

IBM z16 Hardware Overview

—
Lauren Maietti
z/VM CP Developer
lmaiett@us.ibm.com



Contents – Will be updated at the end of presentation creation/review

IBM z16 Highlights	03	Software Support	21
		All Software Support	22
		z/VM Support	23
Hardware Update Overview	09	Summary	25
Models	10	Summary	26
Frame Changes	11	References	27
CPC Drawer Component Changes	12		
I/O Drawer Changes	16		
Capacity & Performance	17		
Capacity Level Changes	18		
Capacity on Demand Enhancements	19		
Performance Information	20		

Overview

IBM z16 Highlights

IBM z16 Built to Build



Accelerated AI

On-chip Integrated Accelerator for AI for **high-speed, real-time** inferencing at scale.

High performance and consistent low-latency inferencing for processing a mix of transactional and **AI workloads** at speed and scale.

Adds more than **6 TFLOPS** of processing power shared by all cores on the chip. **200 TFLOPS** for a 32-chip system.

Cyber Resiliency

Resiliency with flexible capacity to dynamically shift system resources across locations to **avoid disruptions**.

Quantum computing creates threat to today's public key cryptography.

Today's data is in fact at risk for future exposure through **"harvest now, decrypt later"** attacks.

Industry-first **quantum-safe system**

Machine Enhancements

IBM Telum Processor chip using **7nm** technology, has **8** cores each running at **5.2 GHz**.

Per-core performance improvement **11%***.

Maximum configuration capacity increase of **17%** per CEC*.

Redesigned cache structure with **4x** the L2 and **1.5x** the L3 and L4 cache per core*.

A **fully-configured 40TB** system with 4 CPC drawers has up to **25%** increased memory capacity per drawer*.

Modernized Hybrid Cloud

AI and security along with resiliency and capacity enhancements are designed for **mission-critical** workloads in a hybrid cloud environment.

Foundation for **application modernization** and hybrid cloud velocity.

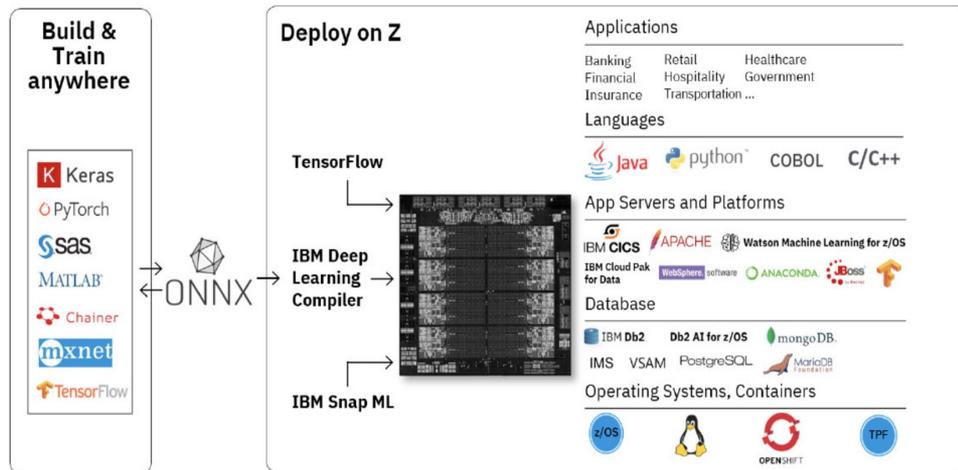
IBM z16 and IBM Z cloud software delivers a broad set of **open and industry-standard tools** (agile DevOps methodology to accelerate modernization)

*compared to the IBM z15

On-Chip AI Acceleration

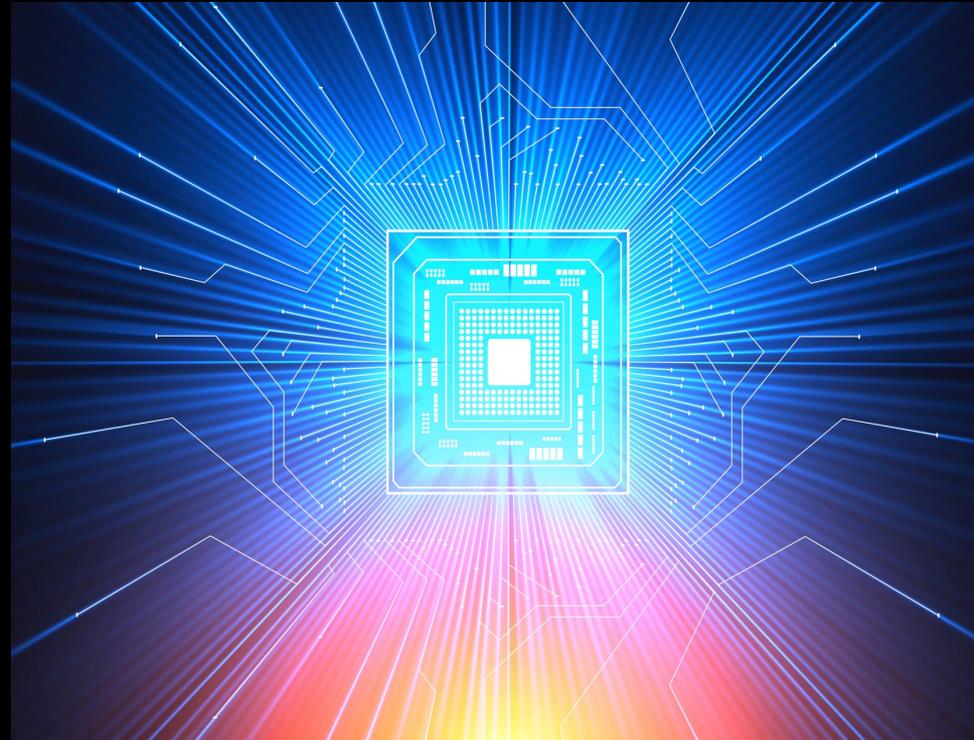
- On-chip Artificial Intelligence Unit (**AIU**) shared by the PU cores
- High-speed, large-scale inferencing for **real-time AI** embedded in transactional and AI workloads. Valuable insights instantly.
 - Ex. fraud prevention and fraud detection, customer behavior predictions and supply chain operations.
- New Instruction: Neural Networks Processing Assist (**NNPA**)
 - Non-privileged CISC memory-to-memory instruction
 - Operates on application memory tensors
- “Train anywhere, Deploy on IBM Z”

Note: Conversion tools and instructions must be used to convert tensor elements to IBM’s Deep Learning Format (DLF)



Quantum-safe Technology

- Quantum-safe capabilities:
 - Key generation
 - Key encapsulation mechanisms
 - Hybrid key exchange schemes
 - End-to-end encryption
 - **APIs** to quantum-safe algorithms in the Crypto Express8S
 - **Secure boot** technology with dual digital signature schemes
- Tools to **discover** where and what crypto used in apps: Compliance-ready CPACF Counters, IBM Application Discovery and Delivery Intelligence (ADDI), Integrated Cryptographic Service Facility (ICSF), and IBM Crypto Analytics Monitor (CAT),



IBM z16 vs. IBM z15

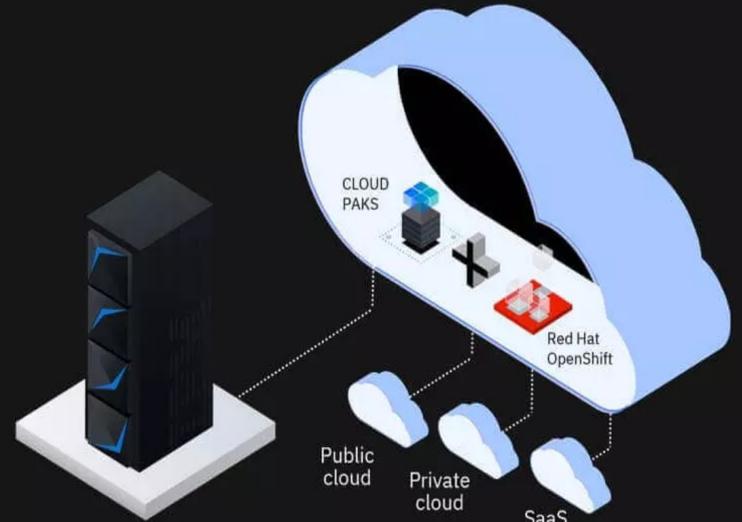
Source: If applicable, describe source origin

	IBM z15 T01	IBM z16 A01
System Capacity and subcapacity settings for CPs	<p>Up to 190 characterizable cores</p> <p>Up to 34 subcapacity settings for CPs</p> <p>Up to 292 total capacity levels</p>	<p>Up to 200 characterizable cores</p> <p>Up to 39 subcapacity settings for CPs</p> <p>Up to 317 total capacity levels</p>
Processor Chip	<p>Multi-core, single-chip modules (SCMs)</p> <p>Up to 12 cores per chip (per SCM)</p> <p>5.2 GHz (14 nm FITFET Silicon-On-Insulator [SOI])</p>	<p>Multi-core, Dual-chip modules (DCMs) with IBM Telum Processor</p> <p>5.2 GHz (using 7nm technology)</p> <p>Up to 16 cores per DCM</p>
Memory Per System	<p>Up to 40 TB of addressable real memory per 5-drawer system</p> <p>8 TB per CPC-drawer</p>	<p>Up to 40 TB of addressable real memory per 4-drawer system</p> <p>10TB per CPC drawer</p>
Hardware System Area (HSA)	256 GB	256GB
Cache	<p>First-level cache (L1 private): 128 KB for instructions, 128 KB for data</p> <p>Second-level cache (L2 private): 4 MB for instructions, 4 MB for data</p> <p>Third-level cache (L3 shared): 256 MB</p> <p>Fourth-level cache (L4 shared): 960 MB</p>	<p>First-level cache (L1 private): 128 KB for instructions, 128KB for data</p> <p>Second-level cache (L2 semi-private): 32 MB</p> <p>Third-level cache (L3 shared-victim): 224 MB</p> <p>Fourth-level cache (L4 shared-victim): 1.7 GB</p>
# of LCSSs and LPARs	<ul style="list-style-type: none"> - 6 Logical Channel Subsystems (LCSS) - 85 LPARs – up to 4TB memory (z/OS and z/VM) - 4 subchannel sets - 64,000 I/O devices per subchannel set 	<ul style="list-style-type: none"> - 6 LCSS - 85 LPARs - up to 4TB memory (z/OS and z/VM) - 4 subchannel sets - 64,000 I/O devices per subchannel set

Hybrid Cloud

Modernize and Integrate applications and data in hybrid cloud

- AI and **Quantum-Safe** Technology
- **Flexible** deployment options to innovate with speed and agility in constantly evolving era
 - Tailor Fit Pricing for IBM Z
 - Flexible Capacity for Cyber Resiliency
 - Unprecedented capacity to meet consolidation needs with innovative I/O features



IBM z16

Hardware Update Overview

Models

IBM z16 A01 processor unit configurations

Feature (Code)	Drawers – Cores	Chips – CPs	Standard SAPs	Standard Spares	IFPs
Max39 (0667)	1 drawer 48 cores	8 0–39	5	2	2
Max82 (0668)	2 drawers 96 cores	16 0–82	10	2	2
Max125 (0669)	3 drawers 144 cores	24 0–125	15	2	2
Max168 (0670)	4 drawers 192 cores	32 0–168	20	2	2
Max200 (0671)	4 drawers 228 cores	32 0–200	24	2	2

IBM z16 – Machine type **3931**

Model number: **A01**

– Configuration:

- **Up to 4** standard "19-inch" frames
- **Up to 4** Central Processor Complex (CPC) drawers built using the **IBM Telum** processor chip design
 - **Up to 200** configurable cores (depending on the configuration)
 - Spare PUs, System Assist Processors (SAPs), and Integrated Firmware Processors (IFPs) are included in the IBM z16 configuration
- **Up to 12** PCIe+ I/O Drawers and 192 adapters with iPDU power (10 drawers/160 adapters with BPA)

Frames

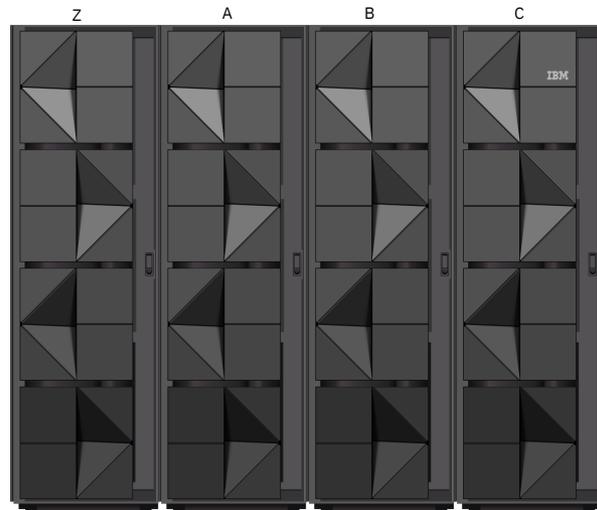
- Up to 4 "19-inch" frames
- Consistent with modern data center layouts.
- Frame Fulfills the requirements for **ASHRAE A3** class environment.

	z15 T01	z16 A01
# of Frames	1-4	1-4
# of CPC Drawers	1-5	1-4
# of I/O Drawers	0-12*	0-12*
Connections	Rear-only	Rear-only
Power Supply	PDU or BPA	PDU or BPA
Cooling	Radiator or water-cooling	Radiator only**

*Maximums vary per power supply
 **Cores cooled with internal water loop

Fully configured PDU-based IBM z16 A01 : 16 drawers (I/O + CPC) and 2 radiator units

Front:



Rear:



CPC DRAWERS

Up to **4 CPC drawers** can be configured (3 in the A Frame and 1 in the B Frame)

All CPC drawers are interconnected with high-speed communications links through the PU chips

Dual-Chip Modules (DCMs):

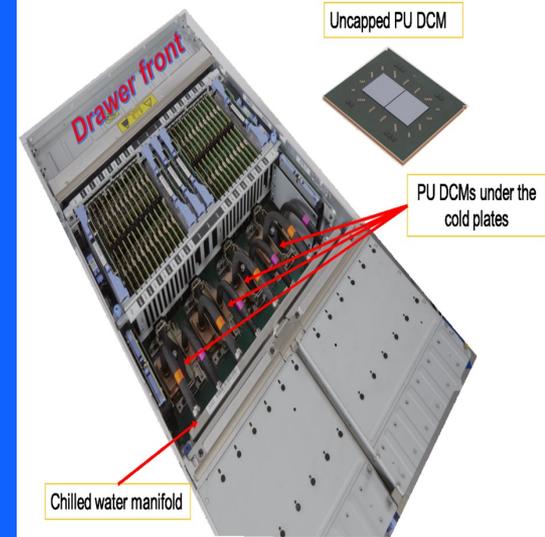
- **4 DCMs** containing **8** processor chips total
 - Up to **8 physical cores** per chip
- Cooled with internal water loop & cold plate

Memory per system
(w/o HSA)

- Minimum of 512 GB and maximum of **40TB per system**

Up to 48 configurable dual-inline memory module (DIMMs) for main memory

Redesigned cache hierarchy



2 dual-function Base Management Cards (BMCs)\Oscillator Cards (OSCs) – redundant interfaces to the internal management network and clock synch to the IBM Z platform.

6 SMP9 connectors for CPC drawer to CPC drawer communication

5 fans at the front of the drawer to cool resources

- PU DCMs are connected to cold plates and internally water-cooled

Fanouts

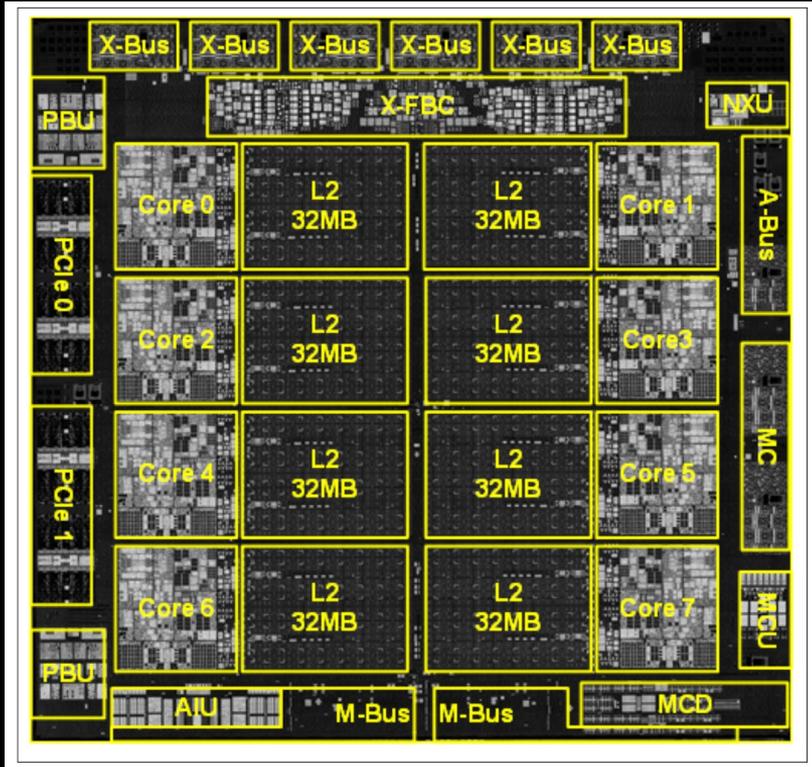
- **Up to 12** PCIe+ fanout adapters to connect to PCIe+ I/O drawers and coupling links.

3 or 4 Power Supply Units (PSUs), depending on the configuration (PDU or BPA)

- Loss of one PSU leaves enough power to satisfy the power requirements of the entire drawer.

CPC Drawer: Processor Chip

Processor Unit Chip



Highlighted Changes:

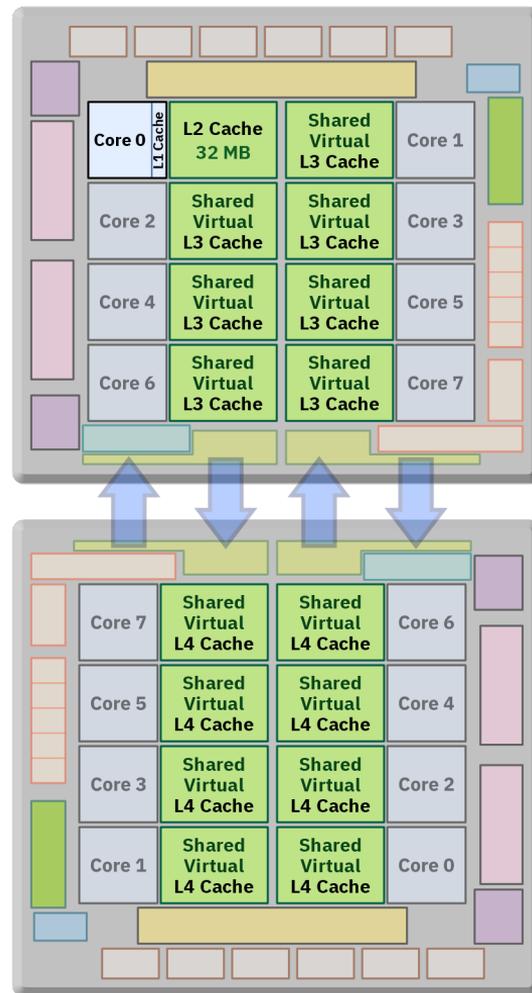
- 7nanometer FinFET technology
 - 18.8 miles of wire
 - 22.5B transistors (vs 9.2B on IBM z15)
- **8** physical cores per chip
 - Max200 is designed to offer **17%** more capacity per box and **40%** more capacity per drawer compared to the IBM z15.
- **Processor Activity Instrumentation** provides counters for cryptography operations executed by the on-core cryptography coprocessor – Central Processor Assist for Cryptographic Function (CPACF).
- Integrated Accelerator for **Artificial Intelligence (AIU)** is implemented on each processor chip and shared among all cores.

CPC Drawer: Cache Structure

Cache structure:

- On-core SRAM L1 Private Cache
 - Divided into **128 KB** cache for instructions and **128 KB** cache for data
- On-core/chip semi-private SRAM L2 cache, implemented as **32 MB** near the core
 - Each core has private 32MB L2 with up to 16 MB of unused cache used towards virtual cache.
- **Up to 224 MB** shared-victim virtual L3 cache comprised of L2s on same chip
- **Up to 1.7 GB** shared-victim virtual L4, consisting of the 'remote' virtual L3 caches of the drawer

4x L2 cache and **1.5x** more L3/L4 cache per core over the IBM z15.



CPC Drawer: Memory

- The amount of memory is directly related to the number of CPC drawers in the system.
- Up to **10 TB** of memory per CPC drawer
- Up to **40 TB** for a 4-CPC drawer system
 - IBM Z platform includes more installed memory than was ordered – part for RAIM design. Included in the base price.
- Offers up to **25% increase** memory capacity per single drawer drawer compared to IBM z15 T01.
- Hardware System Area (HSA) has fixed amount of memory **256GB** managed separately from available memory. No change.
- On IBM z16 platforms, the granularity for memory is in 64, 128, 256, 512, 1024, and 2048 GB increments.

IBM z16 Model A01 memory per feature

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

I/O Drawers

- Maximum IBM z16 configuration can support up to 12 PCIe+ I/O drawers.
- Up to 16 I/O slots over 2 domains for special purpose features such as storage, network, clustering, and cryptography.

New features:

- FICON Express32S – storage connectivity
- OSA-Express7S 1.2 – network connectivity
- RoCE Express3 (Long Reach and Short Reach) – network connectivity
- Coupling Express2 Long Reach – clustering connectivity
- Crypto Express8S

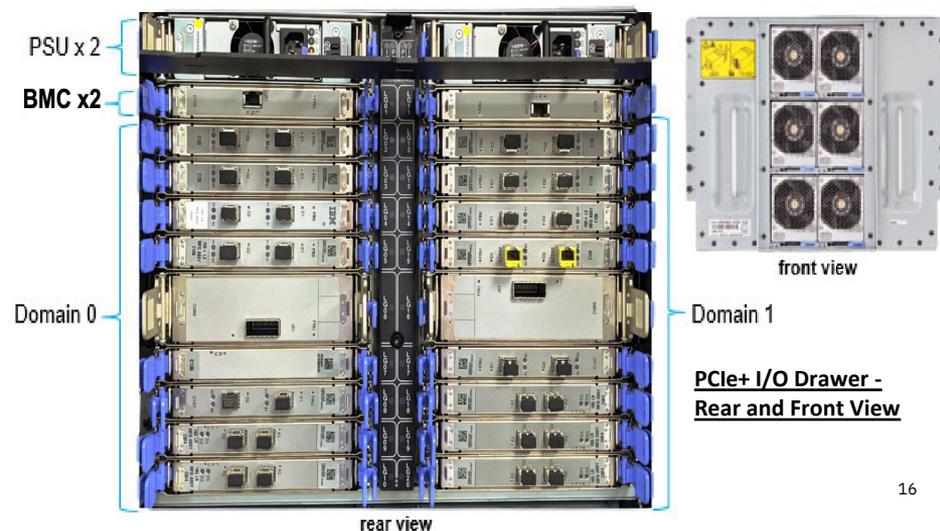
The 2 I/O domains per drawer each contain up to 8 I/O features that support the following types:

- FICON Express32S, FICON Express16SA, or FICON Express16S+
- OSA-Express7S 1.2, OSA-Express7S, or OSA-Express6S
- Crypto-Express8S, Crypto-Express7S, or Crypto-Express6S
- RDMA over Converged Ethernet (RoCE) Express3, RoCE Express2.1, or RoCE Express2
- zHyperLink Express 1.1 and zHyperLink Express
- Coupling Express2 LR

Features no longer available nor carried forward:

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

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.



IBM z16

Capacity and Performance

IBM z16 Capacity Levels

– The IBM z16 Model A01 Feature Max200 is designed to offer up to **17%** more processor capacity in the CEC compared to an IBM z15 Model T01 system

– **4** distinct capacity levels for the first 39 cores:

- 1 full-capacity & 3 sub capacities

– Any cores over 39 all must be full capacity

Granular capacity adds **117** subcapacity settings

200 capacity settings that are available with full-capacity CPs

The **317** distinct capacity settings in the system provide for a range of **1:758** in processing power.

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

Model 7xx = 100% (2253 PCI)
Model 6xx = 66% (1498 PCI)
Model 5xx = 41% (937 PCI)
Model 4xx = 12% (280 PCI)

Compare to z15:

Model 7xx = 100% (2055 PCI)
Model 6xx = 56% (1151 PCI)
Model 5xx = 38% (781 PCI)
Model 4xx = 13% (267 PCI)

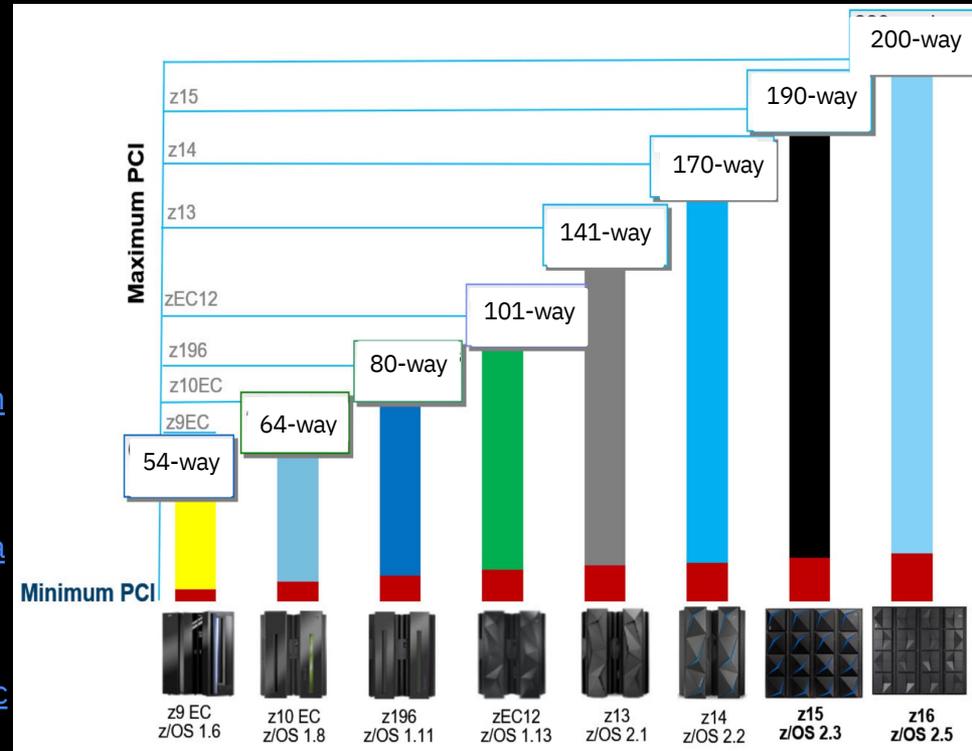
IBM z16 Performance

Performance Highlights (results may vary):

- 11% per-core performance improvement over IBM z15 T01 for typical workload on a typical client configuration

References:

- IBM Presentation and Tools - zPCR:
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS1381>
- Large Systems Performance Reference (LSPR) Report:
<https://www.ibm.com/servers/resourceink/lib03060.nsf/pages/lspindex?OpenDocument>
- Millions of Service Units (MSU) Ratings
<http://www.ibm.com/systems/z/resources/swprice/reference/exhibits/hardware.html>



IBM z16

Software Support

Software Support

IBM and ISV software solution support

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 z/OS Connect Enterprise Edition

The following operating systems are supported on the IBM z16:

z/OS Version 2 Release 5 with program temporary fixes (PTFs)

z/OS Version 2 Release 4 with PTFs

z/OS Version 2 Release 3 with PTFs

z/OS Version 2 Release 2 with PTFs (toleration support only)

z/VM Version 7 Release 3

z/VM Version 7 Release 2 with PTFs

z/VM Version 7 Release 1 with PTFs

12 IBM z16 Technical Introduction

z/VSE Version 6 Release 2 with PTFs

z/TPF Version 1 Release 1 (compatibility support)

IBM plans to support 21st Century Software VSEn V6.3 on IBM z16. For more information,

see this web page <https://www.21stcenturysoftware.com/vse/>

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

SUSE SLES 15 SP3 and SUSE SLES 12 SP5

Red Hat RHEL 8.4 and Red Hat RHEL 7.9

Ubuntu 22.04 LTS and Ubuntu 20.04.1 LTS

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

z/VM Support

PTFs for z/VM 7.1 and 7.2 and in the z/VM 7.3 base

Compatibility support for guest use:

- Embedded Artificial Intelligence Acceleration
- Compliance-ready CPACF Counters
- Breaking-Event-Address-Register Enhancement Facility
- Vector Packed Decimal Enhancements 2
- Reset DAT Protection Facility
- RoCE Express3 Feature
- CEX8S Feature/APIs
- CPU/Core Topology location information within Monitor
- Consolidated Boot Loader for guest IPL from SCSI

Transparent “Support”:

- Flexible Capacity for Cyber Resiliency
- Precision Time Protocol (PTP) Direct Attachment to CEC Enhancements
- N-mode Power STP Imminent Disruption Signal

z/VM PTFs/Prereqs/APARs

Support will be in the base of z/VM 7.3

z/VM 7.1 and 7.2 Pre-req for mixed-machine SSIs PTFs

– [VM66504](#) (available)

z/VM 7.1 and z/VM 7.2 Compatibility PTFs

– [VM66532](#) (available)

For more information on required service see:

<https://www.vm.ibm.com/service/vmreqz16.html>

For more information about IBM z16 migration, see the hardware PSP buckets for 3931DEVICE, and 3931DEVICE z/VM subset.



IBM z16

Summary

Summary

- Model: A01, Machine Type: 3931
- Frame changes:
 - 4 CPC drawers
- CPC Drawer changes:
 - 4 Dual chip modules containing 2 PUs
 - Updated cache structure
 - 10TB memory per drawer, 40TB for the system
- Processor Chip:
 - 7nm technology
 - 8 cores per chip
 - On-chip AI Accelerator
- Cache:
 - Physical private L1 128KB for L1D and L1I, physical semi-private L2 32 MB, virtual shared-victim L3 224MB, virtual shared-victim L4 1.7GB

References

- [IBM z16 Announcement Letter](#)
- [IBM z16 Technical Introduction](#)
- [IBM z16 Technical Guide](#)
- [IBM z15 Technical Guide](#)



Thank you

Lauren Maietti
z/VM CP Software Developer

—

© Copyright IBM Corporation 2020. All rights reserved. The information contained in these materials is provided for informational purposes only, and is provided AS IS without warranty of any kind, express or implied. Any statement of direction represents IBM's current intent, is subject to change or withdrawal, and represent only goals and objectives. IBM, the IBM logo, and ibm.com are trademarks of IBM Corp., 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 at [Copyright and trademark information](#).

