Jim Elliott – Mainframe Consultant
GlassHouse Systems Inc.

IBM z14 (and LinuxONE) Announcement Overview

GlassHouse Systems

2018-04-25

Trademarks and notices

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

Db2*	FICON*	HyperSwap*	IBM Z*	Power Systems	WebSphere*	zEnterprise*	z/OS*
DFSMSdss	Flash Systems	IBM*	LinuxONE	PR/SM	z13*	zHyperLink	z/VM*
DFSMShsm	GDPS*	IBM (logo)*	LinuxONE Emperor II	Storwize*	z13s		z/VSE*
ECKD	HiperSockets	ibm.com	LinuxONE Rockhopper II	System Storage*	z14		

* Registered trademarks of IBM Corporation
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Other product and service names might be trademarks of IBM or other companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

IBM z14 – Designed for trusted digital experiences

- **Pervasive encryption making data your new perimeter**
- **New flexible software pricing offering for modern digital workloads**
- **Machine learning applied directly to your most valuable data**
- **Open and connected with new economics in the cloud**
- **Blockchain on IBM Z where trust meets transparency**

IBM z14

Machine Type: 3906

Models: M01, M02, M03, M04, M05

Model	Customer PUs	Max Memory
M05	170	32 TiB
M04	141	32 TiB
M03	105	24 TiB
M02	69	16 TiB
M01	33	8 TiB

Extending the IBM z14 family

Built on the same IBM z14 technology	Addressing new markets	Standardization and Simplicity	One strong platform and family for the future

IBM z14 Model ZR1 – Built for digital trust, secure cloud

- **Platform simplification**
  - Standardization across many components – including Industry standard 19” rack
  - 16U free space in frame
- **Processor Units (PUs)**
  - 8, 16, 28, 34 PU cores per CPC drawer
  - Feature based sizing – 4, 6/12, 6/24, or 6/30 GCP/PU cores available for characterization
  - 2 SAPs and 1 IFP per system
  - 1 spares designated per system
  - 40 LPARs
- **Memory**
  - RAIM memory design – min of 64 GiB – max to 8 TiB
  - 64 GiB fixed HSA
  - IBM Virtual Flash Memory (replaces Flash Express)
- **I/O**
  - New PCIe Gen 3 IBM zHyperLink™ technology
  - 16 GiBps PCIe Gen 3 I/O interconnects
  - 3 Logical Channel Subsystems (LCSSs) with 3 sub-channel sets per LCSS

IBM z14 Model ZR1

Machine Type: 3907

Model	Machine Type	Customer PUs	Max Memory
M05	3906	170	32 TiB
M04	3906	141	32 TiB
M03	3906	105	24 TiB
M02	3906	69	16 TiB
M01	3906	33	8 TiB
ZR1	3907	4, 12, 24, 30	2 – 8 TiB

Largest z14 ZR1 is expected to provide up to 13% more total z/OS and up to 60% more total Linux® on Z capacity than the largest z13s

IBM Z naming for IBM z14

Brand Name:	**IBM**
Product Class:	**IBM mainframe**
Family Name:	**IBM Z®**
Family Short Name:	**Z**
Product Line Name:	**IBM Z®**
Product Line Short Name:	**Z**
Product Name:	**IBM z14™**
Short Name:	**z14™**
IBM z14 Models:	**ZR1, M01, M02, M03, M04, M05**
Machine Type and Model:	**3907 (z14 Model ZR1), 3906 (z14 Models M01, M02, M03, M04, M05)**

ITSO Redbooks for the IBM z14 model ZR1

- **Available now**
  - New – IBM z14 Model ZR1 Technical Introduction, SG24-8550
  - Updated – IBM Z Connectivity Handbook, SG24-5444-18
  - Updated – IBM Z Functional Matrix, REDP-5157-03
- **April 30, 2018 – Draft Versions**
  - New – IBM z14 Model ZR1 Technical Guide, SG24-8651
- **May 31, 2018 — Draft Versions**
  - IBM z14 Model ZR1 Configuration Setup, SG24-8560

Redbooks
ibm.com/redbooks

IBM Z family

	N-4	N-3	N-2	N-1	N
Dual Frame					
	z10 Enterprise Class (z10 EC) ▪ Ann. 2008-02-26 ▪ 4.4 GHz ▪ Up to 64 cfg cores ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 1.5 TiB	**zEnterprise 196** (z196) ▪ Ann. 2010-07-22 ▪ 5.2 GHz ▪ Up to 80 cfg cores ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 3 TiB	**zEnterprise EC12** (zEC12) ▪ Ann. 2012-08-28 ▪ 5.5 GHz ▪ Up to 101 cfg cores ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 3 TiB	**z13** ▪ Ann. 2015-01-14 ▪ 5.0 GHz ▪ Up to 141 cfg cores ▪ CP, IFL, ICF, zIIP ▪ Up to 10 TiB	**z14 M0x** ▪ Ann. 2017-07-17 ▪ 5.2 GHz ▪ Up to 170 cfg cores ▪ CP, IFL, ICF, zIIP ▪ Up to 32 TiB
Single Frame					
	z10 Business Class (z10 BC) ▪ Ann. 2008-10-21 ▪ 3.5 GHz ▪ Up to 10 cfg cores (5 CP) ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 256 GiB	**zEnterprise 114** (z114) ▪ Ann. 2011-07-12 ▪ 3.8 GHz ▪ Up to 10 cfg cores (5 CP) ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 256 GiB	**zEnterprise BC12** (zBC12) ▪ Ann. 2013/07/23 ▪ 4.2 GHz ▪ Up to 13 cfg cores (6 CP) ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 512 GiB	**z13s** ▪ Ann. 2016-02-16 ▪ 4.3 GHz ▪ Up to 20 cfg cores (6 CP) ▪ CP, IFL, ICF, zIIP ▪ Up to 4 TiB	**z14 ZR1** ▪ Ann. 2018-04-12 ▪ 4.5 GHz ▪ Up to 30 cfg cores (6 CP) ▪ CP, IFL, ICF, zIIP ▪ Up to 8 TiB

IBM Z – Processor roadmap

65 nm	45 nm	32 nm	22 nm	14 nm
z10	**z196**	**zEC12**	**z13**	**z14**
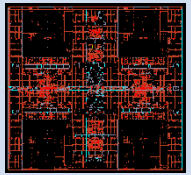	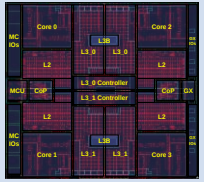	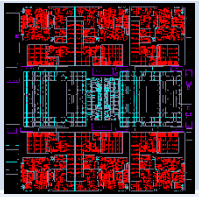	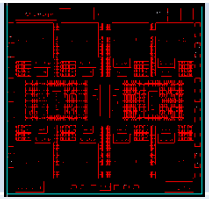	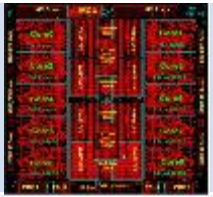
▪ Workload consolidation and integration engine for CPU intensive workloads ▪ Decimal FP ▪ InfiniBand ▪ 64-CP image ▪ Large pages ▪ Shared memory	▪ Top tier single thread performance, system capacity ▪ Accelerator integration ▪ Out of order execution ▪ Water cooling ▪ PCIe I/O fabric ▪ RAIM ▪ Enhanced energy management	▪ Leadership single thread, enhanced throughput ▪ Improved out-of-order ▪ Transactional memory ▪ Dynamic optimization ▪ 2 GiB page support ▪ Step function in system capacity	▪ Leadership system capacity and performance ▪ Modularity and scalability ▪ Dynamic SMT ▪ Supports two instruction threads ▪ SIMD ▪ PCIe attached accelerators ▪ Business analytics optimized	▪ Pervasive encryption ▪ Low latency I/O for acceleration of transaction processing for Db2 on z/OS ▪ Pause-less garbage collection for enterprise scale JAVA applications ▪ New SIMD instructions ▪ Optimized pipeline and enhanced SMT ▪ Virtual flash memory

z14 continues the CMOS mainframe heritage

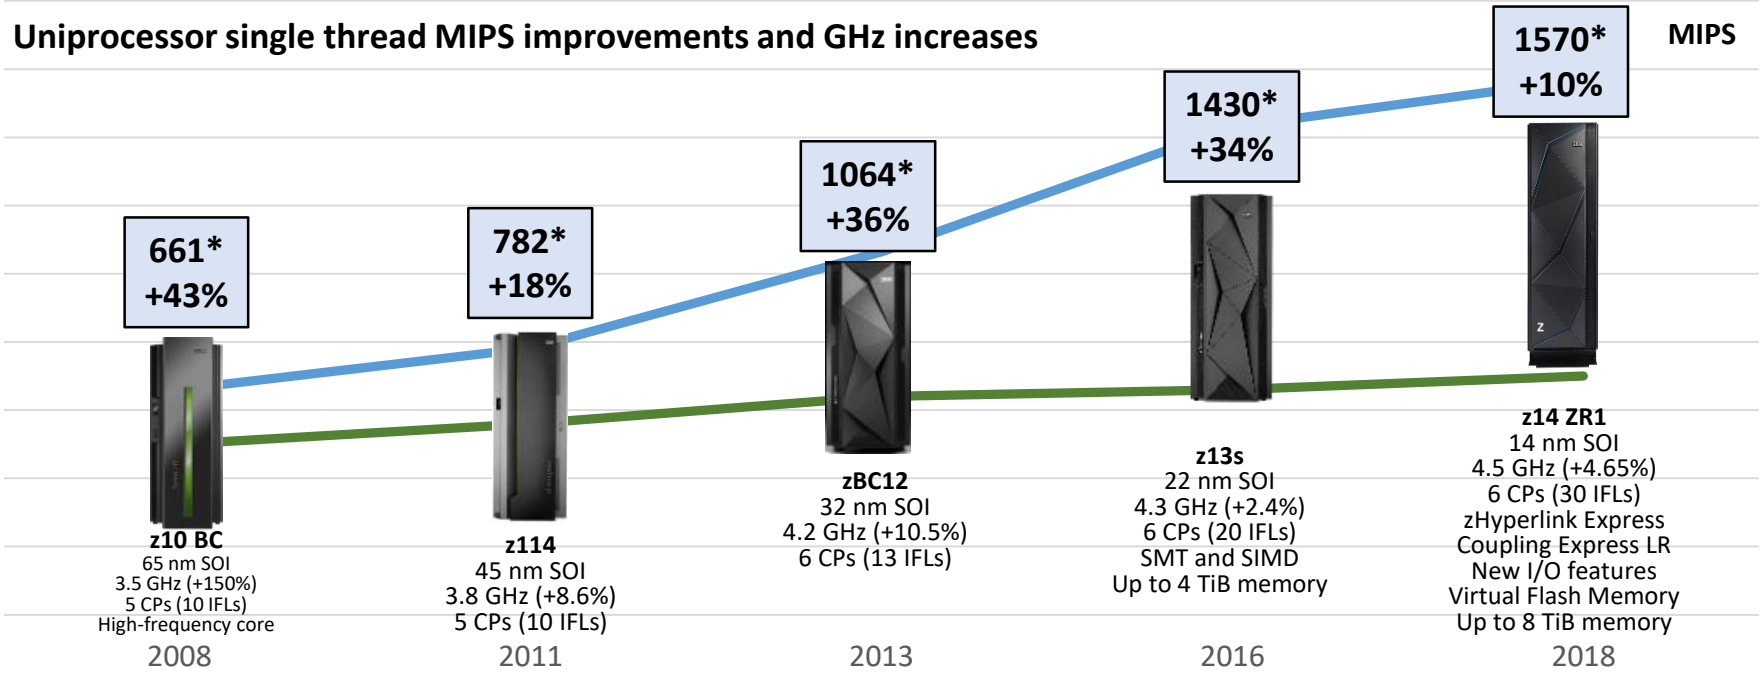

* Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload. MIPS tables are NOT adequate for making comparisons of IBM Z processors. Use IBM Capacity Planning Tools.

z14 single-frame sub-capacity CP granularity

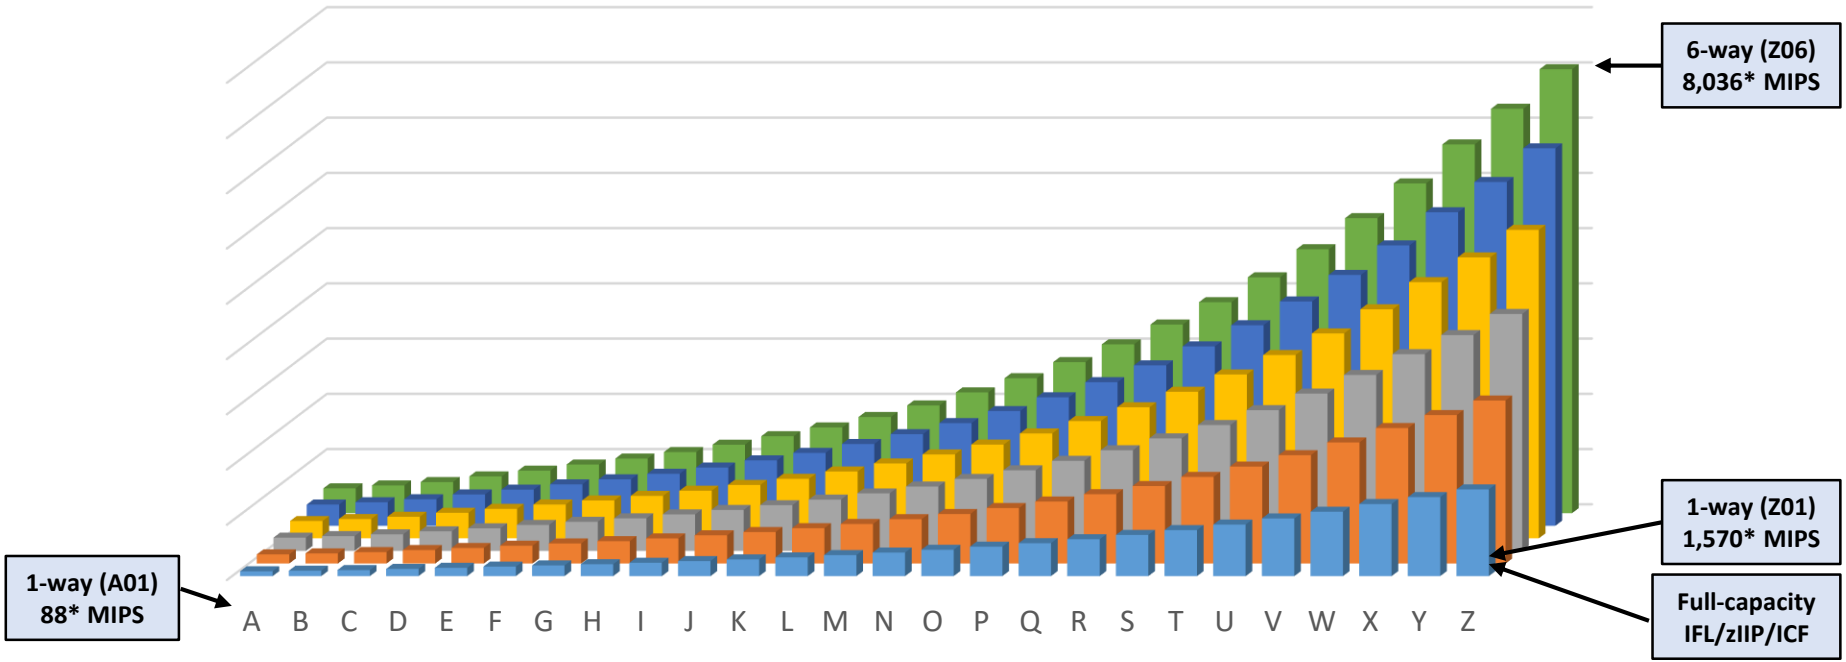

* Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload. MIPS tables are NOT adequate for making comparisons of IBM Z processors. Use IBM Capacity Planning Tools.

z14 ZR1 performance capacity highlights

- **Full speed capacity models (Z0x) capacity ratio to z13s**
  - Up to 54% more total system capacity is expected compared to the z13s (30-way z14 ZR1 vs. 20-way z13s)
  - Up to 1.10x (10%) average performance improvement at equal N-way vs. z13s
- **Full capacity model Z01**
  - Uni-processor capacity (full speed – Z01) 1570 PCI (9.7% increase over z13s)
- **Sub-capacity models**
  - 26 CP capacity levels: 156 sub-capacity settings (26 x 6 CPs)
- **SMT capacity**
  - IFLs and zIIPs can choose to run 2 hardware threads per core
    - Controlled by operating system parameter at the LPAR level
    - Added hardware threads appear as additional processors to the operating systems
    - Default is single thread
  - Likely wide range in capacity improvement per core over single thread: up to 25%
  - SAPs use SMT on z14 ZR1

z14 ZR1 potential sources of variability

- **Expect migrations to z14 ZR1 from z13s to be stable**
- **Workloads migrating to z14 from zBC12 and prior can expect to see slightly less variability than the typical z13s migration experience**
- **zBC12 marked the end of an era while z13s ushered in a new one**
  - Substantial frequency gains from generation to generation are no more across the industry
  - Greater reliance on performance driven by improved IPC in core and nest (e.g., “micro-architecture enhancements”) versus frequency gains
  - Workloads do not all react the same to these changes so there is more variability
  - Micro-benchmarks are particularly susceptible to this effect
- **Moving from MCM to single SCM chip topology also created more variability**
  - Greater reliance on PR/SM to do the right thing in LPAR placement
  - Enabling HiperDispatch is essential in this new era to maximize potential for local cache reuse
- **z14 ZR1 is an evolution of z13s**
  - Reduced variability by consolidating both nodes on a drawer
  - More cores per drawer and each core with higher capacity than z13s means more work can “fit” on a drawer
  - Added architecture and updated z/OS and z/VM to manage locks more intelligently

IBM z14 model ZR1 vs. M01

	z14 ZR1	z14 Model M01 *(new build)*
Uniprocessor Performance[1]	1570 PCI	1695 PCI
z/OS Capacity[1]	88 – 8036 MIPS	250 – 32671 MIPS
Total System Memory	8 TiB	8 TiB
Configurable Engines	30	33
Configurable CPs	0 – 6	0 – 33
LPARS/CSS	40/3	85/6
HiperSockets	32	32
PCIe+ I/O Drawers	0 – 4	0 – 3
I/O slots per PCIe+ I/O drawer	16	32
FICON Channels	128	192
OSA Ports (10GbE/1GbE/1000BASE-T)	32/64/64	48/96/96
IFB host bus Bandwidth	N/A	6.0 GiB/sec
PCIe Gen3 Bandwidth	16.0 GiB/sec	16.0 GiB/sec
PSIFB / ICA-SR (max. features)	0/8	4/10
zIIP Maximum Qty	0 – 12	0 – 22
IFL Maximum Qty[1]	30 @ 4.5 GHz	33 @ 5.2 GHz
ICF Maximum Qty	30	33
Capacity Settings	156	132
Upgradeable	Within z14 Model ZR1 (feature based)	z14 M02, M03, M04

1. For capacity and performance evaluation refer to the latest zPCR https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381
   For Linux, the performance may vary depending on workload characteristics and SMT efficiency for the specific workload

z14 model ZR1 under the covers

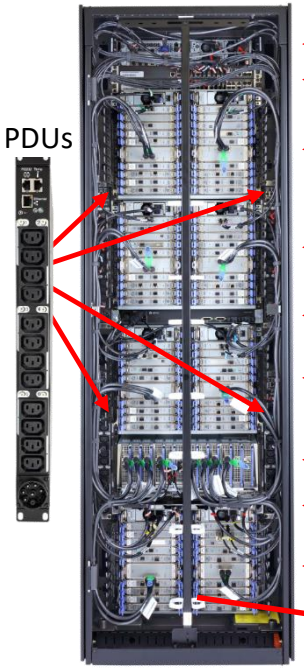

PDUs

Rear View

Support Ethernet Switches

PCIe+ I/O Drawer 4

PCIe+ I/O Drawer 3

KVM Switch

PCIe+ I/O Drawer 1

CPC Drawer

Cable guide (spine)

PCIe+ I/O Drawer 2

Spine detail (with routed I/O cables)

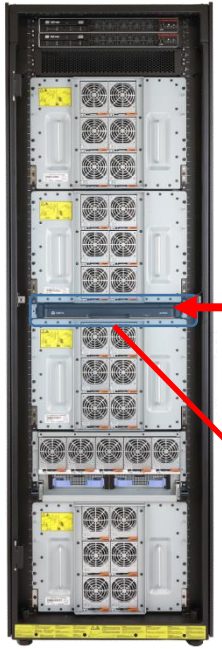

Front View (bezels removed)

Monitor + keyboard for SEs (tray in)

Monitor + keyboard for SEs (tray out)

Specialty processors expand the use of the server while lowering the cost of ownership

Integrated Facility for Linux (IFL) 	Integrated Information Processor[1] (zIIP)	Integrated Coupling Facility (ICF)
▪ Dedicated Linux core on IBM Z ▪ *IT optimization and cloud computing* delivering enhanced economics ▪ Supported by z/VM and KVM virtualization, tooling such as IBM Wave and third parties and the Linux distributions	▪ *Relieves* central processors of running specific workloads ▪ Optimized for strategic web based applications with support for *Java and XML* processing ▪ Focused on data and supporting workloads can help *connect, manage, extend, and protect* data	▪ CF allows multiple processors to access the same data ▪ *PCIe* based long range coupling links ▪ Coupling Facility (CF) traffic is *protected in-flight and at-rest* in the CF by policy driven host-based encryption

1. Supports 2:1 ratio for zIIP to GCP

IBM z14 Model ZR1 – Flexible new way to configure

Feature	Max Memory	Total PUs	Customer PUs	GCPs	IFLs	zIIPs	ICFs	Std. SAPs	Add'l SAPs	Spares	IFP[1]	PCIe I/O Features
Max4	2 TiB	8	4	0-4	0-4	0-2	0-4	2	0-2	1	1	0-16
Max12	4 TiB	16	12	0-6	0-12	0-8	0-12	2	0-2	1	1	0-32
Max24	8 TiB	28	24	0-6	0-24	0-12	0-24	2	0-2	1	1	0-64
Max30	8 TiB	34	30	0-6	0-30	0-12	0-30	2	0-2	1	1	0-64

- ***New feature based sizing*** **(4, 12, 24, 30) – done at configuration time**
- ***New entry level*** **– 88 MIPS for capacity setting A01**
- **Full Z01** ***uniprocessor running up to 1,570 MIPS*** **with maximum z/OS general purpose (GCP) capacity available over 8,000 MIPS**
- **Maximum Linux (IFL) capacity available over 29,000 MIPS**
- ***Same granularity for right sizing*** **– 26 capacity levels x 6 GCPs equals 156 settings**
- **Great economics for** ***standalone Coupling Facility*** **– PCIe Gen3 technology only (no InfiniBand® coupling links)**

IBM z14 Model ZR1
Machine Type: 3907

1. Integrated firmware processor (IFP) used for infrastructure management of PCIe adapters

z14 ZR1 processor drawer (top view)

- **Each PU SCM:**
  - 14nm
  - One Memory Controller per PU Chip
  - Five DDR4 DIMM slots per Memory Controller:
    10 total per logical cluster, 20 total per drawer
- **One CPC drawer / four features (sizes) offered (field upgradeable)**
  - Max4: 1 PU + 1 SC SCMs (8 PUs), 2 PCIe fanouts, up to 2 TiB memory
  - Max12: 2 PU + 1 SC SCMs (16 PUs), 4 PCIe fanouts, up to 4 TiB memory
  - Max24: 4 PU + 1 SC SCMs (28 PUs), 8 PCIe fanouts, up to 8 TiB memory
  - Max30: 4 PU + 1 SC SCMs (34 PUs), 8 PCIe fanouts, up to 8 TiB memory
  - Two logical PU clusters (0 and 1)
  - One SC Chip (672 MiB L4 cache)
  - Two or four PSUs (AC-DC) replace DCAs (DC-DC)
  - Two Oscillator Cards (OSCs)
  - Air cooled PU and SC SCMs
  - Two Flexible Support Processors (FSPs)
  - Up to eight (8) fanout slots for PCIe I/O drawer fanouts or PCIe coupling fanouts

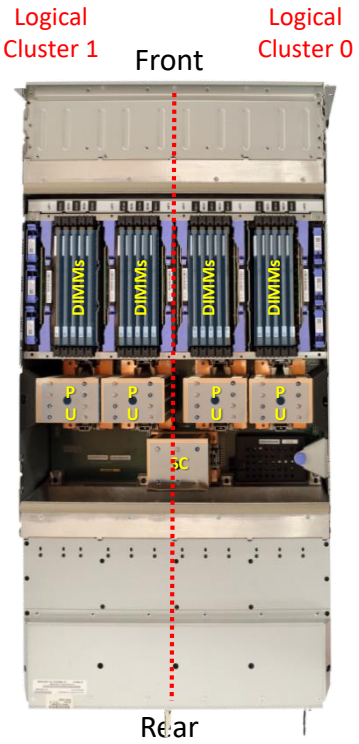

z14 ZR1 CPC drawer layout

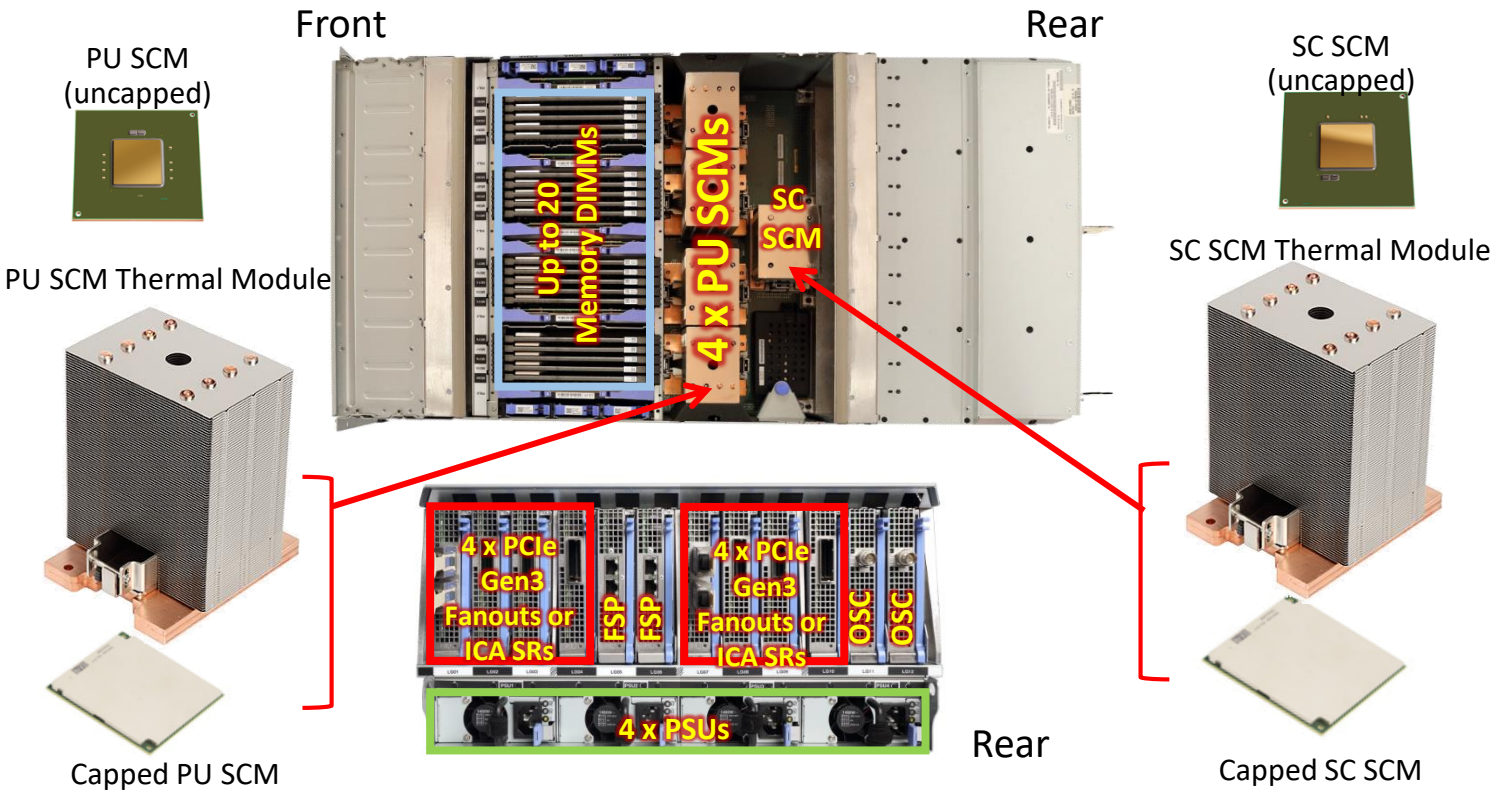

z14 ZR1 10-core processor chip

- **6.1 Billion transistors**
- **25.3 x 27.5 mm chip area**
- **14nm SOI technology,**
- **17 layers of metal**
- **10 cores per CP-chip (5, 6, 7, 8, or 9 active cores per node)**
- **4.5 GHz**

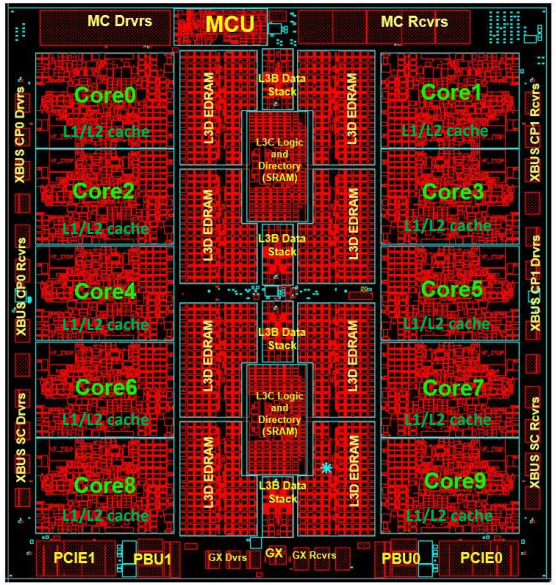

- **Cache Improvements:**
  - New power efficient logical directory design
  - 33% larger L1 Cache (128K)
  - 2x larger L2 Cache (4MB)
  - 2x larger L3 Cache with symbol ECC
- **New Translation/TLB2 design**
  - 4 concurrent translations
  - Reduced latency
  - Lookup integrated into L2 access pipe
  - 2x CRSTE growth
  - 1.5X PTE growth
  - New 64 entry 2gig TLB2
- **Pipeline Optimizations**
  - Improved instruction delivery
  - Faster branch wakeup
  - Reduced execution latency
  - Improved Operand Store Compare avoidance
  - Optimized 2nd generation SMT2
- **Improved Branch Prediction**
  - 33% Larger BTB1 & BTB2
  - New Perceptron Predictor
  - New Simple Call Return Stack

Note: Chip technology also applies to IBM LinuxONE™ Rockhopper™ II (3907 model LR1)

z14 ZR1 system control chip

- SC Chip area: 25.3 x 27.5 mm
- 9.7 Billion transistors
- 14nm SOI technology, 17 layers of metal
- 672 MB shared eDRAM L4 Cache
- System Interconnect
- System Coherency Manager
- X and A[1] bus

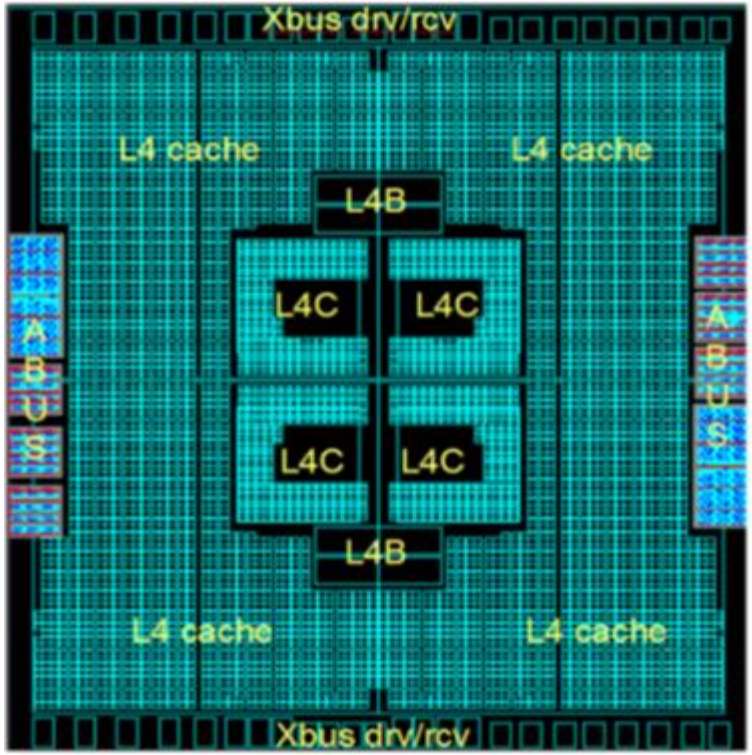

1. A bus not used for z14 ZR1

z14 ZR1 memory ranges and increments

- **Purchased Memory – Memory available for assignment to LPARs**
- **Hardware System Area – Standard 64 GiB of addressable memory for system use outside customer memory**
- **Standard Memory – Provides minimum physical memory required to hold customer purchase memory plus 64 GiB HSA**
- **Preplanned Memory – Provides additional physical memory needed for a concurrent upgrade (LIC CC change only) to a preplanned target customer memory**

Feature	Standard Memory GiB
Max4	64 - 1984
Max12	64 - 4032
Max24	64 - 8128
Max30	64 - 8128

Memory Increment (GiB)	Offered Memory Sizes (GiB)
8	64, 72, 80, 88, 96
32	128, 160, 192, 256, 288, 320, 352, 384
64	448, 512, 576
128	704, 832, 960
256	1216, 1472, 1728, 1984, 2240 … 4032
512	4544, 5056 … 8128

Differentiated value at the core

- **Leader in performance and scale**
  - 10-core processor design in 14nm silicon technology
  - Up to 6 configurable general purpose cores and up to 30 total cores allow the z14 Model ZR1 up to ***13%*** and ***60%*** more capacity for z/OS and Linux respectively than the largest z13s
  - ***1.5x*** more on-chip cache[1] per core optimized for data serving
- **More performance with innovation that helps the full stack**
  - New instructions in the SIMD Facility gives boost for traditional workloads using decimal arithmetic and new applications like analytics
  - Pause-less garbage collection enables enterprise scale Java® applications to run with fewer and shorter pauses for garbage collection on larger and larger heaps
  - Next generation SMT improves performance up to ***25%*** vs. non-SMT for an IFL or zIIP to benefit Linux and zIIP exploiters

1. All comparisons are to the z13s

z14 ZR1 I/O Infrastructure

I/O subsystem internal bus interconnect speeds

PCEe Gen3 z14, z13, z13s

PCIe Gen2 zEC12, zBC12, z196, z114

IFB zEC12, zBC12, z196, z114, z10 EC, z10 BC

eSTI z9 EC, z9 BC

STI z990, z890

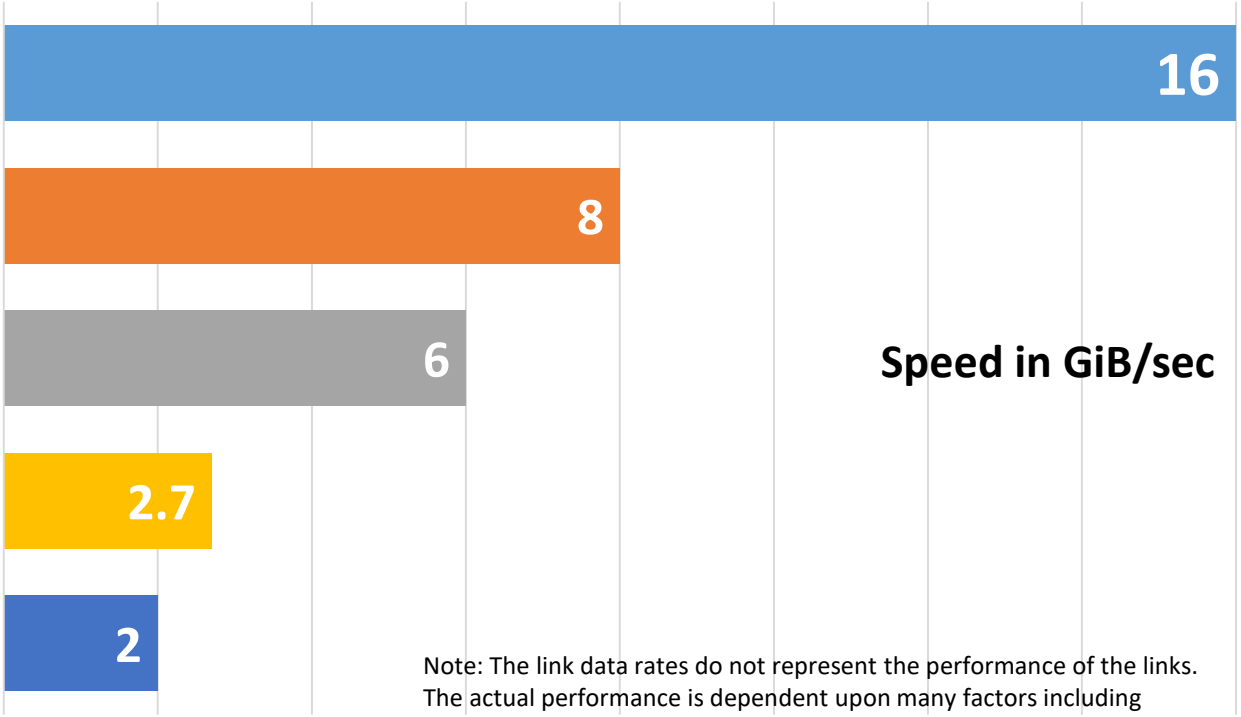

Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

z14 ZR1 “new build” I/O features supported

- **Features in the PCIe+ I/O drawer**
  - FICON Express16S+ LX, SX
  - OSA-Express6S: GbE (LX, SX), 10 GbE (LR, SR), and 1000BASE-T
  - 10GbE RoCE Express2
  - zEDC Express
  - Crypto Express6S
  - zHyperLink Express
  - Coupling Express LR
- **Features in the CPC drawer**
  - Integrated Coupling Adapter (ICA) SR

PCIe+ I/O drawer

16 I/O slots

CPC drawer

8 PCIe Fanout I/O slots

IBM zHyperLink Express

Speed Matters: Breakthrough I/O link technology

- **A direct connect short distance link designed to deliver low latency connectivity between z14 servers and FICON storage systems**
- **zHyperLink improves application response time, cutting I/O sensitive workload response time by up to *50%*[1]**
- **Typical use cases:**
  - Performance improvements are achieved seamlessly without need for application changes
  - Dramatic improvement in data access for OLTP workloads
  - Faster Db2® index splits helps reduce the batch processing window for heavy insert work
  - Better client experience with lower I/O latencies
  - Additional business opportunities for top line growth with more functional applications

1. This response time estimate is based on IBM internal measurements and projections that assume 75% or more of the workload response time is associated with read DASD I/O and the storage system random read cache hit ratio is above 80%. The actual performance that any user will experience may vary.

How does IBM zHyperLink change the game?

- **zHyperLink is fast enough the CPU can just wait for the data**
  - No un-dispatch of the running task
  - No CPU queueing delays to resume it
  - No host CPU cache disruption
  - Very small I/O service time
- **Operating system[1] and middleware[2] (Db2) are changed to keep running over an I/O**
- **Transparently gives Db2[3] apps fundamentally better latency than applications on platforms without zHyperLink**
  - Excluding 100% in memory databases

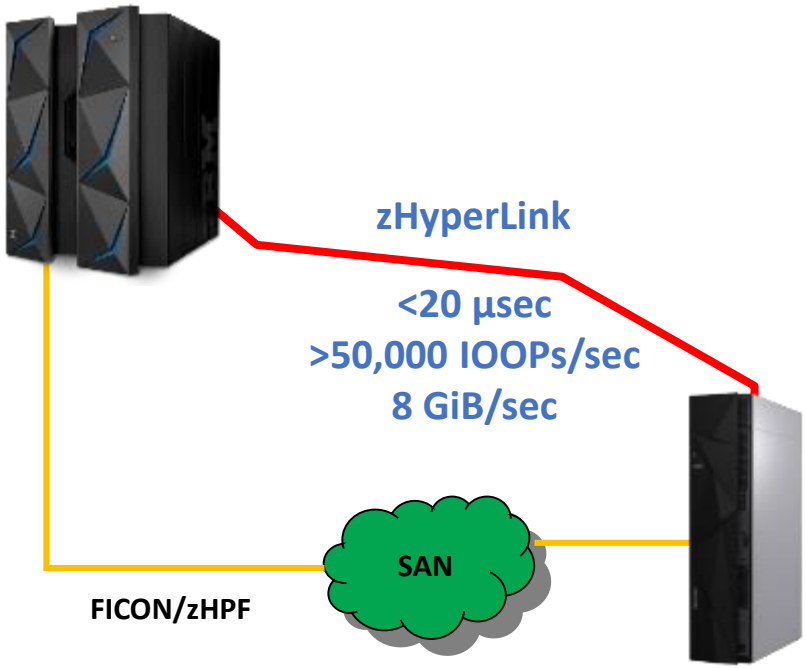

1. z/OS 2.1, 2.2, or 2.3
2. SOD: IBM intends to deliver VSAM exploitation of z14 and DS8880 zHyperLink Express.
3. Db2 for z/OS 12

zHyperLink Express at a glance

- **zHyperLink Express feature**
  - Two ports per feature
  - Maximum of 16 features (32 ports)
  - Function ID Type = HYL
  - Up to 127 Virtual Functions (VFs) per PCHID
  - Point to point connection using PCIe Gen3
  - Maximum distance: 150 meters
- **A standard FICON channel (CHPID type FC) is required for exploiting the zHyperLink Express feature**
- **DS8880 Gen 1, DS8880 Gen2, DS8880F**

I/O performance evolution to IBM zHyperLink with DS8886

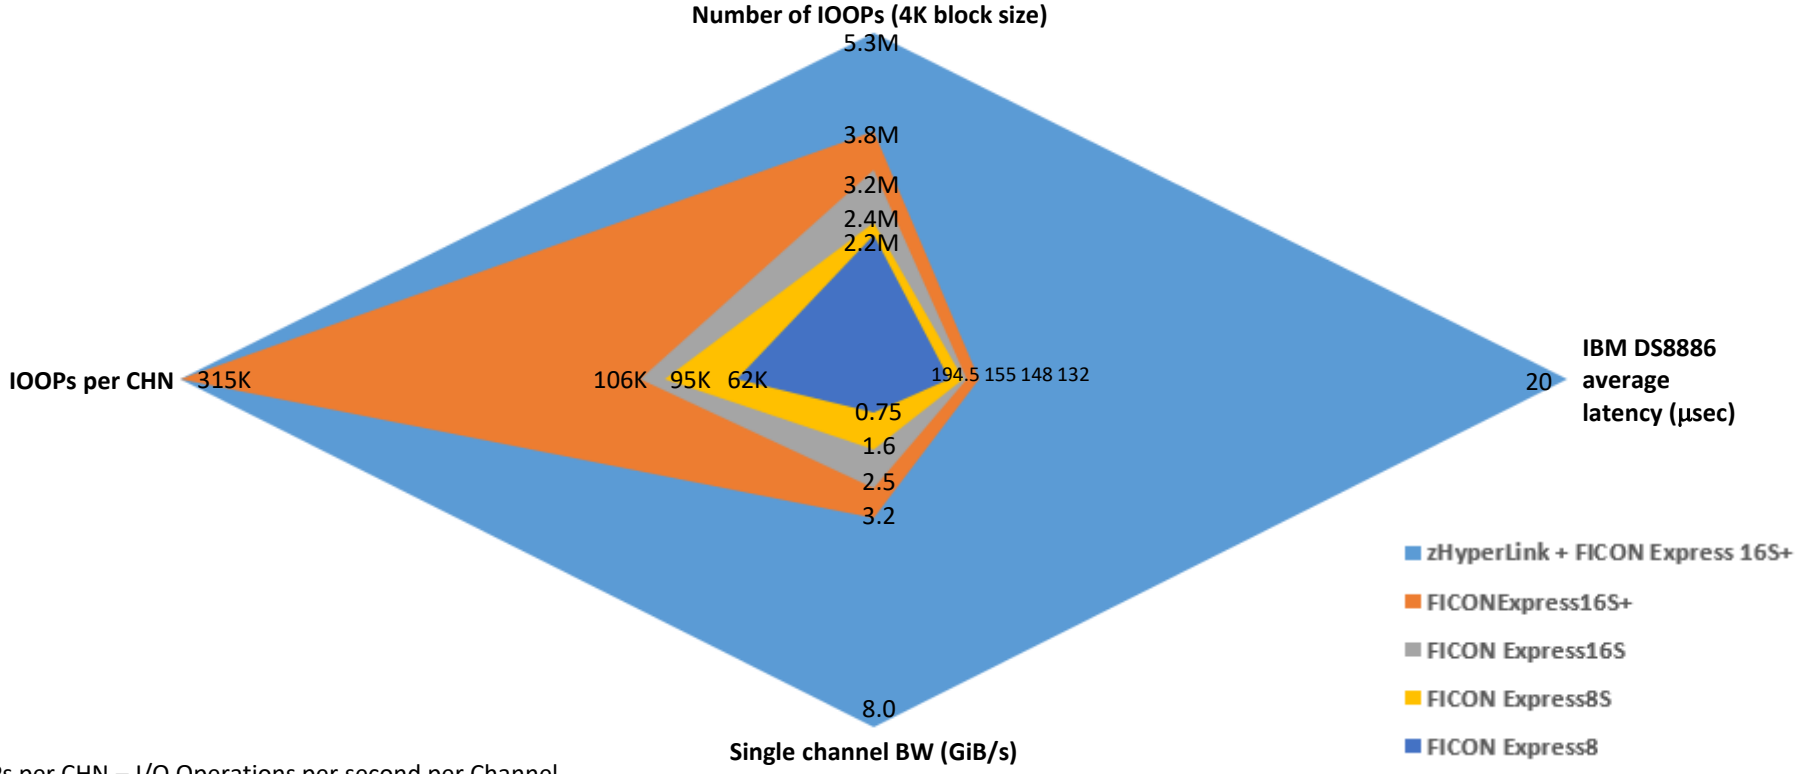

IOOPs per CHN – I/O Operations per second per Channel

Speed and scale getting to the data

- **FICON Express16S+**
  - Up to ***3x***[1] SIO/sec for small data transfer I/O operations and ***25%***[1] SIO/sec increase with mix of large sequential read and write data transfer options
  - Batch Elapsed time improves ***17%***[1] running I/O intensive batch workloads versus same workload using FICON Express16S on a z13s
  - Provides increased scalability by increasing number of devices per channel without degrading performance

1. Performance results are extrapolated from benchmark measured in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed

FICON Express16S+

- **For FICON, zHPF, and FCP**
  - CHPID types: FC and FCP
  - Both ports must be same CHPID type
    - 2 PCHIDs / CHPIDs
- **Auto-negotiates to 4, 8, or 16 Gbps**
  - 2 Gbps connectivity not supported
  - FICON Express8S will be available for 2Gbps (carry forward only)
- **Increased performance compared to FICON Express16S**
- **Small form factor pluggable (SFP) optics**
  - Concurrent repair/replace action for each SFP
  - 10KM LX – 9 micron single mode fiber
    - Unrepeated distance – 10 kilometers (6.2 miles)
  - SX – 50 or 62.5 micron multimode fiber
    - Distance variable with link data rate and fiber type
- **2 channels of LX or SX (no mix)**

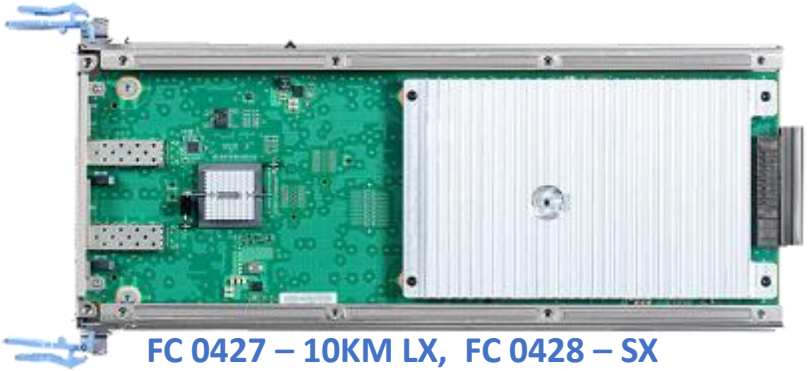

FC 0427 – 10KM LX, FC 0428 – SX

LX/LX or SX/SX

zHPF and FICON performance[1] on IBM Z

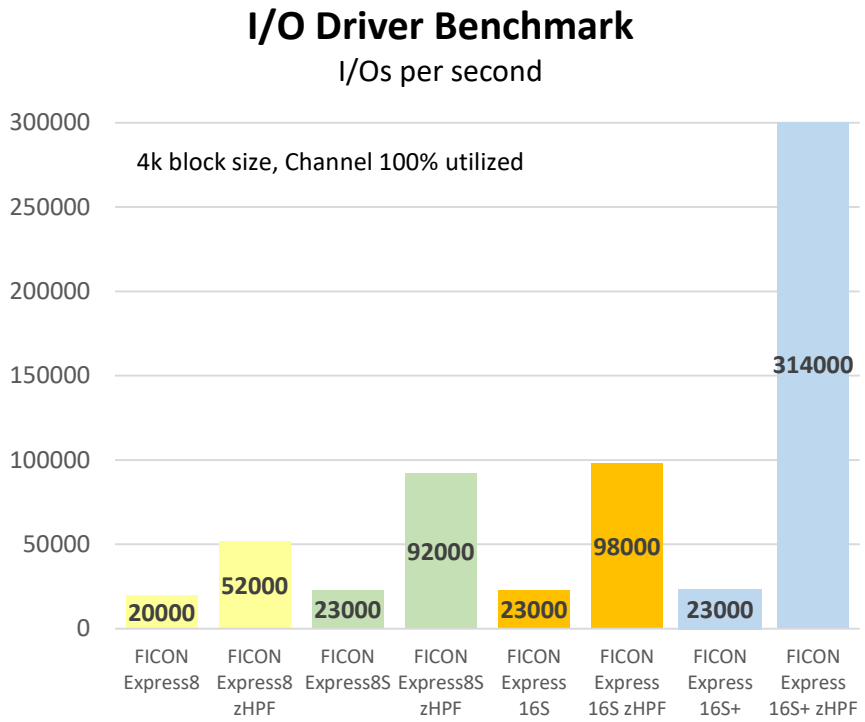

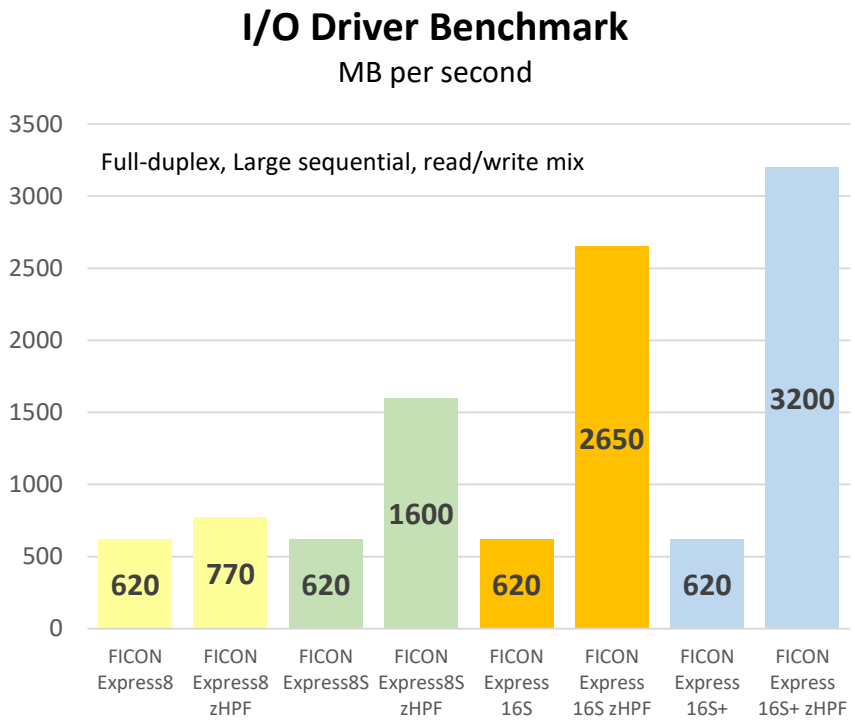

1. This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

FCP performance[1] on IBM Z

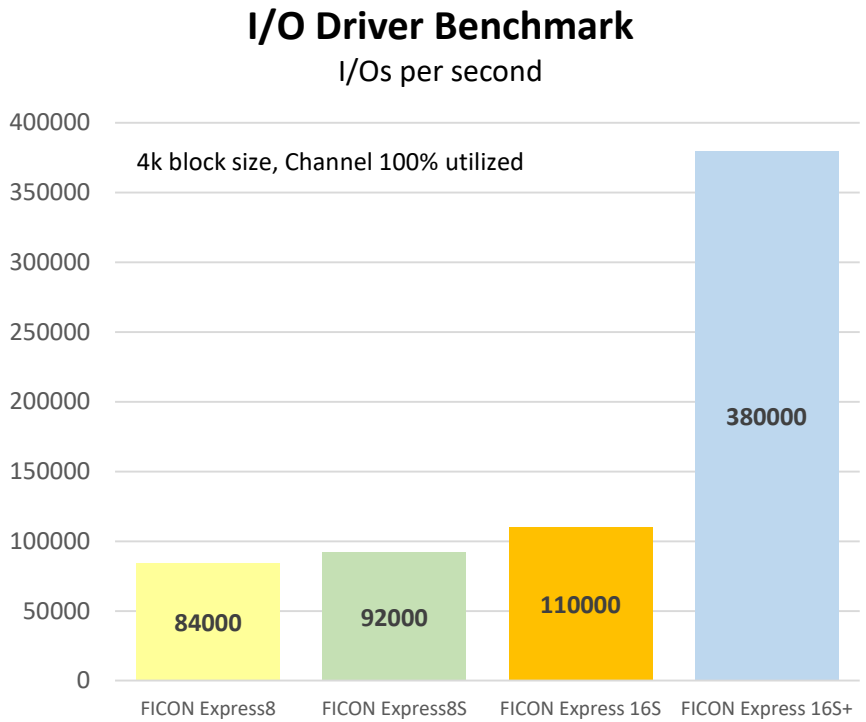

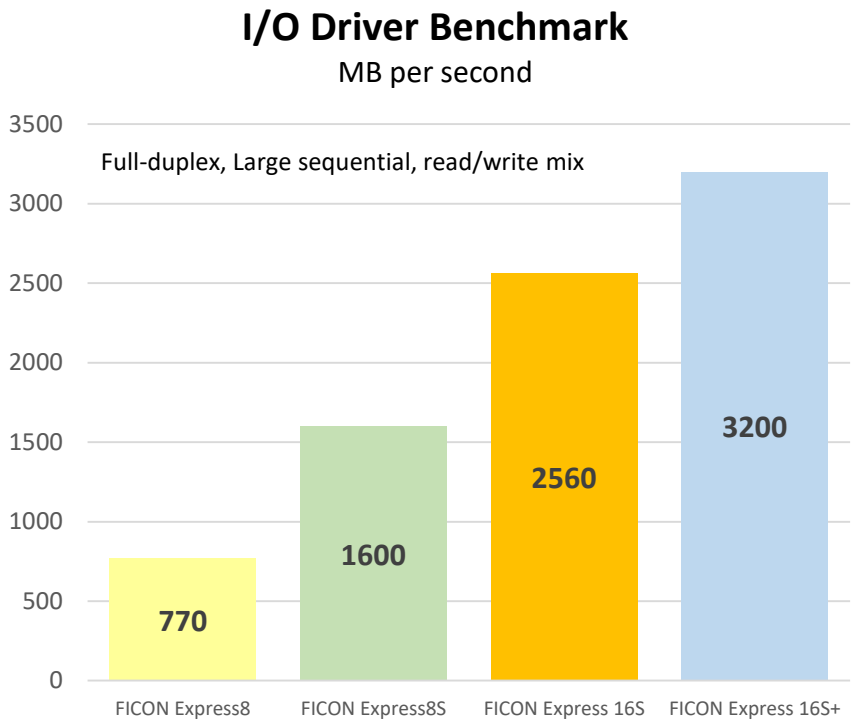

1. This performance was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

OSA-Express6S Fiber Optic Ethernet features

- **10 Gigabit Ethernet (10 GbE)**
  - CHPID types: OSD, OSX
  - Single mode (LR) or multimode (SR) fiber
  - One port of LR or one port of SR
    - 1 PCHID/CHPID
  - Small form factor pluggable (SFP+) optics
  - LC duplex
- **Gigabit Ethernet (1 GbE)**
  - CHPID type: OSD
  - Single mode (LX) or multimode (SX) fiber
  - Two ports of LX or two ports of SX
    - 1 PCHID/CHPID
  - Small form factor pluggable (SFP+) optics
    - Concurrent repair/replace action for each SFP
  - LC Duplex

FC 0424 – 10 GbE LR
FC 0425 – 10 GbE SR

FC 0422 – GbE LX
FC 0423 – GbE SX

OSA-Express6S 1000BASE-T Ethernet feature

- **PCIe form factor feature supported by PCIe I/O drawer**
  - One two-port CHPID per feature
  - Half the density of the OSA-Express3 version
- **Small form factor pluggable (SFP+) transceivers**
  - Concurrent repair/replace action for each SFP
- **Exclusively Supports: Auto-negotiation to 100[1] or 1000 Mbps and full duplex only on Category 5 or better copper**
  - No 10Mbps
  - RJ-45 connector
  - Operates at “line speed”
- **CHPID TYPE Support:**

Operation Mode	Type	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, OS system console operations
QDIO	OSD	TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
Non-QDIO	OSE	TCP/IP and/or SNA/APPN/HPR traffic

1. OSA-Express6S 1000BASE-T adapters (FC 0426) will be the last generation of OSA 1000BASE-T adapters to support connections operating at 100 Mb/second link speed. Future OSA-Express 1000BASE-T adapter generations will support operation only at 1000 Mb/second (1Gb/s) link speed.

Shared Memory Communications architecture

Memory-to-memory communications using high speed protocols and direct memory placement of data for faster communications

Shared Memory Communications Remote Direct Memory Access (SMC-R)	Shared Memory Communications Direct Memory Access (SMC-D)
▪ *New* RoCE Express2 enables SMC-R ▪ Helps to reduce both latency and CPU resource consumption ▪ Up to *50%* CPU savings for FTP file transfers across z/OS systems versus standard TCP/IP[1]	▪ Optimizes LPAR to LPAR inter-system operating system communications ▪ Valuable for communications within a single server without requiring extra hardware ▪ Up to *61%* CPU savings for FTP file transfers for z/OS versus HiperSockets[2]

Any z/OS TCP sockets-based workload can seamlessly use SMC without application changes SMC Applicability Tool (SMCAT) helps assess benefit for your environment

1. Based on internal IBM benchmarks in a controlled environment using z/OS V2R1 Communications Server FTP client and FTP server, transferring a 1.2GiB binary file using SMC-R (10GbE RoCE Express feature) vs. standard TCP/IP (10GbE OSA Express4 feature). The actual CPU savings any user will experience may vary.
2. All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided “AS IS” and no warranties or guarantees are expressed or implied by IBM.

10 GbE RoCE Express2

- **What is it?**
  - The next generation of RoCE technology with the IBM 10GbE RoCE Express2
  - The 10GbE RoCE Express2 provides a technology refresh for RoCE on IBM Z
- **Why?**
  - RoCE is an evolving technology thus it is critical to keep the IBM Z RoCE technology current within the industry
  - Technology currency provides many improvements in the base technology that will provide benefits for application workloads
- **Benefits**
  - Technology currency allows Z customers to benefit from the latest advancements in RoCE architecture, technology and specifications
  - The 10GbE RoCE Express2 provides two physical 10GbE ports (no change)
  - Key difference: RoCE Express2 provides increased virtualization (sharing) capabilities allowing RoCE to be extended to more workloads

Shared Memory Communications over RDMA (SMC-R) concepts

Clustered Systems

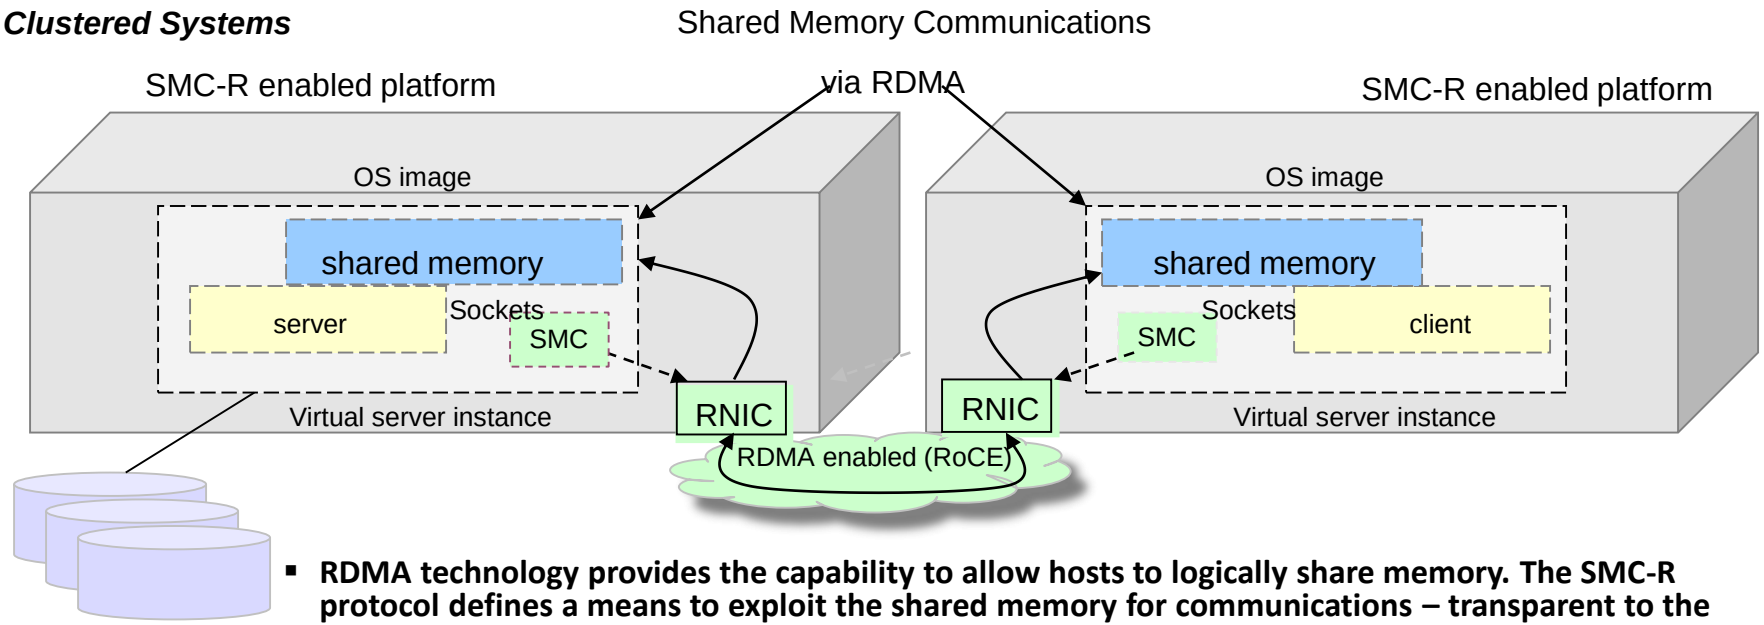

- **RDMA technology provides the capability to allow hosts to logically share memory. The SMC-R protocol defines a means to exploit the shared memory for communications – transparent to the applications! SMC-R RFC: https://datatracker.ietf.org/doc/rfc7609/**
- **This solution is referred to as SMC-R (Shared Memory Communications over RDMA). SMC-R represents a sockets over RDMA protocol that provides a foundation for a complete solution meeting all of the described objectives. SMC-R is an RDMA model exploiting RDMA-writes (only) for all data movement.**

SMC-R exploitation considerations

- **RoCE exploitation in z/OS is provided transparently for applications that exploit TCP sockets using SMC-R**
- **10GbE RoCE Express2 is transparent to:**
  - Socket applications
  - Peer systems connected via RoCE (i.e. the wire flows are unchanged, there are no differences in or awareness of the generation of RoCE)
- **Two (2) RoCE FIDs (unique PCHIDs) are recommended for redundancy, SMC-R Link Groups are formed using the 2 RoCE FIDs (ports)**
- **New consideration: Mixing generations of RoCE adapters on the same stack supported?**
  - 10GbE RoCE Express2 can be mixed with RoCE Express (i.e. provisioned to the same TCP/IP stack or same SMC-R Link Group)
- **Maximum number of RoCE Express and RoCE Express2 features supported per z14 is 8 (combined)**

Improved speed for development and applications on z14

- **Compilers**
  - COBOL v6.2 fully support the Packed Decimal Facility to reduce CPU usage for decimal intensive applications by up to ***38%*** and on average ***19%***
  - Automatic Binary Optimizer v1.3 reduces CPU usage of applications built with COBOL v4 (and below) without source recompilation by up to ***47%***
  - z/OS 2.3 XL C/C++ reduces CPU usage of compute intensive applications on average ***13%***

- **Java SDK 8 SR5**
  - Faster user response times for Java
  - ***4.2x***[1] improvement to AES-GCM crypto to enable best-of-breed security for the API-economy using Java
  - Pause-less garbage collection baked into the processor, reducing pause times by up-to ***3x***[2] for predictable high-perform transaction processing at-scale
  - 50+ new instructions on the z14 co-designed and exploited by Java

1. z14 z/OS WAS Liberty with GCM Mode Encryption: Customers running WebSphere Liberty for z/OS, using clear key encryption AES_128_GCM cipher, can get up to 4.2X improvement in throughput per core with IBM Java 8 SR5 on z14 compared to Java 8 SR3 on z13)
2. Performance results based on IBM internal tests running Java Store Inventory and Point-of-Sale in COMPOSITE mode application benchmark on standalone dedicated IBM z14 and z13 machines using z/OS 2.2 2 with APAR OA51643 and no other workloads running in the LPAR under test. Both z14 and z13 were configured with 1CP and 8 SMT zIIPs with total 17 hardware threads. Capacity projections were done to estimate the benefit of moving from z13 z/OS 2.2 Java 8 SR3 to z14 z/OS 2.2 Java 8 SR5 with Pause-less garbage collection enabled by java option -Xgc:concurrentScavenge. The response-time constrained Service Level Agreements (SLAs) metric used for this claim was based on geometric mean of (throughput @ 10ms, 25ms, 50ms, 75ms and 100ms response time SLAs). Hardware instrumentation data was collected and analyzed on all benchmarks to verify performance results. IBM 64-bit SDK for z/OS Java Technology Edition, Version 8 SR3 on z13 was used as the baseline. IBM 64-bit SDK for z/OS Java Technology Edition, Version 8 SR5 is scheduled to GA September 2017.

Innovation for analytics with Linux on IBM Z

- **Massive scale up of JVMs with higher capacity IFLs, pause-less garbage collection, and 50+ new instructions co-designed and exploited by Java**
- **New SIMD instructions improve running analytics on IBM Z with increased parallelism**
- **A comprehensive portfolio of cognitive and analytics solution is available for Linux, allowing Linux to become the analytics hub in the enterprise**
- **Up to *2x* more memory for greater processing scale and performance, enables more in-memory workloads and in-line analytics for delivering richer transactional experiences**
- **Performance, networking and efficiencies running Linux side-by-side with z/OS**

Capacity and features to manage more data

- **Up to 8 TiB memory**
  - Support new workloads, in-memory databases and efficiently process huge amounts of information for real-time business insights
- **IBM Virtual Flash Memory**
  - Next generation of Flash Express to provide higher levels of availability and performance during workload transitions and spikes
- **Single Instruction Multiple Data**
  - Enhanced math libraries provide performance improvements for analytical workloads
- **Shared Memory Communications – Direct Memory Access (SMC-D)**
  - Up to ***61%*** CPU savings for FTP file transfers for z/OS versus HiperSockets™[1]

1. All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

Designed for data
I/O options that protect, access, share

Pervasive Encryption		Getting to Data		Accessing the Web		Clustering to Protect	
New	CF Encryption	*New*	zHyperLink Express	*New*	10 GbE RoCE Express2	*New*	Coupling Express LR
New	Crypto Express6S	*New*	IBM Virtual Flash Memory	*New*	OSA-Express6S		ICA SR
New	Speed of CPACF	*New*	FICON Express16S+		HiperSockets	*Plus*	Improved CF scalability, constraint relief, and diagnostic enhancements
New	Payment Card Industry (PCI) HSM		zHPF – Extended Distance II		SMC-D		
New	TKE 9.0		zEDC Express				

Virtual Flash Memory

What is VFM? – High level

- The “storage class memory” provided by Flash Express adapters is replaced with Virtual Flash Memory (VFM) feature which uses part of the main memory
- Customer can buy one to four 512 GiB “units” of VFM
- Much simpler management of VFM resource (HMC task)
- No hardware repair and verify (no cables, no adapters)
- Better performance since no I/O to a PCIe attached adapter takes place
- RAS: Memory protected by RAIM and ECC (internal / main memory)

What is VFM? – The details

- Data persistence is not supported (no change from Flash Express)
- No change to the existing customer operating system interface (i.e. existing architecture is unchanged)
- Allocation of VFM storage moves to LPAR activation since PR/SM "owns" management of partition memory
- Customer specifies initial and maximum amount of VFM
- VFM can be added or deleted by operating systems using existing "storage class memory" after partition activation
- Simplification in management of VFM since there is no hardware adapter to manage
- VFM allocations and definitions for all partitions can be viewed through "Storage Information" panel

VFM activation profiles

IBM **Hardware Management Console**

SEARCH FAVORITES

Home | Customize/Delete Activatio...

Customize Image Profiles: MUSCA:MUSCA11 : MUSCA11 : Storage

- MUSCA:MUSCA11
  - MUSCA11
    - General
    - Processor
    - Security
    - Storage
    - Options
    - Load
    - Crypto

Central Storage

Amount in: Gigabytes (GB)

Initial: 8.0

Reserved: 8.0

Storage origin:
- (●) Determined by the system
- ( ) Determined by the user

Origin: 0.0

Virtual Flash Memory (GB)

Choose memory amounts in 16GB increments up to a maximum of 1024GB.

Initial: * 16

Maximum: * 32

Cancel | Save | Copy Profile | Paste Profile | Assign Profile | Help

z14 Model ZR1 Capacity and Performance Planning

z14 ZR1: Primary performance drivers

- **Memory subsystem**
  - Focused on keeping data “closer” to the processor unit
    - Larger L1, L2, and L3 caches
    - One unified L4 shared by L3s
  - Twice (2x) configurable memory (8 TiB vs. 4 TiB for z13s)
- **Processor**
  - Improved IPC with microarchitecture enhancements
    - Pipeline optimizations and improved branch prediction
    - Cache redesigned to use virtual TLB1 and reduce TLB2 miss latency
    - Four faster dynamic address translation engines vs. one for z13
  - Improved SMT for zIIPs, IFLs, and SAPs (new for z14)
  - SIMD architecture extensions for analytics and register-based decimal operations
  - New Guarded Storage Facility enables near “pause-less” garbage collection for Java
  - Improved crypto co-processor with 4x to 6x performance improvement for AES
  - Up to 10 processor units (cores) per chip vs. 8 on z13
- **Up to 30 configurable processor units (cores) vs. 20 on z13s**
- **156 sub-capacity settings (same as z13s)**
- **HiperDispatch**
  - Exploits new chip configuration
  - Required for SMT on zIIPs
- **PR/SM**
  - 40 customer partitions (same as z13s)
  - Up to 30 LCPs and 8 TiB memory per partition (actual limits are operating system dependent)
  - Improved resource allocation algorithms based on z13s experience

z14 ZR1: Additional details

- **Processor microarchitecture improvements**
  - 4.5 GHz z14 SF vs 4.3 GHz z13s (+4.65%)
  - Merged L1/TLB1 – eliminates TLB1 miss penalty, inlined TLB2 access on L1 miss mitigates TLB2 access penalty
  - Four HW-implemented translation engines on z14 SF vs one picocoded engine on z13s
  - 2x CRSTE (combined region segment table entry) and 1.25x PTE (page table entry) growth
  - Branch prediction improvements; 33% BTB1-and-2 growth, new perceptron predictor and simple call-return stack
  - Pipeline optimization; improved instruction delivery, faster branch wakeup, reduced execution latency, improved OSC prediction, …
  - Optimized 2nd generation SMT; improved thread balancing, multiple outstanding translations, optimized hang avoidance mechanisms
  - Improved Hot Cache Line handling; dynamic throttling – e.g., "XI strong-arming" and a Hot Line Table
- **Cache Improvements**
  - L1 I-Cache increased from 96K to 128K per Core (1.33x)
  - L2 D-Cache increased from 2MB to 4MB per Core (2x)
  - L3 Cache increased from 64MB to 128MB per CP (2x)
  - L4 Cache changed from 480MB L4 + 224MB NIC to 672MB L4, however only one L4 per drawer on z14 SF vs two on z13s
  - L4 sequential prefetch
- **Storage Hierarchy Improvements**
  - Single SC Topology (Reduced system latencies)
  - Redesigned System protocols (Reduce contention points)
  - New L3 fetch miss resumption on L4 sequential pre-fetch conflict
  - Bus Feeds and Speeds
  - Wider processor store bus
  - Improved Hot Cache Line handling; contention affinity, single SC
- **Software / Firmware Improvements**
  - PR/SM: Improved Memory Affinity, improved logical partition placement algorithms based on z13s experience
  - New Translation Management Facility (avoid expensive ops)
  - New per-work-unit dispatch cache footprint metrics (cache efficiency)
  - Improved Hot Cache Line handling (new instructions)

Efficiency managing data movement to improve access time

- **On-chip compression coprocessor**
  - Enhancements enable further compression of data including Db2® indices, improving memory, transfer and disk efficiency
  - In the future[1] Db2 plans to enable new order-preserving compression for Db2 indices using compression coprocessor to support index compression
- **zEDC**
  - Compression further reduces cost to pervasively encrypt data with less data to encrypt
  - More data active and effective compression with a dedicated compression accelerator
  - Disk savings with improved utilization of storage tiers with DFSMSdss™ use of compression

1. IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Sharing the data

- **OSA-Express6S**
  - Technology refresh of OSA-Express6S
- **10GbE RoCE Express2**
  - Technology refresh used with SMC-R
  - Provides increased virtualization / sharing capability allowing RoCE to be extended to more workloads
- **Coupling technology enhancements**
  - Ethernet based coupling links using 10GbE RoCE technology – Coupling Express LR
  - Better utilization of Coupling Facility (CF) processors with scalability improvements
  - Faster problem resolution with additional CF Request Diagnostics
  - Additional physical and logical coupling links offers Coupling Link Constraint Relief

Enterprise service agility with Linux on Z
Improved IT economics

- **Linux cores (IFLs) will benefit from next generation SMT for improved virtualization performance with up to *25%* vs. non-SMT on z14**
- **Up to 30 configurable Linux cores provide efficiency at scale and improve price/performance**
- **z14 has a redesigned cache architecture with *1.5x* more on-chip cache per core – compared to the z13**
- **Up to *2x* memory to support large VM consolidations, provide a higher ceiling for vertical scale needs and to support data-in-memory applications**
- **FICON Express16S+ with FCP protocol for small data transfer I/O operations achieved a greater than 3x improvement over FICON Express16S[1]**

1. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

Simplicity with Linux on IBM Z
Efficient and economic operation

- **Easier system administration using IBM Dynamic Partition Manager with simplified hardware configuration and provisioning tools enables greater time to value**
- **Comprehensive open source software portfolio with such products as Docker, Go, Chef, Puppet, MongoDB, PostgreSQL, Apache Spark, or Node.js to modernize your applications**
- **Improved on chip and coprocessor cryptographic performance to ensure your data can be protected from external and internal breaches**
- **Massive scale up of JVMs with pause-less garbage collection**
- **8 TiB of memory for greater processing scale and performance, enables more in-memory workloads and in-line analytics for delivering richer transactional experiences**
- **New DBaaS reference architectures to support cloud deployments**

HMC Highlights

Next generation Hardware Management Console

Empowering users by providing them with a modern workspace that equips them to securely and confidently manage system hardware from anywhere

Tree structure

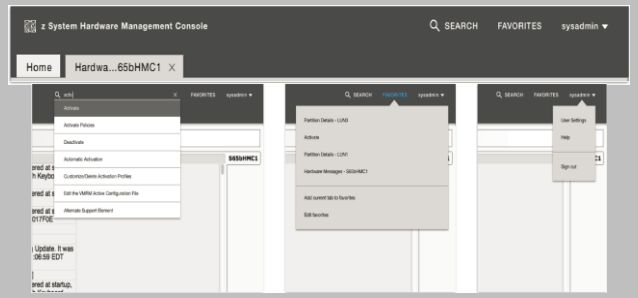

Multi-factor authentication

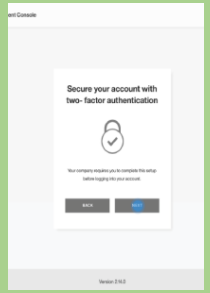

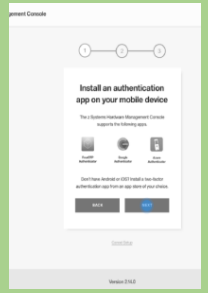

Mobile capabilities

Workspace enhancements

- **Classic UI Style NO longer available on HMC/SE 2.14.0 and later**
- **Workspace Enhancements to aid in transition to Tree Style UI**
  - A new masthead for the HMC and SE will help users quickly find and launch tasks
  - Tasks will now open in tabs within the user interface (instead of separate browser windows) to make finding and managing running tasks easier

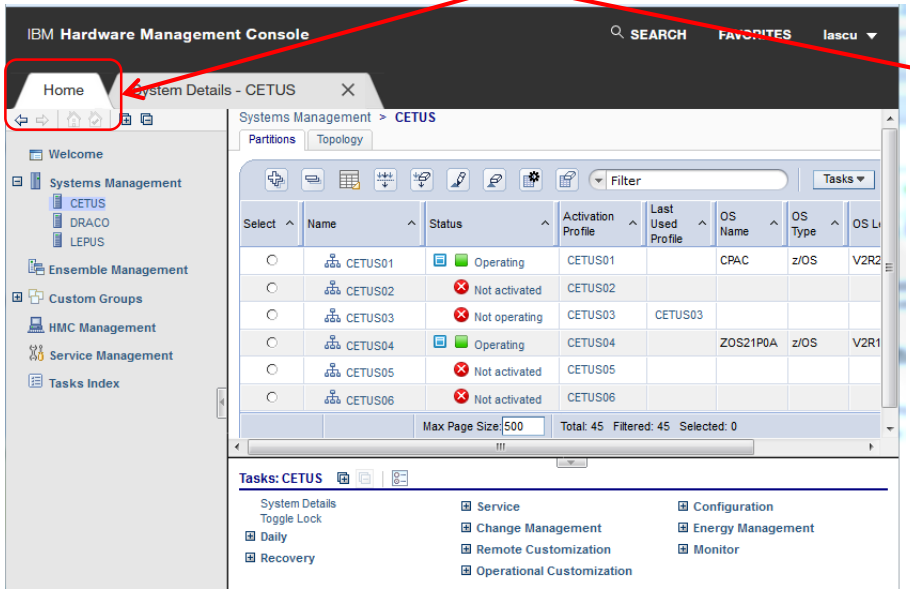

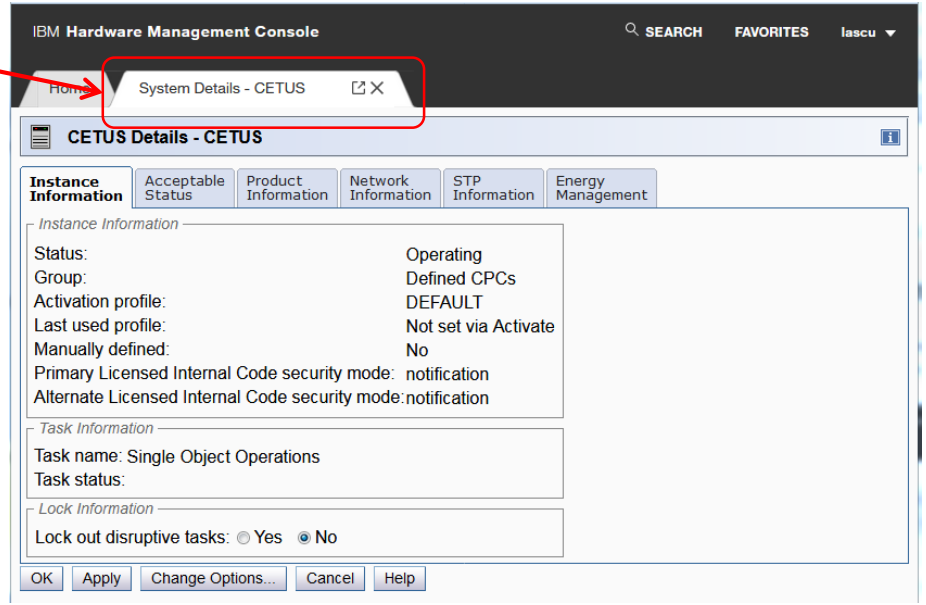

Java applet removal

- **2.13.1: HMC Tasks no longer Java Applet based implementations**
  - Operating System Messages
  - Integrated 3270 Console
  - Integrated ASCII Console
  - Text Console
- **2.14.0: Open IOCDS Source option on Input/Output Configuration task**
  - With HMC/SE 2.14.0 can now edit IOCDS source directly on HMC console
  - Alternative with remote browsing still available:
  - Use Export Source File option on Input/Output Configuration task
  - Use editor on your own workstation
  - Use Import Source file option to put back onto Support Element

Simplified workflow for system time management

- **Improved help tools to compliment and improve system admin skills**
  - Inline definition of technical terms
- **Improved user experience with visual representation of configuration panels**
  - Guidance provided within the workflows
  - Topology displays of system time networks
  - Errors surfaced in visualization for easier problem resolution of setup errors
- **Single point of system time management for multiple systems**

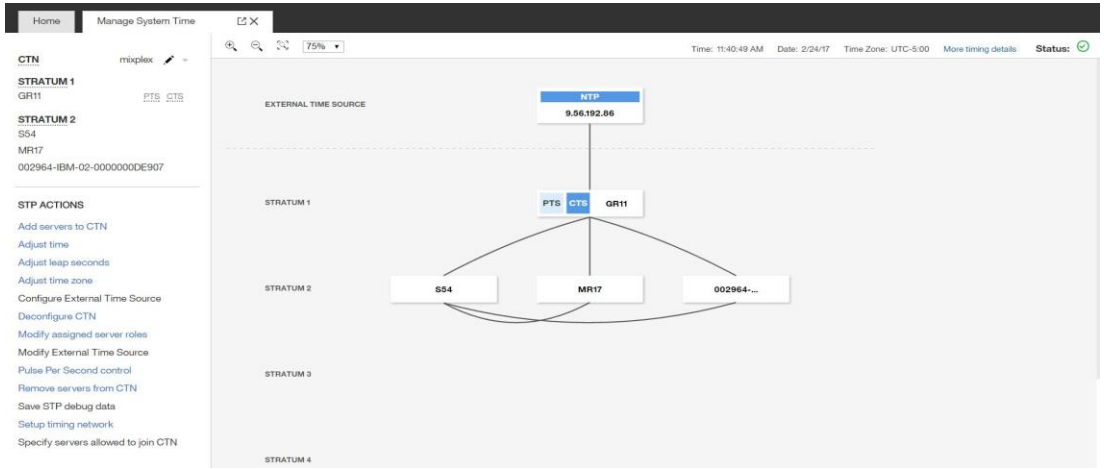

HMC Multi-Factor Authentication

- **What is Multi-Factor Authentication?**
  - Method of user authentication that requires multiple pieces of evidence
  - Factor – component used to prove identity
    - Something you know e.g. username, password, PIN
    - Something you have e.g. badge, smartphone, keyfob, USB stick, key, bank card
    - Something you are (biometrics) e.g. fingerprint, eye iris, voice, typing speed/pattern
  - Online banking, ATMs, Google, Facebook, etc.

- **HMC Multi-Factor Authentication:**
  - Optional
  - Configurable on a per-user, per-template basis
  - Local and remote GUI logon
  - LDAP-authenticated
  - Web Services APIs
  - Locally-authenticated users
  - users
  - Pattern/Template users
  - Standalone solution required
    - No network connectivity
    - Partitions not running (e.g., LDAP, RACF)
  - User-supplied smartphone
  - HMC, SE and TKE

z14 ZR1 Power and Cooling

Data center planning and service updates

- **Standardization across many components – including Industry standard 19” rack**
- **16U free space in frame**
  - Save space in the data center with 1U rack-mounted HMC and TKE
- **Non-raised floor option, overhead power and cabling**
- **2 or 4 30A single-phase 200V-240V power cables**
- **New – lowering costs and raising RAS with ASHRAE A3 envelope**

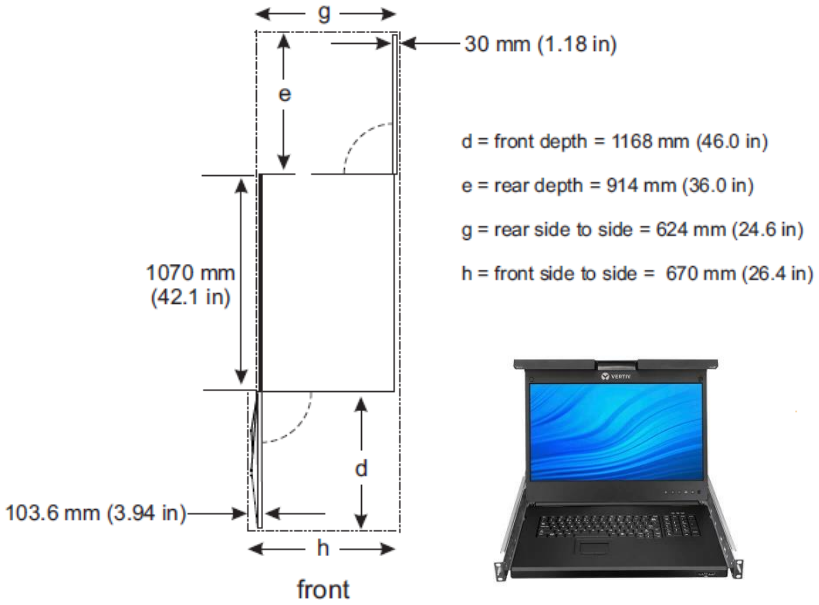

Power infrastructure design

- **IBM z14 ZR1 power characteristics:**
  - Packaged in a 42U 19” rack
  - Single Phase (200 – 240 VAC)
  - Intelligent Power Distribution Units (PDUs) – two or four PDUs per system, depending on configuration, switchable (Ethernet controlled)
    - Rating per power cord: 200 – 240 VAC, 30A
  - Redundant power supply units (PSUs) for:
    - Support Elements (1+1)
    - PCIe+ I/O Drawers (1+1)
    - CPC Drawer (two or four PSUs, redundant, configuration dependent)
  - Redundant Ethernet Switches
    - Each Ethernet switch has one built in power supply, the entire switch is a FRU
  - Max. supported rack power – 9600W
  - Rack Airflow – front to rear
  - All cabling in the rear of the rack

- **IBM z14 ZR1 does not provide:**
  - Bulk Power Regulator (BPR)
  - External EPO switch
  - Internal Battery Feature (IBF)
  - Three phase power
  - 480 VAC
  - HVDC (high voltage DC)
  - Balanced power
- **Notes:**
  - Power is also provided for hosting non-Z equipment in the z14 ZR1 rack (16U feature)
  - Power specifications for “hosted” equipment are provided in the IBM z14 ZR1 Installation manual for Physical Planning, GC28-6974
  - When servicing equipment hosted in the z14 ZR1 rack an assessment must be carried out to avoid power outages to the other equipment installed in the same rack

Utility power consumption

CPC Drawer Feature	Number of PCIe+ I/O drawers / PDUs				
	0/2	1/2	2/4	3/4	4/4
0636 (Max4)	1.36 kW	2.26 kW	-	-	-
0637 (Max12)	1.77 kW	2.67 kW	3.58 kW	-	-
0638 (Max24)	2.59 kW	3.48 kW	4.39 kW	5.29 kW	6.12 kW
0639 (Max30)	2.59 kW	3.48 kW	4.39 kW	5.29 kW	6.12 kW

Notes:

- 16U Reserved (FC 0617) limits the number of PCIe+ I/O drawers to two (2)
- With 16U reserved maximum 3400 W in addition to the installed Z equipment can be used
- With 16U Reserved there are always four (4) PDUs installed

Cooling considerations

- **Airflow direction is front-to-back (i.e. other airflow orientation (for non-Z equipment) could lead to hot exhaust air being recirculated back to the front of the drawers)**
- **Cabling must be dressed appropriately such it does not block the front or back (i.e. would result in reduced airflow and increased component temperatures)**
- **All unused rack locations must be blocked with front and rear filler plates**
- **No cabling in the front of the rack (i.e. if power cables for non-Z equipment are present, these should be routed to the sides and towards the rear side of the rack)**

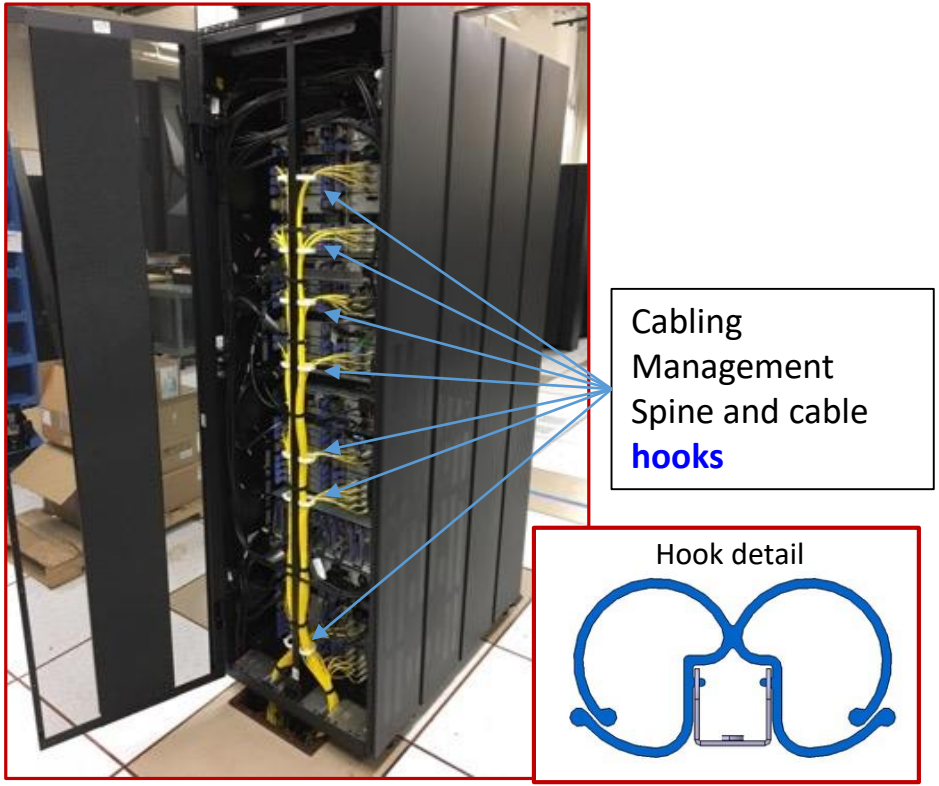

16U Reserved Feature

Imagine the possibilities of building an “all-in-one” solution

- **With smaller I/O configurations – a new 16U Reserved feature code can be added that tags 16U of space in rack as “Available”**
- **This creates *a new opportunity* to customize a comprehensive solution that fixes your requirements[1]**
- **Populate with your choice of server, switch or storage elements[2]**
- **Available only on the IBM z14 ZR1**

1. Hardware service to these IBM or non IBM options may be provided by a 3rd party – but must be able to manage analysis of power, thermal, air flow and other requirements as listed in the IMPP guide. Training may be offered via IBM as negotiated.
2. Requirements for physical structures as well as interactions with the ‘mainframe server’ will be provided

IBM Z and IBM Storage synergy

Storage Networking

Other examples of uses for 16u Reserved include IBM 1u HMC and TKE, Power Systems™, NVMe

Flash and Hybrid Storage Systems

z/VM®, z/VSE®, Linux on Z (FCP only)	z/OS, z/VM, z/TPF, z/VSE, Linux on Z (FCP and FICON)

FlashSystem™ A9000[1]

Storwize® V7000 / V7000F

FlashSystem V9000

FlashSystem FS900[1]

DS8880

DS8880F Storage[2]
Statement of Direction

1. FlashSystem A9000 and FS900 are not supported by z/VSE
2. IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Dynamic Partition Manager

IBM Dynamic Partition Manager

Simplified configuration of logical partitions

- **Standardizes configuration and management of all system resources from a single management endpoint**
- **Developed for servers with z/VM 6.4, KVM, and/or Linux as a partition-hosted operating system**
- **Ease Linux installation with auto configuration of devices[1]**
- **Guided storage setup, provisioning and management – new FICON ECKD support**
- **Secure FTP through HMC for booting and installing an operating system via FTP**

1. Where Linux distro installers exploit function

DPM functions available via graphical and API interfaces

IBM **Hardware Management Console**

SEARCH FAVORITES octavian

Home | Partition Details - M90-L...

Partition Details - M90-LNXT02

General
Status
Controls
Processors
Memory
Network
Storage
Accelerators
Cryptos
Boot

Processors

Processor mode: Shared Dedicated

Processors: 7

1 18 35 52 69 Processors
2 36 70 104 138 Threads

Processing weight: 100

999 Very High
900
700 High
500 Medium
300 Low
100 Very Low
1

Enforce weight capping
Enforce absolute processor capping

Number of processors (0.01-255.0): 1

Manage Processor Sharing

Processors

70 60 50 40 30 20 10 0

Entitled
Dedicated
M90-LNXT02
Shared

Shared Processors
Virtual/Physical:58.82%

M90-LNXT02
M90-KVMP10
M90-LNXT01
M90-LNXP11
M90-LNXP10

Active Processing Weights

M90-LNXT02
M90-LNXT01
M90-KVMP10
M90-LNXP11
M90-LNXP10

Memory

Related Tasks
Stop
System Details
Manage Adapters
Monitor System

OK Apply Cancel Help

Coupling Technology Changes

Coupling Facility enhancements

- **Ethernet based Coupling Links**
  - ***New*** Coupling Express LR using 10GbE RoCE technology
  - Integrated Coupling Adapter (ICA SR)
- **Better utilization of Coupling Facility (CF) processors with scalability improvements**
- **Faster problem resolution with additional CF Request Diagnostics**
- **Additional physical and logical links offers Coupling Link Constraint Relief**

Parallel Sysplex Coupling Links – ICA SR

- **IBM Integrated Coupling Adapter (ICA SR) – FC 0172**
  - Coupling connectivity into the future (short distance)
    - ICA SR is recommended for short distance coupling z13/z13s to z13/z13s and beyond
  - Coupling channel type: CS5
    - Performance similar to coupling over InfiniBand 12X IFB3 protocol
  - PCIe Gen3, fanout in the CPC drawer, 2-ports per fanout, up to 150m
    - 8 GigaBytes per second (GiBps)[1]
  - z13/z13s/z14 to z13/z13s/z14 and up connectivity
  - Maximum configurations supported:
    - 40 links per z13 CPC; Up to 4 CHPIDs per port, 8 buffers (i.e. 8 subchannels) per CHPID
    - 80 links per z14 CPC; Up to 4 CHPIDs per port, 8 buffers (i.e. 8 subchannels) per CHPID
  - ICA requires new cabling for single MTP connector; cables - 150m: OM4; 100m OM3
  - Differs from 12X Infiniband split Transmit/Receive connector
  - Available as of z13 GA1

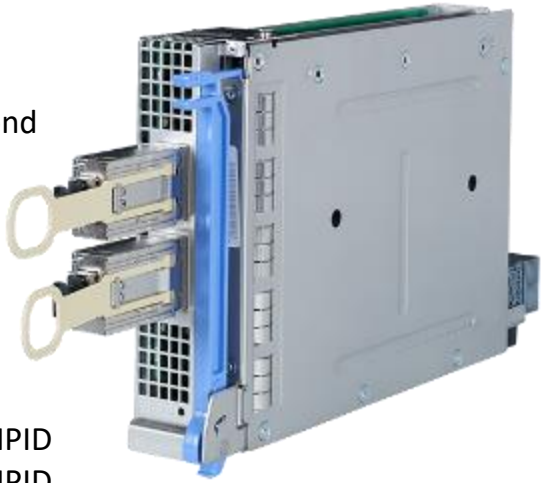

1. Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

Parallel Sysplex Coupling Links – CE LR

- **Coupling Express Long Reach (CE LR) – FC 0433**
  - Coupling connectivity into the future (long distance)
    - Coupling Express LR is recommended for long distance coupling z13/z13s to z13/z13s and up
  - New coupling channel type: CL5
  - Performance is similar to coupling over InfiniBand 1x
  - PCIe I/O drawer required for CL5 adapter – even for standalone CF usage
  - Feature (2-port card) with Coupling Optics and firmware
  - 10 Gbps[1], Up to 4 CHPIDs per port, 32 buffers (i.e. 32 subchannels) per CHPID
  - Distance: 10 km unrepeated; up to 100 km with a qualified DWDM
    - RPQ 8P2197 is required for 20 km support, while more than 100 km requires RPQ 8P2981.
  - Maximum configuration supported:
    - 64 ports (32 features) for z14 and z13 , 32 ports (16 features) for z13s
  - Point-to-point just like InfiniBand 1X and ISC-3
    - Can not be utilized in a switched environment
  - Cabling: utilizes same 9µ, single mode fiber type as 1X IFB and ISC-3
  - z13 GA2+, z13s GA1+ and z14 availability
  - No Going Away Signal for STP

1. Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

Coupling link roadmap (dual-frame systems)

z10	z196	zEC12	z13		z14 M0x	zFuture[1]	
			GA1	GA2+			
				Coupling Express LR[2]	Coupling Express LR	Coupling Express LR	
			ICA SR	ICA SR	ICA SR	ICA SR	
	12X HCA3-O and 1X HCA3-O LR	12X HCA3-O and 1X HCA3-O LR	12X HCA3-O and 1X HCA3-O LR		12X HCA3-O and 1X HCA3-O LR		Initial offering
12X HCA2-O and 1X HCA2-O LR	12X HCA2-O and 1X HCA2-O LR	12X HCA2-O and 1X HCA2-O LR					New build or carry forward
ISC-3	ISC-3	ISC-3					Last generation to order
ICB-4							Last generation to support

1. All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only
2. Coupling Express LR is the future Long Distance Ethernet Coupling Link which will reside in the PCIe I/O drawer. Like ICA SR, it will require z13 to z13(+) connectivity.

Coupling link roadmap (single-frame systems)

z10 BC	z114	zBC12	z13s		z14 ZR1	zFuture[1]
			GA2	GA2+		
				Coupling Express LR[1]	Coupling Express LR	Coupling Express LR
			ICA SR	ICA SR	ICA SR	ICA SR
	12X HCA3-O and 1X HCA3-O LR	12X HCA3-O and 1X HCA3-O LR	12X HCA3-O and 1X HCA3-O LR			
12X HCA2-O and 1X HCA2-O LR	12X HCA2-O and 1X HCA2-O LR	12X HCA2-O and 1X HCA2-O LR				
ISC-3	ISC-3	ISC-3				
ICB-4						

Initial offering
New build or carry forward
Last generation to order
Last generation to support

1. Coupling Express LR is the future Long Distance Ethernet Coupling Link which will reside in the PCIe I/O drawer. Like ICA SR, it will require z13 to z13(+) connectivity.

z14 Coupling Connectivity

Z14 DF, z13, and z13s
12x IFB, 12x IFB3, 1x IFB, ICA SR, CE LR

zEC12 and zBC12
12x IFB, 12x IFB3, 1x IFB

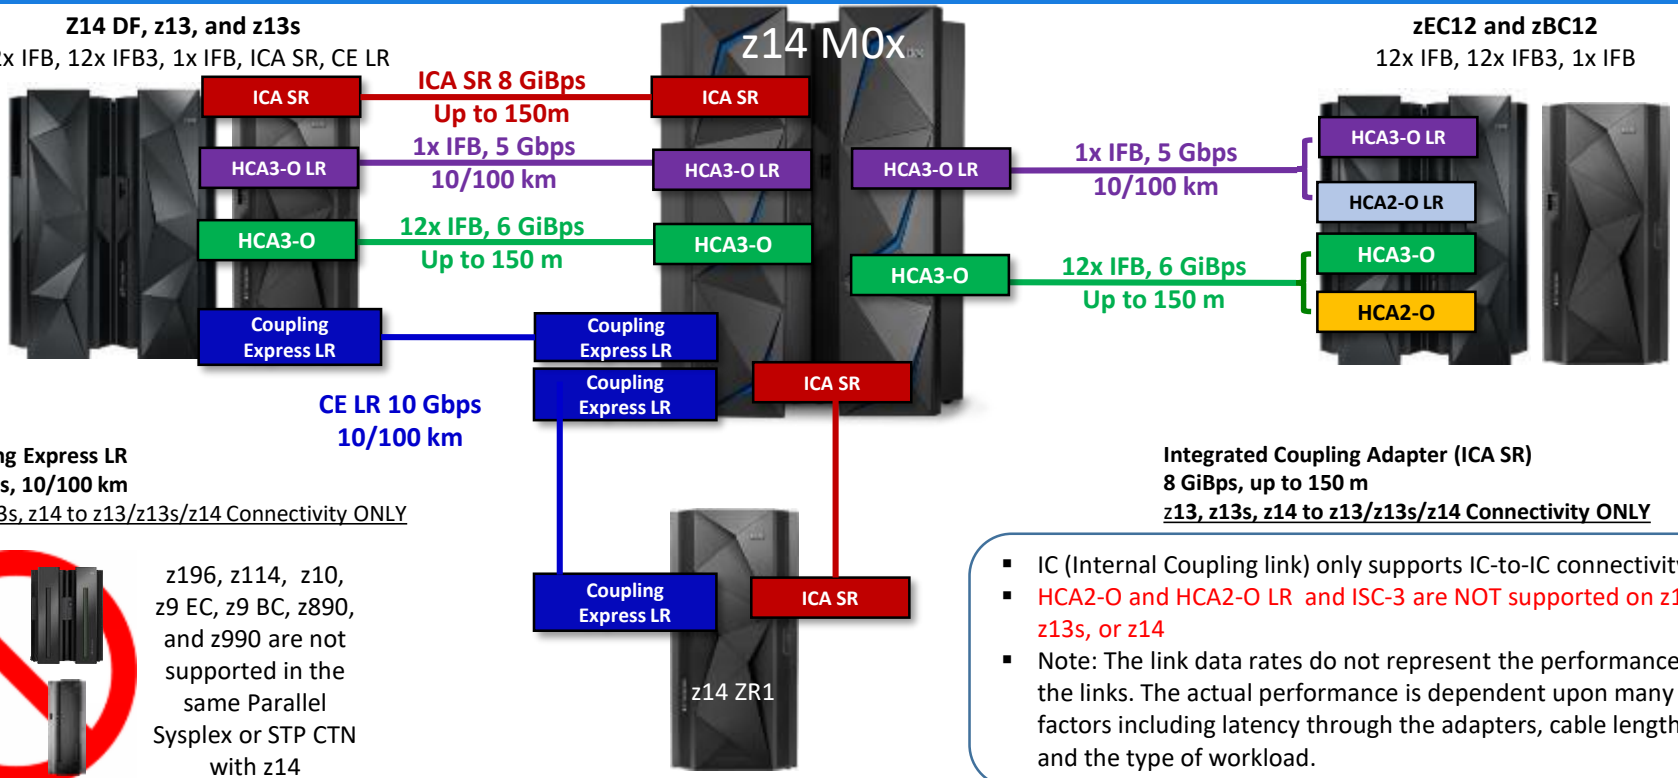

Coupling Express LR
10 Gbps, 10/100 km
z13, z13s, z14 to z13/z13s/z14 Connectivity ONLY

z196, z114, z10, z9 EC, z9 BC, z890, and z990 are not supported in the same Parallel Sysplex or STP CTN with z14

Integrated Coupling Adapter (ICA SR)
8 GiBps, up to 150 m
z13, z13s, z14 to z13/z13s/z14 Connectivity ONLY

- IC (Internal Coupling link) only supports IC-to-IC connectivity
- HCA2-O and HCA2-O LR and ISC-3 are NOT supported on z13, z13s, or z14
- Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

Coupling link history (z196/z114 to z14)

Sysplex Coexistence	N-3	N-2	N-1	N		
	z196 / z114	zEC12 / zBC12	z13 / z13s	z14 M0x	z14 ZR1	Comments
ISC3	NB & CF	CF	Not Supported	Not Supported	Not Supported	Long Distance Lower Bandwidth
HCA2-O 12x	NB & CF	CF	Not Supported	Not Supported	Not Supported	Short Distance High Bandwidth
HCA2-O 1x LR	NB & CF	CF	Not Supported	Not Supported	Not Supported	Long Distance Lower Bandwidth
HCA3-O 12x	NB	NB & CF	NB & CF	NB & CF	Not Supported	Short Distance High Bandwidth
HCA3-O 1x LR	NB	NB & CF	NB & CF	NB & CF	Not Supported	Long Distance Lower Bandwidth
ICA SR	Not Available	Not Available	NB	NB & CF	NB & CF	Short Distance High Bandwidth
CE LR	Not Available	Not Available	NB	NB & CF	NB & CF	Long Distance Lower Bandwidth
Internal Coupling	Yes	Yes	Yes	Yes	Yes	Internal to server only High Bandwidth

NB = New Build, Technology Exchange, Migration Offering
CF = Carry Forward, available via MES upgrade from pervious technology only

Coupling link considerations for z14 ZR1

- **z14 ZR1 supports only PCIe-based coupling (ICA SR and CE LR)**
  - Maximum number of physical ICA SR coupling links (ports) is 16 per CPC
  - However, long distance coupling requires PCIe+ I/O Drawer for hosting CE LR features, thus reducing the maximum number of ICA SR features.
- **While physical coupling links can be (and most often are) shared across images and sysplexes, there are some customers who configure dedicated physical CF links**
  - "Stacked" consolidated sysplexes, on a set of physical CECs, with dedicated connectivity provided for each sysplex
  - These kinds of dedicated/isolated configurations drive requirements for higher maximum limits on both physical links and logical CHPIDs
- **CPCs that host standalone CFs tend to have the highest per-CPC consumption of coupling connectivity resources**
  - CFs are "focal points" for both z/OS-to-CF and CF-to-CF sysplex connectivity, as well as STP timing roles
- **z14 M0x (3906) is the last machine supporting the Infiniband coupling**
  - Additional coupling link configuration complexity is expected for transitioning to ICA SR and CE LR coupling

z14 ZR1 Pervasive Encryption and Cryptography

The IBM z14 makes pervasive encryption achievable

- **Pervasive encryption gives Z clients a simplified way to protect data at a much coarser scale at industry best performance – even with no change it runs faster!**
- **Pervasive encryption provides the ability to encrypt data by policy without application change**
- **Pervasive encryption greatly simplifies audit and makes it easier for clients to pass compliance audits**

Encrypting as much as possible of your data and transactional pipeline helps reduce potential data breach risks and financial losses

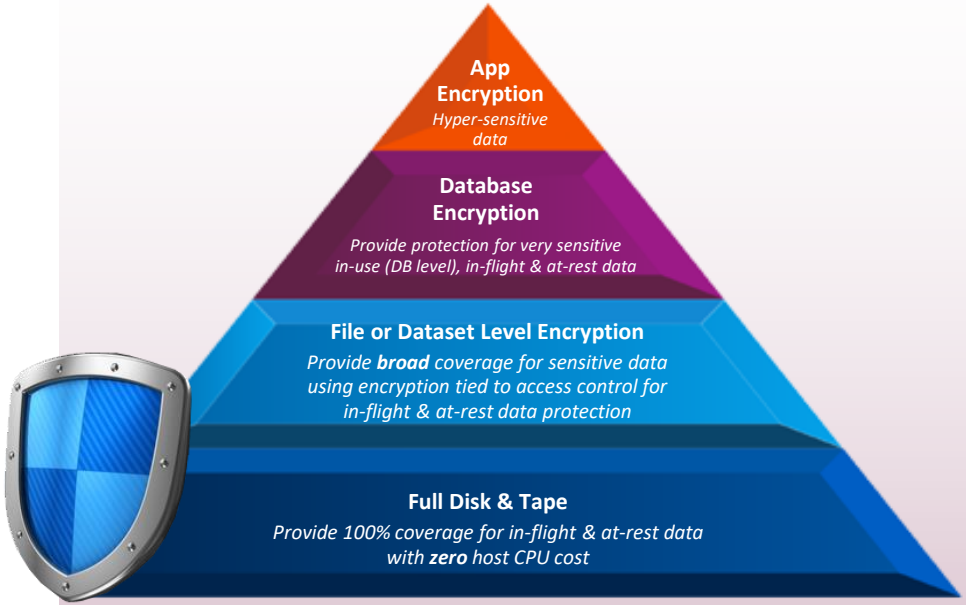

IBM z14 – Performance that changes the game for security

- **Performance with integrated cryptographic hardware**
  - ***6x*** faster encryption for like modes and data sizes with enhanced on-chip (CPACF) cryptographic performance compared to IBM z13s™[1]
  - ***2X*** the SSL handshake performance on z14 ZR1 with Crypto Express6S compared to z13s with Crypto Express5S[2]
- **Datasets automatically protected with z/OS dataset encryption**
  - Protect z/OS data sets automatically throughout their life cycle
  - Enforce consistent policy over access to encrypted content
- **Protection in the sysplex**
  - Data is encrypted/decrypted at a host and is protected in flight and at rest inside the Coupling Facility (CF)

1. This performance claim is extrapolated from IBM internal tests comparing CPACF throughput rates on z14 to z13. It is based on the assumption that the only factor that has a significant influence on CPACF throughput is cycle time. Further, that the cycle time ratio between z14 and z13 is within 1% of the cycle time ratio between z14 ZR1 and z13s.
2. This performance claim is extrapolated from IBM internal tests comparing SSL handshake throughput rates on z14 with CEX6S to z13 with CEX5S. It is based on the fact that the CEX6S in a z14 ZR1 is the very same device as the CEX6S in a z14, with no difference in components, clock speed, features or anything else that would have an impact on the CEX6S throughput.

Encryption for data at rest and inflight for Linux on Z

- **Fast encryption of Linux workloads delivered with near-zero overhead via hardware accelerated encryption of Central Processor Assist for Cryptographic Function (CPACF) and new Crypto Express6S**
  - All encryption functions within the Linux kernel and the openSSL, open Cryptoki and GSKIT libraries are benefiting
    and transparently delivering the performance to the applications and middleware
- **Forthcoming Linux on Z[1] enhanced security via "protected key" encryption for data at-rest**
  - Protected key encryption is processed in the CPACF for high speed and stored in a hardware security module (HSM) and enables fast encrypting and decrypting of complete disks (volumes) or selected partitions
- **Forthcoming Linux on Z[1] enablement to create true unique cryptographic data using the new true random number generator with CPACF**
- **Performance boost for Java with new Galois Counter Mode (GCM) encryption for minimum latency and operation overhead**

1. IBM is working with the Linux distribution partners to get the functionality included in their distribution for Linux on Z

Note: All claims noted on this slide are based on IBM Internal measurements. Results may vary. Additional information is available upon request.

IBM z14 Crypto Express6S feature

- **General availability in z14**
  - Based on IBM PCIe Cryptographic Coprocessor (PCIeCC), available as machine type 4768
- **Enhanced card performance**
  - Main PPC 476 (qty 2 in lock-step) @ 1.2Ghz (vs. Crypto Express5S at 800 MHz)
  - PCIe Gen 2 (vs. Crypto Express5S at PCIe Gen 1)
  - DDR3 1600
  - Persistent Memory Management for faster boot time (FPGA)
  - New Miniboot implementation for certification compliance
    - Allows easier transition of algorithms when certification requirements change
- **Enhanced Public Key Cryptography algorithms performance**
- **Upgraded secure module tamper detection technology**
  - Improved thermal capabilities for increased performance
  - Continued support for temperature and voltage detection
- **Physical Security Standards in progress/planned:**
  - FIPS 140-2 level 4
  - Common Criteria EP11 EAL4
  - Payment Card Industry (PCI) HSM
  - German Banking Industry Commission (GBIC, formerly DK)
  - Note: PCI-HSM certification is new for Crypto Express6S, the others also apply to Crypto Express5S

Crypto Express6S internal view

IBM z14 CPACF

- **CPACF – Central Processor Assist for Cryptographic Functions**
  - CPACF enhancements are available to all system products/components. ICSF WD17 is required for exploitation via ICSF callable services
  - GCM and SHA-3 improvements will be included in the crypto performance whitepaper
- **z14 CPACF Performance**
  - CPACF encryption rates for like modes and data sizes on z14 are up to six times faster than z13
  - Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13 (results may vary)

IBM z14 CPACF – SHA-3, SHAKE algorithms

- **SHA-3 was standardized by NIST in 2015 (FIPS Pub. 202) as an alternative to SHA-2**
- **SHA-2 is still acceptable and there is no indication that SHA-2 is vulnerable or that SHA-3 is more or less vulnerable than SHA-2.**
- **CPACF is introducing support for the four SHA-3 hashing algorithms**
  - SHA3-224, SHA3-256, SHA3-384, SHA3-512
- **And the two extendable output functions as described by the standard**
  - SHAKE-128, SHAKE-256
- **The SHAKE variants are known as an extendable output functions and allow generating hash values of user-specified lengths**
  - Note however that they are not approved as hash functions
- **These algorithms are exposed to applications via the ICSF One-Way Hash Generate callable service (CSNBOWH)**

IBM z14 CPACF – True Random Number Generation (TRNG)

- **CPACF is adding support for TRNG, an improvement over Deterministic RNG in the sense that the numbers generated are more random**
  - However, this comes with a performance penalty
- **ICSF will use the best of both worlds and take a hybrid approach to random number generation – TRNG will be used to seed a DRNG which will then be used to generate random numbers**
- **Random numbers may be obtained/used via the following ICSF callable services:**
  - PKCS #11 Pseudo-Random Function (CSFPPRF)
  - Random Number Generate (CSNBRNG, CSNBRNGL)
  - PKCS#11 services when generating a clear key

IBM z14 CPACF – GCM improvements

- **CPACF introduced a hardware instruction to perform Galois Counter Mode encryption**
- **Previously, multiple hardware instructions had to be invoked to achieve GCM**
- **A single hardware instruction enables better performance**
- **CPACF GCM is exposed to applications via the following ICSF callable services:**
  - Symmetric Key Encipher (CSNBSYE)
  - Symmetric Key Decipher (CSNBSYD)
  - PKCS #11 Secret Key Encrypt (CSFPSKE)
  - PKCS #11 Secret Key Decrypt (CSFPSKD)

Summary

IBM z14 ZR1 Functional Comparison to IBM zBC12

Performance and Scale	▪ Uni Performance ▪ System Capacity ▪ Models ▪ Processing cores ▪ Granular Capacity ▪ Memory ▪ SIMD	▪ New up to 50% performance improvement over IBM zEnterprise® Business Class (zBC12[1]) ▪ New up to 240% (30-way to 13-way) more capacity for z/OS system total z/OS capacity performance improvement over zBC12[1] ▪ Six CPC feature combinations versus zBC12 has five CPC models ▪ Up to six (10-core) chips and two (6-core) chips on zBC12 ▪ Up to 156 capacity settings on the z14 ZR1 and also on zBC12 for the precise capacity you need ▪ New up to 8 TB RAIM memory vs. 512 GB RAIM memory on zBC12, ideal for analytics ▪ z14 vector processing provides richer, complex analytics models, faster analytics to traditional workloads
Virtualization	▪ LPAR virtualization ▪ RoCE adapter ▪ Simplified LPAR mgt	▪ 40 partitions on z14 ZR1 versus 30 on zBC12 to allow for more customer workloads ▪ New 10 GbE RoCE Express2 with additional virtual functions per port vs. dedicated zBC12 10GbE Express ports ▪ Enhanced IBM Dynamic Partition Manager configuration and management of system resources – not on zBC12
Infrastructure Efficiency	▪ Networking ▪ FICON ▪ IBM zHyperLink ▪ WWPN ▪ HMC ▪ JAVA ▪ IBM Virtual Flash Memory	▪ New OSA-Express6S with ASIC improvements over zBC12 using OSA-Express4S ▪ New FICON Express16S+ versus FICON Express 8S on zBC12 ▪ New IBM zHyperLink – New short distance z14 channel that on IBM DS8880 System Storage® for lower latency ▪ I/O serial number migration allows keeping same serial number on replacement server not on zBC12 ▪ New next gen HMC on z14 ZR1– simplified panels, new mobile capabilities, security enhancements, easier help ▪ Enterprise Java applications run with reducing pause from garbage collection activities on heaps – not on zBC12 ▪ New memory for server side availability and support for large pages – not available on zBC12
Sysplex-Coupling	▪ Coupling Express LR ▪ ICA-SR ▪ STP	▪ Coupling with CE LR in PCIe I/O drawer versus HCA-3 InfiniBand in processor drawer, helps free up slots ▪ Short distance coupling with PCIe-based links (ICA SR) ▪ New Simplified STP management with HMC simplification not available on zBC12
Data Center Efficiency	▪ Accessible frame ▪ Environmental ▪ DC Footprint	▪ New accessible frame to support private cloud and customer equipment for efficient data center use ▪ More choice in top and bottom exit cabling; New ASHRAE A3 rating vs. ASHRAE 2 on zBC12 ▪ New 19 inch rack takes up only 2 floor tiles, 40% less space than z13s
Security	▪ Cryptographic Coprocessor ▪ Crypto Express adapter ▪ IBM SSC ▪ Secure Console Access	▪ CPACF for improved performance and new true Random Number Generator ▪ Crypto Express6S high performance + new algorithms for ECC, SHA, VISA Format preserving encryption ▪ Secure deployment of software virtual appliances not available on zBC12 ▪ Protect sensitive data with TLS support in the OSA-ICC; this is not available on zBC12

1. Based on preliminary internal measurements and projections and compared to the z13s and zBC12. Results may vary by customer based on individual workload, configuration and software levels. Visit LSPR website for more details at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex

IBM z14 ZR1 Functional Comparison to IBM z13s

Performance and Scale	▪ Uniprocessor Performance ▪ System Capacity ▪ SMT ▪ Cache ▪ Processing cores ▪ Granular Capacity ▪ Memory ▪ Compression ▪ SIMD	▪ New up to 10% performance improvement over IBM z13s (z13s)[1] ▪ New up to 60% system (30-way to 20-way) total z/OS capacity performance improvement over z13s[1] ▪ Next gen SMT delivers up to 25% performance improvement over z13s and newly available for SAP processors ▪ New z14 ZR1has 1.5x more on-chip cache per core versus z13s ▪ Up to 30 configurable cores to configure, up to 20 on z13s ▪ New up to 156 capacity settings (same as z13s) ▪ Up to 8 TB RAIM memory vs. 4 TB RAIM memory on z13s ▪ Improved CMPSC compression and Huffman Coding compression ratio using zEDC Express versus on z13s ▪ New instructions for perform boost to traditional workloads and new analytics workloads versus SIMD on z13s
Virtualization	▪ LPAR virtualization ▪ RoCE adapter Virtualization ▪ Simplified LPAR management	▪ 40 partitions – same as z13s ▪ New 10 GbE RoCE Express2 with additional virtual functions (31 VFs) per physical port (same as 10 GbE Express on z13s) ▪ Enhanced IBM Dynamic Partition Manager for config and mgt of system resources – new z/VM and ECKD disk support
Infrastructure Efficiency	▪ Networking ▪ FICON ▪ zHPF ▪ IBM zHyperLink ▪ LCSS/Subchannel sets ▪ HMC ▪ Pause-less garbage collection ▪ IBM Virtual Flash Memory	▪ New OSA-Express6S with improvements over OSA-Express5S on z13s ▪ FICON Expresss16S+ on z14ZR1 offers up to three times the I/Os / second compared to FICON16S ▪ zHPF extended distance II – faster remote site recovery through improved I/O service time for remote data writes ▪ New IBM zHyperLink short distance channel for IBM DS8880 System Storage for low latency; not on z13s ▪ Up to3 LCSS and 3 Subchannel sets – same as z13s ▪ Simplified HMC with new panels, mobile capabilities, security enhancements replacing classic UI on z13s ▪ New enterprise scale Java applications support for larger heaps with less pause delay for garbage collection ▪ New VFM replacement for Flash Express helping improve availability – available only on z14 family
Sysplex-Coupling	▪ Coupling Express LR ▪ Coupling – ICA SR ▪ STP	▪ Coupling with Coupling Express LR Links vs. z13s HCA3 IFB links (CE LR now available on z13s) ▪ Short distance coupling with PCIe-based links (ICA SR) – same as z13s ▪ New Simplified STP management with HMC enhancements and improved user interface not available on z13s
Data Center Efficiency	▪ Accessible frame ▪ Environmental ▪ DC Footprint	▪ New accessible frame to support private cloud and efficiently use data center space ▪ More choice in top exit and bottom exit cabling on z14 ZR1. New ASHRAE A3 rating ▪ New 19 inch rack takes up only 2 floor tiles, 40% less space than z13s
Security	▪ Cryptographic Coprocessor ▪ Crypto Express ▪ Firmware Integrity monitoring ▪ IBM Secure Service Container	▪ CPACF for improved performance (~2-6X faster) and new true Random Number Generator versus z13s ▪ Crypto Express6S performance increase plus new algorithms for elliptic curve, SHA, VISA FPE vs. on Crypto Express5S ▪ New Optional integrity monitoring on Support Element and HMC to protect against tampering ▪ Secure deployment of software virtual appliances

1. Based on preliminary internal measurements and projections and compared to the z13s and zBC12. Official performance data will be available upon announce. Results may vary by customer based on individual workload, configuration and software levels. Visit LSPR website for more details at: https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex

IBM z14
Machine Type: 3906
Models:
M01, M02, M03, M04, M05

Model	Customer PUs	Max Memory
M05	170	32 TiB
M04	141	32 TiB
M03	105	24 TiB
M02	69	16 TiB
M01	33	8 TiB
ZR1	4, 12, 24, 30	8 TiB

IBM z14
Machine Type: 3907
Model ZR1

IBM z14 Model ZR1
IBM LinuxONE Rockhopper II

Extending the IBM z14 and LinuxONE Families

Building on the breakthrough technologies and strong 2017 launches

Designed for the Secure Cloud

IBM LinuxONE Emperor™ II
Machine Type: 3906
Models:
LM1, LM2, LM3, LM4, LM5

Model	Customer PUs	Max Memory
LM5	170	32 TiB
LM4	141	32 TiB
LM3	105	24 TiB
LM2	69	16 TiB
LM1	33	8 TiB
LR1	4, 12, 24, 30	8 TiB

IBM LinuxONE Rockhopper II
Machine Type: 3907
Model LR1

IBM z14 functions available across all z14 models

I/O
FICON Express16S+, zHyperLink Express, OSA-Express 6S, Crypto Express6S, zEDC Express, RoCE Express2, zHPF, including IFP to support the PCIe IO features

RAIM memory

HiperSockets

Specialty processors
IFL, zIIP, ICF

Security
CPACF, Crypto Express6S, GCM Encryption (Java), TKE

Pause-less garbage collection

SIMD

SMC-R and SMC-D

IBM Virtual Flash Memory (VFM)

IBM Dynamic Partition Manager (DPM)

IBM Secure Service Container

On chip compression

Coupling
- ICA SR
- CE LR

HMC Mobile app

Physical planning
- Overhead cabling and power
- ASHRAE A3

IBM z14
Machine Type: 3907
Model: ZR1

IBM z14
Machine Type: 3906
Models: M01, M02, M03, M04, M05

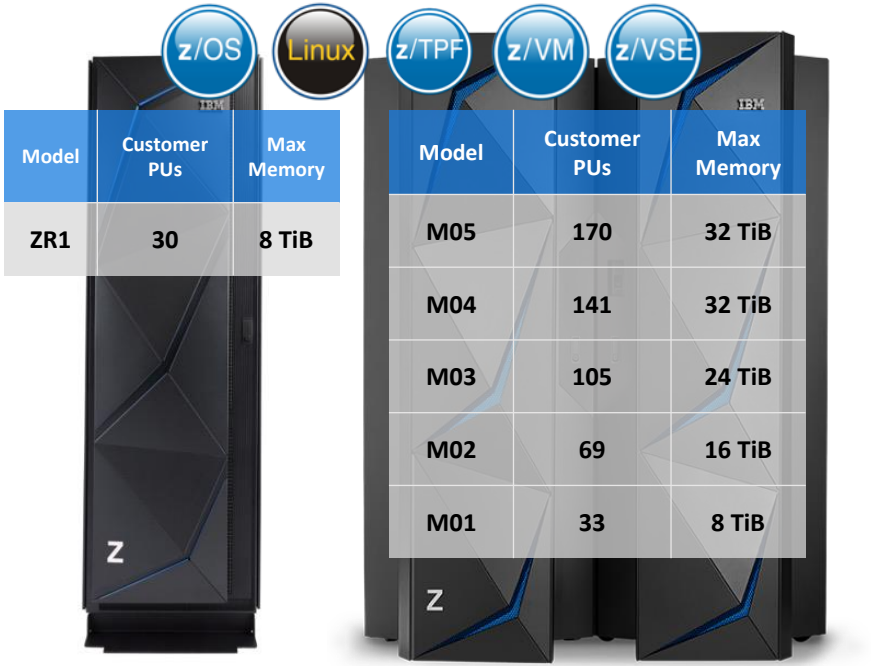

Model	Customer PUs	Max Memory
ZR1	30	8 TiB

Model	Customer PUs	Max Memory
M05	170	32 TiB
M04	141	32 TiB
M03	105	24 TiB
M02	69	16 TiB
M01	33	8 TiB

An IBM z14 for everyone

"Right size" your mainframe to fit your needs

IBM z14 Model ZR1 – if you …	IBM z14 Models M01-M05 – if you …
▪ Currently have an IBM z13s, zBC12, z114, z10 BC ▪ Have a requirement to scale up to 8,000 GCP MIPS for current and future growth ▪ Require 30 or less specialty processors – IFLs, zIIPs, ICFs ▪ Have smaller Coupling and/or I/O attachment requirements – including zHyperLink Express ▪ Have a data center strategy that is PDU-based with 200v-240v power with air cooling ▪ Want a 19" industry standard form factor allowing clients to lower power costs and have a 40% smaller footprint cost that fits freely in any data center ▪ Want to customize servers by adding storage, server or switch options in optional 16u of available frame space	▪ Currently have an IBM z13, zEC12, z196, z10EC ▪ Have a requirement for current and future growth for capacity that scales up to 146,000 GCP MIPS ▪ Have a large disk installment with large I/O requirements – including zHyperLink Express ▪ Have a data center strategy that is bulk power based on 480v with either air or water cooling ▪ Need new ways to address your 'green' requirements – i.e. water cooling ▪ Have a large Capacity Back Up (CBU) requirement – and like the control of having your disaster recovery site right in your own shop

IBM z14

Extending the IBM z14 family

Breakthrough technologies

Designed for the Secure Cloud

Jim Elliott
Mainframe Consultant

885 Don Mills Road, Suite 300, Toronto, ON M3C 1V9
Office: 416-229-2950 x341 Toll Free: 1-877-877-3757 x341
Mobile: 416-527-0666 jelliott@ghsystems.com **ghsystems.com**
jlelliotton @jlelliotton Jim.Elliott.Canada jlelliotton.blogspot.com