteristics of the Z family, with the following RAS improvements:

- System Recovery Boost

  System Recovery Boost is new with the z15. It offers more CP capacity during particular system recovery operations to accelerate system (operating system and services) start when the system is being started or shutdown. System Recovery Boost is operating system-dependent. No other hardware, software, or maintenance charges are required during the boost period for the base functions of System Recovery Boost.

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

  System Recovery Boost provides the following options for the capacity increase:

  – Subcapacity Boost: During the boost period, subcapacity engines are being 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 information about the optional temporary capacity upgrade, see “System Recovery Boost upgrade (optional, z15 T01 only)” on page 86.

  System Recovery Boost can also be used for recovery processes, such as IBM HyperSwap[5], 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.

[5] IBM HyperSwap is a high availability 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.

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 System Recovery Boost users are z/OS, z/VM, z/VSE, and z/TPF. z/VM uses System Recovery Boost if it runs on 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 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.

System Recovery Boost support is available with GDPS V4.2 by way of 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.

- Level 3 and Level 4 cache enhancements that use powerful symbol ECC to extend the reach of prior z14 cache and memory improvements for improved availability. The level 3 and level 4 cache powerful symbol ECC is designed to make it resistant to more failure mechanisms. Preemptive dynamic random access memory (DRAM) marking was added to the main memory to isolate and recover failures more quickly.
- On-chip compression co-processor for accelerating operations on a core level, for all LPARs. The technology replaces the zEDC feature, improving reliability and availability.
- A PU single-chip module (SCM) uses 14 nm SOI technology and consists of 17 layers of metal, 9.1 billion transistors (with up to 12 active cores per chip that run at 5.2 GHz with for z15 T01, or with up to 11 active cores per chip running at 4.5 GHz for z15 T02)), which enhances thermal conductivity and improves reliability.
- VFM: Flash Express PCIe feature replacement with memory dual inline memory modules (DIMMs), which is more robust solution that uses RAIM protection against memory faults.

z15 T01 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 or repair. To ensure that the z15 T01 configuration supports removal of a CPC drawer with minimal effect to the workload, consider the flexible memory option (see “Flexible memory” on page 76).

New for z15 T02, Concurrent Drawer Add (for Max31) is now supported[6], while EDA enables Concurrent Drawer Repair for Max65 feature.

z15 also continues to have RAIM[7] 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 is used in storage systems for several years. For more information about RAS features, see *IBM z15 (8561) Technical Guide*, SG24-8851.

z15 T01 can be configured with of a maximum of five CPC drawers that are designed as a field replaceable unit (FRU). Connections among the CPC drawers are established through symmetric multiprocessor (SMP) cables. Each CPC drawer consists of two logical PU clusters and contains four CP SCMs and one Storage Controller (SC) SCM. In addition to

[6] Concurrent drawer add can be performed if Feature Code 2271 (CPC1 Reserve) is part of the initial system order.

[7] 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.

SCMs, CPC drawers host memory DIMMs, connectors for I/O, oscillator interface, and manifolds.

z15 inherits I/O infrastructure reliability improvements from z14, including Forward Error Correction (FEC) technology that enables better recovery of FICON channels. FICON Express16SA and FICON Express16S+ features continue to provide a new standard for transmitting data over 16 Gbps links by using 64b/66b encoding.

The z15 air-cooled configuration includes a newly designed front-to-rear radiator cooling system. The radiator pumps, blowers, controls, and sensors are N+2 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 reliability of the pumps. The replacement of pumps or blowers is concurrent with no affect on performance.

A water-cooling system is also an option with the z15 T01, with water-cooling unit (WCU) technology. Two redundant WCUs run with two independent chilled water feeds. One WCU and one water feed can support the entire system load. The water-cooled configuration is backed up by the rear door heat exchangers in the rare event of a problem with the chilled water facilities of the customer.

RAS also includes the following enhancements:

- Integrated sparing
- Error detection and recovery improvements in caches and memory
- 25 GbE RoCE Express2 (Optics as FRU)
- PCIe coupling links (improved diagnostics)
- Enhanced channel logging
- OSA firmware changes to increase the capability of concurrent maintenance change level (MCL) updates
- System power cycle management (servicing capability)
- CFCC level 24 (various enhancements for improving CF resiliency)
- IBM RMF reporting improvements

The z15 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 a z15 (for example, as a result of a DR situation), the SAPs can be concurrently added to the CPC configuration.

It is possible to concurrently add CP, zIIP, IFL, and ICF processors to an LPAR. This function is supported by z/VM, and (with appropriate program temporary fixes (PTFs)) by z/OS, and z/VSE. It is possible to concurrently add memory to an LPAR as well. This feature is supported by z/OS and z/VM.

The z15 supports adding Crypto-Express features to an LPAR dynamically by changing the cryptographic information in the image profiles. Users can also dynamically delete or move Crypto-Express features. This enhancement is supported by z/OS, z/VM, and Linux on IBM Z.

5.4.1 RAS capability for the Support Element

Enhancements are made to the Support Element (SE) design for z15. Two 1U trusted servers are inside the z15 A frame: one is always the primary SE and the other is the alternative SE.

The primary SE is the active SE. The alternative acts as the backup. Information is mirrored once per day. The SE servers include N+1 redundant power supplies.

More powerful SEs offer RAS improvements, such as ECC memory, redundant physical networks for SE networking requirements, redundant power modules, and better thermal characteristics.

5.4.2 RAS capability for the Hardware Management Console

Enhancements are also made to the HMC designs for z15. The new1U HMC design offers the same RAS improvements as those improvements of the 1U SE. The 1U HMC option can be mounted in a customer-supplied rack. The MiniTower design that is used before z15 is still available.

Back up options are available for the HMC and SE. The HMC can host an FTP server with scheduled backups.

5.5 High availability with Parallel Sysplex

The Parallel Sysplex technology is an IBM Z clustering technology that allows users to build highly resilient, highly scalable, dynamic, and robust 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.

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 Z, or within Z, running alongside other LPARs. PUs that are characterized as CPs or ICFs are used to configure CF LPAR. ICFs are preferred because of software licensing reasons. Two or more Coupling Facilities (CFs) are recommended for high availability.

- 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, CICS TS, and WebSphere MQ for synchronizing the access to the shared data and resources.

  z/VM can emulate CFCC as a guest VM, which allows 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, easy-to-use mechanism for duplexing structure data held in CFs. This function provides a robust recovery mechanism for failures, such as loss of a single structure on CF or loss of connectivity to a single CF. The recovery is done through rapid failover to the other structure instance of the duplex pair.

For more information about deploying system-managed CF structure duplexing, see the technical paper *System-Managed CF Structure Duplexing,* ZSW01975USEN. You access the paper 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 Z platforms, 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 Z I/O virtualization capabilities, such as MIF. MIF allows several LPARs to share I/O paths in a secure way, which maximizes use and greatly simplifies the configuration and connectivity.

A properly configured Parallel Sysplex cluster is designed to maximize availability *at the application level*. Rather than a quick recovery of a failure, the Parallel Sysplex design objective is *zero 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 of the data. Subsystems in z/OS declare resources to the CF. The CF accepts and manages lock and unlock requests on those resources, which helps to ensure the data integrity. A duplicate CF further enhances the availability. Key users of the data-sharing capability are Db2, WebSphere MQ, WebSphere ESB, IMS, and CICS.

  Because these 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. z/OS WLM component is a key workload distributor, ensuring that the required SLA (Service Level Agreement) 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

  Global Resource Serialization (GRS) component of z/OS manages the access to the shared resources (such as CPU, memory, network, and storage) across all members of 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. A single-system image ensures reduced complexity from 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[8]

  Three hardware generations (the current and the two previous generations) often are supported in the same Parallel Sysplex. z15 supports the following coupling features:

  - Integrated Coupling Adapter Short Reach (ICA SR 8x link, up to 150 m [492 feet])
  - Coupling Express Long Reach (CE LR) 1x link, up to 10 km [6.2 miles] unrepeated)

  Software support for multiple releases or versions is provided.

- CF encryption support

  Supports encrypted data while it is being transferred to and from the CF because it is 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-Express5 Coprocessor (CEX5C), the Crypto-Express6 Coprocessor (CEX6C), or the Crypto-Express7 Coprocessor (CEX7C) feature.
  - Support provided can be enabled only when all systems are z/OS 2.3 or higher. Toleration support with reduced functionality is provided for z/OS 2.2 and z/OS 2.1.

- 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 a z15 or z14 is driven by an updated HCD instance that is running in a z/OS LPAR on a remote z15 or z14.

- Improved performance and resilience

  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.

- z15 is providing new resiliency mechanism

  The CF provides more information to z/OS about every message path that appears active. Namely, the current SYID (system ID) 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.

[8] Provided coupling and timing links are not InfiniBand.

The components of a Parallel Sysplex as implemented within the IBM z/Architecture are shown in Figure 5-4. The configuration is one of many possible Parallel Sysplex configurations.

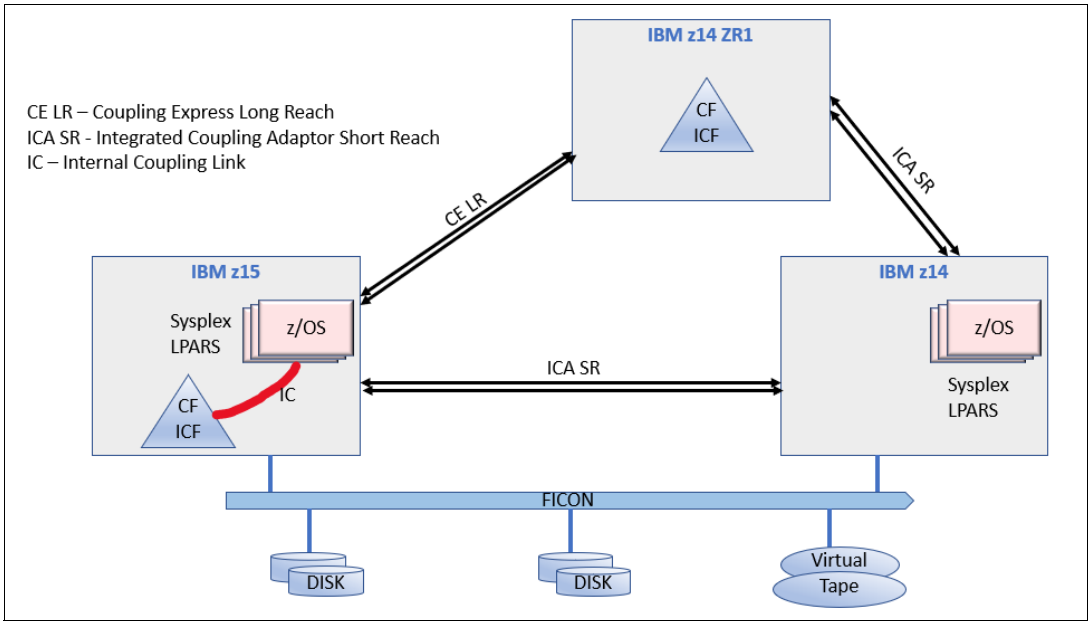

Figure 5-4 A sample Parallel Sysplex deployment

Server Time Protocol (STP) over coupling links provides time synchronization to all systems in the sysplex.

Note: IBM z15 can be coupled with other systems only by way of ICA SR or CE LR. zEC12, zBC12, z196, z114, and older systems cannot coexist in the same sysplex with z15.

For more information about coupling link options, see 4.4, “Clustering connectivity” on page 64.

5.6 Pervasive encryption

Cryptography and corporate security always were a fundamental aspect of IBM Z platform development. Z 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 as well, 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 are becoming more demanding too, such as 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 the data-centric one: the data is becoming the new perimeter, and the encryption applies to all of the data, regardless of its origin and location. At the same time, this approach does not require costly application changes, being transparent to the applications and their service consumers.

Pervasive encryption provides 360 degrees of data encryption, for both 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 the standards and compliance.

Pervasive encryption uses the hardware cryptography acceleration in the z15, which is proven to be more effective, performant, and stable compared to the software encryption.

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

- Integrated cryptographic hardware:
  - Central Processor Assist for Cryptographic Functions (CPACF) is a co-processor on every PU that accelerates symmetric encryption operations.
  - Crypto-Express features are hardware security modules (HSMs)[9] with these features:
    - Complying with Federal Information Processing Standards (FIPS) 140-2 Level 4 (achieving the highest level of compliance within this standard).
    - Accelerating various cryptographic algorithms (Digital signature sign/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, providing highest encryption performance.
- Data set and file (also known as volume) encryption: Linux on 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 Z is a new capability that is introduced with the z15 T02, available for all models of the z15 generation. Engineered to help protect against insider and outsider threats in multitenanted cloud environments, it ensures that users and system administrators cannot access sensitive data in Linux based virtual environments.

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

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 to data in the clear.

- Secure Boot: This new enhancement targets securing the booting process of an open source operating system, for example, the family of Linux distributions. With the increased number of public open source repository attacks, the extra step of verifying the operating system kernel version was introduced. Secure boot integrity checks validate that an operating system kernel is from an official provider and is not compromised. This new enhancement can be used by Linux on 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 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, switch logs, or technicians who are using Fibre Channel (FC) analyzers to examine the packets during problem determination to name a few.

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 the IBM z15 T01 and the IBM Storage subsystem (DS8900F). 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 Express16SA (Feature Code 0436 or Feature Code 0437), turn on the Fibre Channel Endpoint Security panels on the HMC so setup can be done.

IBM Secure Key Lifecycle Manager (SKLM) acts as a trusted authority for key generation operations and authentication server. It provides shared secret key generation in a relationship between an FC initiator (IBM Z) and the IBM Storage target. The solution implements authentication and key management called IBM Secure Key Exchange (SKE).

Before establishing the connection, each link must be authenticated, and if successful, then becomes a trusted connection. A policy sets the rules, for example, enforcing the 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) is encrypted when it leaves either endpoint (source) and decrypted at the destination. Encryption/decryption is done at the FC adapter level. There is no involvement of the operating system that is running on the host (IBM Z) in Endpoint Security related operations. Tools are provided at the operating system level for displaying information about encryption status.

5.6.2 IBM Data Privacy Passports

Data is almost never static in the system. It flows between various applications and system components, completes an extract, transform, and load (ETL) process that moves the data to another platform for further analysis and reporting, and other actions. Today, data privacy (controlling the usage of the data) is as important as data protection (the use of secure protocols or other means to encrypt and protect the data). Data protection ensures a secure flow of the data from a starting point to an endpoint with data remaining protected and encrypted always. However, data protection capabilities are not intended to provide data flow auditing and access management for the data, which is where *IBM Data Privacy Passports* can help.

Data Privacy Passports provides a data-centric audit and protection (DCAP) approach for the protection of eligible data across the enterprise. It is the next logical step from the IBM Z pervasive encryption strategy to protect data that is on IBM Z and as it moves throughout the enterprise and beyond.

The concept of DCAP is a change from the current model. Before the data is moved around the enterprise, it is repackaged into a secure object. In Data Privacy Passports, this object is called the Trusted Data Object (TDO). The Data Privacy Passports feature does this protection at a field-level, which means that there is a level of granularity to this protection that cannot be obtained from more broad protection techniques.

Important: The Data Privacy Passports feature supports data sources that can be accessed through a JDBC connection.

Data Privacy Passports is composed of a component that is known as the *Passport Controller* that provides all the protection, enforcement, policy, and key management for the solution. The Passport Controller provides an intercept point to transform *raw* eligible data to enforce data protection. The data is protected at the point of extraction and enforced at the point of consumption. The flexible policy defines the pool of trusted users and their access to the data based on their role. The key lifecycle management (generation, revocation, and other actions) is done within Passport Controller.

The data is encrypted with standard AES-256 keys, which are managed by the Passport Controller. Future access to sensitive data can be revoked remotely by using Data Privacy Passports and can even be made unusable by destroying its encryption key.

For more information, see *Protecting Data Privacy Beyond the Trusted System of Record*, REDP-5567.

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 way of tamper protection during installation and runtime, restricted administrator access, and encryption of data in-flight and at-rest.

A Secure Service Container is an integrated IBM Z appliance and hosts the most sensitive workloads and applications. They act as a highly protected and secured digital vault and enforce 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 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 (TLS, 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, 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 verified booting 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 that is certified by EAL5+ Common Criteria Standard, which is the highest level that ensures the workload separation and isolation.
- Horizontal workload isolation - separation from the rest of the host environment.

A Secure Service Container is a powerful IBM technology for providing the extra protection of 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.

6

Operating system support

This chapter describes the operating system requirements and support considerations for the IBM z15 and its features.

Support and use 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 8561DEVICE (z15 T01) and 8562DEVICE (z15 T02) at Preventive Service Planning buckets for mainframe operating environments.

This chapter includes the following topics:

6.1 Software support summary

The software portfolio for the z15 includes various operating systems and middleware that support the most recent and significant technologies. The following major operating systems are supported:

- z/OS
- z/VM
- z/VSE
- z/TPF
- Linux on IBM Z and the kernel-based virtual machine (KVM) hypervisor

6.1.1 Operating system summary

The current and minimum operating system levels that are required to support the z15 are listed in Table 6-1 on page 105. Operating system levels that are no longer in service are not covered in this publication. These older levels can support certain features.

Program temporary fixes (PTFs) and PSP buckets: The use of several features depends on a particular operating system. In all cases, PTFs might be necessary with the operating system level indicated.

Review the IBM Z hardware fix categories before the operating system is installed.

PTFs for z/OS, z/VM, and z/VSE 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.

Table 6-1 Operating system requirements

Operating system[a]	End of service	Notes
z/OS V2R5[b]	Not announced	See the z/OS, z/VM, z/VSE, and z/TPF subsets of the 8561DEVICE (z15 T01) and 8562DEVICE (z15 T02) Hardware PSP buckets and Fix Categories before installing the z15.
z/OS V2R4	Not announced.	
z/OS V2R3	September 2022.[c]	
z/OS V2R2	September 2020.	
z/OS V2R1	September 2018.	
z/VM V7R2	Not announced.	
z/VM V7R1	Not announced.	
z/VM V6R4	March 2021.	
z/VSE V6R2	Not announced.	
z/TPF V1R1	Not announced.	
Linux on IBM Z[d]	Support information is available for SUSE[e], Red Hat[f], and Canonical.[g]	
KVM hypervisor	For more information about minimal and recommended distribution levels, see the Linux distributors' websites.	

a. Only z/Architecture mode is supported.
b. Preview announcement at the time of this writing - per IBM United States Software Announcement 221-057
c. 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.
d. For more information, see the Linux on IBM Z page of the IBM Z website.
e. For more information, see the Support page of the SUSE website.
f. For more information, see the Red Hat Enterprise Linux Life Cycle page of the Red Hat website.
g. For more information, see the Ubuntu for IBM Z page of the Ubuntu website.

z/Architecture mode: As announced on January 14, 2015 with Announcement letter 115-001, all systems support operating systems that are running in z/Architecture mode only. This support applies to operating systems that are running native on PR/SM and operating systems that are running as second-level guests.

IBM operating systems that run in ESA/390 mode are no longer in service or currently available with only extended service contracts.

All 24-bit and 31-bit problem-state application programs that were originally written to run on the ESA/390 architecture are unaffected by this change.

6.1.2 Application development and languages

Software developers can take the advantage of having the multiple programming language environments that run on IBM Z. Various teams at IBM constantly work on extending the compilers and development tool support for Z platform to provide developers with everything that is needed for the agile and other modern development methods.

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

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

z/OS platform is actively embracing open source projects, extending z/OS support (in partnership with Rocket Software) for Anaconda stack (used by Python developers for business analytics) and Apache Spark. For more information, see this IBM z Systems Development Blog entry.

IBM recently announced the z/OS Container Extensions (zCX) feature, which allows you to run Docker containers natively under z/OS. At the time of this writing, this feature requires z/OS 2.4 and a z14 platform (and higher). For more information, see this web page.

The extensive set of advanced integrated development environments (IDEs) and integration tools are available for continuous development, testing, and deployment of application code.

IBM Z embraces emerging concepts, such as DevOps. DevOps is the process of bringing Development and Operations together to share processes and procedures to reduce the risk of change and improve the speed of deployment. Organizations are embarking on their journey with digital transformation and entering the API economy. So it becomes essential to connect business-critical applications that run on Z platforms 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, for example, IBM Application Delivery Foundation for z and IBM Rational® Team Concert®, which can be coupled with IBM Application Discovery and Delivery Intelligence (ADDI) application discovery technology, enable developers to understand the applications, gain cognitive insights into the process, and evolve those valuable older assets at the speed of business with reduced risk to the enterprise.

The use of modern development practices is part of the transformation. IBM Rational Team Concert and open source-based Git Version Control Tools for IBM z/OS (ported by Rocket Software) are modern source code managers that run on and support z/OS.

For more information about software for Z platforms, see the Products catalog website.

Note: The use of the most recent versions of the compilers is of utmost importance. The compilers take advantage of the latest technologies on Z platform and of the performance benefits that are introduced. Examples of benefits include new cache structures, new machine instructions, and instruction execution enhancements.

The IBM z15 inherits and further enhances the features and functions from its predecessor, the IBM z14, such as single-instruction, multiple-data (SIMD) architectural notation, which provides more efficient vector processing of the data.

Java applications benefit from Guarded Storage Facility (GSF), which enables pause-less garbage collection.

IBM z15 brings on-chip compression for the acceleration of the data compression that can be used by all applications.

Operating systems that run on z15 can use System Recovery Boost (a new feature to accelerate the recovery after an outage).

The following new security functions were introduced to complement the 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.
- Secure Boot: A feature for verifying the open source operating systems' kernel to ensure that it comes from the trusted provider.

6.1.3 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 uses the most recent IBM z/Architecture and performance optimization, enhanced XML parsing support, and capability of programming with Unicode. It also supports Java 8 SDKs for Java interoperability and delivers new COBOL statements, new and changed compiler options, and changed APIs.

- IBM Automatic Binary Optimizer for z/OS

  The Automatic Binary Optimizer for z/OS improves the performance of compiled COBOL programs. The Optimizer does not require source code, source code migration, or the tuning of performance options. It uses modern optimization technology to target Z platforms and accelerate the performance of compiled COBOL applications.

- Enterprise PL/I

  The latest version of Enterprise PL/I provides web interoperability, which includes web services, XML parsers, and Java Platform, Enterprise Edition (Java EE). The compiler also includes the expanded support for UTF-16. Enterprise PL/I for z/OS allows you to capitalize on IT investments while modernizing your infrastructure.

- z/OS XL C/C++

  The 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 in concert with z/OS Problem Determination Tools.

IBM Enterprise COBOL and Enterprise PL/I support are strategic components (separately orderable products) for IBM Rational Developer for IBM Z software. These features provide a robust, integrated development environment (IDE) for COBOL and PL/I and connecting web services, Java Platform, Enterprise Edition (Java EE) applications, and traditional business processes.

z/OS XL C/C++ programmers can also tap into *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.

6.2 Support by operating system

This section lists the support by in-service operating systems of selected functions of the IBM z15.

For more information about the IBM z15 and its features, see *IBM z15 (8561) Technical Guide*, SG24-8851 and *IBM z15 (8562) Technical Guide*, SG24-8852.

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 all z15 enhancements. However, z15 capabilities differ depending on the z/OS release.

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 z15 is V2.1 with PTFs (IBM Extended Software Support Services offering must be purchased).

z/OS supports the following select (but not limited to) functions:

- Coupling Facility Control Code (CFCC) Level 24
- HiperDispatch Enhancements
- z/OS SLIP enhancement
- HIS support
- Exploitation of new hardware instructions – XL C/C++ ARCH(13) and TUNE(13)
- Cryptographic support
- OSA Express support
- RDMA over Converged Ethernet (RoCE) Express support
- On-chip compression acceleration
- On-chip sort acceleration
- System Recovery Boost

Before the z15 migration process is started, see Migration Considerations, Hardware PSP Buckets, and Fix Categories at this web page.

For more information, see the IBM Support website.

For more information about z/OS downloads, see the z/OS website.

6.2.2 z/VM

z/VM (V6R4[1], V7R1, and V7R2) supports the following functions on the z15:

- z/Architecture support.
- On-chip compression acceleration.
- On-chip sort acceleration.
- System Recovery Boost.
- ESA/390-compatibility mode for guests.

[1] End of support as of March 31st, 2021

- Crypto-Express7S adapter and cryptographic enhancements.
- Message-Security-Assist Extension (MSA-9/Elliptic Curve Cryptography (ECC) authentication).
- Dynamic I/O enhancements: Dynamic I/O support is provided for managing the configuration of OSA-Express7S[2] OSD CHPIDs, IBM Fibre Connection (FICON) Express16SA[2] Fibre Channel (FC) and Fibre Channel Protocol (FCP) CHPIDs, RoCE Express2, and Coupling Express Long Reach (CE LR) features.

> **z/VM logical partitions (LPARs):** z15 central processors (CPs) and Integrated Facilities for Linux (IFLs) feature increased capacity over the capacity of their predecessors. Therefore, we suggest that your 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.

IBM z/VM V7.2, in addition to V7.1 enhancements, brings new functions and management capabilities, such as:

- Centralized Service Management for non-Single System Image (SSI) environments to deploy service to multiple systems, regardless of geographic location, from a centralized primary location.
- 4TB real memory support: With the available PTF for APAR VM66173, z/VM delivers support for up to 4 TB of real memory, enabling z/VM systems to address a full 4 TB of first-level (real) memory, doubling the previously supported limit of 2 TB.
- Multiple Subchannel Set (MSS)[3] Multi-Target Peer-to-Peer Remote Copy (MT-PPRC) z/VM support for the GDPS environment. This feature allows a device to be the primary to up to three secondary devices, each defined in a separate alternative subchannel set (supporting up to three alternative subchannel sets). Also provides the CP updates necessary for virtual machine (VM) and HCD support of alternative subchannel sets.
- New Architecture Level Set of z13, z13s (LinuxONE Emperor / Rockhopper), or newer processor families.

IBM z/VM V7.1 offers a new release delivery and new function support that provides predictability for lifecycle management of z/VM systems through continuous delivery. z/VM V7.1 includes the following enhancements:

- Integration of the SSI function and Live Guest Relocation into the base, at no extra cost.
- Systems Recovery Boost support for CPs when configured as subcapacity.
- Upgrades to a new Architecture Level Set (for IBM zEC12 and higher).
- Extended address volume (EAV) paging.
- Eighty logical processors per LPAR.
- Dynamic Crypto.

For more information about PTF availability, see the z/VM Continuous Delivery News web page.

For more information about for z15 migration, see the hardware PSP buckets for 8561DEVICE, and 8562DEVICE z/VM subset.

---

[2] Certain I/O features are not available on IBM z15 Model T02 (machine type 8562). For the supported I/O features, see Chapter 3, “IBM z15 Model T02 hardware overview” on page 41.

[3] Supported by IBM z/VM 7.2 only.

For more information about all z15 features and functions that are supported by the z/VM releases, see *IBM z15 (8561) Technical Guide*, SG24-8851 and *IBM z15 (8562) Technical Guide*, SG24-8852.

6.2.3 z/VSE

The z15 support is provided by z/VSE 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.
- System Recovery Boost (subcapacity CP speed boost only).

For more information about all z15 features and functions that are supported by the z/VSE releases, see *IBM z15 (8561) Technical Guide*, SG24-8851 and *IBM z15 (8562) Technical Guide*, SG24-8852.

6.2.4 z/TPF

The z15 support is provided by z/TPF V1R1 with PTFs. For more information about all z15 features and functions that are supported by the z/TPF, see *IBM z15 (8561) Technical Guide*, SG24-8851 and *IBM z15 (8562) Technical Guide*, SG24-8852.

6.2.5 Linux on IBM Z

The Red Hat, SUSE, and Ubuntu releases that are supported on the z15 are listed in Table 6-2.

Table 6-2 Linux on IBM Z distributions

Linux on IBM Z distribution[a]	Version and release
SUSE Linux Enterprise Server	SLES 15 SP1 with service
SUSE Linux Enterprise Server	SLES 12 SP4 with service
Red Hat Enterprise Server	RHEL 8.0 with service
Red Hat Enterprise Server	RHEL 7.7 with service
Red Hat Enterprise Server	RHEL 6.10 with service
Canonical	Ubuntu 20.04 LTS with service
Canonical	Ubuntu 18.04.1 LTS with service
Canonical	Ubuntu 16.04.6 LTS with service

a. Only z/Architecture (64-bit mode) is supported. IBM testing identifies the “recommended levels” of the tested distributions.

For more information about certified and tested Linux distributions on Z, see the Linux on IBM Z page of the IBM Z website.

For more information about all z15 features and functions that are supported by the Linux on IBM Z distributions, see *IBM z15 (8561) Technical Guide*, SG24-8851 and *IBM z15 (8562) Technical Guide*, SG24-8852.

6.2.6 Kernel-based virtual machine hypervisor

For the z15, the KVM is delivered and supported by the Linux distribution partners. For more information about KVM support for the Z platform, see the following resources:

- The documentation for your distribution.
- *The Virtualization Cookbook for IBM Z Volume 5: KVM*, SG24-8463.

6.3 Software licensing

The Z software portfolio includes operating system software (that is, z/OS, z/VM, z/VSE, and z/TPF) and middleware that runs on these operating systems. The portfolio also includes middleware for Linux on IBM Z environments. For the z15, 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 the 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, and z/TPF operating systems. Charges are based on processor capacity, which is measured in millions of service units (MSU) per hour.

- IBM Z Tailored Fit pricing for software[4]

  Tailored Fit Pricing is a new, flexible software pricing model that dramatically simplifies the existing 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 can also 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 Z Tailored Fit Pricing for software

[4] Tailored Fit Pricing for IBM Z Hardware Consumption Solution is also available for z15 systems - see the following web page: https://www.ibm.com/it-infrastructure/z/pricing

Subcapacity pricing terms for z/VM and select z/VM-based programs

Subcapacity pricing for the z/VM V6 operating environment is available when you are running z/VM, Version 6, Release 3, or higher. Software pricing at less than full machine capacity can provide more flexibility and improved cost of computing as you manage the volatility and growth of new workloads.

For more information about subcapacity pricing terms for z/VM and z/VM-based programs, see announcement letter 217-267, dated July 17, 2017.

For more information about software licensing options that are available for z15, see *IBM z15 (8561) Technical Guide*, SG24-8851and *IBM z15 (8562) Technical Guide*, SG24-8852.

IBM Secure Service Container

The IBM Secure Service Container was expanded to integrate with the IBM Cloud Private platform for hybrid and private cloud deployments on IBM Z. You can deploy containerized IBM Middleware applications and use common management tools for deploying homegrown or other third-party Docker and Kubernetes-based applications.

Secure Service Container supports the deployment of software container technology without requiring application changes. The Secure Service Container is an appliance that contains hardware, software, and middleware. It provides a trusted execution environment for the most sensitive applications by using end-to-end encrypting the entire stack for both data-at-flight and data-at-rest, memory, and network.

Other applications can access Secure Service Container services transparently by way of REST APIs. Secure Service Container protects the applications that run inside them from various cyber attack vectors from internal and external threats.

6.4 References

For current planning information, see the following operating system web pages:

- z/OS
- z/VM
- z/VSE
- z/TPF
- Linux on IBM Z

Redbooks

IBM z15 Technical Introduction

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8850-02

ISBN 0738459178

Printed in U.S.A.