GLASSHOUSE SYSTEMS

IBM z14™ Technical Overview for CMG Canada

Jim Elliott
Mainframe Consultant
GlassHouse Systems Inc.

GLASSHOUSE SYSTEMS

Presentation date: 2017-10-24

Trademarks

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

BladeCenter*	DFSMShsm	HiperSockets	Ibm.com	LinuxONE Emperor	z13*	z/OS*
CICS*	ECKD	HyperSwap*	IBMZ*	PR/SM	z14	z/VM*
DB2*	FICON*	IBM*	Infiniband*	System Storage*	zEnterprise*	z/VSE*
DFSMSdss	GDPS*	IBM (logo)*	LinuxONE	WebSphere*	zHyperLink	

* Registered trademarks of IBM Corporation
Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.
ITIL is a Registered Trade Mark of AXELOS Limited.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
UNIX is a registered trademark of The Open Group in the United States and other countries.
VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.
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.

Agenda

- IBM z14 Overview
- IBM Z Pervasive Encryption
- IBM z14 Designed for Data Serving
- IBM Z Coupling Facility Enhancements
- IBM z14 Capacity and Performance Planning
- IBM z14 Designed for Simplification
- IBM z14 Summary

GlassHouse Systems Inc.

- **Background**
  - Established in 1993 as an enterprise IT infrastructure provider
  - HQ in Toronto, operating throughout Canada and the USA
  - Enterprise clients, public and private sector, cross industry
  - Preeminent IBM Business Partner in Canada
- **Value to clients**
  - Our commitment – to every customer success with long term relationships
  - Our team – senior technical and sales professionals with deep technical skills, and an understanding of IT industry directions and business drivers
  - Our technology partners – leading IT providers, committed and working with GHS to execute and deliver solutions that address specific challenges
  - Our experience and execution – assessment, design, implementation, on-going support and ease of use

GlassHouse Systems Inc. offerings

Systems	Software	Services
▪ IBM Z, IBM Power, IBM Storage ▪ Lenovo Enterprise Systems ▪ VMware Hybrid Cloud ▪ CISCO® Networking, UCS™, VersaStack™	▪ Security – SIEM, Patch Management, Data Security ▪ Analytics – Descriptive, Predictive, Cognitive ▪ Recovery – Software Defined Storage Solutions, Enterprise BC	▪ GHS ***insight***™ – Capacity & Performance Analysis; Solution Design; Configuration Development & Validation ▪ GHS ***advantage***™ – On-site Professional Services; SoW Driven and PM managed; Assessments, Migrations & Implementation ▪ GHS ***manage***™ – IBM AIX Cloud; IBM i Cloud; Linux Cloud; RIMS; Security Managed Services

IBM z14 Overview

IBM z14 – The cornerstone for a trust economy

- **Pervasive Encryption**
  - On-chip cryptographic coprocessor performance for ***6x*** faster encryption of data in-flight and at-rest
  - Crypto Express6S with ***2x*** performance increase
  - No application changes using DFSMS
  - Protection of database data according to secure policy with z/OS dataset encryption, CF encryption, z/VM encrypted hypervisor paging and z/TPF transparent database encryption
  - Secure deployment of software appliances including tamper protection with IBM Secure Service Container

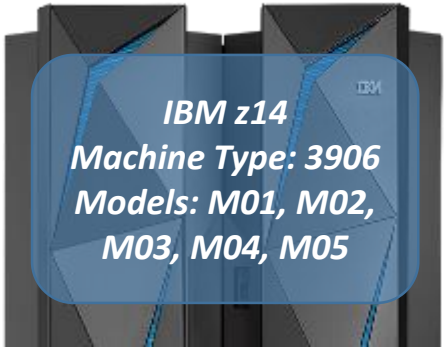

- **Intelligence and Insight**
  - Support data-in-memory applications and new workloads using ***3x*** (32 TB) memory
  - Enhanced math libraries with SIMD
  - Massive scale up of JVMs with pause-less garbage collection
  - zIIP SMT performance for Java-based analytics (e.g. Spark)

- **Open and connected**
  - Scaling improvements with 170 configurable cores
  - Improved virtualization for Linux with next generation SMT
  - Data serving enhanced with ***1.5x*** more on-core cache
  - Reduced latency for workloads with FICON Express16S+
  - Dramatic latency reduction with direct link between server and DS8880 using IBM zHyperLink
  - OSA-Express6S
  - Data center planning enhancements improve floor space and raise RAS
  - New PCIe based long range coupling links

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 TB Memory	**zEnterprise 196** (z196) ▪ Ann. 2010-07-22 ▪ 5.2 GHz ▪ Up to 80 cfg cores ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 3 TB Memory	**zEnterprise EC12** (zEC12) ▪ Ann. 2012-08-28 ▪ 5.5 GHz ▪ Up to 101 cfg cores ▪ CP, IFL, ICF, zAAP, zIIP ▪ Up to 3 TB Memory	**z13** ▪ Ann. 2015-01-14 ▪ 5.0 GHz ▪ Up to 141 cfg cores ▪ CP, IFL, ICF, zIIP ▪ Up to 10 TB Memory	**z14** ▪ Ann. 2017-07-17 ▪ 5.2 GHz ▪ Up to 170 cfg cores ▪ CP, IFL, ICF, zIIP ▪ Up to 32 TB Memory
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 GB Memory	**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 GB Memory	**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 GB Memory	**z13s** ▪ Ann. 2016-02-16 ▪ 4.3 GHz ▪ Up to 20 cfg cores (6 CP) ▪ CP, IFL, ICF, zIIP ▪ Up to 4 TB Memory	

IBM z14 – Differentiated value at the core

- **Leader in performance and scale**
  - New 10-core processor design in 14nm silicon technology
  - Up to 170 configurable cores for ***35%*** total capacity improvement over the largest IBM z13
  - ***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 z13

IBM z14 specifications

- **Processor Units (PUs)**
  - 41 (49 for M05) PU cores per processor drawer
  - 5 (6 for M05) chips per processor drawer
  - 33, 69, 105, 141, or 170 PU cores available for characterization
  - Up to 23 System Assist processors (SAP) per system (standard) plus SAPs are SMT
  - 1 Integrated Firmware Processor (IFP) for PCIe I/O
  - 2 spares designated per system
  - 85 Logical Partitions (LPAR)
  - Sub-capacity available for up to 33 CPs – 3 sub-capacity points
  - Enhanced performance for compression and crypto coprocessor
- **Memory**
  - RAIM memory design
  - System minimum of 128 GB – up to 8 TB / drawer
  - 192 GB fixed HSA from 1st processor drawer
  - Up to 32 TB for system and up to 16 TB per LPAR
    - z/OS supports up to 4 TB, Linux on Z supports up to 16 TB, z/VM supports up to 2 TB
  - IBM Virtual Flash Memory (replaces Flash Express)
- **I/O**
  - New PCIe Gen 3 IBM zHyperLink™ technology
  - Last high-end server to support InfiniBand® coupling features
  - 16 GBps PCIe Gen 3 I/O Interconnects
  - 6 Logical Channel Subsystems (LCSSs) with 4 Sub-channel sets per LCSS

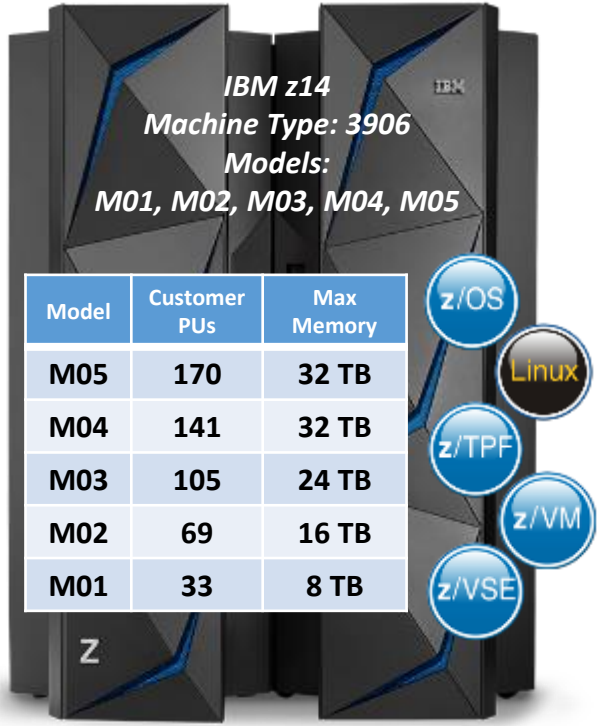

Model	Customer PUs	Max Memory
M05	170	32 TB
M04	141	32 TB
M03	105	24 TB
M02	69	16 TB
M01	33	8 TB

IBM Z – Processor roadmap

65 nm	45 nm	32 nm	22 nm	14 nm
z10	**z196**	**zEC12**	**z13**	**z14**
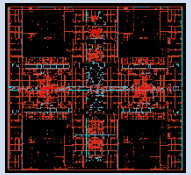	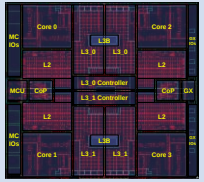	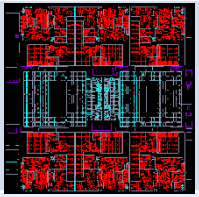	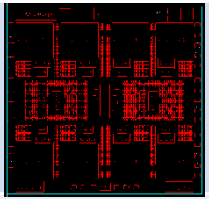	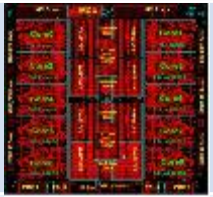
▪ Workload consolidation and integration engine for CPU intensive workloads ▪ Decimal FP ▪ InfiniBand ▪ 64-CP image ▪ Large pages ▪ Shared memory	▪ Top tier single thread performance, system capacity ▪ Accelerator integration ▪ Out of order execution ▪ Water cooling ▪ PCIe I/O fabric ▪ RAIM ▪ Enhanced energy management	▪ Leadership single thread, enhanced throughput ▪ Improved out-of-order ▪ Transactional memory ▪ Dynamic optimization ▪ 2 GB 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

IBM z14 continues the CMOS mainframe heritage

- z14 ~10% more than equal z13 n-way
- Up to 35% max capacity 170-way (z14) vs. 141-way (z13)
- SMT vs. Single Thread ~10–40% (average 25%)
- SMT z14 vs. z13 ~ 15% (z/VM Guests)

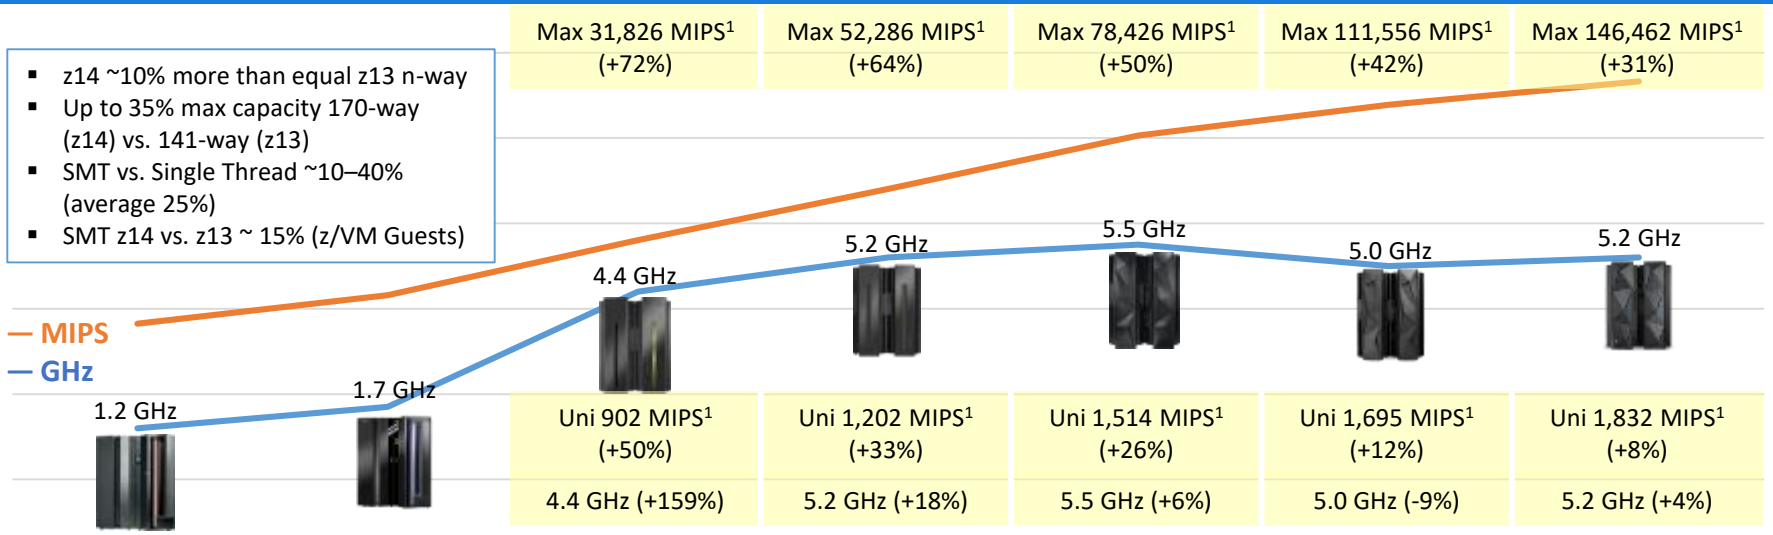

		Max 31,826 MIPS[1] (+72%)	Max 52,286 MIPS[1] (+64%)	Max 78,426 MIPS[1] (+50%)	Max 111,556 MIPS[1] (+42%)	Max 146,462 MIPS[1] (+31%)
		Uni 902 MIPS[1] (+50%)	Uni 1,202 MIPS[1] (+33%)	Uni 1,514 MIPS[1] (+26%)	Uni 1,695 MIPS[1] (+12%)	Uni 1,832 MIPS[1] (+8%)
		4.4 GHz (+159%)	5.2 GHz (+18%)	5.5 GHz (+6%)	5.0 GHz (-9%)	5.2 GHz (+4%)
2003 – z990	**2005 – z9 EC**	**2008 – z10 EC**	**2010 – z196**	**2012 – zEC12**	**2015 – z13**	**2017 – z14**
▪ 130 nm SOI ▪ 32 Cores[2] (42) ▪ Superscalar ▪ Modular SMP	▪ 90 nm SOI ▪ 54 Cores[2] (64) ▪ System level scaling	▪ 65 nm SOI ▪ 64 Cores[2] (77) ▪ High-freq. core ▪ 3-level cache	▪ 45 nm SOI ▪ 80 Cores[2] (96) ▪ OOO core ▪ eDRAM cache ▪ RAIM memory	▪ 32 nm SOI ▪ 101 Cores[2] (120) ▪ OOO and eDRAM cache improvements ▪ PCIe Flash ▪ Arch. extensions for scaling	▪ 22 nm SOI ▪ 141 Cores[2] (168) ▪ SMT and SIMD ▪ Up to 10TB of Memory	▪ 14 nm SOI ▪ 170 Cores[2] (196) ▪ Enh. SMT and SIMD ▪ Up to 32 TB of Memory

1. MIPS are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
2. Number of PU cores for customer use (total cores)

IBM Z – Systems designed to meet workload requirements

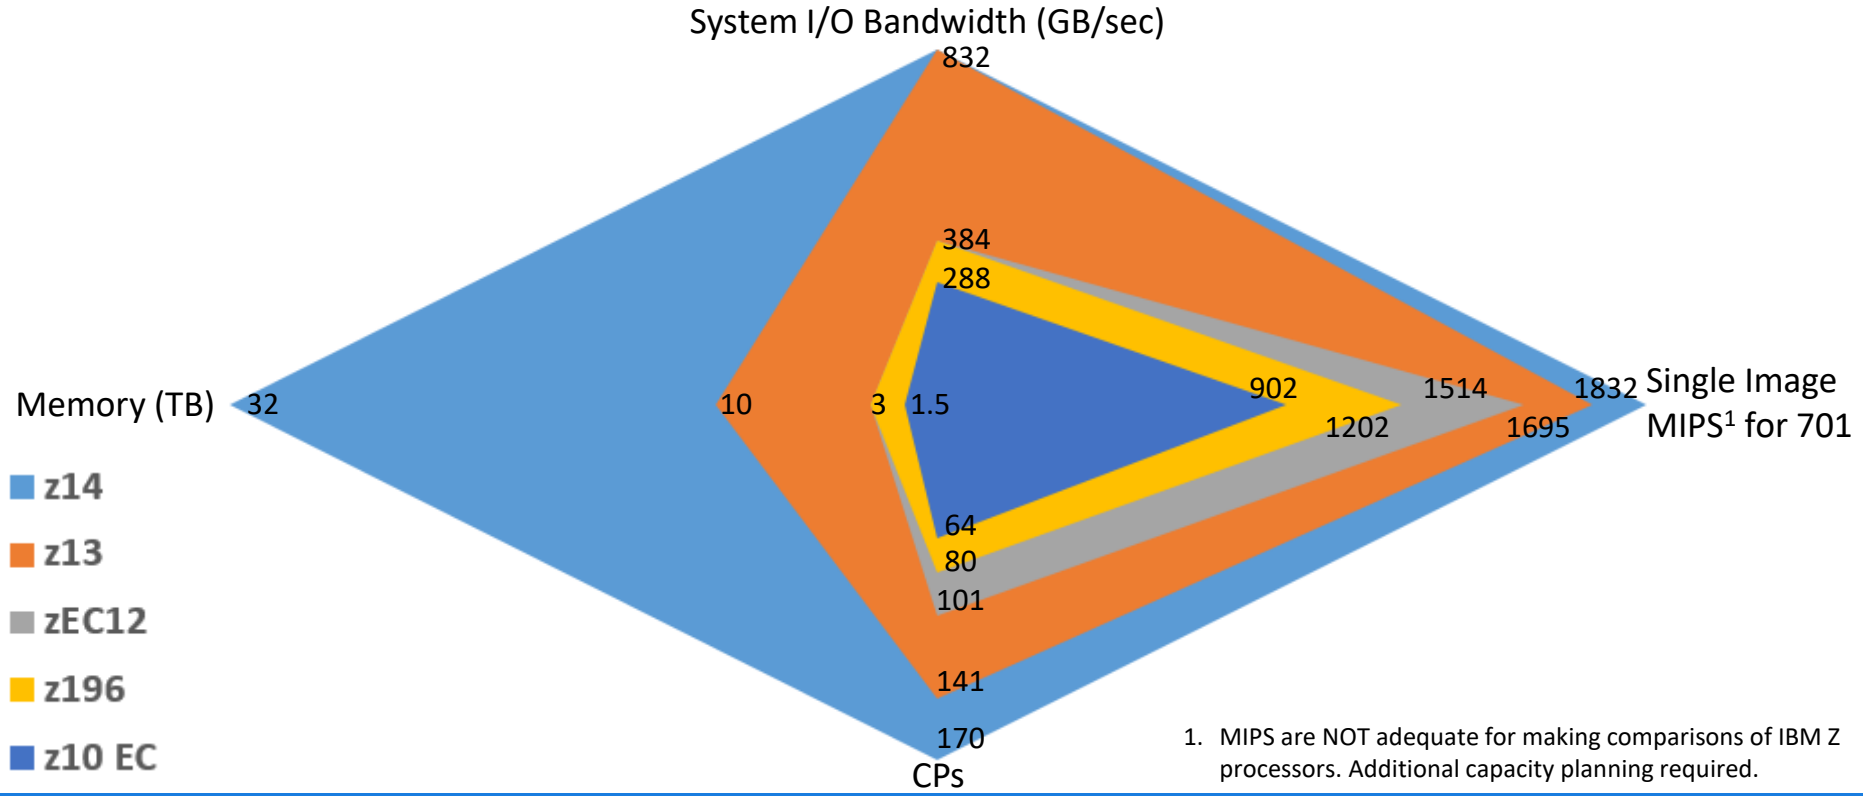

1. MIPS are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required.

IBM z14 processor drawer

- **Each PU SCM:**
  - 14nm
  - One memory controller per CP chip
  - Five DDR4 DIMM slots per memory controller: 15 total per logical cluster
- **Each drawer:**
  - Two logical CP clusters (0 and 1)
  - Five PU chips: 41 active PUs – M01 – M04
  - Six PU chips: 49 Active PUs – M05
  - One SC chip (672 MB L4 cache)
  - Populated DIMM slots: 25 DIMMs to support up to 8 TB of addressable memory (10 TB RAIM)
  - Water cooling for PU SCMs, air cooled SC SCM
  - Two Flexible Support Processors (FSP)
  - Ten fanout slots for PCIe I/O drawer fanouts or PCIe coupling fanouts
  - Four fanout slots for PSIFB coupling link fanouts

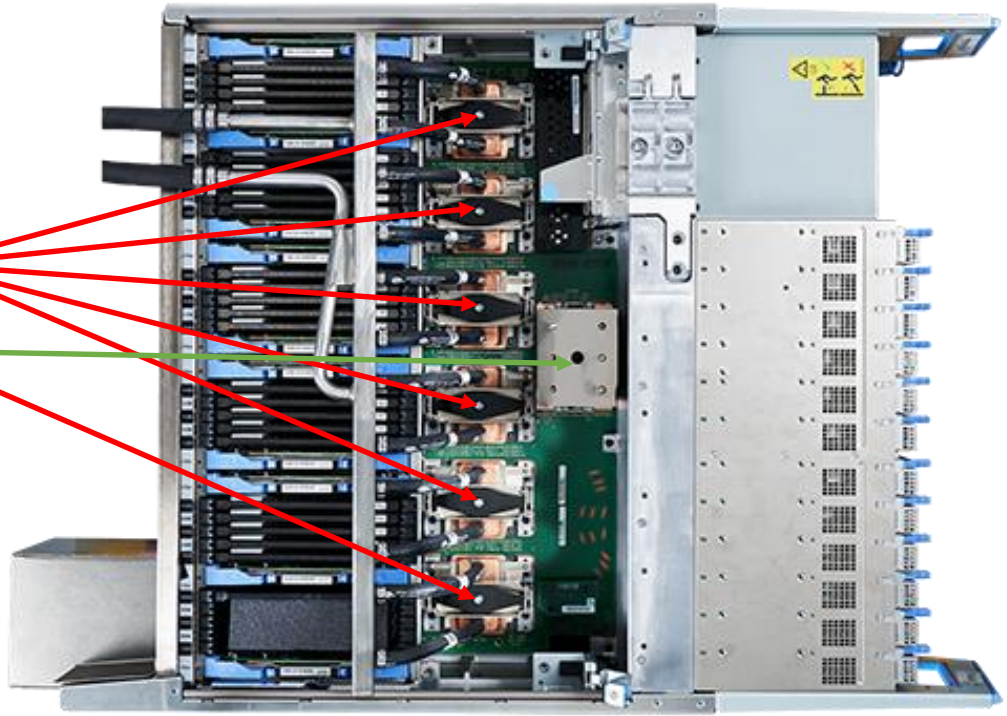

IBM z14 processor drawer

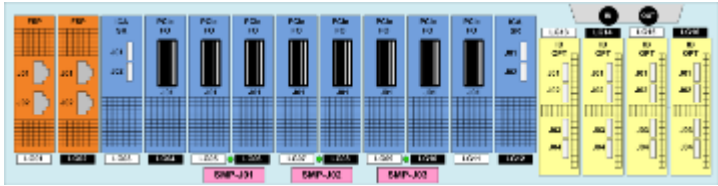

- **The z14 can support up to five PCIe I/O drawers**
- **Flexible Support Processor (FSP):**
  - Two per CPC drawer (LG01, LG02)
- **PCIe Fanout Slots:**
  - Total of 10 per CPC Drawer (LG03 … LG12)
  - Used for either connectivity of CPC Drawer to PCIe I/O Drawer or ICA SR coupling links
  - For every ICA SR configured, the number of PCIe fanouts available for the PCIe I/O drawer are reduced by the number of ICA SRs configured
    - Only Models M04/M05 can support the maximum numbers of PCIe I/O drawers and ICA SRs
    - For all other model (M01, M02, M03), the number of PCIe I/O drawers supported will vary with the number of ICA SRs configured (this may result in ‘half’ drawer configurations)
- **IFB Fanout Slots:**
  - Total of four per CPC Drawer (LG13 … LG16)
  - Used for HCA3-O (12x) or HCA3-O LR (1x) PSIFB coupling links

IBM z14 processor chip design

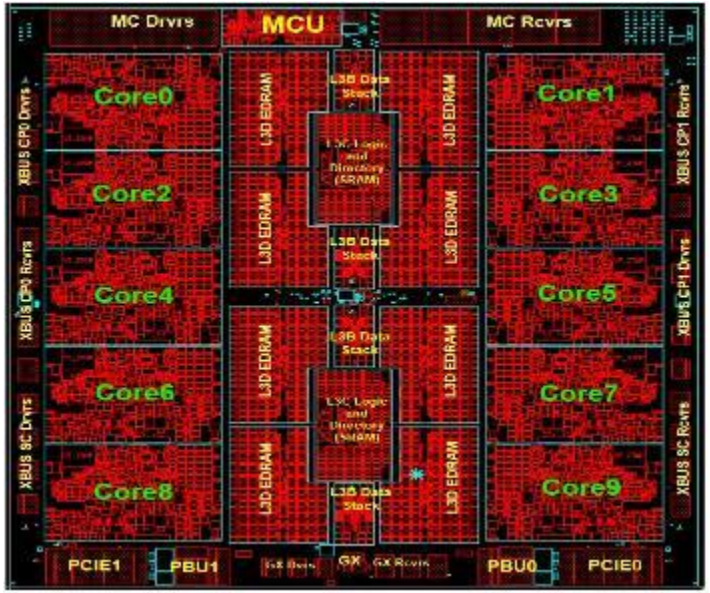

- **14nm SOI technology**
  - 17 layers of metal
  - 10 cores per CP-chip
- **Chip area**
  - 26.5 x 27.8 mm

- **Up to ten active cores (PUs) per chip**
  - 5.2 GHz (versus 5.0 GHz with the IBM z13®)
  - 33% larger L1 Instruction Cache (128KB) on core
  - 2x larger L2 Data Cache (4MB) on core
  - 2x larger L3 Cache on chip with symbol ECC – shared by on-chip cores and communicates with cores, memory, I/O and system controller single chip module
- **New instructions for Single Instruction/Multiple Data (SIMD)**
- **Single thread or next generation 2-way simultaneous multithreading (SMT) operation**
- **Guarded Storage Facility (GSF)**
- **Compression enhancements – Huffman Coding – in future[1] DB2® indices can take advantage of the compression coprocessor**
- **Improved instruction execution bandwidth:**
  - Pipeline optimization – Improved instruction delivery, faster branch wakeup, optimized hardware/millicode interaction, reduced execution latency, improved OSC prediction
  - Translator / TLB enhancements
  - New instructions for old and new workloads – i.e. Vector BCD arithmetic (COBOL)
- **I/O buses**
  - One InfiniBand I/O bus
  - Two PCIe I/O buses
- **Memory Controller (MCU)**
  - Interface to controller on memory DIMMs, supports RAIM design

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

IBM z14 system controller chip design

- SC chip area: 25.3 x 27.5 mm
- 14nm SOI technology, 17 layers of metal
- 672 MB shared eDRAM L4 Cache
- System Interconnect
- System Coherency Manager
- X and A Bus

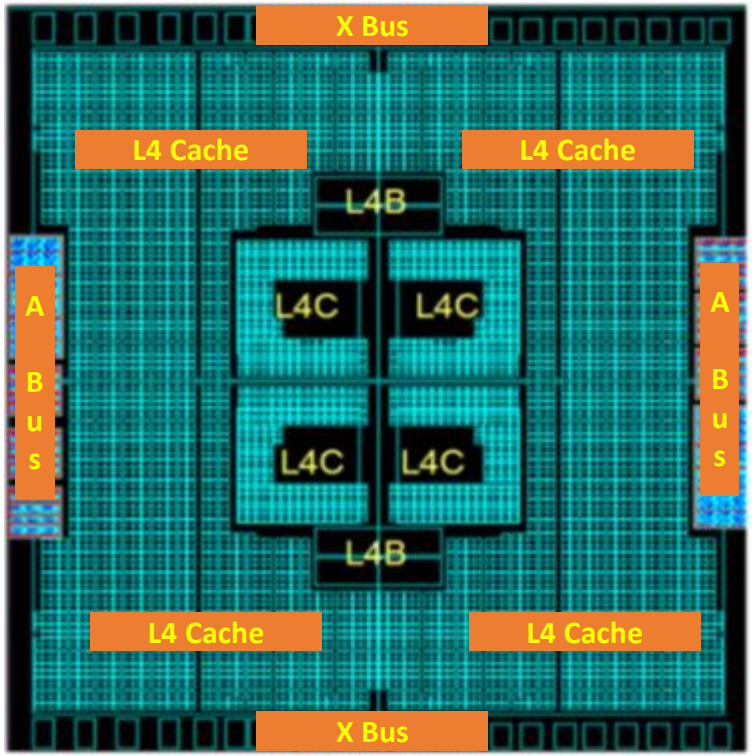

z14 processor allocation / usage

Model	CPs	IFLs	Unassigned IFLs	zIIPs	ICFs	IFPs	Standard SAPs	Additional SAPs	Spares
M01	0-33	0-33	0-32	0-22	0-33	1	5	0-4	2
M02	0-69	0-69	0-68	0-46	0-69	1	10	0-8	2
M03	0-105	0-105	0-104	0-70	0-105	1	15	0-12	2
M04	0-141	0-141	0-140	0-94	0-141	1	20	0-16	2
M05	0-170	0-170	0-169	0-112	0-170	1	23	0-16	2

- **z14 Models M01 to M04 use drawers with 41 cores. The Model M05 has 4 drawers with 49 cores.**
- **The maximum number of logical ICFs or logical CPs supported in a CF logical partition is 16**
- **The integrated firmware processor (IFP) is used for PCIe I/O support functions**
- **Concurrent Drawer Add is available to upgrade in steps from model M01 to model M04**
- **Field upgrade to model M05 NOT supported (M05 is Factory built only)**

1. **At least one CP, IFL, or ICF must be purchased in every machine**
2. **Two zIIPs may be purchased for each CP purchased if PUs are available. This remains true for sub-capacity CPs and for "banked" CPs.**
3. **On an upgrade from zEC12, installed zAAPs are converted to zIIPs by default. (Option: Convert to another engine type)**
4. **The IFP is conceptually an additional, special purpose SAP**

IBM z14 memory for the digital enterprise

- Up to 32TB of memory to support new workloads, data-in-memory applications, efficiently process huge amounts of information for real-time business insights
- In memory databases are critical to enable faster insightful analysis by enabling correlations and other analyses not otherwise made possible
- Large memory can aid compression by providing large buffers to stage processing
- Shift in new types of applications, that perform random access on data versus sequential access, require more data in memory to sustain SLAs
- Large VM consolidations can use memory to provide a higher ceiling for vertical scale needs

IBM z14 purchased memory offering ranges

Model	Standard Memory GB	Flexible Memory GB
M01	320 - 8000	NA
M02	320 - 16192	320 - 8000
M03	320 - 24384	320 - 16192
M04	320 - 32576	320 - 24384
M05	320 - 32576	320 - 24384

- **Purchased Memory – Memory available for assignment to LPARs**
- **Hardware System Area – Standard 192 GB of addressable memory for system use outside customer memory**
- **Standard Memory – Provides minimum physical memory required to hold customer purchased memory (minimum of 128 GB) plus 192 GB HSA**
- **Flexible Memory – Provides additional physical memory needed to support activation base customer memory and HSA on a multiple CPC drawer z14 with one drawer out of service.**
- **Plan Ahead Memory – Provides additional physical memory needed for a concurrent upgrade (LIC CC change only) to a preplanned target customer memory**

IBM z14 memory granularity

- **Memory granularity for assignment to an LPAR and for configuration On and Off is the LPAR's Increment size**
  - z14 physical Increment size is fixed at 1 GB (increased from 512 MB on z13)
  - Central memory granularity is virtualized as a multiple of 1024 MB (1 GB) for each LPAR based on the size of the larger of its two Elements: Initial Central and Reserved Central
  - Increment size increases with Element size because operating systems that support memory reconfiguration (z/OS and z/VM) support only up to 512 Increments per Element
- **Action: Review MVS RSU parameter. Change RSU (if non-zero) to MB or GB values, not as a number of Increments to avoid problems if LPAR memory grows. A MB or GB value will be rounded up if not equal to the partition's Increment size.**

Note: For z/OS V2R3 with z14, a minimum of processor storage of 8GB for z/OS LPAR and 2GB for z/OS on z/VM is recommended.

IBM Virtual Flash Memory cuts away at availability lapses

- Next generation of Flash Express to provide higher levels of availability and performance
- Moved to RAIM storage – doesn't require PCIe slots
- Slashes latency for critical application processing such as diagnostics collection

Typical Client Use Cases:

- Improve availability and performance during workload transition and spikes
- Faster, less disruptive diagnostics with faster first failure data capture time
- Less paging with use of pageable large pages for Java or DB2
- Cost effective, resilient solution for overflow of MQ shared queues in Coupling Facility

IBM Z Virtual Flash Memory

- **The "storage class memory" provided by Flash Express (FC 0402 and FC 0403) adapters is replaced with Virtual Flash Memory (VFM) which is part of the main memory**
- **VFM is offered as a priced hardware feature**
  - Customer can buy one to four "units" of VFM initial purchase
- **A "unit" will be 1.5 TB (1536 GB) on z14**
  - Approximately same size as a Flash Express pair of adapters
- **Much simpler management of VFM resource (HMC task)**
- **No hardware repair and verify (no cables, no adapters)**
- **Better performance since no "I/O" to attached adapter takes place**
- **RAS: Memory protected by RAIM and ECC (internal / main memory)**

Note: Use cases and exploitation for VFM have not been changed (e.g. z/OS paging, CF shared queue overflow), they just transparently benefit from the changes in the hardware implementation.

Quick summary of physical planning considerations

- Floor space – No change unless ordering Thin Covers
- Overhead I/O or Power – No change
- Power – No change to typical power consumption
- Environment – New ASHREA A3 Classification
- Customer Water – No change
- Weight – Sight increase in weight depending on configuration
- Airflow – No change
- New feature – Thin Covers

Max power for z14 (vs. z13)

- **z14 max power slightly higher, at least for max configuration**
  - 5.2 Ghz vs 5.0 Ghz
  - 48 additional physical cores on max system
  - max memory = 32TB vs 10TB
- **Estimated maximums:**
  - z13 – 28.3kW
  - z14 – 30kW
- **Typical configurations will have approximately the same power as the equivalent z13**
- **Customer chilled water saves approximately 2KW (6.6%) on a maximum configuration system**
- **Nameplate rating unchanged:**

– 200V 3 phase AC 50/60hz	50 amps
– 208-240V 3 phase AC 50/60hz	48 amps
– 380-480V 3 phase AC 50/60hz	26 amps
– 480V 3 phase AC 60hz	21 amps
– 380-520V DC	44 amps

z14 dimensions and weights

- Base overall system dimensions are identical to z13
- z14 has a thin cover option (non-acoustic) which will shorten the system depth by approximately 15.5" and reduce the system weight by 108 lbs. when this option is selected
- System weights for z14 are close to the same as for z13

z14 system weights (lbs.)

Model	Min.	Typ.	Max.	Water add	Top exit I/O add	Battery feature add	Balanced power add	Earthquake feature add	Thin door add
M01	2772	3378	3677	127	156	447*	239	180	-108
M02	2951	3885	4742	127	156	671	109	180	-108
M03	3228	4549	5020	104	156	671	55	180	-108
M04	3486	4800	5248	104	156	671	0	180	-108
M05	3499	5126	5290	104	156	671	0	180	-108

IBM Z Pervasive Encryption

Designed for data protection

- **Market view**
  - Nearly 4 million records stolen per day, 157,364 per hour, and 2,623 per minute[1]
  - Of the 9 billion records breached since 2013 only 4% were encrypted[1]
  - 1 in 4 companies is likely to experience a breach[2]
  - The greatest security mistake organizations make is failing to protect their networks and data from internal threats[3]

The world's premier system for enabling data as the new security perimeter

- **Pervasive encryption**
- **No application changes**
- **Protect from internal and external threats**

1. http://breachlevelindex.com/
2. https://www.ibm.com/security/infographics/data-breach/
3. Steve Marsh in article: https://digitalguardian.com/blog/expert-guide-securing-sensitive-data-34-experts-reveal-biggest-mistakes-companies-make-data

The IBM z14 makes pervasive encryption achievable

- **Pervasive encryption gives IBM 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 of your data and transactional pipeline helps reduce potential data breach risks and financial losses**

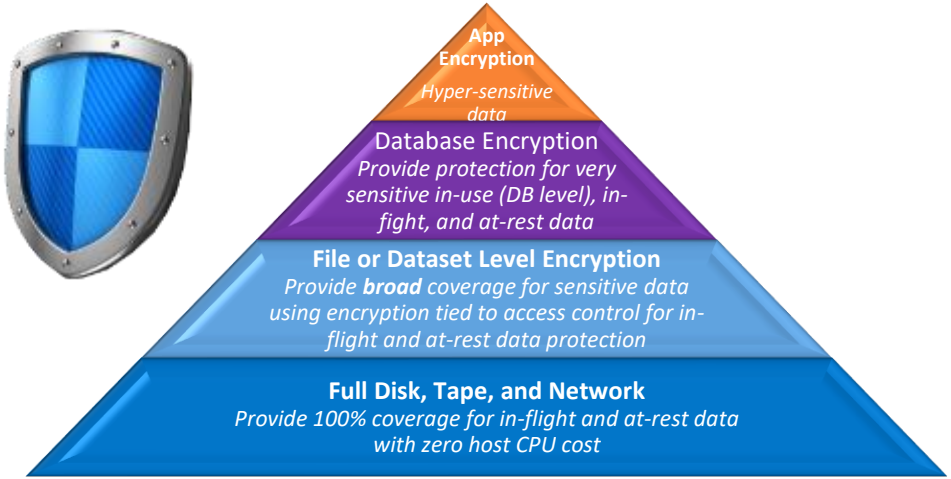

IBM z14: Performance that changes the game for security

- **Performance with integrated cryptographic hardware**
  - ***6x*** faster encryption of data in-flight and at-rest with enhanced on-chip cryptographic performance compared to z13
  - ***2x*** the SSL handshake performance on z14 with Crypto Express6S compared to z13 with Crypto Express5S[1]
- **Datasets automatically protected with z/OS Dataset Encryption**
  - Protect z/OS data sets[2] and zFS file systems[3] 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)[3]

1. Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13. Results may vary.
2. On October 4th, 2016 IBM announced a Statement of Direction to deliver z/OS dataset encryption capability in z/OS V2.2 (Announcement Letter)
3. IBM z/OS 2.3 Announcement Letter

Pervasive encryption with Linux on z14

- **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
- **Linux[1] expects to get enhanced security with the "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
- **Linux[1] is expected to be enabled to create true unique cryptographic data using the new true random number generator with CPACF**

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

More security enhancements on IBM z14

- *Ensure stronger security in the payment card industry* using Crypto Express6S compliance with security standards
- *Stronger cryptographic computation* using True Random Number generator support on CPACF
- *Performance boost for Java* with new Galois Counter Mode (GCM) encryption for minimum latency and operation overhead
- *New audit log application* and other performance improvements on TKE 9.0
- *Audit network encryption attributes* within z/OS network traffic using new z/OS Encryption Readiness Technology (zERT)[1] tool

1. zERT does not require a z14

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)

IBM z14 Designed for Data Serving

Designed for data serving

- **Market view**
  - Cognitive businesses uses insight from all data to enhance their digital intelligence and disrupt industries
  - Over 90% of the world's data cannot be Googled[1]
  - Accelerating time to market is essential – success is measure in days, not weeks
  - Data gravity – analyze data where it resides

Designed for data serving in a cognitive world

- **Speed, scale and reduced latency**
- **Efficiency for managing data**
- **Secure and flexible access to data**

1. Source: http://www.tennessean.com/story/money/tech/2014/05/02/jj-rosen-popular-search-engines-skim-surface/8636081/

I/O subsystem internal bus interconnect speeds

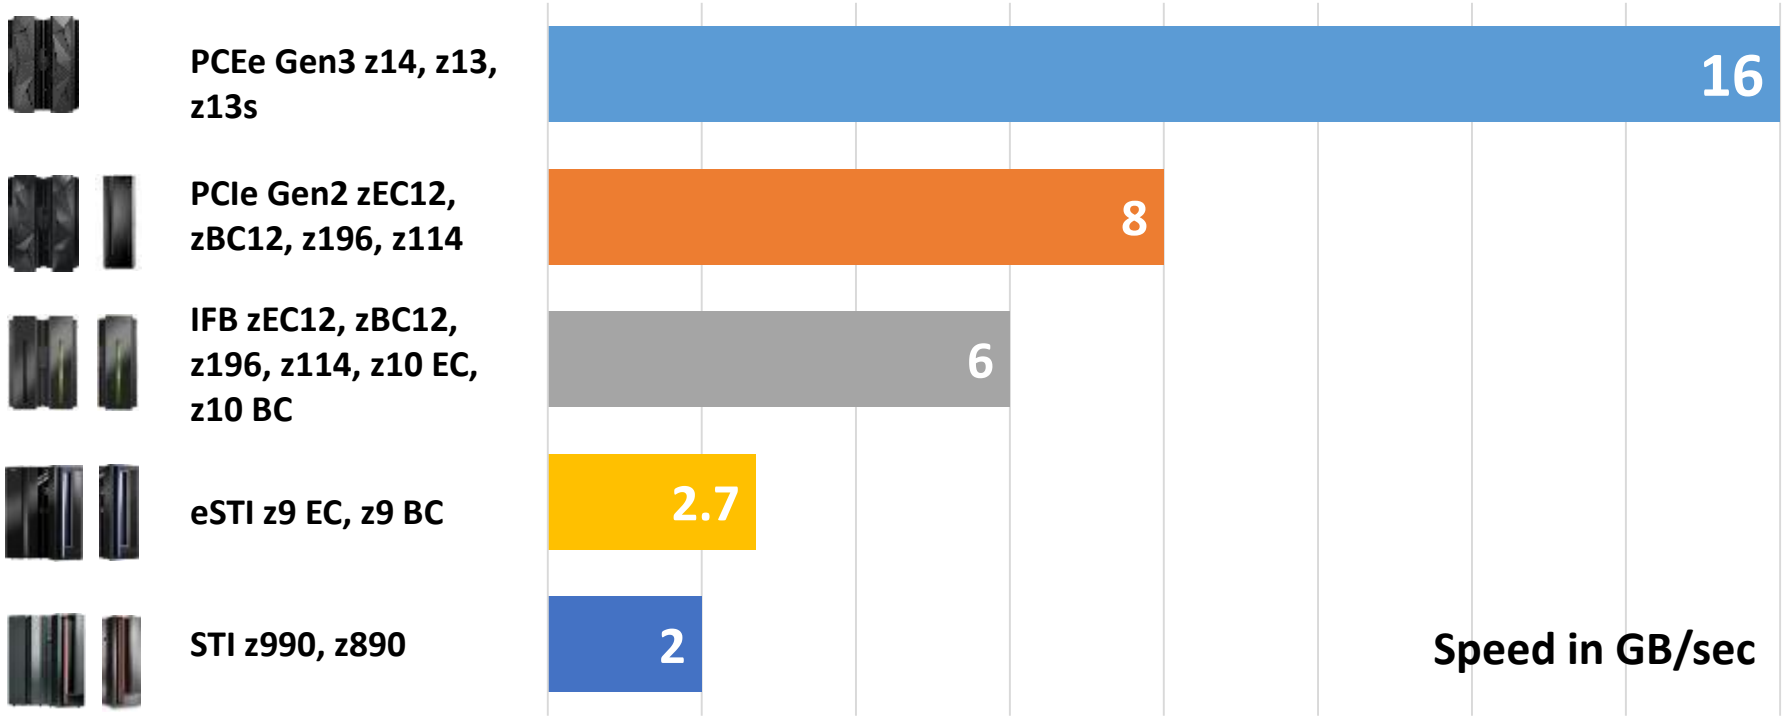

IBM zHyperLink Express

Speed matters: Breakthrough I/O link technology

- **A new direct connect short distance link designed to deliver low latency connectivity between z14 and FICON storage systems**
- **zHyperLink improves application response time, cutting I/O sensitive workload response time by *up to 50%* without requiring application changes[1]**

Typical Client 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

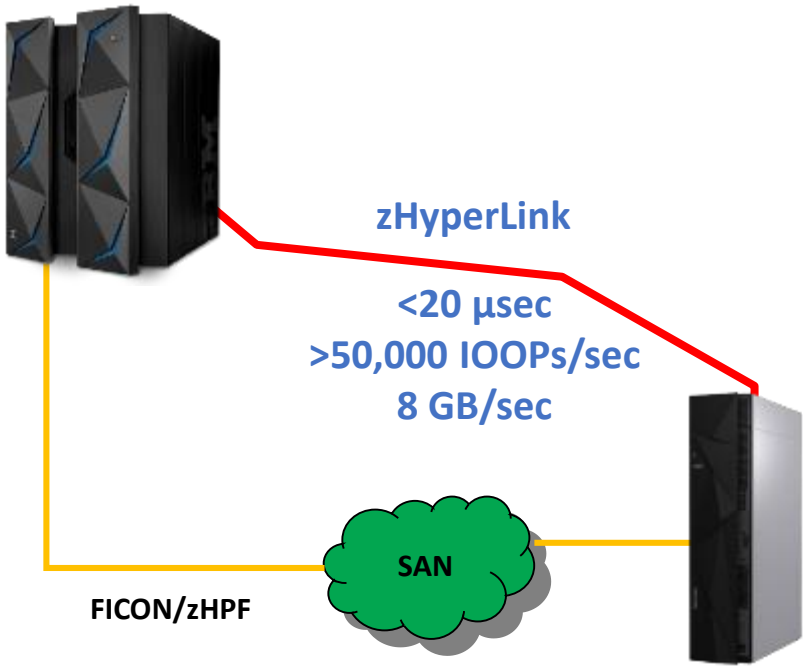

1. z/OS 2.1, 2.2, or 2.3
2. SOD: IBM intends to deliver VSAM exploitation of z14 and DS8880 zHyperLink Express.
3. Db2 for z/OS 11 or 12

zHyperLink Express at a glance

- **zHyperLink Express feature (FC 0431)**
  - 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 models 984, 985, 986 and 988[1]**

1. http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=ca&infotype=an&supplier=897&letternum=ENUS117-005

I/O performance evolution to IBM zHyperLink with DS8886

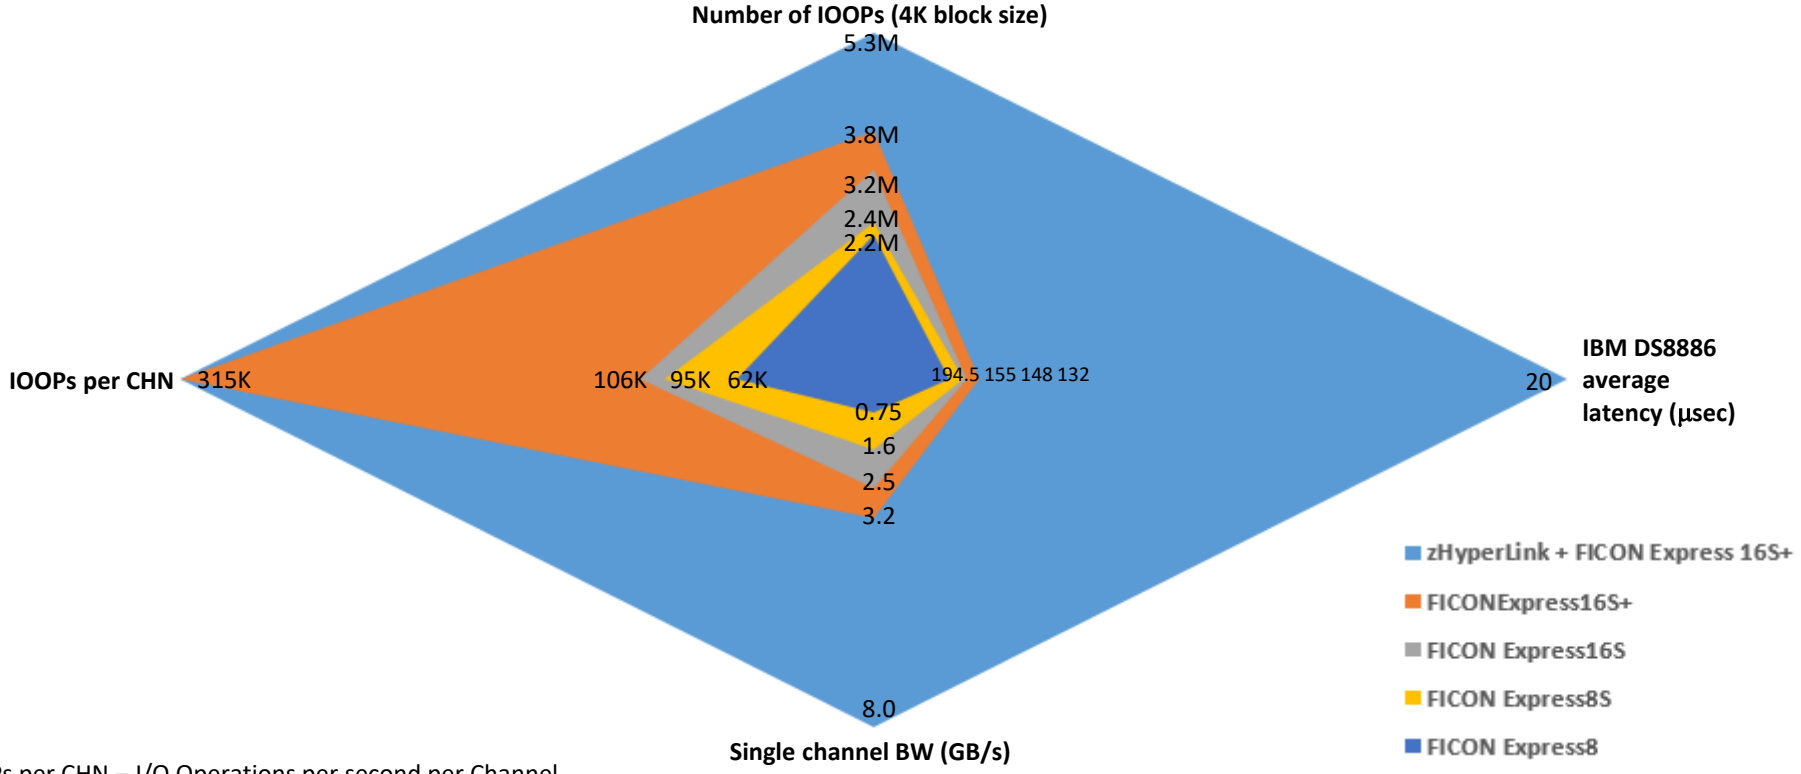

IOOPs per CHN – I/O Operations per second per Channel

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 1.2GB 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.

Shared memory architecture

Faster communications that preserve TCP/IP qualities of service

- **Shared Memory Communications – Direct Memory Access (SMC-D) optimizes z/OS for improved performance in '*within-the-box*' communications versus standard TCP/IP over HiperSockets or Open System Adapter**

Typical Client Use Cases:

- **Valuable for multi-tiered work co-located onto a single IBM Z server without requiring extra hardware**
- **Any z/OS TCP sockets based workload can seamlessly use SMC-D without requiring any application changes**
- **With z/VM 6.3 guest exploitation, you can understand the value for your z/OS SMC-R and SMC-D workloads before going into production**

SMC Applicability Tool (SMCAT) is available to assist in gaining additional insight into the applicability of SMC-D (and SMC-R) for your environment

- *Up to **61%** CPU savings for FTP file transfers across z/OS systems versus HiperSockets[1]*
- *Up to **9x** improvement in throughput with more than a **88%** decrease in CPU consumption and a **90%** decrease in response time for streaming workloads versus using HiperSockets[1]*
- *Up to **91%** improvement in throughput and up to **48%** improvement in response time for interactive workloads versus using 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.

Optimize server to server networking – transparently

Exploitation of RDMA over Converged Ethernet (RoCE)

- **Designed to take advantage of high speed protocols and direct memory placement of data for faster communications**
- **Increased sharing / virtualization with new 10Gbe RoCE Express2 feature**

Typical Client Use Cases:

- **Helps to reduce both latency and CPU resource consumption over traditional TCP/IP for communications across z/OS systems**
- **Any z/OS TCP sockets based workload can seamlessly use SMC-R without requiring any application changes**
- **With z/VM 6.3 or higher guest exploitation, you can understand the value for your z/OS workloads before going into production**

- *Up to **50%** CPU savings for FTP file transfers across z/OS systems versus standard TCP/IP[1]*
- *Up to **48%** reduction in response time and **10%** CPU savings for a sample CICS® workload exploiting IPIC using SMC-R versus TCP/IP[2]*
- *Up to **40%** reduction in overall transaction response time for WAS workload accessing z/OS DB2[3]*
- *Up to **3X** increase in IBM MQ messages delivered across z/OS systems*

1. Based on internal IBM benchmarks in a controlled environment using z/OS V2R1 Communications Server FTP client and FTP server, transferring a 1.2GB 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. Based on internal IBM benchmarks using a modeled CICS workload driving a CICS transaction that performs 5 DPL (Distributed Program Link) calls to a CICS region on a remote z/OS system via CICS IP interconnectivity (IPIC), using 32K input/output containers. Response times and CPU savings measured on z/OS system initiating the DPL calls. The actual response times and CPU savings any user will experience will vary.
3. Based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.
4. Based on internal IBM benchmarks using a modeled WebSphere MQ for z/OS workload driving non-persistent messages across z/OS systems in a request/response pattern. The benchmarks included various data sizes and number of channel pairs The actual throughput and CPU savings users will experience may vary based on the user workload and configuration.

Shared Memory Communications over RDMA (SMC-R) concepts

Clustered Systems

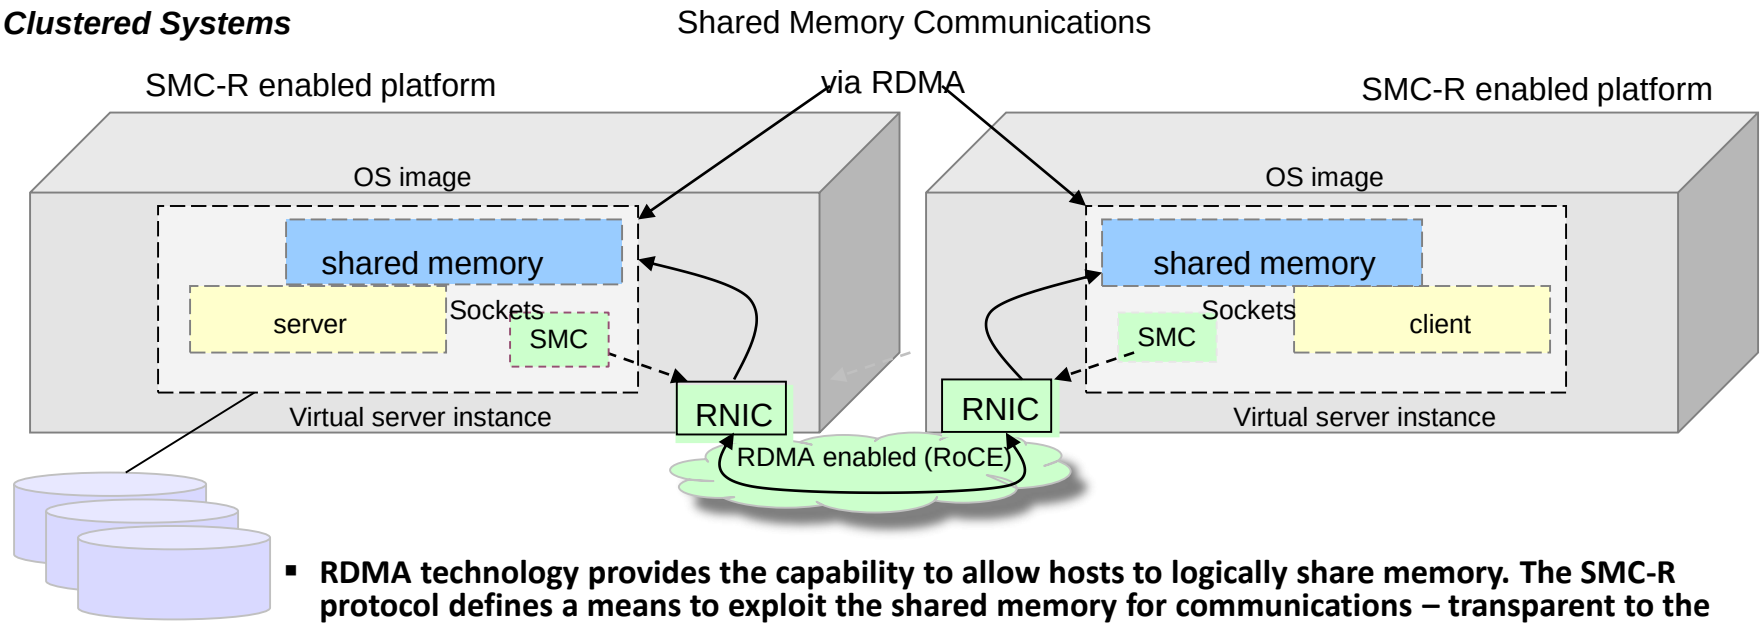

- **RDMA technology provides the capability to allow hosts to logically share memory. The SMC-R protocol defines a means to exploit the shared memory for communications – transparent to the applications! SMC-R RFC: https://datatracker.ietf.org/doc/rfc7609/**
- **This solution is referred to as SMC-R (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)**

Other RoCE considerations

- **Linux: Linux can exploit RoCE Express2 as a standard NIC (Network Interface Card) for Ethernet**
  - A specific Linux distribution level is required (reference PSP bucket for additional details).
    - SLES 11 SP4, SLES 12 SP2
    - RHEL 6.8, RHEL 7.3
    - Ubuntu 16.04 (+ additional patches)
  - Note: Linux does not currently support SMC-R.
- **Configuration or deployment issues that should be considered**
  - In addition to the existing RoCE Express hardware installation procedures, when the FID is configured in HCD the RoCE Express2 port number is also required
  - Port number must be specified, there is no default
  - RoCE Express2 will support a greater number of Virtual Functions per physical port (63) which will benefit the Linux shared RoCE environment
- **z/VM guest support for both SMC-R and SMC-D**

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)**

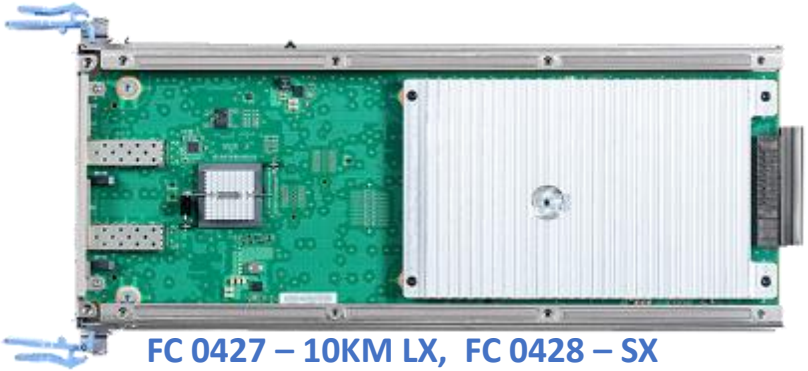

FC 0427 – 10KM LX, FC 0428 – SX

LX/LX or SX/SX

zHPF and FICON performance[1] on IBM Z

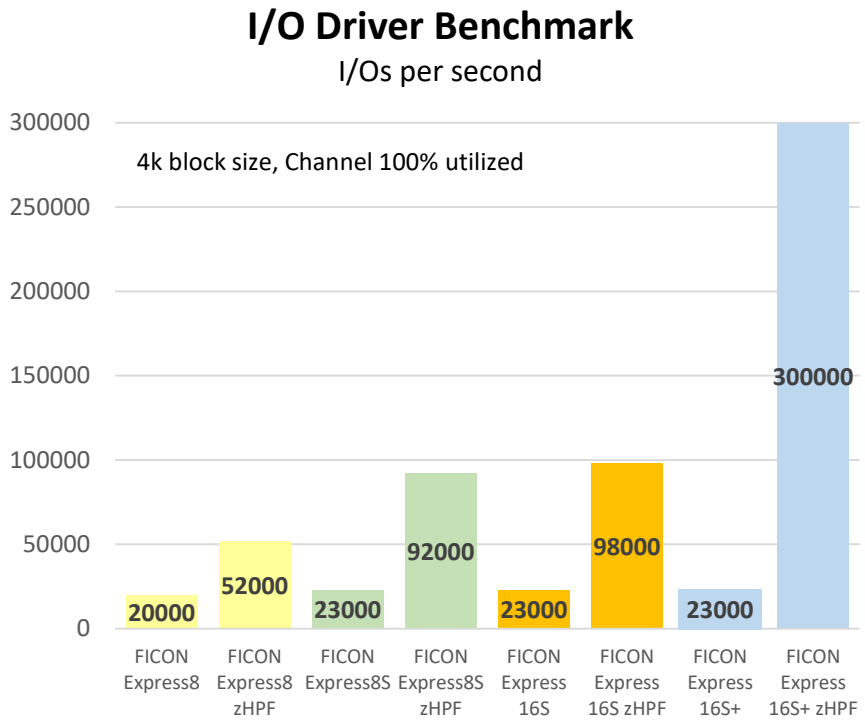

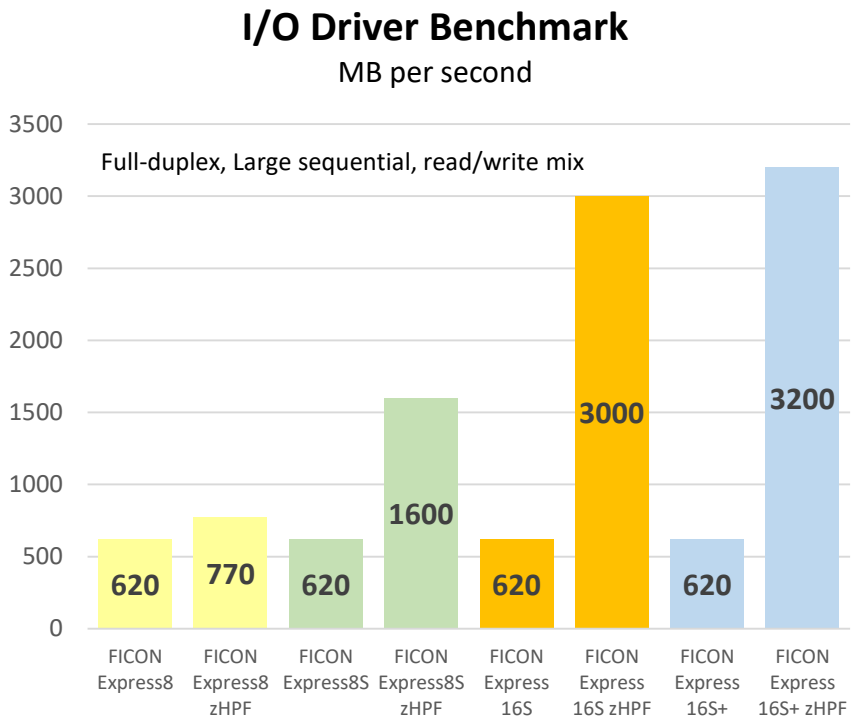

1. This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration,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

FCP performance[1] on IBM Z

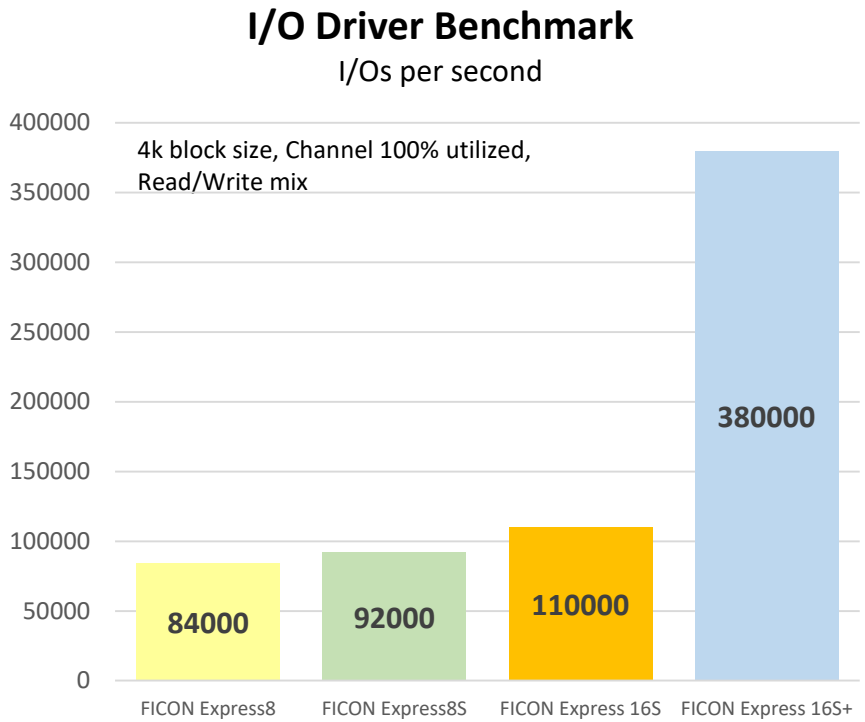

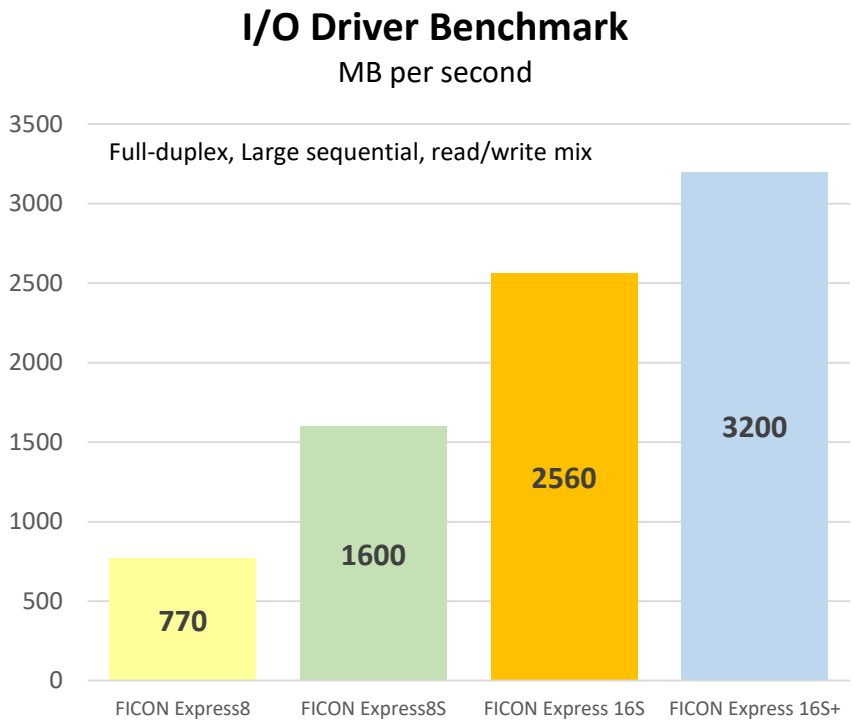

1. This performance was measured in a controlled environment running an I/O driver program. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, 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
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN)

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.

z14 new build I/O features supported

- **Features – PCIe I/O drawer (FC 4013) provides 32 I/O slots**
  - FICON Express16S+: LX (FC 0427), SX (FC 0428)
  - OSA-Express6S: 1 GbE LX (FC 0422), 1 GbE SX (FC0423), 10 GbE LR (FC 0424), 10 GbE SR (FC 0425), and 1000BASE-T (FC 0426)
  - 10GbE RoCE Express2 (FC 0412)
  - zEDC Express (FC 0420)
  - Crypto Express6S (FC 0893)
  - zHyperLink Express (FC 0431)
  - Coupling Express LR (FC 0433)
- **PCIe Coupling Link feature (CPC drawer PCIe fanout)**
  - ICA SR - two 8GBps PCIe Gen3 Coupling Links (FC 0172)
- **InfiniBand Coupling features (CPC drawer HCA fanout)[1]**
  - HCA3-O two 12x 6GBps[2] InfiniBand DDR Coupling Links (FC 0171)
  - HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links (FC 0170)

CPC drawer

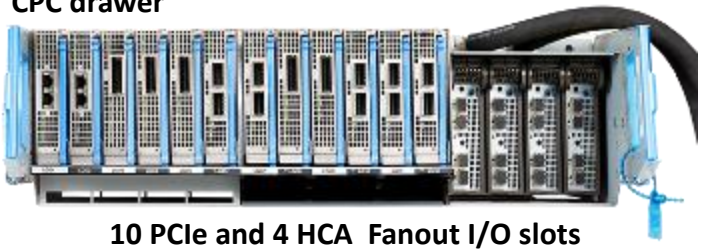

10 PCIe and 4 HCA Fanout I/O slots

1. z14 is the **LAST** IBM Z server to support InfiniBand features
2. 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 carry-forward (MES) I/O features supported

- **Features – PCIe I/O drawer**
  - FICON Express16S: LX (FC 0418), SX (FC 0419)
  - FICON Express8S: 10 KM LX (FC 0409), SX (FX 0410)
  - OSA-Express5S: 1 GbE LX (FC 0413), 1 GbE SX (FC0414), 10 GbE LR (FC 0415), 10 GbE SR (FC 0416), and 1000BASE-T (FC 0417)
  - OSA-Express4S: 1000Base-T (FC 0408)
  - 10GbE RoCE Express (FC 0411)
  - zEDC Express (FC 0420)
  - Crypto Express5S (FC 0890)
  - Coupling Express LR (FC 0433)
- **PCIe Coupling Link feature (CPC drawer PCIe fanout)**
  - ICA SR - two 8GBps PCIe Gen3 Coupling Links (FC 0172)
- **InfiniBand Coupling features (CPC drawer HCA fanout)**[1]
  - HCA3-O two 12x 6GBps[2] InfiniBand DDR Coupling Links (FC 0171)
  - HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links (FC 0170)

1. z14 is the **LAST** IBM Z server to support InfiniBand features
2. 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.

IBM Z Coupling Facility Enhancements

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 (GBps)[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

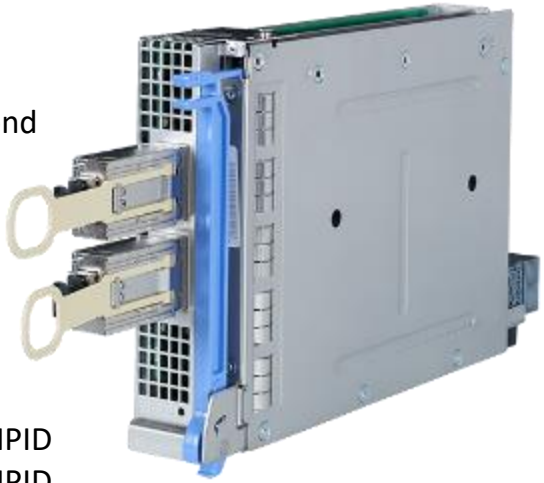

1. Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

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	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		zNext SF[1]	
			GA2	GA2+		
				Coupling Express LR[2]	Coupling Express LR	
			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			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.

z14 Coupling Connectivity

z14, z13, and z13s
12x IFB, 12x IFB3, 1x IFB, ICA SR, CE LR

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

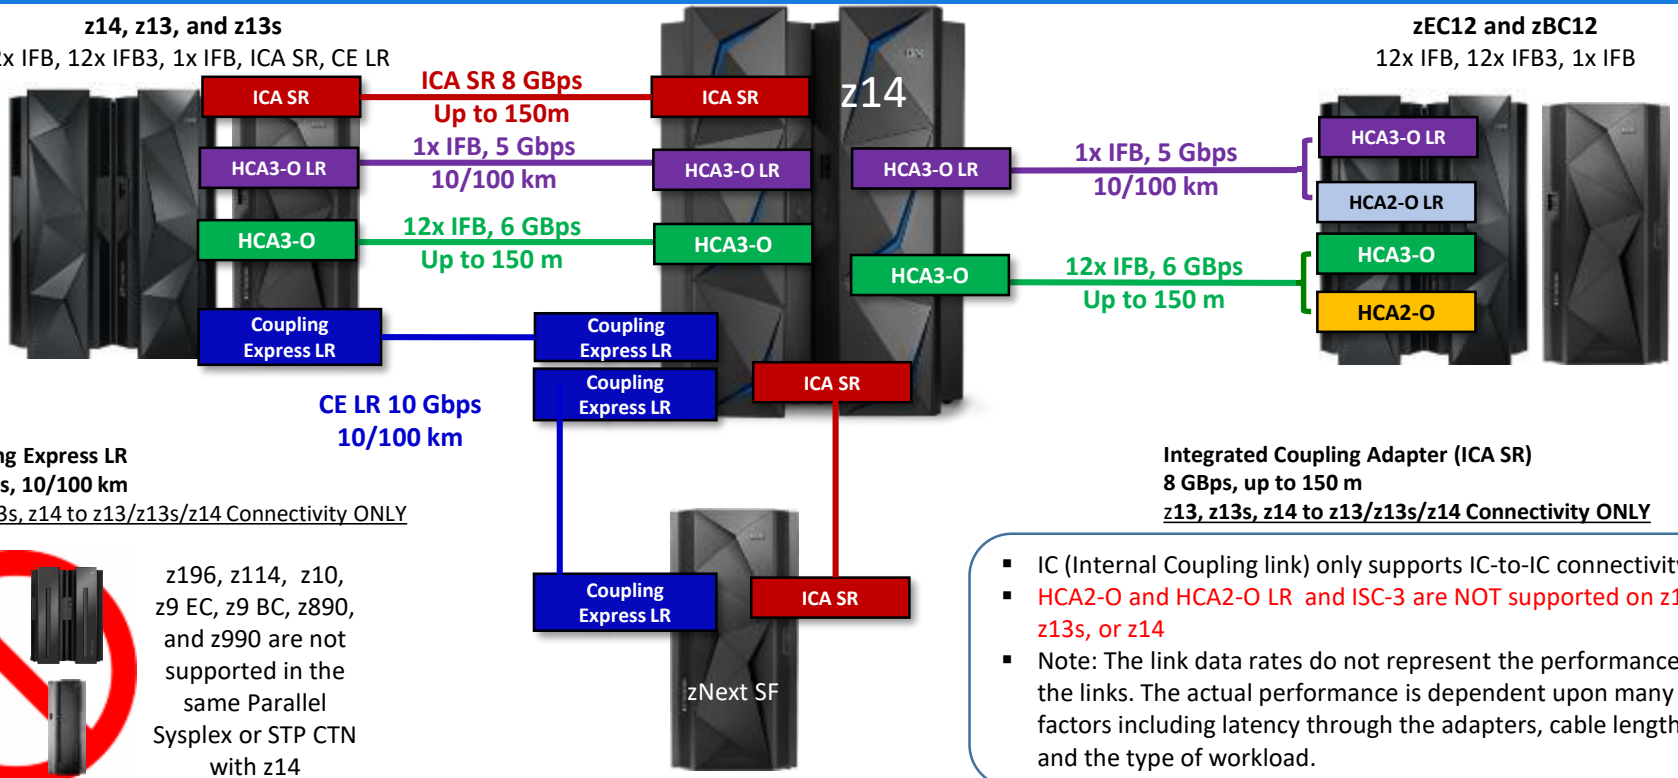

Coupling Express LR
10 Gbps, 10/100 km
z13, z13s, z14 to z13/z13s/z14 Connectivity ONLY

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 GBps, 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.

IBM z14 Capacity and Performance Planning

Performance drivers with z14

- **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 versus 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 versus 8 on z13
- **Memory subsystem**
  - Focused on keeping data "closer" to the processor unit
    - Larger L1, L2, and L3 caches
    - One unified L4 cache per drawer shared by L3s
  - 3.2x configurable memory (32 TB versus 10 TB for z13)
- **Up to 170 configurable processor units (cores) versus 141 on z13**
- **3 sub-capacity settings (same as z13)**
- **HiperDispatch**
  - Exploits new chip configuration
  - Required for SMT on zIIPs
- **PR/SM**
  - 85 customer partitions (same as z13)
  - Up to 170 LCPs and 16 TB memory per partition (actual limits are OS dependent)
  - Improved resource allocation algorithms based on z13 experience

z14 vs. z13 hardware SMP topology

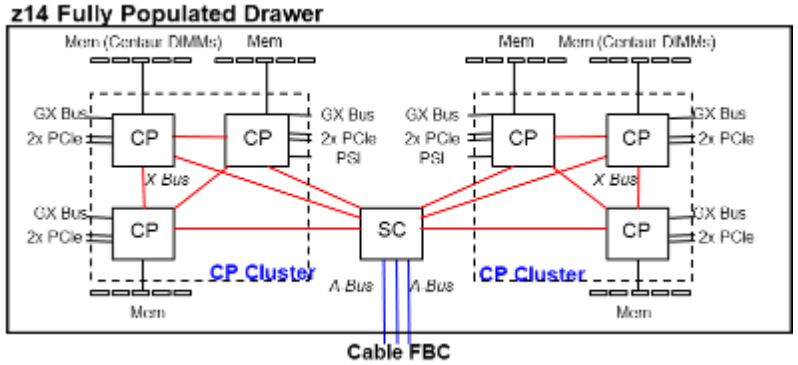

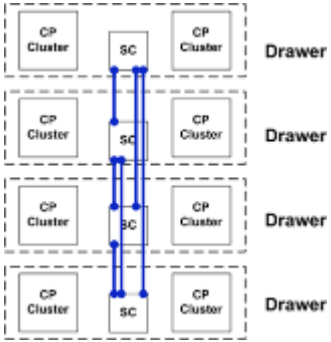

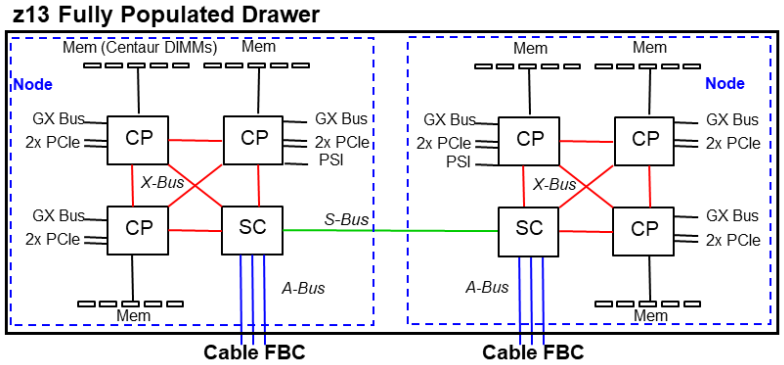

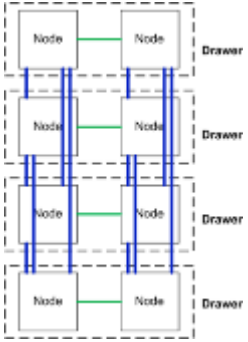

z14 vs. z13 cache topology comparison

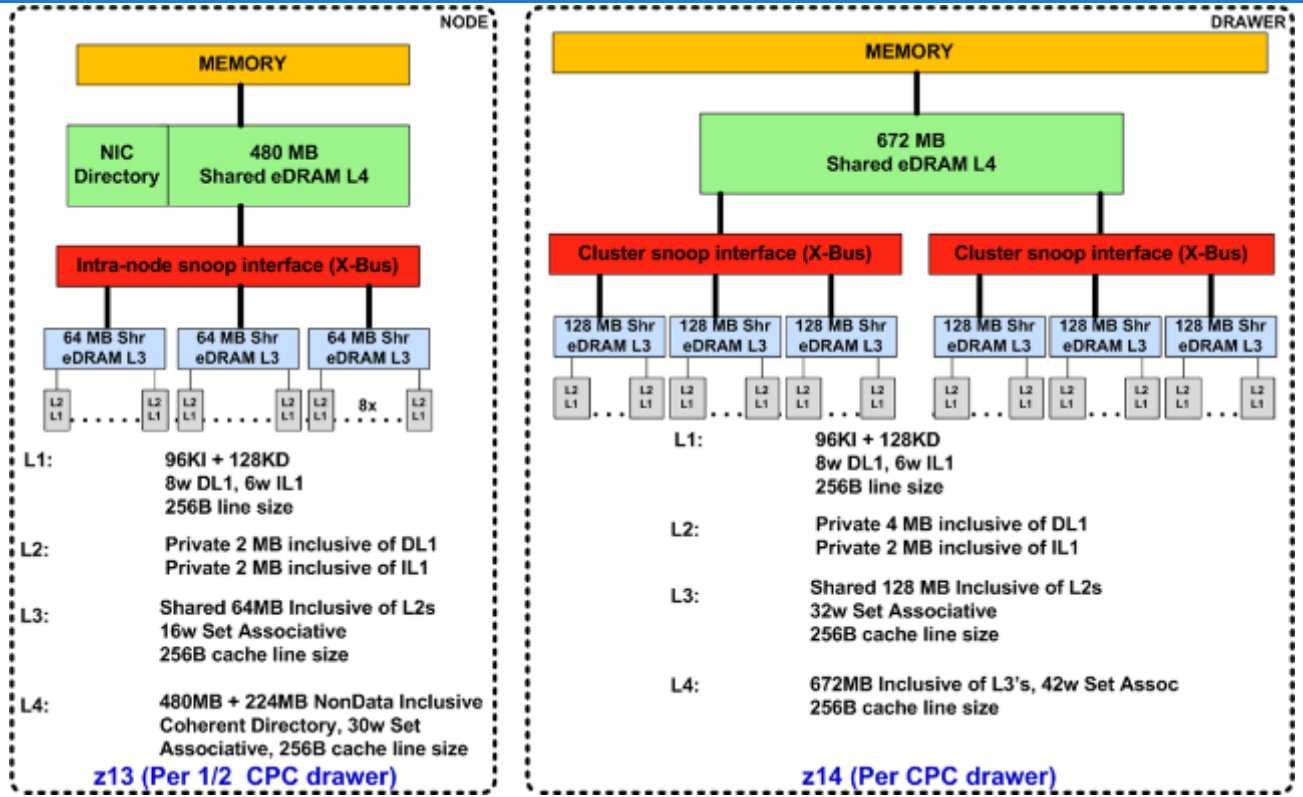

z14 vs. z13 vs. zEC12 vs. z196 vs. z10 EC MIPS comparison

- z14 up to 170-way: 256 – up to 146,462 MIPS
- z13 up to 141-way: 250 – up to 111,556 MIPS
- zEC12 up to 101-way: 240 – 78,426 MIPS
- z196 up to 80-way: 240 – 52,286 MIPS
- z10 EC up to 64-way: 214 – 31,826 MIPS

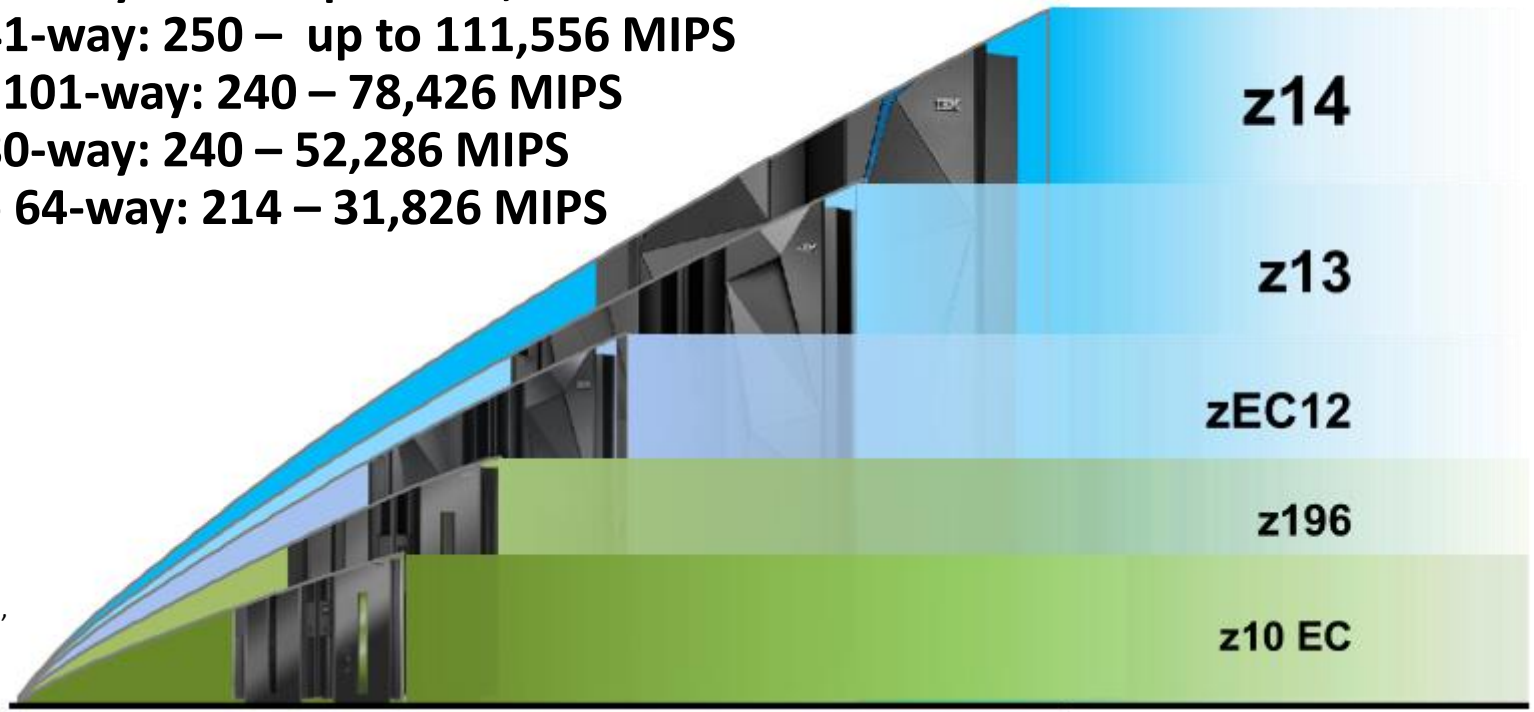

Note: For z14 – z/OS V2.2, for z13 and zEC12 - z/OS 2.1, and for z196 and z10 EC z/OS V1.13 LSPR numbers quoted.

z14 performance capacity highlights

- **Full speed capacity models (7xx) capacity ratio to z13**
  - Up to 35% more total system capacity is expected compared to the z13 (170-way z14 vs. 141-way z13)
  - Up to 1.10x (10%) average performance improvement at equal N-way vs. z13
- **Full capacity model 701**
  - Uni-processor capacity (full speed – 701) 1832 MIPS (8% increase over z13)
- **Sub-capacity models**
  - Uni-processor capacity ratio – sub-capacity models:
    - 401: 0.14x (256 MIPS)
    - 501: 0.41x (751 MIPS)
    - 601: 0.59x (1081 MIPS)
  - Up to 33 CPs (general purpose) for each sub-capacity model
- **SMT capacity**
  - IFLs and zIIPs can choose to run 2 hardware threads per core
    - Controlled by an operating system parameter at the LPAR level
    - Added hardware threads appear as additional processors to the operating systems
    - Default is single thread
  - SAPs use SMT active z14 by default (this cannot be altered by the customer)
  - Likely wide range in capacity improvement per core over single thread: up to 25%

z14 performance – potential sources of variability

- **Expect migrations to z14 from z13 to be stable**
- **Workloads migrating to z14 from zEC12 and prior can expect to see slightly less variability than the typical z13 migration experience**
- **zEC12 marked the end of an era while z13 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”) vs. 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 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 is an evolution of z13**
  - Reduced variability by consolidating both nodes on a drawer
  - More cores per drawer and each core with higher capacity than z13 means more work can “fit” on a drawer
  - Added architecture and updated z/OS to manage locks more intelligently

LSPR: Performance showcase for IBM Z processors

- **IBM Z provides capacity comparisons among processors based on a variety of measured workloads which are published in the Large System Performance Reference (LSPR)**
  - https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex
  - Now based on z/OS 2.2, DB2 11, CICS 5.3, IMS 14, Cobol 6.1, WAS 8.5.5.9
    Note: Capacity ratios may vary if using older versions of these subsystems
  - Minor tweaks to the three workload categories based on customers' CPU MF data for zEC12 to z13 migrations
- **Old and new processors are measured in the same environment with the same workloads at high utilizations**
- **Over time, workloads and environment are updated to stay current with customer profiles**
  - Old processors measured with new workloads/environment may have different average capacity ratios compared to when they were originally measured
- **LSPR presents capacity ratios among processors**
- **Single number metrics include PCI/MIPS, MSUs, and SRM Constants**
  - MIPS are based on the ratios for the "Average" workload category with a median customer LPAR configuration

zPCR support for z14

- **Supports 85 partitions included in a capacity study**
- **zAAPs are not supported on z13, z13s, and z14**
  - If changing host to z14, zAAP partitions will be invalid and excluded
  - Replace zAAP partitions with zIIP partitions using Partition Definition window
- **Absolute Capping (zEC12, zBC12, z13, z13s, and z14 only)**
  - User can specify an absolute capping limit in 1/100th of a processor
- **SMT support (z13, z13s, and z14 only)**
  - z/OS 2.1 or later on zIIPs
  - z/VM 6.3 or later, KVM 1.1.1 or later, and Linux on IFLs
  - Default capacity benefit for zIIPs and IFLs is now 25% (20% for z13/z13s IFLs)
  - User can set estimated capacity benefit for IFL and zIIP partitions from 0% to 60% (1% increments)
  - Measured capacity benefit read from EDF (z/OS or z/VM) or RMF CPU Activity Report (z/OS)
    - Increase measured benefit by 5% for z/VM when changing LPAR host from z13 to z14, otherwise use the same measured benefit for z14
  - Margin-of-error using default SMT capacity benefit can be significantly higher
  - Margin-of-error is indeterminate if user overrides default SMT capacity benefit
  - New extract program versions (CP3KEXTR, CP3KVMXT) will be necessary to assess z14 upgrade

Critical migration action

- **Before and after CPU MF Counters data will be critical to determine the source of variation for workloads that do encounter it**
  - Ensure the CPU MF data is captured and kept for analysis
  - For z/VM “Before” and “After” peak hour data must be written to disk
- **Critical migration action for every z14 candidate (z/OS and z/VM)**
  - CPU MF Counters must be enabled on your current processor
  - CPU MF Counters must be enabled on your new z14 processor
- **Take action to validate CPU MF is implemented or get a plan started to implement CPU MF**
- **Minimal overhead – we encourage customers to run continuously**
- **CPU MF Counters enablement resources**
  - CPU MF Webinar Replays and Presentations
    - http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4922
  - z/OS CPU MF - “Detailed Instructions” Step by Step Guide
    - http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TC000066
  - z/VM Using CPU Measurement Facility Host Counters
    - http://www.vm.ibm.com/perf/tips/cpumf.html

IBM z14 Designed for Simplification

Designed for simplification

- **Market view**
  - 38% of HR managers surveyed said training and developing employees is their greatest staffing concern[1]
  - Simplified management of mainframe systems
  - Majority of large enterprises have multiple virtualization technology silos, each managed with its own administrative solutions and skills, which may increase overall operational cost and complexity[2]

The best infrastructure to support an open and connected world

- **'From anywhere' mobile access**
- **Simplified sys admin of z/OS**
- **Standardization for skills transfer**

1. Business Management Daily – Dec 21, 2015
2. IDC – July 2015 #257441

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

Also available for z13 and z13s!

Tree structure

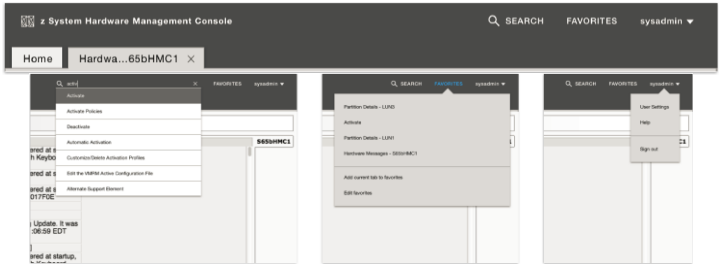

Multi-factor authentication

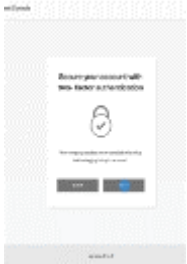

New mobile capabilities

New HMCs/TKEs for z14 and z13/z13s

Description	New Feature	Old Feature(s)
HMC	0082	0091, 0092, 0095
HMC Rack Mount	0083	0094, 0096
TKE Rack Mount w/4768	0085	0097
TKE w/4768	0086	0098
TKE 9.0 LIC	0879	0877 TKE 8.0 LIC, 0878 TKE 8.1 LIC

- **Effective September 13, 2017 the "old" features can no longer be ordered on z13 and z13s**

HMC workspace enhancements

- **Classic UI Style no longer available on HMC/SE 2.14.0**
- **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

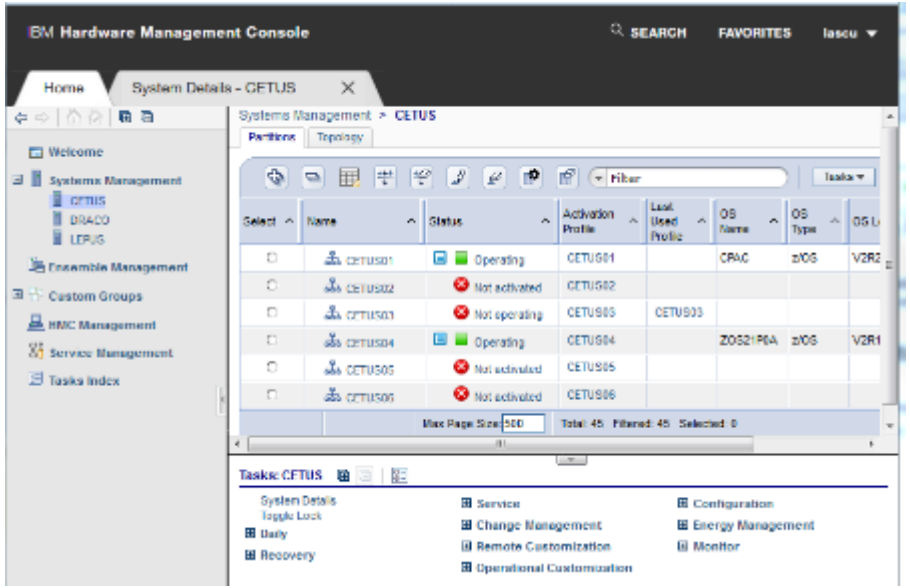

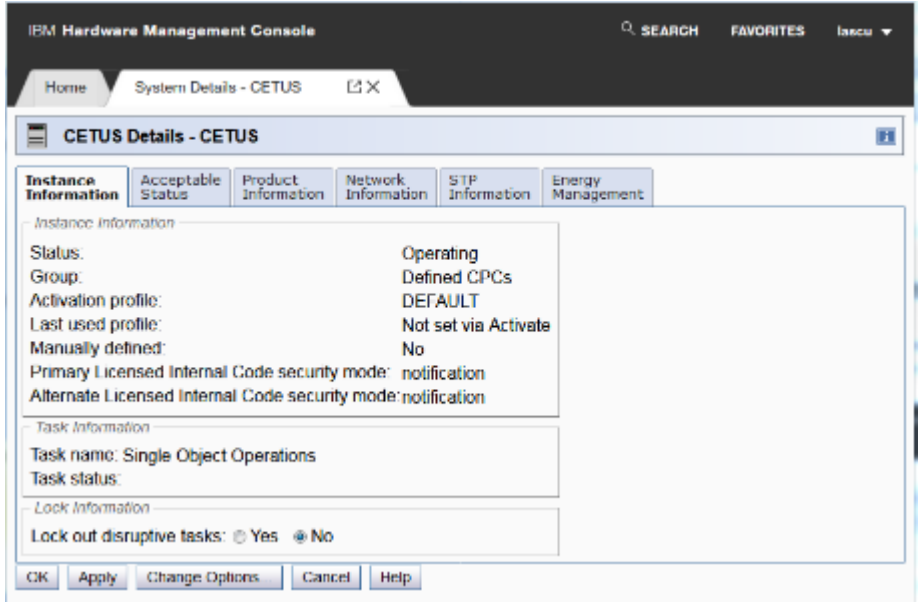

HMC multi-factor authentication

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

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**

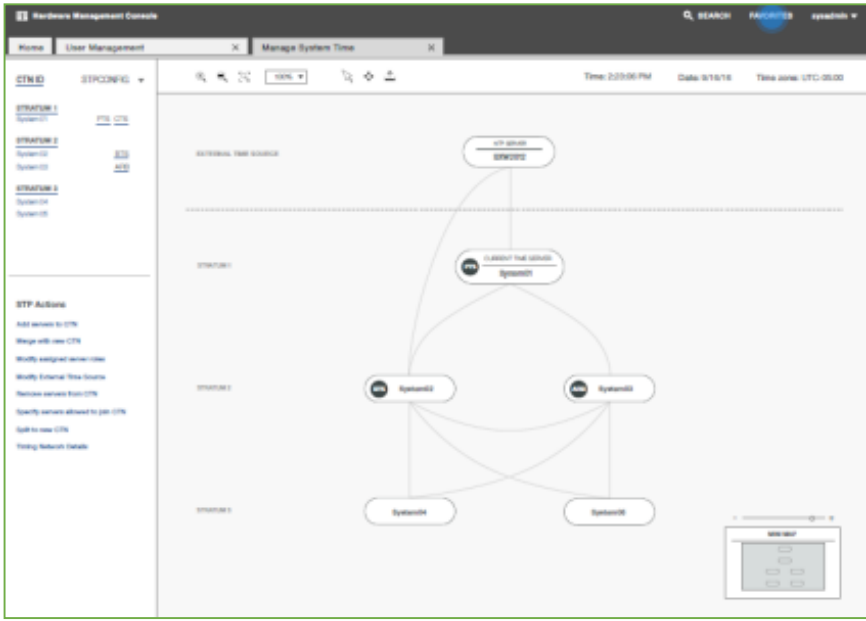

IBM z/OS Management Facility (z/OSMF)

- **Modernization: Support for new browser-based management console**
- **Productivity: Embedded active user assistance (such as wizards)**
- **Always on: z/OS V2.3 change so z/OSMF is available all the time**
- **Standardize: Software installation and configuration experience enhanced in V2.3**
- **Sysplex Management: New z/OSMF plug-in provides detailed views of sysplex infrastructure resources such as CFs and CF structures, CF Structure Connectors, couple datasets and policies, and Coupling links**

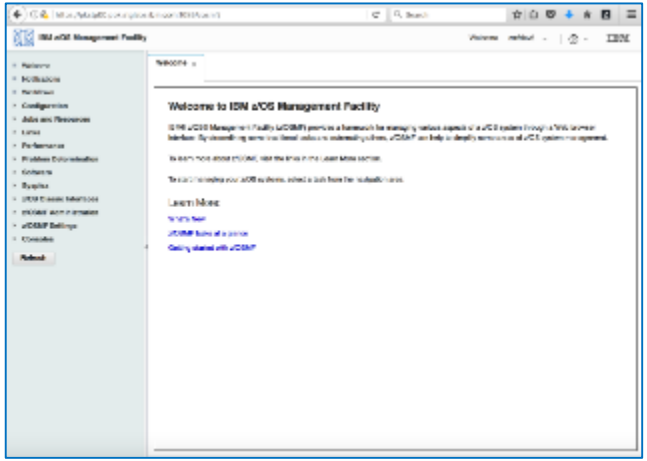

IBM Dynamic Partition Manager
Simplified configuration of logical partitions

- **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 – SOD: FICON ECKD support**
- **Secure FTP through HMC for booting and installing an operating system via FTP**

Benefits for users new to IBM Z[2]:

- **Modify system resources without disrupting running workloads**
- **Create alarms for events, conditions, and state changes**
- **Update individual partition resources to adjust capacity, redundancy, availability, or isolation**

1. Where Linux distro installers exploit function
2. Not available for z/OS, z/VSE, or z/TPF

Simplicity with Linux on 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**
- **32TB 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**

Simplicity with Linux on Z

Efficient and economic operation

- **Run WebSphere Liberty on z14 using clear key encryption can get up to *2.6x* improvement in throughput per core with IBM Java 8 SR5 compared to x86[1]**
- **Scale-up a MongoDB instance to 17 TB in a single system without database sharding and get *2.4x more throughput* on z14 leveraging additional memory compared to z13[2]**
- **Scale-out to 2 million Docker containers in a single system, no application server farms necessary[3]**

1 Performance results based on IBM internal tests running DayTrader 3 with WebSphere Liberty 8.5.5.9 using SSL clear key and TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 cipher.
2 Performance result based on IBM internal tests comparing MongoDB performance in native LPAR on z14 using additional memory versus z13 driven by YCSB 0.11.0 (write-heavy, read-only). Results may vary. z14 configuration: LPAR with 12 dedicated IFLs and 20 TB memory running on SLES 12 SP2 (SMT mode) a MongoDB Enterprise Release 3.4.1 instance (no sharding, no replication) with a 17 GB database. The database was located on an 18 TB LUN on an IBM FlashSystem 900. z13 configuration: LPAR with 12 dedicated IFLs and 10 TB memory running on SLES 12 SP2 (SMT mode) a MongoDB Enterprise Release 3.4.1 instance (no sharding, no replication) with a 17 GB database. The database was located on an 18 TB LUN on an IBM FlashSystem 900.
3 Performance result is extrapolated from IBM internal tests running in a z14 LPAR with 10 dedicated IFLs and 16 GB memory 1000 BusyBox Docker containers with ApacheHTTP. Results may vary. Operating system was SLES 12 SP2 (SMT mode). Docker 1.12 was

Operating systems exploiting hardware innovation

z/OS Version 2.3		▪ Provide a simple, transparent, and consumable approach to enable extensive encryption of data in-flight and at-rest to substantially reduce the people and hardware costs associated with protecting data and achieving compliance mandates. ▪ Simplify the overall management of the z/OS ecosystem, increasing the productivity and value of system administrators and easing the on-boarding of new team members ▪ Provide a simple, consumable approach for self-service provisioning and rapid delivery of aaS, while enabling for the API economy.
z/VSE Version 6		▪ Performance and functional enhancements for online processing plus faster I/O with FICON Express16S+ with link rate of 16 Gbps ▪ Improved network performance with OSA Express6S and security with firewall functionality, including z/VSE Network Appliance deployable in the IBM Secure Service Container ▪ Better and faster HW Encryption with Crypto Express6S ▪ Wide portfolio using z/VSE Connectors and Linux on Z ▪ More capabilities will be available with the upcoming z/VSE V6.2 (see the z/VSE Announcement Preview)
Linux on Z		▪ Multithreading with next generation SMT may allow for per core software savings ▪ Ability to host and manage more workloads for service agility and improved IT economics ▪ Hardware accelerated encryption improve cryptographic functions inside Linux, openSSL, openCryptoki ,and GSKIT ▪ Secure Service Container, based on Linux, can be used for isolated partitions to protect data and applications ▪ Linux distributions from Canonical, Red Hat and SUSE to your selection
z/TPF		▪ Management of extreme transaction volumes up to hundreds of thousands of transactions per second ▪ Fast / consistent response across predictable and unpredictable peaks ▪ Low cost per transaction for large applications ▪ Centralized database handling routines to effectively manage databases ▪ HiperDispatch workload balancing to optimize processor utilization

Hypervisors and virtualization for IBM Z

PR/SM LPAR and IBM DPM

- Virtualization is ***built into the DNA of IBM Z***
- PR/SM™ ***manages and virtualizes all*** the installed and enabled system resources as a single large SMP system
- ***Full sharing/partitioning of the installed resources*** with the highest levels of efficiency and utilization
- ***Scale up or scale out on demand*** with support for up to 85 partitions
- IBM Dynamic Partition Manager ***simplifies provisioning and management experience***
- Assured ***workload isolation*** with the highest EAL5+ security certification
- New ***dynamic optimization and scalability*** enhancements

z/VM Version 6.4

IBM Wave for z/VM

- ***Enables extreme scalability, security and efficiency*** creating cost savings opportunities by taking by exploiting Guest Enhanced DAT to allow guests to take advantage of large (1MB) pages
- ***Ease Migration*** with upgrade in place infrastructure provides a seamless migration path from previous z/VM releases (z/VM 6.2 and z/VM 6.3) to the latest version
- ***Operational improvements*** by enhancing z/VM to provide ease of use
- ***SCSI*** for guest attachment of printers, disk drives, scanners, and other peripherals, and host or guest attachment of disk drives
- IBM Wave for z/VM ***simplifies the management*** of virtual Linux servers from a single user interface

KVM for z Systems

- Support new analytics workloads with ***Single Instruction Multiple Data (SIMD)*** for competitive advantage
- High compute capacity with support for ***Simultaneous Multithreading (SMT)*** to meet new business requirements
- Improved problem determination and high availability setup to ***reduce down time*** and quickly react to business needs
- Secure and protect business data with ***Crypto exploitation***
- ***Technology developed by IBM … product offered by Linux Distribution partners***

z/OS support summary

Release	z9 EC z9 BC	z10 EC z10 BC	z196 z114	zEC12 zBC12	z13 z13s	z14	End of Service	Extended Defect Support[1]
z/OS 1.13	X	X	X	X	X	X[1]	9/16	9/19[3]
z/OS 2.1	X	X	X	X	X	X	9/18[3]	9/21[3]
z/OS 2.2		X	X	X	X	X	9/20[3]	9/23[3]
z/OS 2.3[2]				X	X	X	9/22[3]	9/25[3]

- **Notes:**
  1. The IBM Software Support Services for z/OS V1.13, offered as of October 1, 2016, provides the ability for customers to purchase extended defect support service for z/OS V1.13
  2. Planned to be Generally Available in September 2017
  3. Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

z/OS support for z14

Release	IBM.Device.Server.z14-3906.RequiredService						IBM.Device.Server.z14-3906.Exploitation										
	Base Support	CPU Measurement Facility (HIS)	Crypto Express6S Toleration	FICON Express 16S+	z14 Assembler Support	OSA-Express6S	Guarded Storage Facility (GSF)	Instruction Execution Protection (IEP)	IBM Virtual Flash Memory	Crypto Express6S	RoCE Express2	z14 XL C/C++	CF Level 22	Coupling Express LR	HiperDispatch Enhancements	Data Set Encryption	zHyperLink
z/OS V1.13[S]	P		W,P[AO]	P	P	P			W								
z/OS V2.1	P	P	P	P	P	P			Y	W	P		P	P	P	C	P
z/OS 2.2	P	P	P	P	P	P	P	P	Y	W	P		P	P	P	P	P
z/OS 2.3	Y	Y	P	Y	Y	Y	Y	Y	Y	W	Y	Y	Y	P	Y	Y	Y

Notes

S	IBM Software Support Services required for extended z/OS V1.13 support.	C	Coexistence support is required, if exploited. Dependent upon the specific function. There could be partial support on lower levels. Full support in z/OS V2.3.
P	PTF is required, use SMP/E FIXCAT for identification.	AO	Requires the ICSF web deliverable for FMID HCR77A0 minimally, with PTF.
Y	Support is in the base.	W	A web deliverable is required, available at ibm.com/systems/z/os/zos/downloads/

z/VM support summary

z/VM Release	GA	End of Service	End of Marketing	Minimum Processor Level	Maximum Processor Level	Security Level
6.4	2016-11-11			IBM System z196 & z114®	-	
6.3	2013-07-26	2017-12-31[1]	2016-11-10	IBM System z10®	-	EAL 4+ OSPP-LS
6.2	2011-12-02	2017-06-30[2]	2013-07-23	IBM System z10®	z13	-
5.4	2008-09-12	2017-12-31[3]	2012-03-12	IBM eServer zSeries 800& 900	zEC12	-

4 releases in service on January 1, 2017

1 release in service on January 1, 2018

1. EOS announced 2015-02-03
2. EOS announced 2016-02-02
3. EOS announced 2016-08-02

z/VSE support summary

z/VSE Release	z800 / z900 z890 / z990	z9	z10	z196 / z114 / zEC12 zBC12 / z13 / z13s	z14	VSE EoM	VSE EoS
z/VSE V6.2	No	No	No	Yes	Yes	tbd	tbd
z/VSE V6.1	No	No	Yes	Yes	Yes	tbd	tbd
z/VSE V5.2	No	Yes	Yes	Yes	Yes	2017-03-13	2018-10-31
z/VSE V5.1	No	Yes	Yes	Yes	Yes	2014-05-23	2016-06-30

- **z/VSE 5.2 and 6.1 support z14 at GA, PTFs may be required**
- **Only z/VSE 5.1, 5.2, 6.1 and 6.2 can IPL on z14**
- **z/VSE provides toleration support**
  - Compilers will not change, no additional exploitation for z14
- **Supported z14 adapters**
  - FICON Express16S+ (CHPID type FC) - utilizing FICON or channel to channel
  - FICON Express16S+ (CHPID type FCP) for support of SCSI devices
  - OSA-Express6S features (e.g. 1000BASE-T Ethernet, GbE)
  - Crypto Express6S toleration

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%***[1]
  - 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%***[2]

- **Java SDK 8 SR5**
  - Faster user response times for Java
  - ***4.2x*** 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***[4] for predictable high-perform transaction processing at-scale
  - 50+ new instructions on the z14 co-designed and exploited by Java

1. On z14 over same applications built with Enterprise COBOL for z/OS V6.1 on z14
2. Over the same general compute intensive applications built with Enterprise COBOL for z/OS V4.2 on z14
3. IBM Java 8 SR5 on z14 compared to Java 8 SR3 on z13
4. Using new IBM Java 8 SR5 Pause-less garbage collection feature on z14 compared to using Java 8 SR3 on z13.

Summary

Additional information

- **IBM z14 product information** web site
- **IBM Redbooks**
  - IBM Z Functional Matrix, REDP-5157-02
  - IBM Z Connectivity Handbook, SG24-5444-17
  - IBM z14 Technical Introduction, SG24-8450-00
  - IBM z14 Technical Guide, SG24-8451-00
- **Education**
  - IBM Z YouTube channel

Protecting your investment in technology

- **Designed to protect your investment**
  - Offering upgrades from IBM z13 and IBM zEnterprise® EC12 (zEC12) to the z14
- **Full upgradeability within Models M01 – M04**
  - No upgrade to Model M05
  - No upgrade from any IBM LinuxONE Emperor™
- **On demand offerings offer temporary or permanent growth when you need it**

Leading in the trust economy

- Designed with pervasive encryption for piece of mind that data and privacy is always protected
- Designed with improvements in speed, efficiency and access improvements to both serve up data to build services and new offerings and to perform machine learning and insight on the data where it resides
- Designed to be open and industry standard to bridge the skills gap and make the infrastructure easier to manage

IBM z14 Functional Comparison to IBM z13

Performance and Scale	▪ Uniprocessor Performance ▪ System Capacity ▪ SMT ▪ SIMD ▪ Cache ▪ Models ▪ Processing cores ▪ Granular Capacity ▪ Memory ▪ Compression	▪ New up to 10% performance improvement over IBM z13 (z13)[1] ▪ New up to 35% system total z/OS capacity performance improvement over z13[1] ▪ New 2nd generation SMT delivers virtualization benefits for Linux and up to 25% performance improvement for Linux on Z and zIIP workloads vs non-SMT on z14 ▪ New instructions for perform boost to traditional workloads and new analytics workloads versus SIMD on z13 ▪ New z14 has 1.5x more on-chip cache per core versus z13 ▪ Five models with up to four CPC drawers (z13 also has five CPC models and four drawers) ▪ New up to 170 cores to configure, up to 141 on z13 ▪ New up to 269 capacity settings versus 231 on the z13 ▪ New up to 32 TB RAIM memory versus 10 TB RAIM memory on z13 ▪ New improvement in CMPSC compression and Huffman Coding compression ratio using zEDC Express versus z13
Virtualization	▪ LPAR virtualization ▪ RoCE adapter ▪ Simplified LPAR management	▪ 85 partitions – same as z13 ▪ New 10 GbE RoCE Express2 with additional virtual functions (VFs) per physical port (10GbE Express on z13) ▪ Enhanced IBM Dynamic Partition Manager allows for config and management of system resources – new support for z/VM and for ECKD disk

1. For average LSPR workloads running z/OS 2.1. Official performance data are available and can be obtained online at LSPR (Large Systems Performance Reference) website at: https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocumen. Actual performance results may vary by customer based on individual workload, configuration and software levels.

IBM z14 Functional Comparison to IBM z13

Infrastructure Efficiency	▪ Networking ▪ HiperSockets and SMC-D ▪ FICON ▪ zHPF ▪ IBM zHyperLink ▪ Forward Error Correction ▪ FICON dynamic routing ▪ LCSS/Subchannel sets ▪ WWPN ▪ HMC ▪ Pause-less garbage collection ▪ Upgradeability ▪ IBM Virtual Flash Express	▪ New OSA-Express6S with improvements over z13 using OSA-Express5S ▪ Up to 32 HiperSockets (same as z13) and memory-to-memory communications with SMC-D offers within-the-box communications for z/OS ▪ New FICON Express16S+ will provide an increase in I/O rates over FICON Expresss16S ▪ zHPF extended distance II offers faster remote site recovery with improved I/O service time improvement when writing data remotely (GDPS® HyperSwap®) same as z13 ▪ New IBM zHyperLink – New short distance z14 channel that can be installed on IBM DS8880 System Storage® for lower latency not on z13 ▪ Industry standard FEC for optical connections for substantially reduced I/O link errors same as z13 ▪ Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches same on z13 ▪ Up to six LCSS and 4 Subchannel sets – same as z13 ▪ I/O serial number migration allows keeping same serial number on replacement server same as z13 ▪ New next generation HMC with simplified panels, new mobile capabilities, security enhancements (including multi-factor authentication), easier help panels – not on z13. (No Classic Style User Interface on z14) ▪ New enterprise scale JAVA applications to run without periodic pause for garbage collection on larger & larger heaps ▪ Upgradeable from z13 and zEnterprise EC12 (zEC12) ▪ New memory replacement for Flash Express helping improve availability – available only on z14

IBM z14 Functional Comparison to IBM z13

Resiliency and Availability	▪ Coupling – HCA-3 ▪ Coupling – ICA SR ▪ Coupling Express LR ▪ STP ▪ Sparing ▪ Rack Mounted Accessories ▪ Environmentals	▪ Coupling with HCA-3 InfiniBand Coupling Links – long and short distance – same as z13 ▪ Short distance coupling with PCIe-based links (ICA SR) – same as z13 ▪ New Coupling Express LR – Coupling Express LR will be available on z13 ▪ New Simplified STP management with HMC enhancements not available on z13 ▪ Enhanced integrated sparing on z14 and z13 reducing the number of on site service and maintenance events ▪ Rack-mounted HMC and TKE options to save space in the data center ▪ Optional non raised floor, overhead cabling, water cooling and DC power plus New ASHRAE A3 ration – (not ASHRAE rating on z13)
Security	▪ Cryptographic Coprocessor ▪ Crypto Express ▪ IBM Secure Service Container ▪ Secure Console Access	▪ CPACF for improved performance and new true Random Number Generator versus z13 ▪ New Crypto Express6S with performance increase plus new algorithms for elliptic curve, SHA, VISA FPE versus z13 Crypto Express5S ▪ Secure deployment of software virtual appliances ▪ Protection of sensitive data by using Transport Layer Security (TLS) support in the Open Systems Adapter-Integrated Console Controller (OSA-ICC)

IBM z14 Functional Comparison to IBM zEC12

Performance and Scale	▪ Uniprocessor Performance ▪ System Capacity ▪ SMT ▪ SIMD ▪ Cache ▪ Models ▪ Processing cores ▪ Granular Capacity ▪ Memory ▪ Compression	▪ New up to 22% performance improvement over IBM zEnterprise EC12[1] ▪ New up to 89% system total z/OS capacity performance improvement over zEC12[1] ▪ SMT delivers up to 55% price performance improvement for Linux on IBM z Systems and up to 55% price performance improvement for zIIP workloads versus single threaded only on zEC12 ▪ Vector processing (SIMD) model provides construction of richer, complex analytics models, increased programmer productivity, faster analytics to traditional workloads versus no SIMD on zEC12 ▪ New z14 has over 60% more on-chip cache per core versus zEC12 ▪ Five models with up to four CPC drawers (zEC12 also has five CPC models and four books) ▪ New up to 170 cores to configure, up to 101 on zEC12 ▪ New up to 269 capacity settings versus 161 on the zEC12 ▪ New up to 32 TB RAIM memory versus 3 TB RAIM memory on zEC12 ▪ New improvement in CMPSC compression and Huffman Coding compression ratio using zEDC Express versus zEC12
Virtualization	▪ LPAR virtualization ▪ RoCE adapter ▪ Simplified LPAR management	▪ 85 partitions versus 60 on zEC12 ▪ New 10 GbE RoCE Express2 with additional virtual functions (VFs) per physical port (versus dedicated 10GbE Express on zEC12) ▪ Enhanced IBM Dynamic Partition Manager allows for configuration and management of system resources– new support for z/VM and for ECKD disk versus no DPM on zEC12

1. For average LSPR workloads running z/OS 2.1. Official performance data are available and can be obtained online at LSPR (Large Systems Performance Reference) website at: https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocumen. Actual performance results may vary by customer based on individual workload, configuration and software levels.

IBM z14 Functional Comparison to IBM zEC12

Infrastructure Efficiency	▪ Networking ▪ HiperSockets and SMC-D ▪ FICON ▪ zHPF ▪ IBM zHyperLink ▪ Forward Error Correction ▪ FICON dynamic routing ▪ LCSS/Subchannel sets ▪ WWPN ▪ HMC ▪ Pause-less garbage collection ▪ IBM Virtual Flash Express	▪ New OSA-Express6S with improvements over zEC12 using OSA-Express5S ▪ Up to 32 HiperSockets (same as zEC12) and memory-to-memory communications with SMC-D offers within-the-box communications for z/OS (SMC-D not on zEC12) ▪ New FICON Express16S+ and increased FICON subchannels to 32K versus FICON Express 8S and 24K on zEC12 ▪ zHPF extended distance II offers faster remote site recovery with improved I/O service time improvement when writing data remotely (GDPS HyperSwap) versus zHPF only on zEC12 ▪ New IBM zHyperLink - New short distance z14 channel that can be installed on IBM DS8880 System Storage® for lower latency not on zEC12 ▪ Industry standard FEC for optical connections for substantially reduced I/O link errors not available on zEC12 ▪ Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches not on zEC12 ▪ Up to six LCSS and 4 Subchannel sets – versus four LCSS and 3 subchannel sets on zEC12 ▪ I/O serial number migration allows keeping same serial number on replacement server not on zEC12 ▪ New next generation HMC with simplified panels, new mobile capabilities, security enhancements (including multi-factor authentication), easier help panels – not on zEC12. (No Classic Style User Interface on z14) ▪ New enterprise scale JAVA applications to run without periodic pause for garbage collection on larger & larger heaps not on zEC12 ▪ New memory replacement for Flash Express helping improve availability – available only on z14

IBM z14 Functional Comparison to IBM zEC12

Resiliency and Availability	▪ Coupling – HCA-3 ▪ Coupling – ICA SR ▪ Coupling Express LR ▪ STP ▪ Sparing ▪ Rack Mounted Accessories ▪ Environmentals	▪ Coupling with HCA-3 InfiniBand Coupling Links – long and short distance – same as zEC12 ▪ Short distance coupling with PCIe-based links (ICA SR) – not available on zEC12 ▪ New Coupling Express LR – Coupling Express LR – not available on zEC12 ▪ New Simplified STP management with HMC enhancements not available on zEC12 ▪ Enhanced integrated sparing on z14 reducing the number of on site service and maintenance events not on zEC12 ▪ Rack-mounted HMC and TKE options to save space in the data center not available on zEC12 ▪ Optional non raised floor, overhead cabling, water cooling and DC power plus New ASHRAE A3 ration – (not ASHRAE rating of zEC12)
Security	▪ Cryptographic Coprocessor ▪ Crypto Express ▪ IBM Secure Service Container ▪ Secure Console Access	▪ CPACF for improved performance and new true Random Number Generator not on zEC12 ▪ New Crypto Express6S with performance increase plus new algorithms for elliptic curve, SHA, VISA FPE versus zEC12 Crypto Express4S ▪ Secure deployment of software virtual appliances not on zEC12 ▪ Protection of sensitive data by using Transport Layer Security (TLS) support in the Open Systems Adapter-Integrated Console Controller (OSA-ICC) not on zEC12

Jim Elliott
Mainframe Consultant

885 Don Mills Road, Suite 300, Toronto, ON M3C 1V9 t: (416) 229-2950 x341
c: (416) 527-0666 jelliott@ghsystems.com **ghsystems.com**
jlelliotton @jlelliotton Jim.Elliott.Canada jlelliotton.blogspot.com

LinkedIn Twitter Facebook Blog