IBM z14

Franco Pinto
Client Technical Specialist
IBM Z Server Solutions

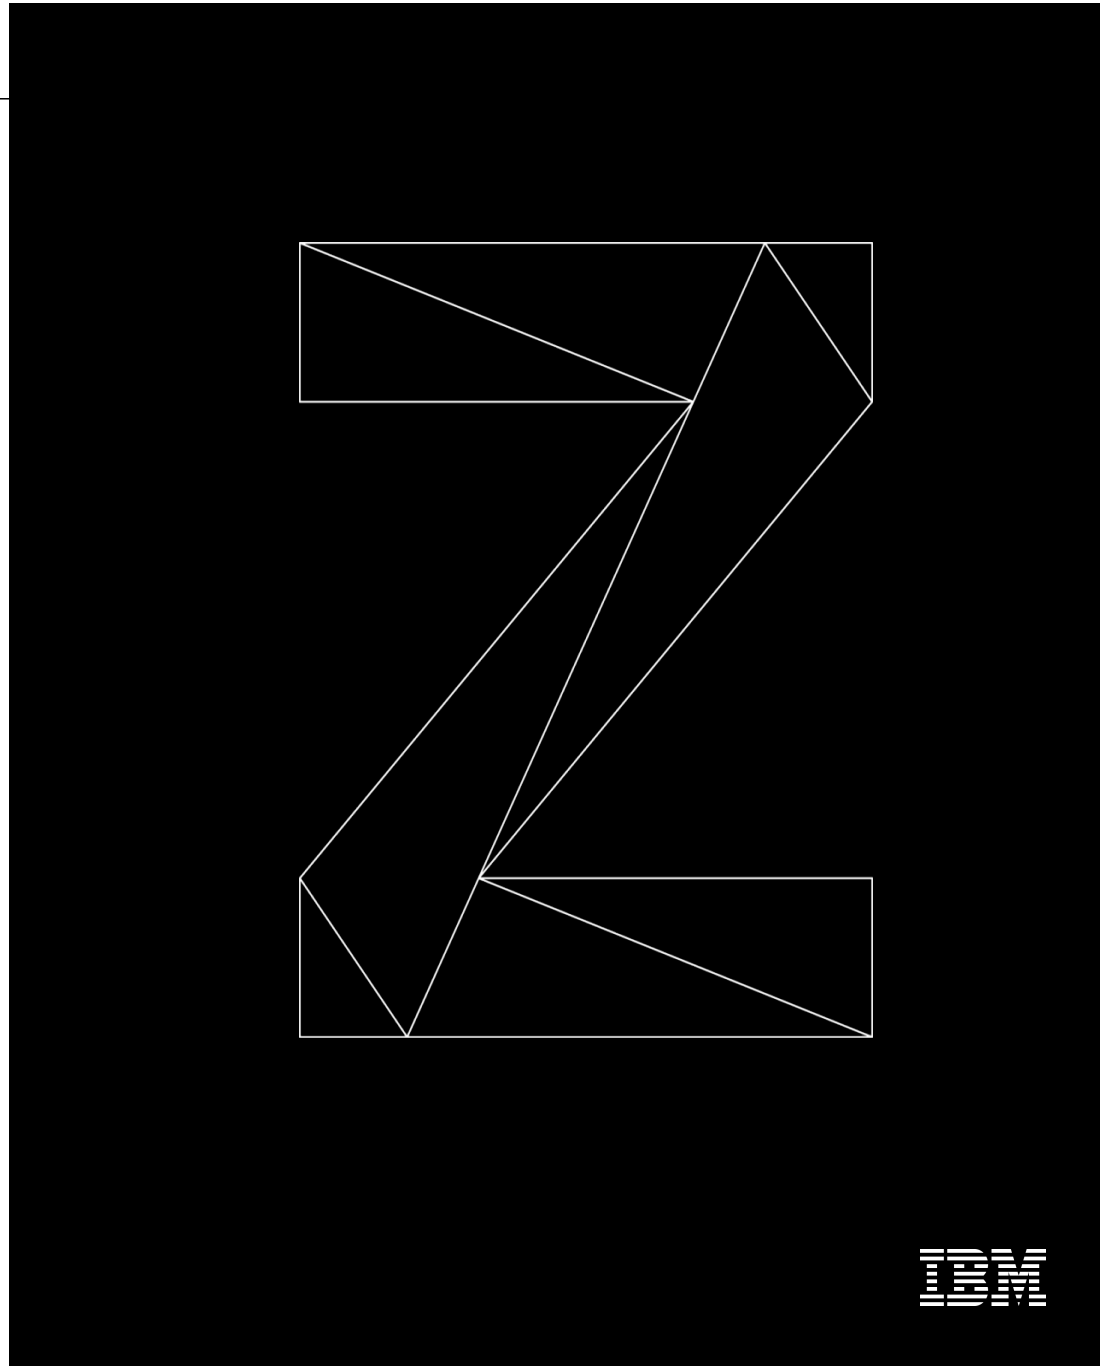

Trademarks, Thanks and Contact

This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:
http://www.ibm.com/legal/copytrade.shtml

Thanks to

- **Octavian Lascu**, IBM ITSO
- **Gregory Hutchison**, IBM Washington Systems Center
- **Riaz Ahmad**, IBM Washington Systems Center
- **Kenneth J Stine**, IBM Washington Systems Center
- **Michael Grossmann**, IBM Deutschland

Franco Pinto

Client Technical Specialist
IBM Global Markets, Systems HW Sales
Systems Technical Sales for IBM Z

Phone: +41-79-6255769
Mail: franco.pinto@ch.ibm.com

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 Sync 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 Sync 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.

Table of Content

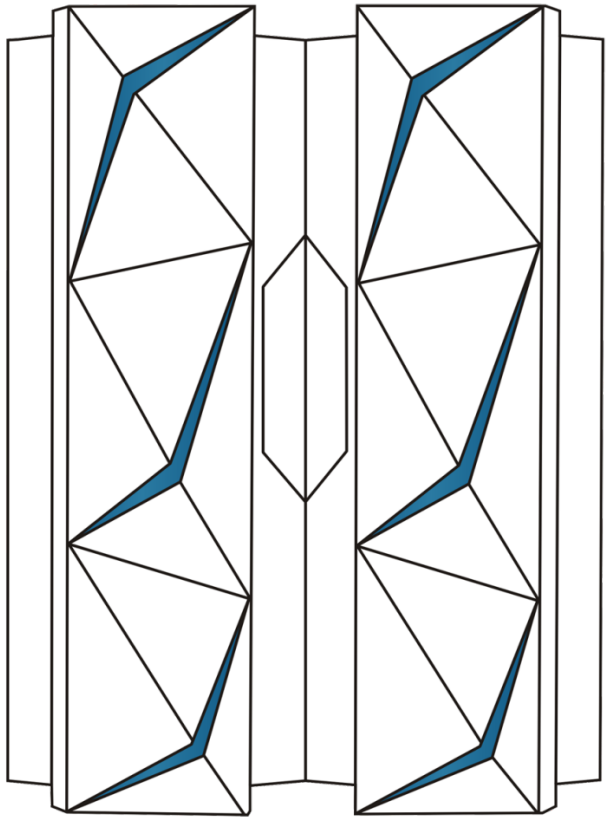

IBM Z: Designed for Trusted Digital Experiences

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

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

Designed for data serving in a cognitive world

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

The best infrastructure to support an open and connected world

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

IBM z14 at a glance

System, Processor, Memory
Five hardware models: M01, M02, M03, M04, M05
10 core 5.2GHz 14nm PU SCM
1 - 170 PUs configurable as CPs, zIIPs, IFLs, ICFs, up to 196 Pus
Increased Uniprocessor capacity
Up to 33 sub capacity CPs at capacity settings 4, 5, or 6
CPC Drawers and backplane Oscillator
Enhanced SMT and new instructions for SIMD
Enhanced processor/cache design with 1.5x more on-chip cache sizes
Up to 32 TB DRAM, protected by Redundant Array of Independent Memory (RAIM)
Virtual Flash Memory (VFM)
192 GB HSA
Improved pipeline design and cache management

Announce: July 17, 2017

I/O Subsystem, Parallel Sysplex, STP, Security
PCIe Gen3 I/O fanouts with 16 GBps Buses
6 CSS, 4 Subchannel sets per CSS
0 – 5 PCIe I/O Drawer Gen3 (no I/O Drawer)
Next generation FICON Express16S+
10 GbE RoCE Express2
Integrated Coupling Adapter (ICA SR) and Coupling express LR for coupling links
Support for up to 256 coupling CHPIDs per CPC
CFCC Level 22
Crypto Express6S and CMPSC compression and Huffman Coding compression
STP configuration and usability enhancements (GUI)
IBM zHyperLink Express
OSA-Express6S
Secure Service Container

RAS, simplification and others	
L3 Cache Symbol ECC	Acoustic and thin covers (space saving)
N+1 radiator design for Air Cooled System	Drop “Classic” HMC UI
ASHRAE Class A3 design	Enhanced SE and HMC Hardware (security)
Support for ASHRAE Class A3 datacenter	TKE 9.0 LICC
Largesum TCP/IP hardware Checksum (OSA-Express6S)	Pause-less garbage collection
Universal Spare SCM s (CP and SC)	Simplified and enhanced functionality for STP configuration
Enhanced Dynamic Memory Relocation for EDA and CDR	Virtual Flash Memory (replaces IBM zFlash Express)

PR/SM
Up to 170 CPUs per partition
IBM Dynamic Partition Manager updates
Up to 85 LPARs
16 TB Memory per partition

IBM z14 – At a Glance

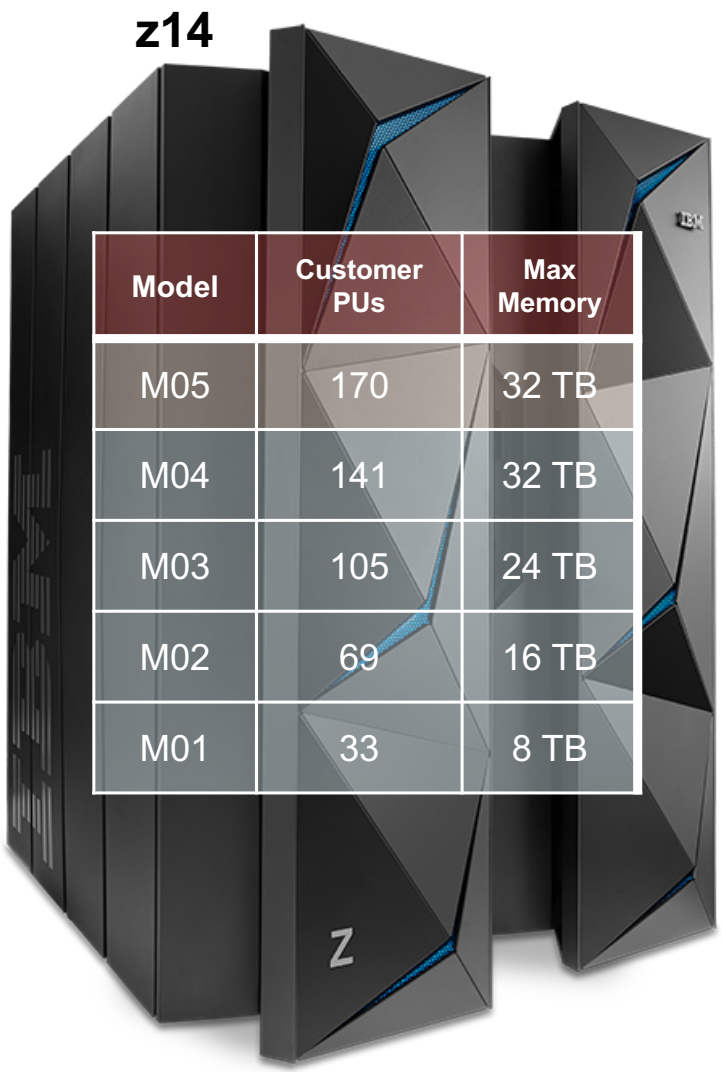

Model	Customer PUs	Max Memory
M05	170	32 TB
M04	141	32 TB
M03	105	24 TB
M02	69	16 TB
M01	33	8 TB

- Machine Type
  - 3906
- 5 Models
  - M01, M02, M03, M04 and M05
- Processor Units (PUs)
  - 41 (49 for M05) PU cores per CPC drawer
  - Up to 39 SAPs per system
  - 2 spares designated per system
  - Dependent on the H/W model - up to 33, 69, 105, 141,170 PU cores available for characterization
    - Central Processors (CPs), Internal Coupling Facility (ICFs), Integrated Facility for Linux (IFLs), IBM Z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs) and Integrated Firmware Processor (IFP)
    - 85 LPARs,
  - Sub-capacity available for up to 33 CPs
    - 3 sub-capacity points (4xx, 5xx, 6xx)
- Memory
  - RAIM Memory design
  - System Minimum of 64 GB
  - Up to 8 TB per drawer
  - Up to 32 TB for System and up to 16 TB per LPAR (OS dependent)
    - 192 GB Fixed HSA, standard
    - 32/64/96/128/256/512 GB increments
  - Virtual Flash Memory
- I/O
  - Up to 40 PCIe Gen3 Fanouts @ 16 GBps each and Integrated Coupling Adapters @ 2 x 8 GBps per System
  - Up to 16 InfiniBand Fanouts for HCA3-O and HCA3-O LR features
  - 6 Logical Channel Subsystems (LCSSs)
    - 4 Sub-channel sets per LCSS
  - Server Time Protocol (STP) Optional (No ETR)

IBM z14 – At a Glance (2)

- Scalable, robust, efficient and secure platform
  - Powerful core infrastructure
  - Unparalleled virtualization capabilities
  - Ease of use and flexibility
- Pervasive encryption – Built into and around the platform: data centric
  - Protecting the data throughout the journey: from data inception to data in flight to data at rest
  - Data access, data privacy, and data integrity: Applications, Databases, Bulk encryption
  - Special purpose hardware (specialized accelerators, co-processors), efficient use of CPU resources
  - Certifications – HSM, PCI, FIPS etc.,
- Workloads – production and development – time to value – IaaS, SaaS, Business Process as a service (Blockchain, SSC as infrastructure)
- Analytics and Cognitive, Machine learning
  - Keep data on the platform avoid the costs associated with movement of data , keep data current by continuous streaming of ingested data (as the data is created onto the platform it becomes available for analytics / cognitive / machine learning)
  - Exploit the integrated facilities (SMT, SIMD Decimal FPU, and so on)

z14 Continues the CMOS Mainframe Heritage

~ 10% for equal z13 n-way
Up to 35% max capacity 170-way vs 141-way (z13)
SMT vs Single Thread ~ 10–40% (average 25%)
- both zIIP & IFL
SMT z14 vs z13 ~ 15% (z/VM Guests)

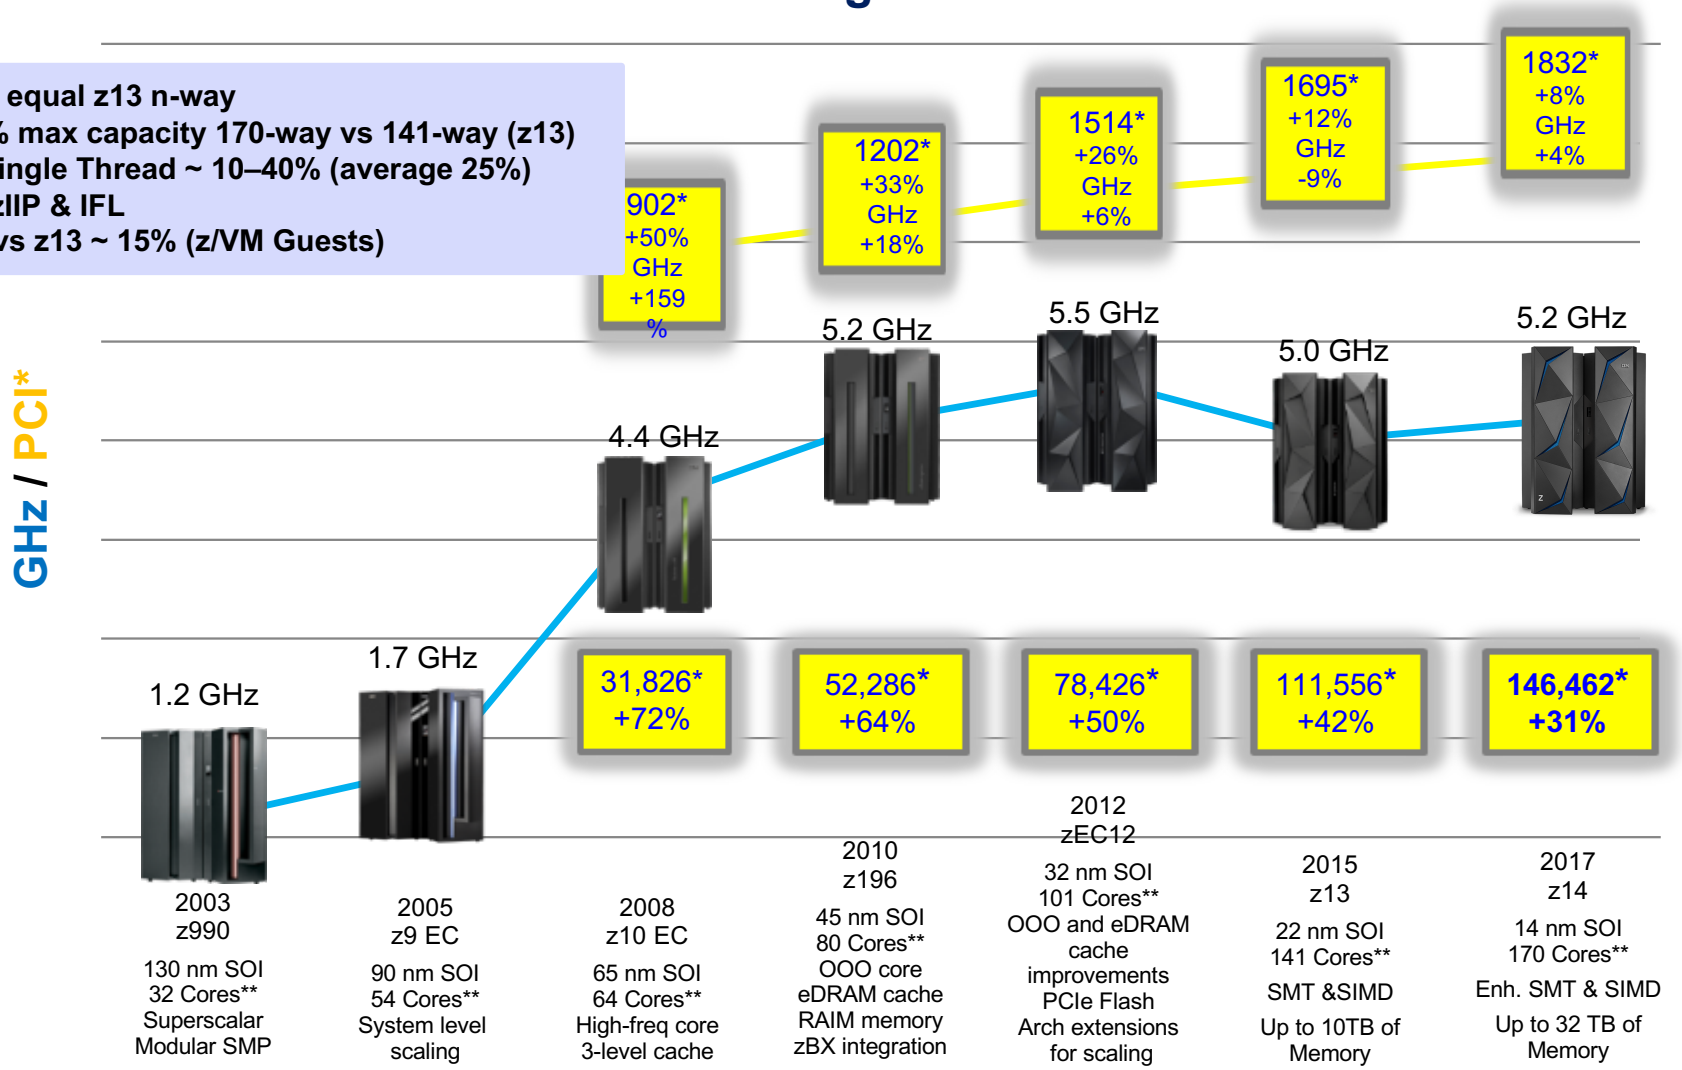

* MIPS Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
** Number of PU cores for customer use

Balanced system design

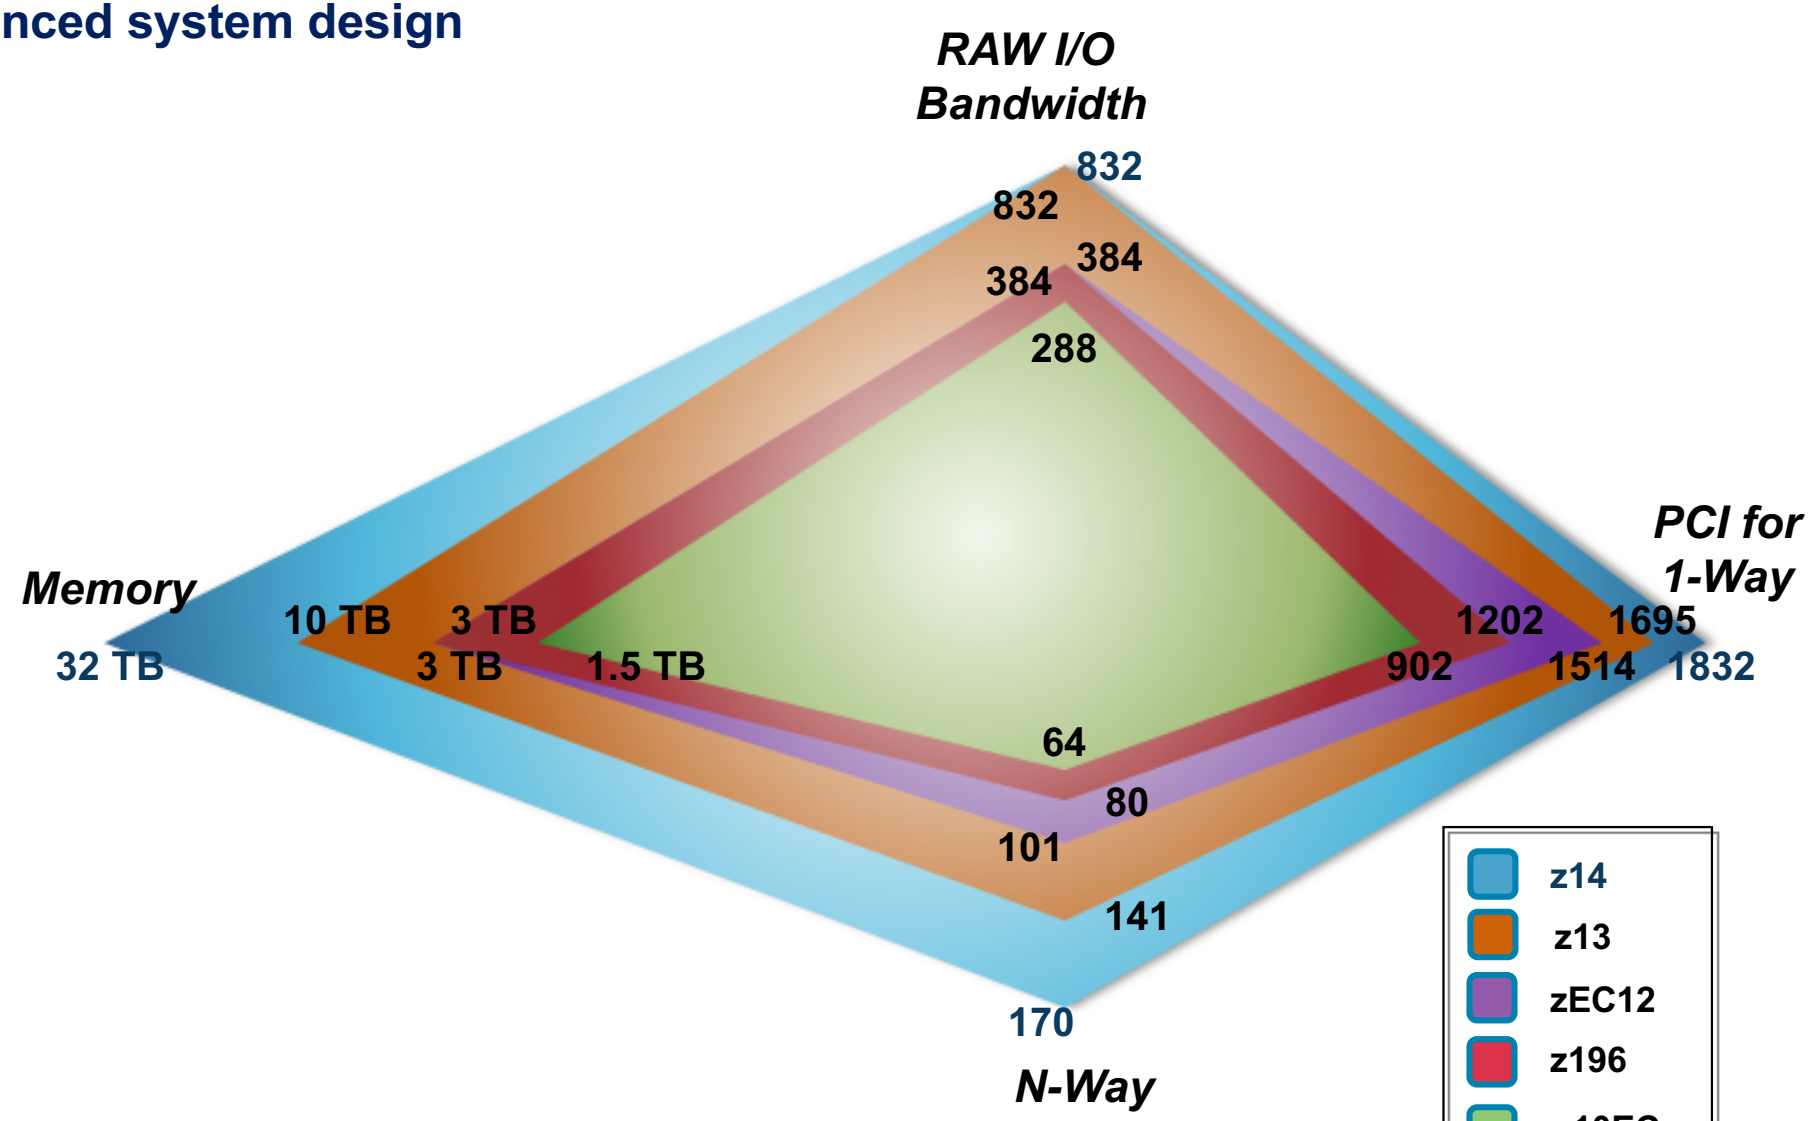

z14 System Design Changes

- 14 nm Processor with improved SIMD, SMT
- 10 Cores per CP SCM design
- 5 or 6 CP SCMs per Drawer
- Integrated I/O with PCIe Direct Attach
- Single System Controller Chip
- Simplified CPC Drawer SMP Fabric

- Crypto Express6S
- OSA-Express6S
- FICON Express16S+
- RoCE Express2
- IBM zHyperLink Express
- Coupling Express Long Reach
- Radiator Design improvements
- Expanded operating environment (ASHRAE Class A3)
- Thin doors (optional)

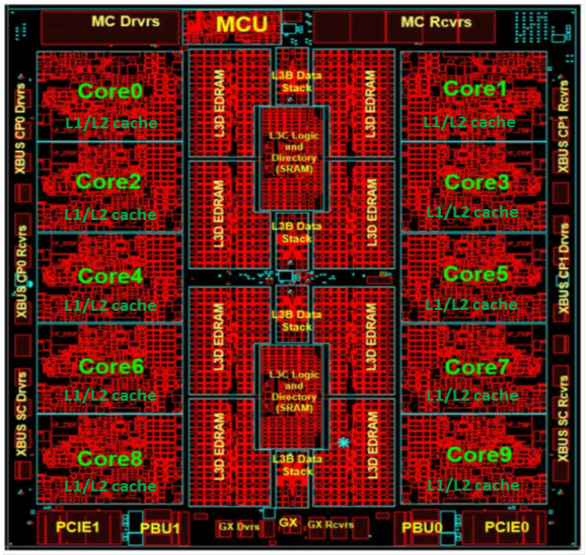

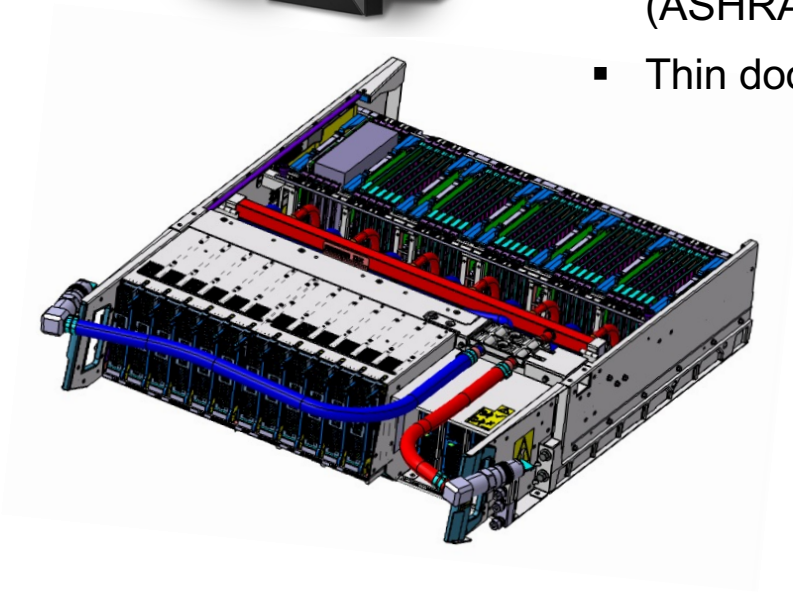

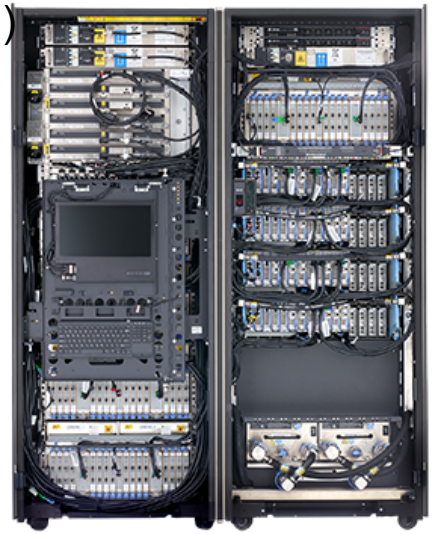

IBM z14 (M/T 3906)

- Announced – 07/2017, GA – 09/2017
- 5 models – M05, M04, M03, M02, M01
  - Up to 170 customer configurable engines
- Sub-capacity Offerings for up to 33 CPs
- PU (Engine) Characterization
  - CP, IFL, ICF, zIIP, SAP, IFP (No zAAPs)
- Enhanced SIMD and SMT
- On Demand Capabilities
  - CoD: CIU, CBU, On/Off CoD, CPE
- Memory – up to 32 TB
  - Up to 16 TB per LPAR (OS dependent)
  - 192 GB Fixed HSA
  - Virtual Flash Memory (zFlash Express replacement
- Channels
  - PCIe Gen3 16 GBps channel buses
  - Six LCSSs, up to 85 LPARs
  - Four Subchannel Sets per LCSS
  - FICON Express16S+ (New build) or 16S and 8S (Carry forward)
- Channels (cont.)
  - IBM zHyperLink Express
  - OSA-Express6 (5S carry forward)
  - HiperSockets™ – up to 32
  - zEnterprise Data Compression (zEDC)
  - Enhanced 10 GbE RoCE Express2 (10 GbE RoCE Express Carry forward)
  - Shared Memory Communications - Direct Memory Access (SMC-D)
- Crypto Express6S (5S Carry forward)
- Parallel Sysplex clustering:
  - ICA SR (PCIe) Coupling,
  - InfiniBand Coupling (HCA3-O, HCA3-O LR)
  - Coupling Express Long Reach
  - Coupling Facility Control Code Level 22
- Operating Systems
  - z/OS®, z/VM®, z/VSE, z/TPF, Linux on z Systems , KVM for IBM Z
- IBM Dynamic Partition Manager (DPM)
- IBM Secure Service Container

Table of Content

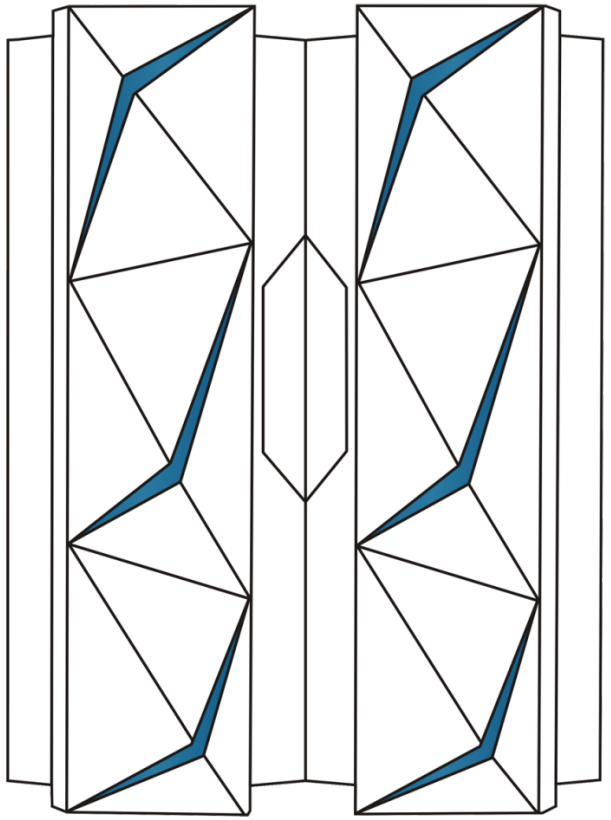

z14 Radiator-based Air cooled – Front View (Model M04 or M05)

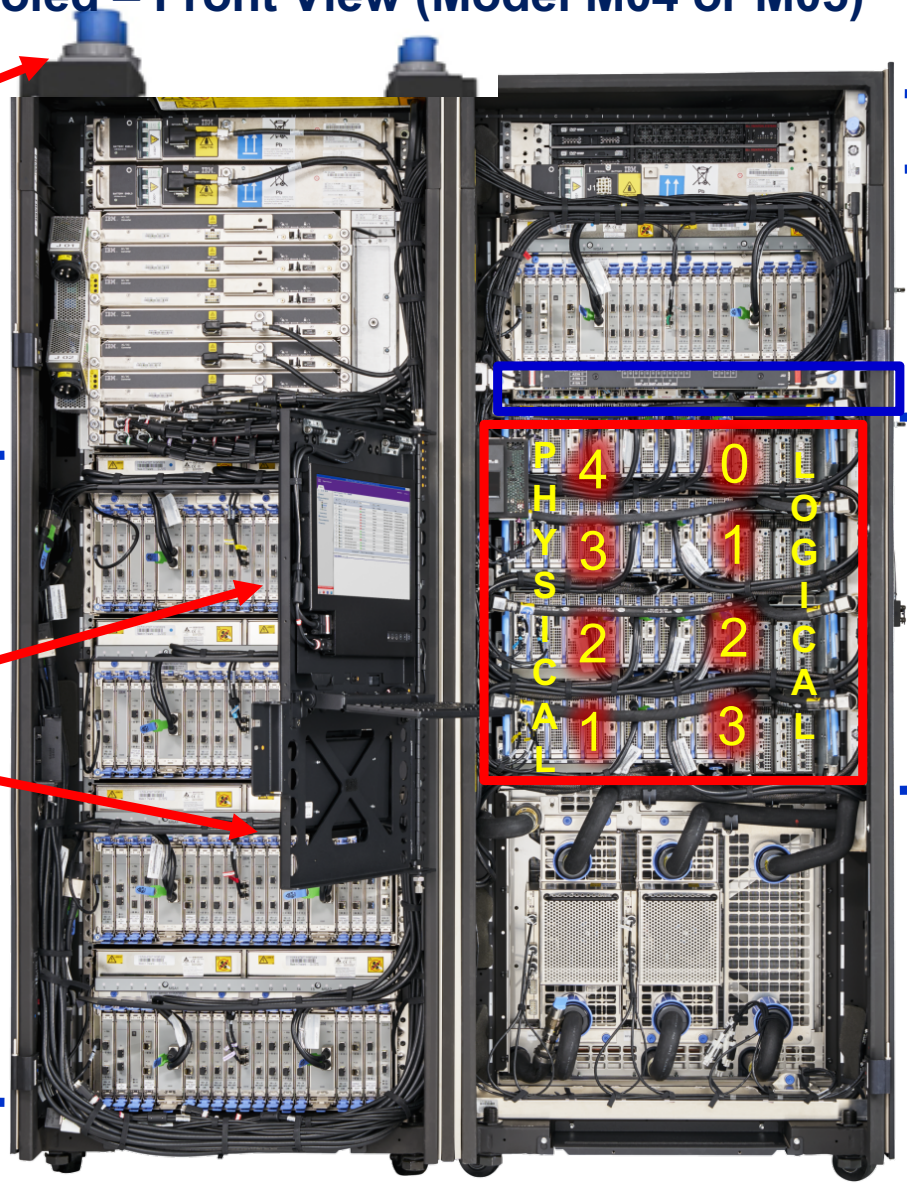

Overhead Power Cables (option)

Internal Batteries (optional)

Power Supplies

Displays and keyboards for Support Elements

PCIe I/O drawers numbers 1 to 4

2 x 1U Support Elements

PCIe I/O drawer number 5

System Control Hubs (used to be BPHs)

CPC Drawers, PCIe Fanouts, Cooling water manifold and pipes, PCIe I/O interconnect cables, FSPs and Ethernet cables

Note: CPC Drawer plugging numbers are on the left and logical numbers on the right

Radiator Pumps

Overhead I/O feature is a co-req for overhead power option

z14 Processor Drawer (Top View)

- 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
  - Ten fanout slots for PCIe I/O drawer fanouts or PCIe coupling fanouts
  - Four fanout slots for PSIFB coupling link fanouts

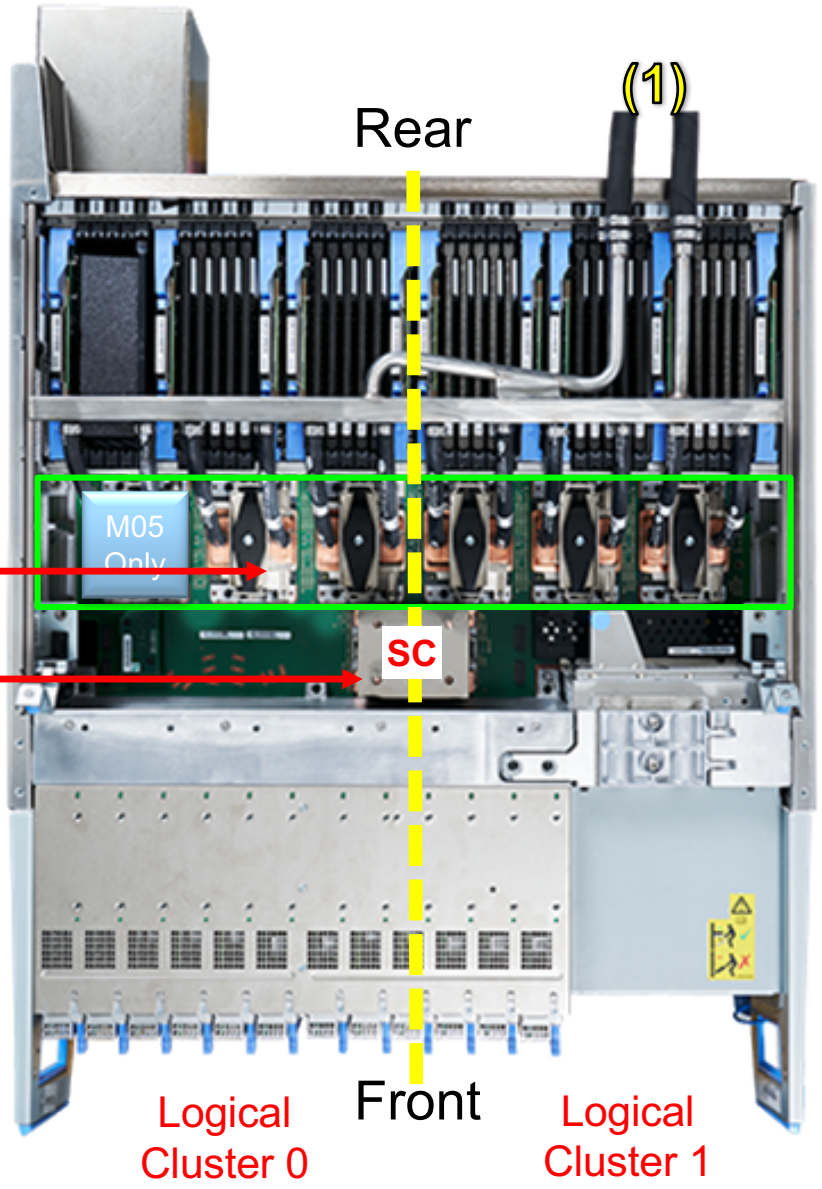

z14 PU and SC SCM assembly

PU SCM (uncapped)

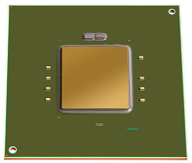

6x PU SCMs

Capped PU Chips

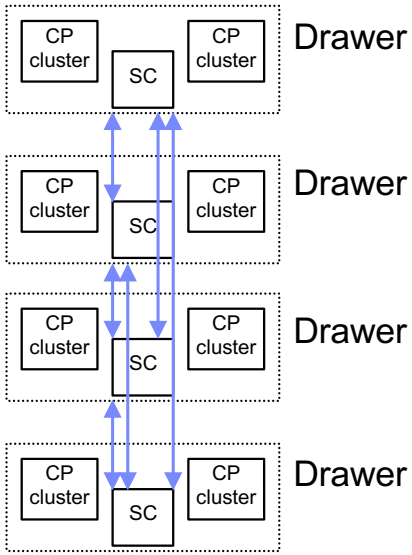

4 Drawer System
Fully Interconnected

1x SC SCM (Air Cooled)

6x PU SCMs under the cold-plates

Capped SC

SC SCM with Thermal Module

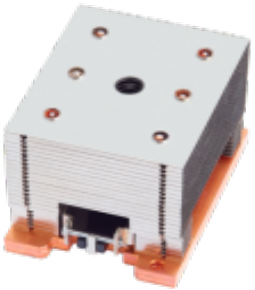

SC Chip

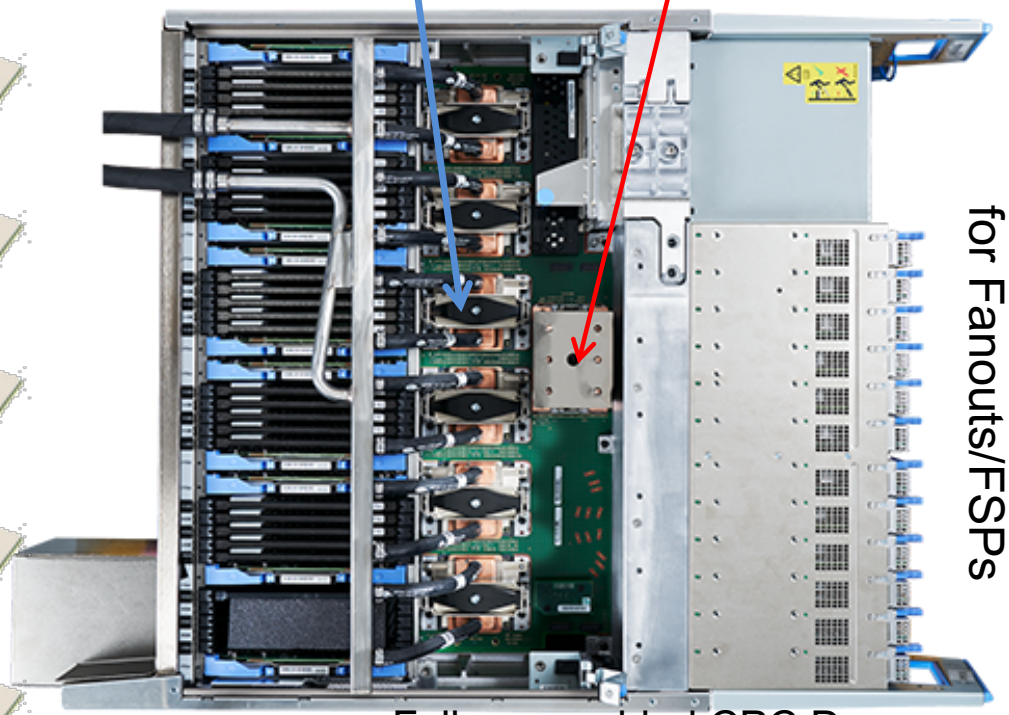

Front of CPC Drawer for Fanouts/FSPs

Fully assembled CPC Drawer with the chilled water supply manifold lifted to the left

IBM Z – Processor Roadmap

65 nm
z10
2/2008

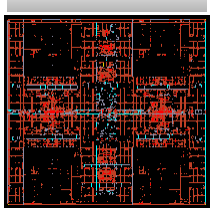

- Workload Consolidation and Integration Engine for CPU Intensive Workloads
- Decimal FP
- Infiniband
- 64-CP Image
- Large Pages
- Shared Memory

45 nm
z196
9/2010

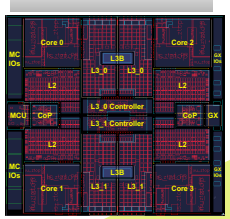

- Top Tier Single Thread Performance,System Capacity
- Accelerator Integration
- Out of Order Execution
- Water Cooling
- PCIe I/O Fabric
- RAIM
- Enhanced Energy Management

32 nm
zEC12
8/2012

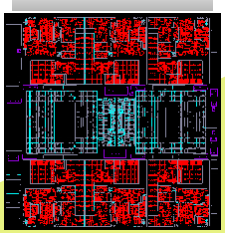

- Leadership Single Thread, Enhanced Throughput
- Improved out-of-order
- Transactional Memory
- Dynamic Optimization
- 2 GB page support
- Step Function in System Capacity

22 nm
z13
1/2015

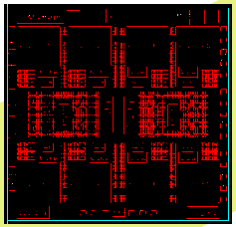

- Leadership System Capacity and Performance
- Modularity & Scalability
- Dynamic SMT
- Supports two instruction threads
- SIMD
- PCIe attached accelerators
- Business Analytics Optimized

14 nm
z14
7/2017

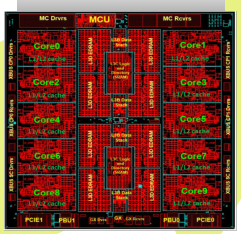

- Pervasive encryption
- Low latency I/O for acceleration of transaction processing for DB2 on z/OS
- Pause-less garbage collection for enterprise scale JAVA applications
- New SIMD instructions
- Optimized pipeline and enhanced SMT
- Virtual Flash Memory

z14 processor design summary

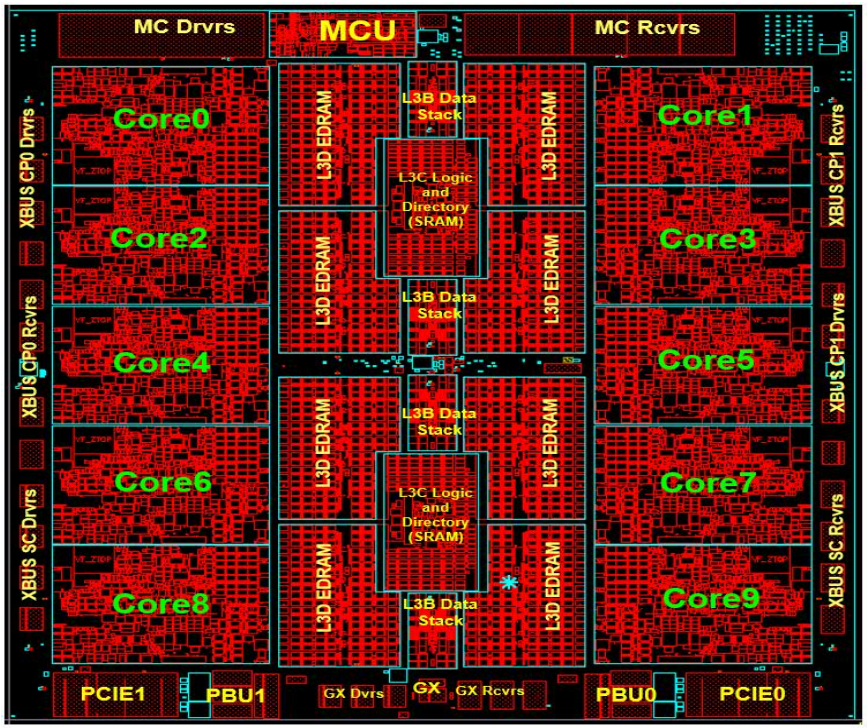

- 14nm SOI technology,
- 17 layers of metal
- 10 cores per CP-chip,
- 5.2GHz

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

Performance and Capacity

- ~ 10% for equal z13 n-way (on average)
- ~ 31% max capacity 170-way vs 141-way (z13)
- SMT vs Single Thread ~ 10-40% (average 25%)
  - for both zIIP and IFL
- SMT z13 vs z14 ~ 15% (z/VM Guests)

Where does the performance come from?

- It is not just Cache
- It is not just Cycle Time
- 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

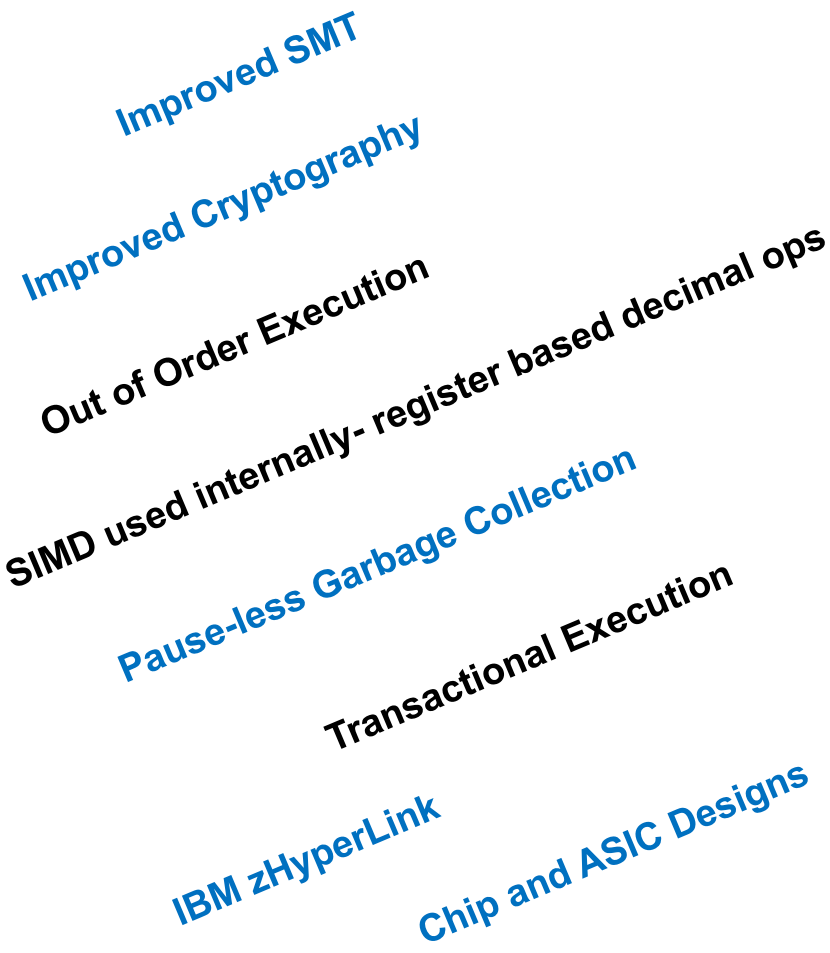

Performance drivers with z14

- 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)
- 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
- 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

Instruction Execution Protection Facility (no-execute memory)

Linux guest optimization

- Through enhanced hardware features and explicit software requests to obtain memory areas as non-executable it will be possible to protect areas of memory from unauthorized execution. Attempts to execute instructions from these areas will result in protection exception errors.

- The intent is to mark memory regions as non-executable, an attempt to execute instruction code in these regions will cause an exception

- LE (Language Environment) will mark the stack and heap as non-executable memory to prevent certain buffer overflow attacks from succeeding.
  - Will prevent malicious users from being able to insert executable code into an area of memory (e.g. a heap or stack) and cause it to get control in an unauthorized manner.

- New keyword on virtual memory obtains to indicate whether memory will be used to contain executable code
  - EXECUTABLE=YES|NO on STORAGE OBTAIN or IARV64

z14 Processor Enhancements: Guarded Storage Facility (GSF) for Pause-less Garbage Collection

- Problem:
  - When garbage collection occurs today, all threads running under a JVM must stop
  - Customers are consolidating from multiple to single JVM environments to increase productivity and save money.
  - The consolidation effort generates heap sizes >100GB where garbage collection pauses can take minutes!
  - Long pause times cause transactional application failures and SLA violations.
- Solution:
  - Define flexible new architecture that provides hardware assisted read barriers for guarded storage involved in a garbage collection/compaction event.
  - Whenever a Pointer is loaded from memory, the pointer is checked against a pending GC, and in case of a “hit”, the control flow is redirected
  - The Dynamic Runtime can then assist in GC-ing the pointed-to object, before resuming the SW thread.
  - Software exploitation of fast hardware barrier detection and acceleration will allow application threads to run concurrently during the majority of garbage collection events
- Impact:
  - Reduces worst case latency impacts for critical applications like financial trading platforms
  - Maintains SLAs, and keeps IBM Z servers in our customers modernization roadmaps.

Guarded Storage Facility (GSF) for Pause-less Garbage Collection

- Designed to improve the performance of Java while garbage collection is active
- Up to 64 regions of memory can be marked as guarded
- Reading a pointer with the new LGG or LLGFSG instruction will do a range check on the loaded value and automatically invoke a user space handler if one of the guarded regions is affected

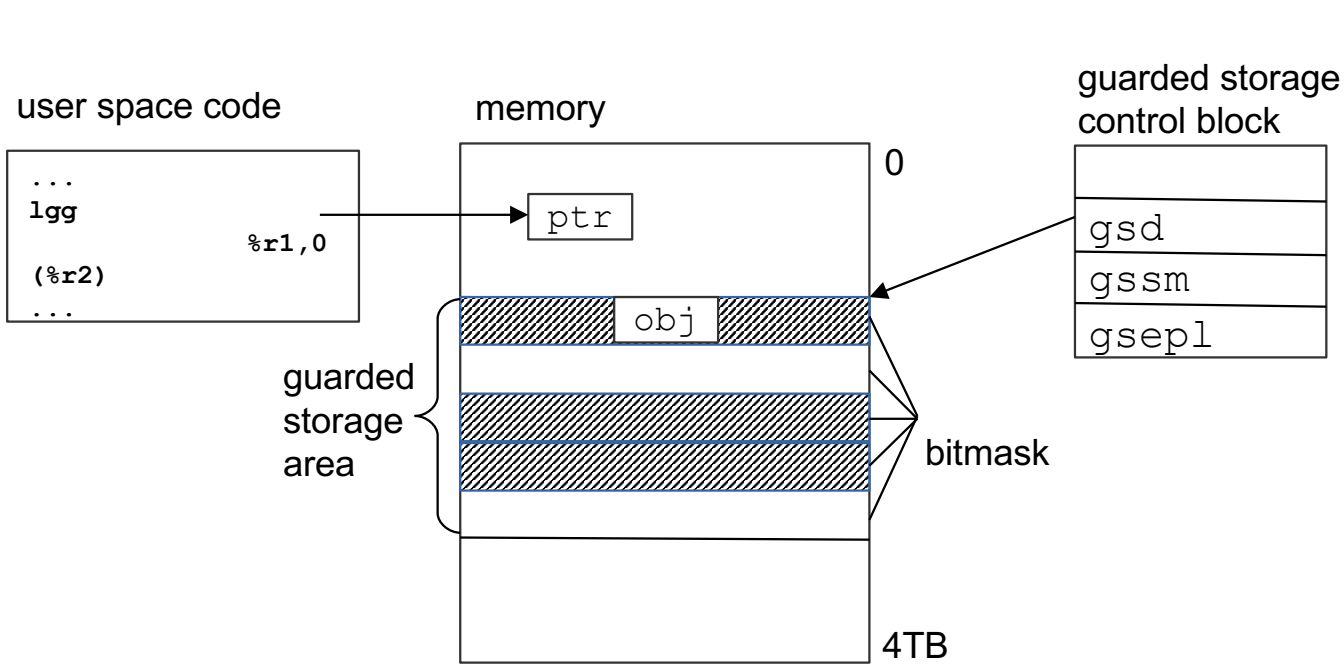

implicit branch to the guarded storage event handler

guarded storage event parm. list

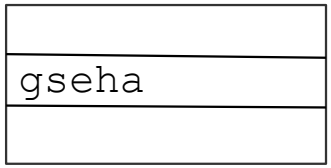


```
lgr   %r14,%r15
aghi  %r15,-320
stmg  %r0,%r14,192(%r15)
stg   %r14,312(%r14)
la    %r2,160(%r15)
stgsc 160(%r15)
lg    %r14,24(%r2)
lg    %r14,40(%r14)
la    %r14,6(%r14)
stg   %r14,304(%r15)
brasl %r14,gs_handler
lmg   %r0,%r15,192(%r15)
br      %r14
```

z14 On-Drawer and System Topology

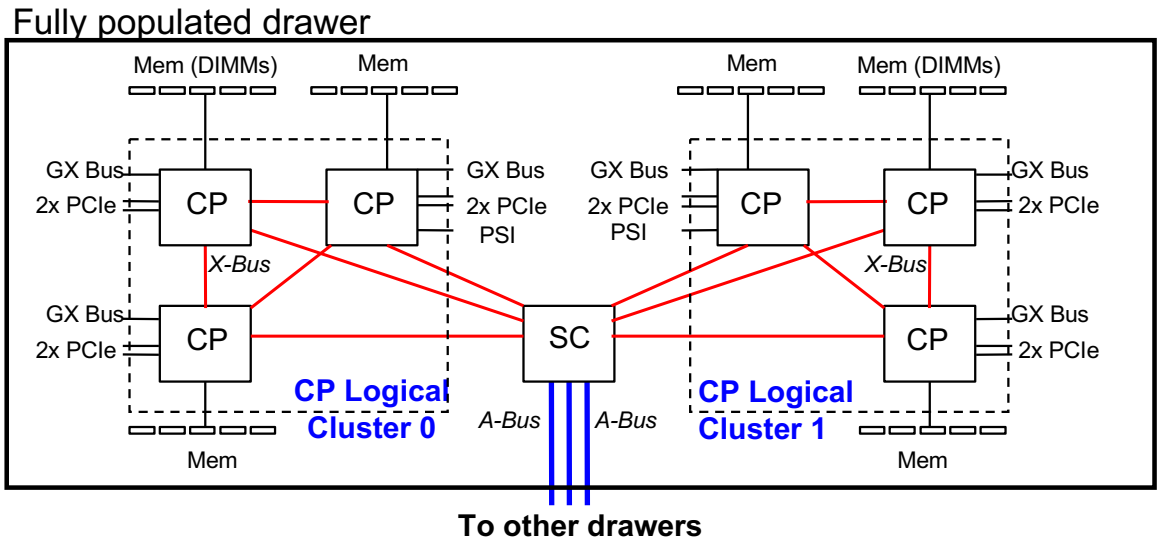

- Chips
  - Five or six PU chips
  - One SC chip (672 MB L4 cache)
- RAIM Memory
  - One Memory Controller per CP Chip
  - Five DDR4 DIMM slots per Controller, 25 total per drawer
- SC and CP Chip Interconnects
  - X-bus: SC and CPs to each other
  - A-bus: SC to SC chips in the remote drawers

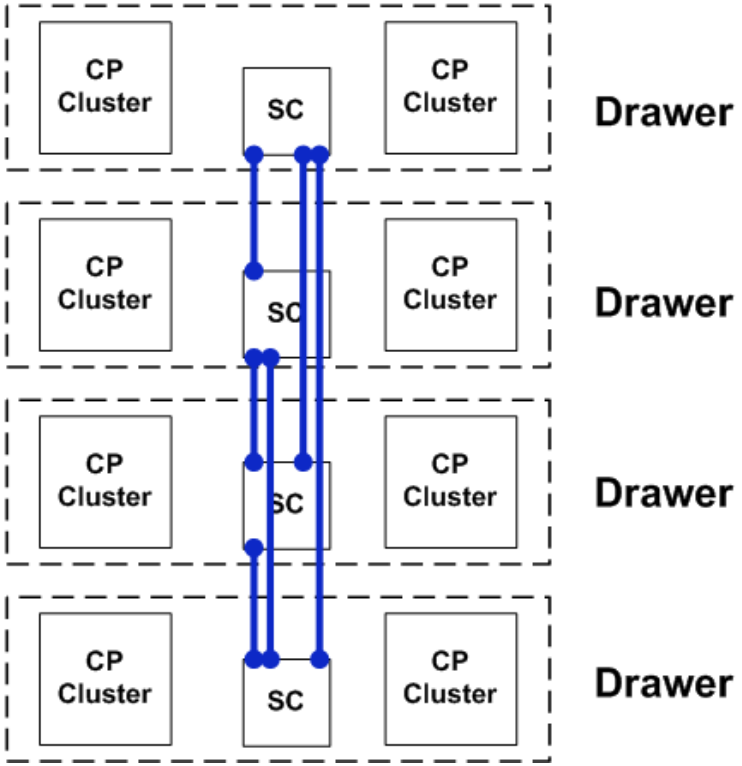

4 Drawer System
Fully Interconnected

z13 vs z14 Cache Topology Comparison

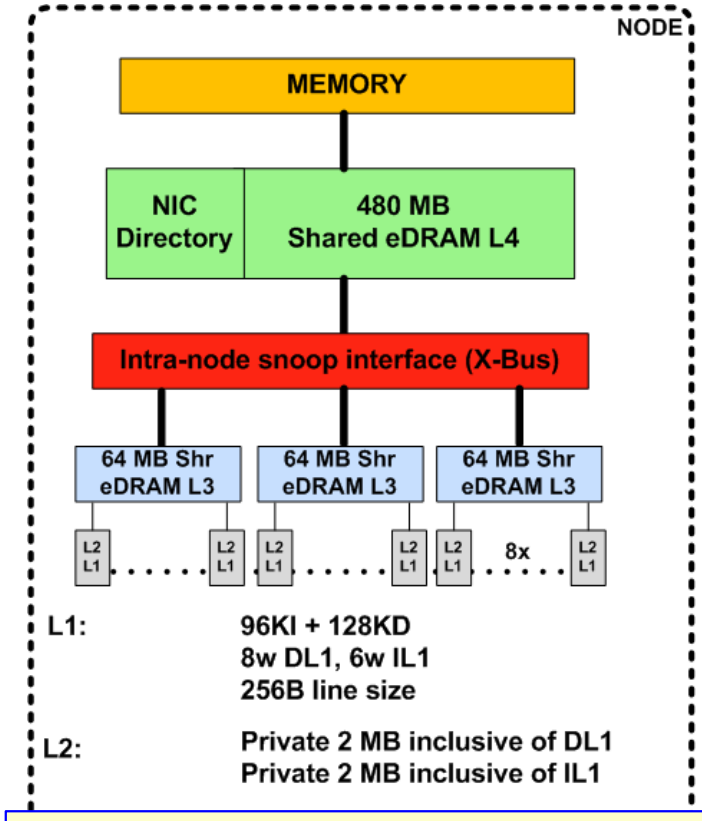

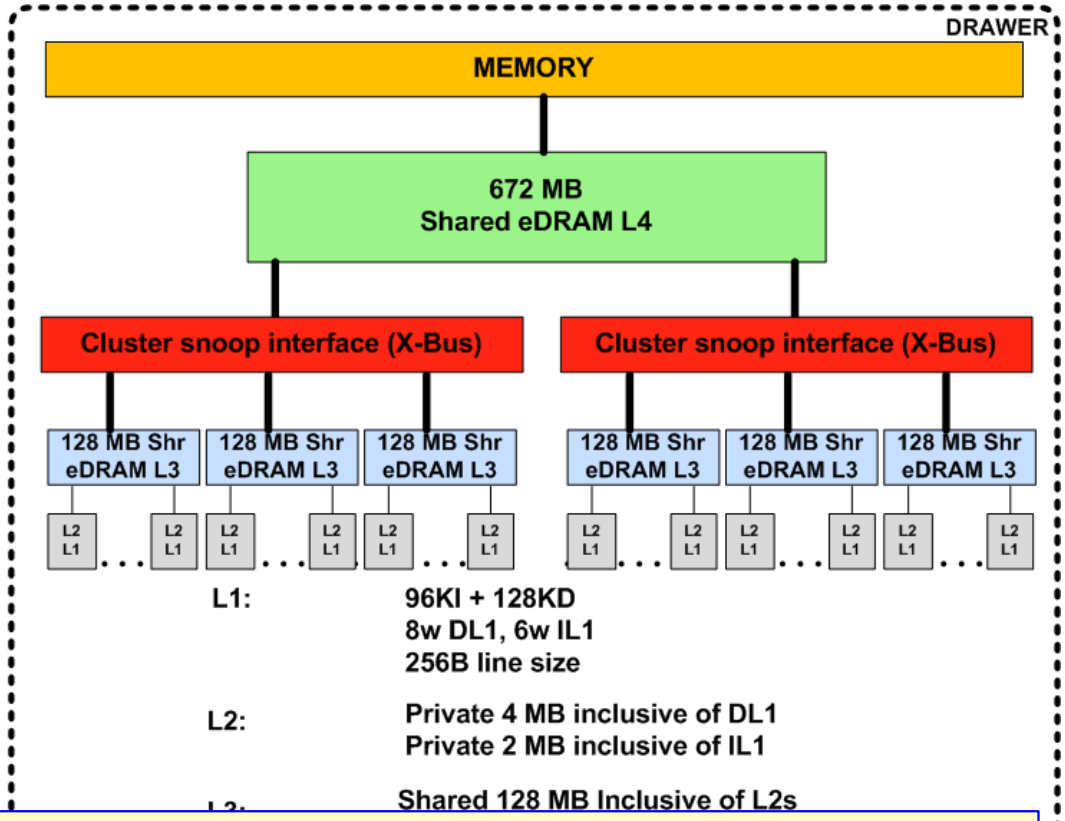

L1 (Instruction) increased by 33%. 96K[i] 128K[d] -to- 128K[i] 128K[d]
L2 (Data) increased by 2x. 2MB[i] 2MB[d] -to- 2MB[i] 4MB[d]
L3 increased by 2x. 64MB -to- 128MB
L4 from 960MB -to- 672MB per drawer
- Single System Controller (SC) latency Improvements, thus less L4 cache required.

Co-processor Compression (CMPSC) Improvements

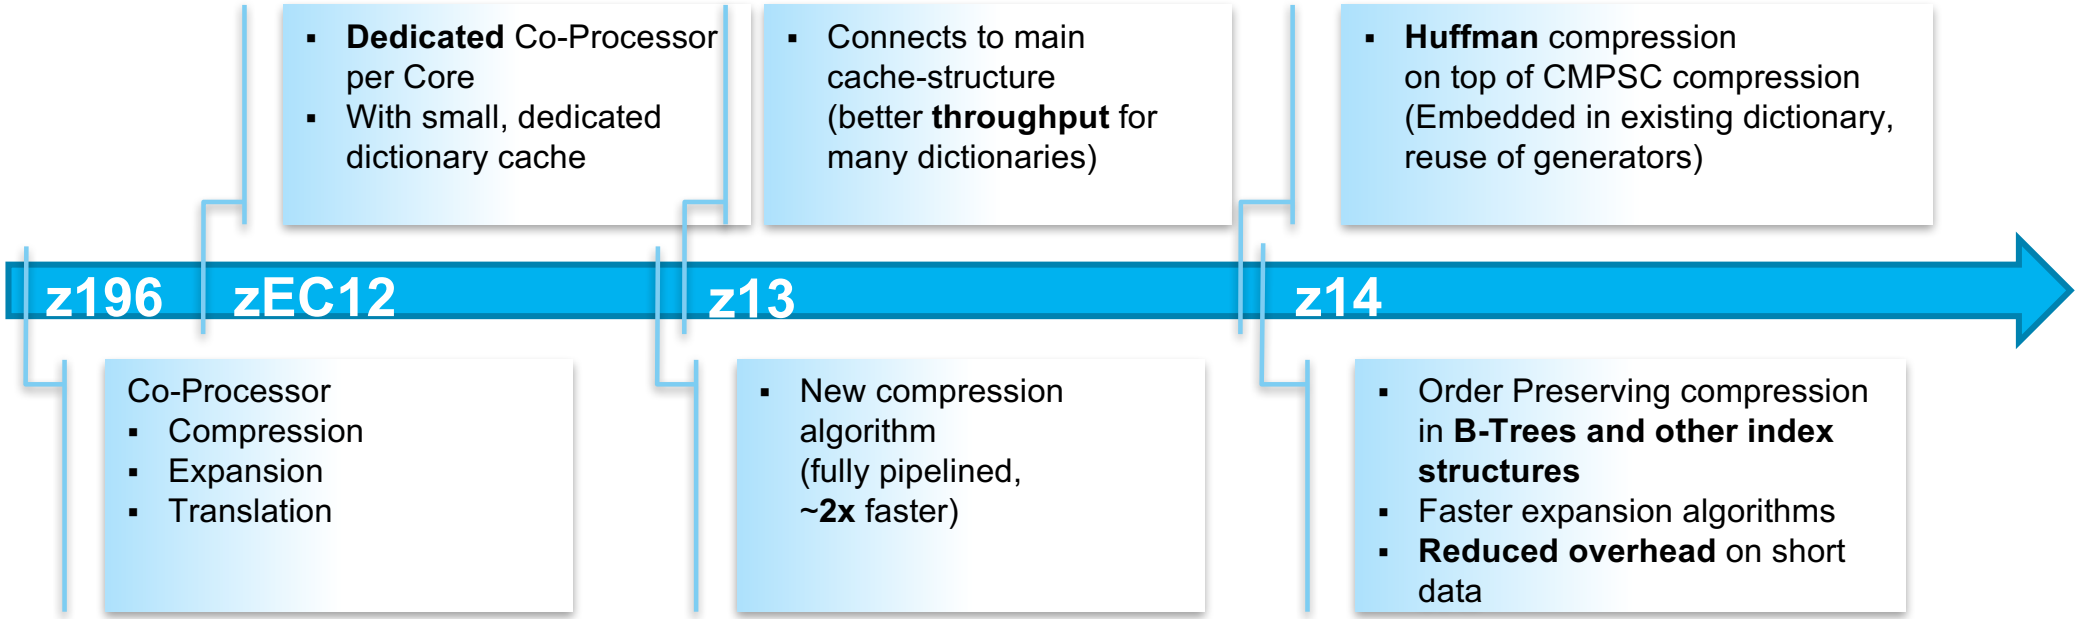

Compression/Expansion Strategy (z14 and beyond)

- Improved compression/expansion performance, in particular reduced overhead for short compression
  → disk/memory savings by exploiting more compression with low CPU consumption
- Improved compression ratio (e.g. Huffman coding in z14)
  → additional disk/memory savings even where compression is already in use today
- Order-preserving compression for search trees, sort files, etc.
  → disk/memory savings for large parts of data that have not been practical to compress previously

Crypto Functions in the Dedicated Co-processor – new for z14

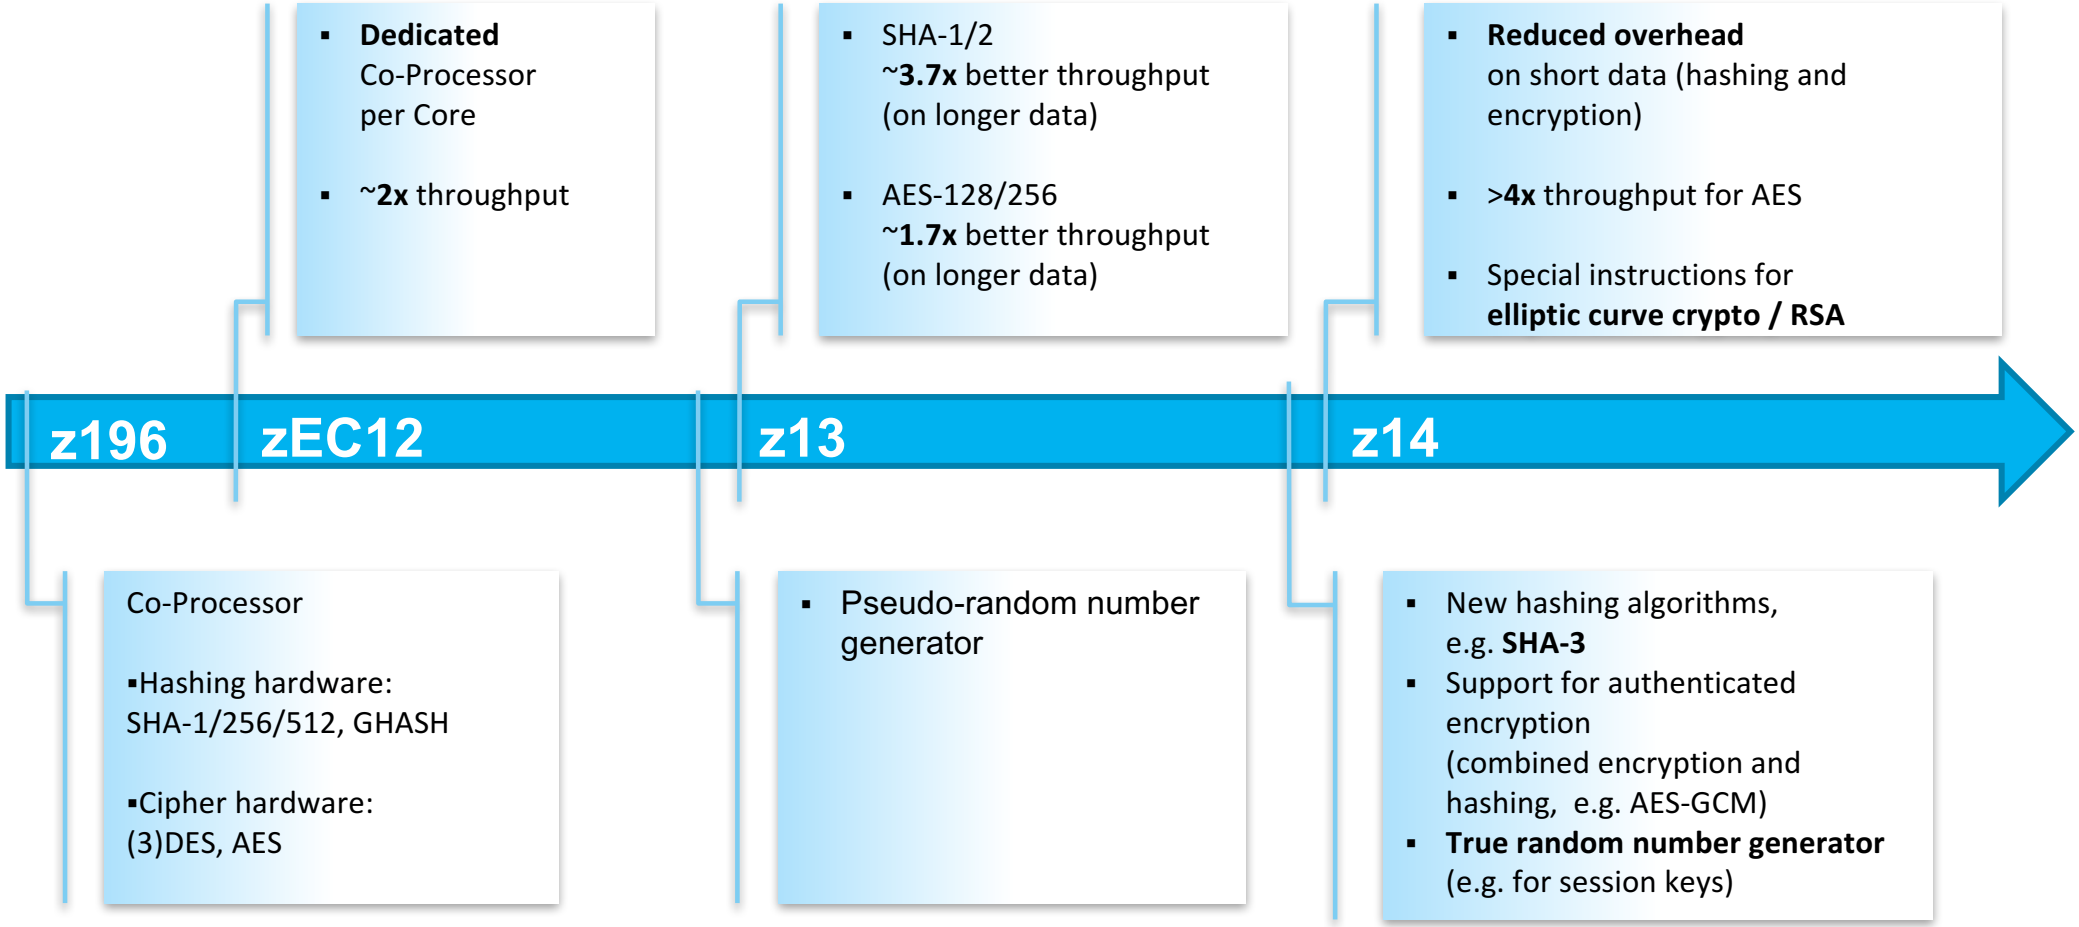

- New instruction added : KMGCM for end to end implementation of NIST GCM standard. (800-38D)
- KIMD / KLMD extended to implement SHA-3 standard. (FIPS 202)
- AES throughput improved to 3.5 to 4B/cycle

z14 Basics of Capacity on Demand (function unchanged)

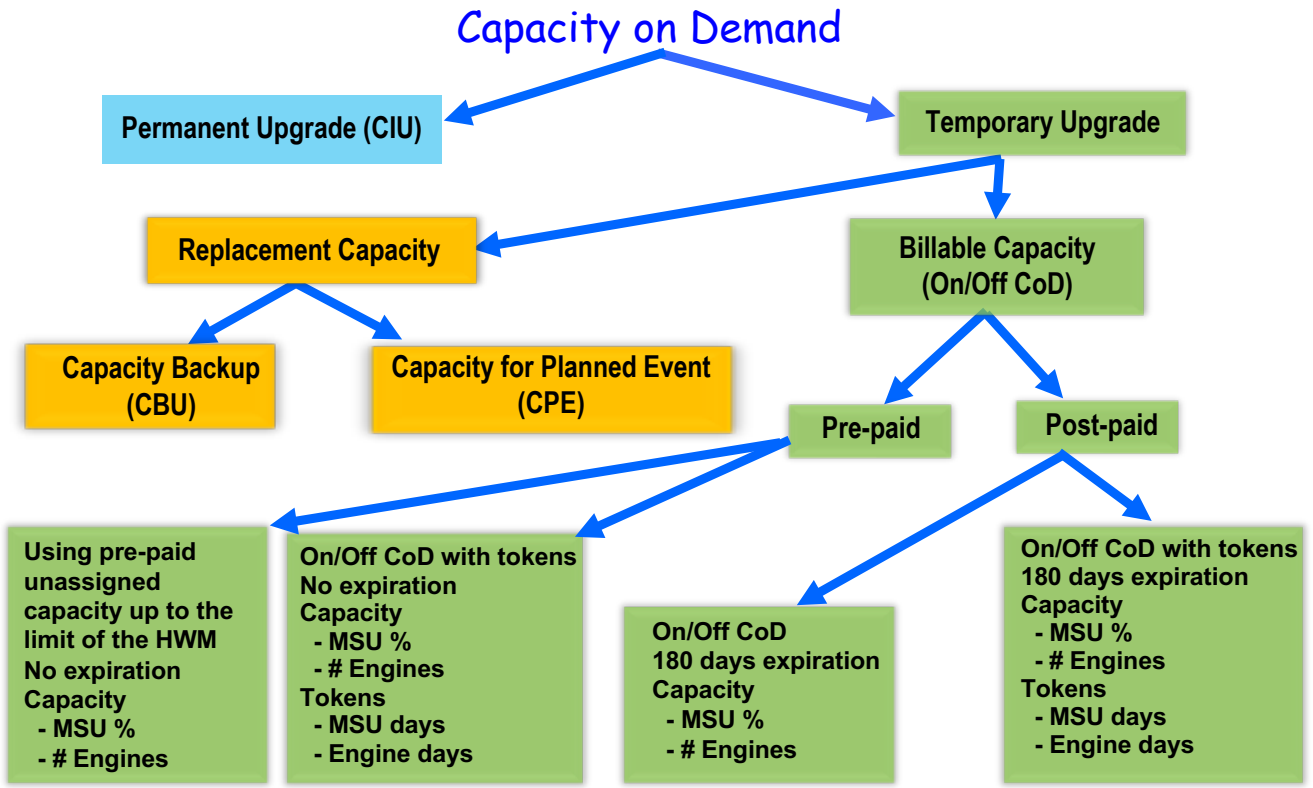

Table of Content

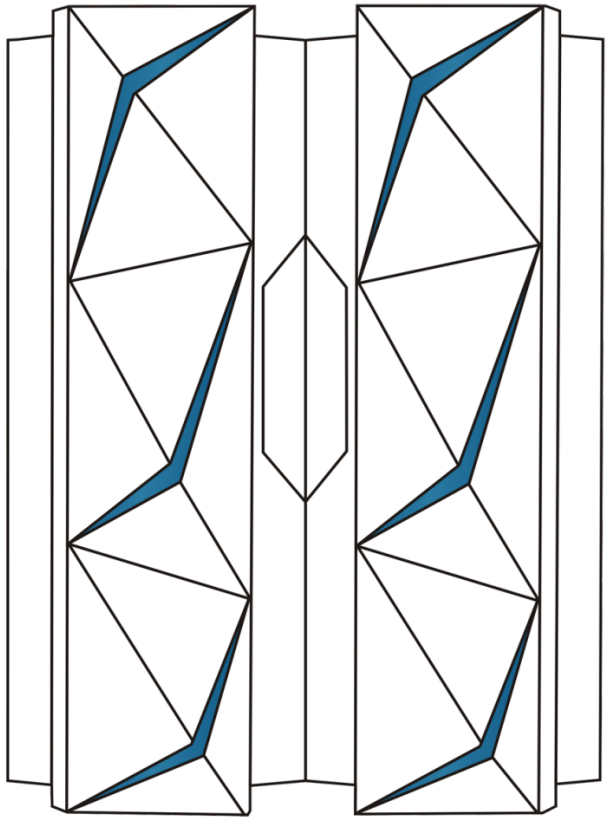

z13 z/Architecture Extensions

- Simultaneous multithreading (SMT) operation
  - Up to two active execution threads per core can dynamically share the caches, TLBs and execution resources of each IFL and zIIP core. SMT is designed to improve both core capacity and single thread performance significantly.
  - PR/SM online logical processors to dispatches physical cores; but, an operating system with SMT support can be configured to dispatch work to a thread on an IFL or zIIP core in single thread or SMT mode so that HiperDispatch cache optimization is considered. (Zero, one or two threads can be active in SMT mode). Enhanced hardware monitoring support will measure thread usage and capacity.
- Core micro-architecture radically altered to increase parallelism
  - New branch prediction and instruction fetch front end to support SMT and to improve branch prediction throughput.
  - Wider instruction decode, dispatch and completion bandwidth: Increased to six instructions per cycle compared to three on zEC12
  - Larger instruction issue bandwidth: Increased to up to 10 instructions issued per cycle (2 branch, 4 FXU, 2 LSU, 2 BFU/DFU/SIMD) compared to 7 on zEC12
  - Greater integer execution bandwidth: Four FXU execution units
  - Greater floating point execution bandwidth: Two BFUs and two DFUs; improved fixed point and floating point divide
- Single Instruction Multiple Data (SIMD) instruction set and execution: Business Analytics Vector Processing
  - Data types: Integer: byte to quad-word; String: 8, 16, 32 bit; binary floating point
  - New instructions (139) include string operations, vector integer and vector floating point operations: two 64-bit, four 32-bit, eight 16-bit and sixteen 8-bit operations.
  - Floating Point Instructions operate on newly architected vector registers (32 new 128-bit registers). Existing FPRs overlay these vector registers.

z Systems SMT Exploitation -- Core, Thread dispatching

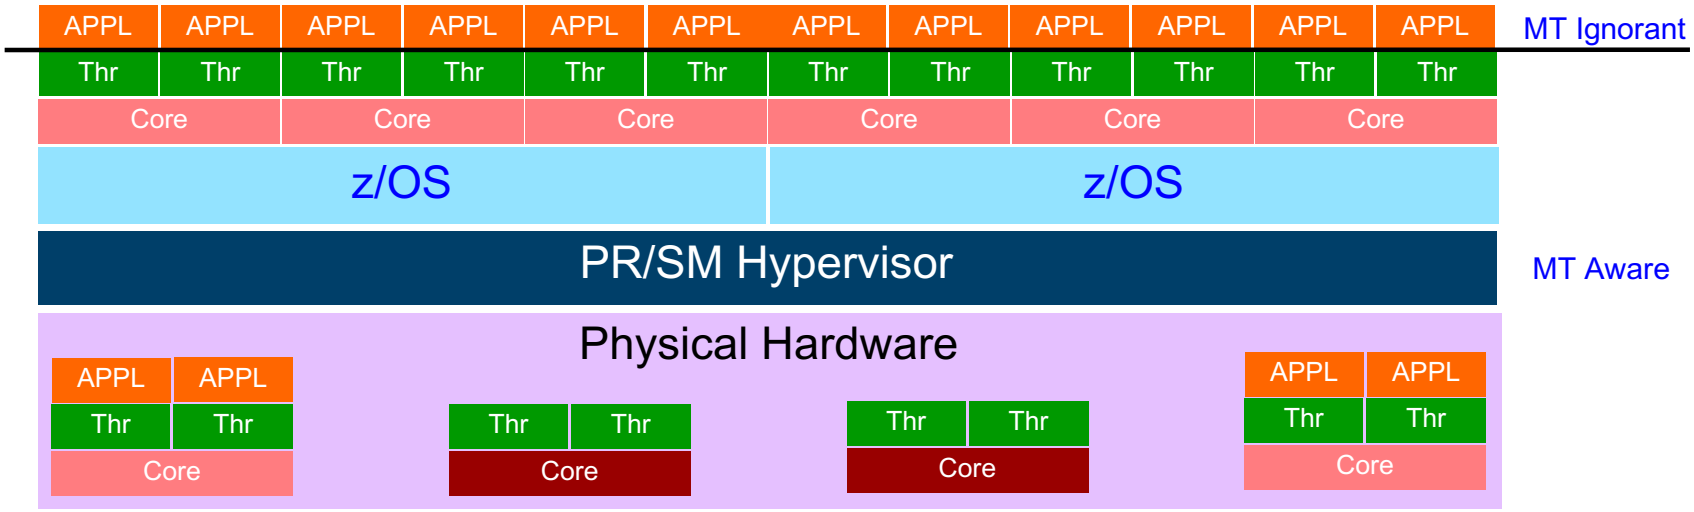

- PR/SM supports SMT for SMT aware OS like z/OS via core dispatching
- z/OS controls and manages whole core (all threads) to:
  - Maximize core throughput (fill running cores, spill to waiting cores)
  - Maximize core availability (meet goals using fewest cores)
- Limits SMT variability to a single z/OS workload
  - Makes capacity, accounting, latency, response time more predictable and repeatable
- SMT is transparent to applications

Simultaneous Multithreading (SMT) on z13

- Simultaneous multithreading allows instructions from one or two threads to execute on a zIIP or IFL processor core.
- SMT helps to address memory latency, resulting in an overall **capacity*** (throughput) improvement per core
- Capacity improvement is variable depending on workload. For **AVERAGE** workloads the **estimated** capacity* of a z13 zIIP/IFL with exploitation of the SMT option is**:**
  - zIIP is 38% greater than a zEC12 zIIP
  - IFL is 32% greater than a zEC12 IFL
  - zIIP is 72% greater than a z196 zIIP
  - IFL is 65% greater than a z196 IFL
- SMT exploitation: z/VM V6.3 + PTFs for IFLs and z/OS V2.1 + PTFs in an LPAR for zIIPs
- The use of SMT mode can be enabled on an LPAR by LPAR basis via operating system parameters.
  - When enabled, z/OS can transition dynamically between MT-1 (multi thread) and MT-2 modes with operator commands.
- Notes:
  1. SMT is designed to deliver better overall capacity (throughput) for many workloads. Thread performance (instruction execution rate for an individual thread) may be faster running in single thread mode.
  2. Because SMT is not available for CPs, LSPR ratings do not include it

*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. 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 .

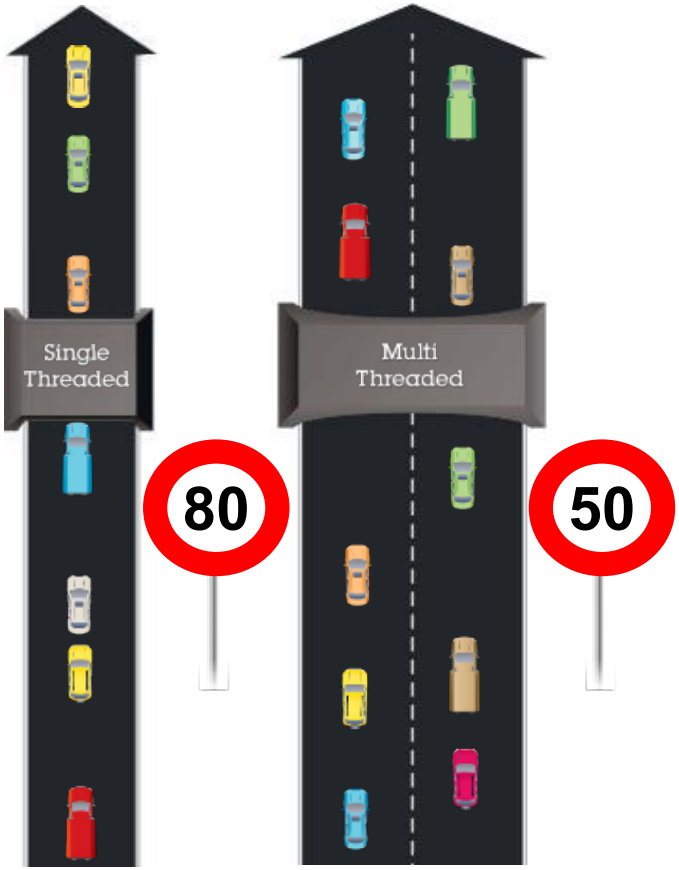

*Which approach is designed for the highest volume** of traffic? Which road is faster?*

***Two lanes at 50 carry 25% more volume if traffic density per lane is equal*

z/OS SMT Support (for zIIP Only)

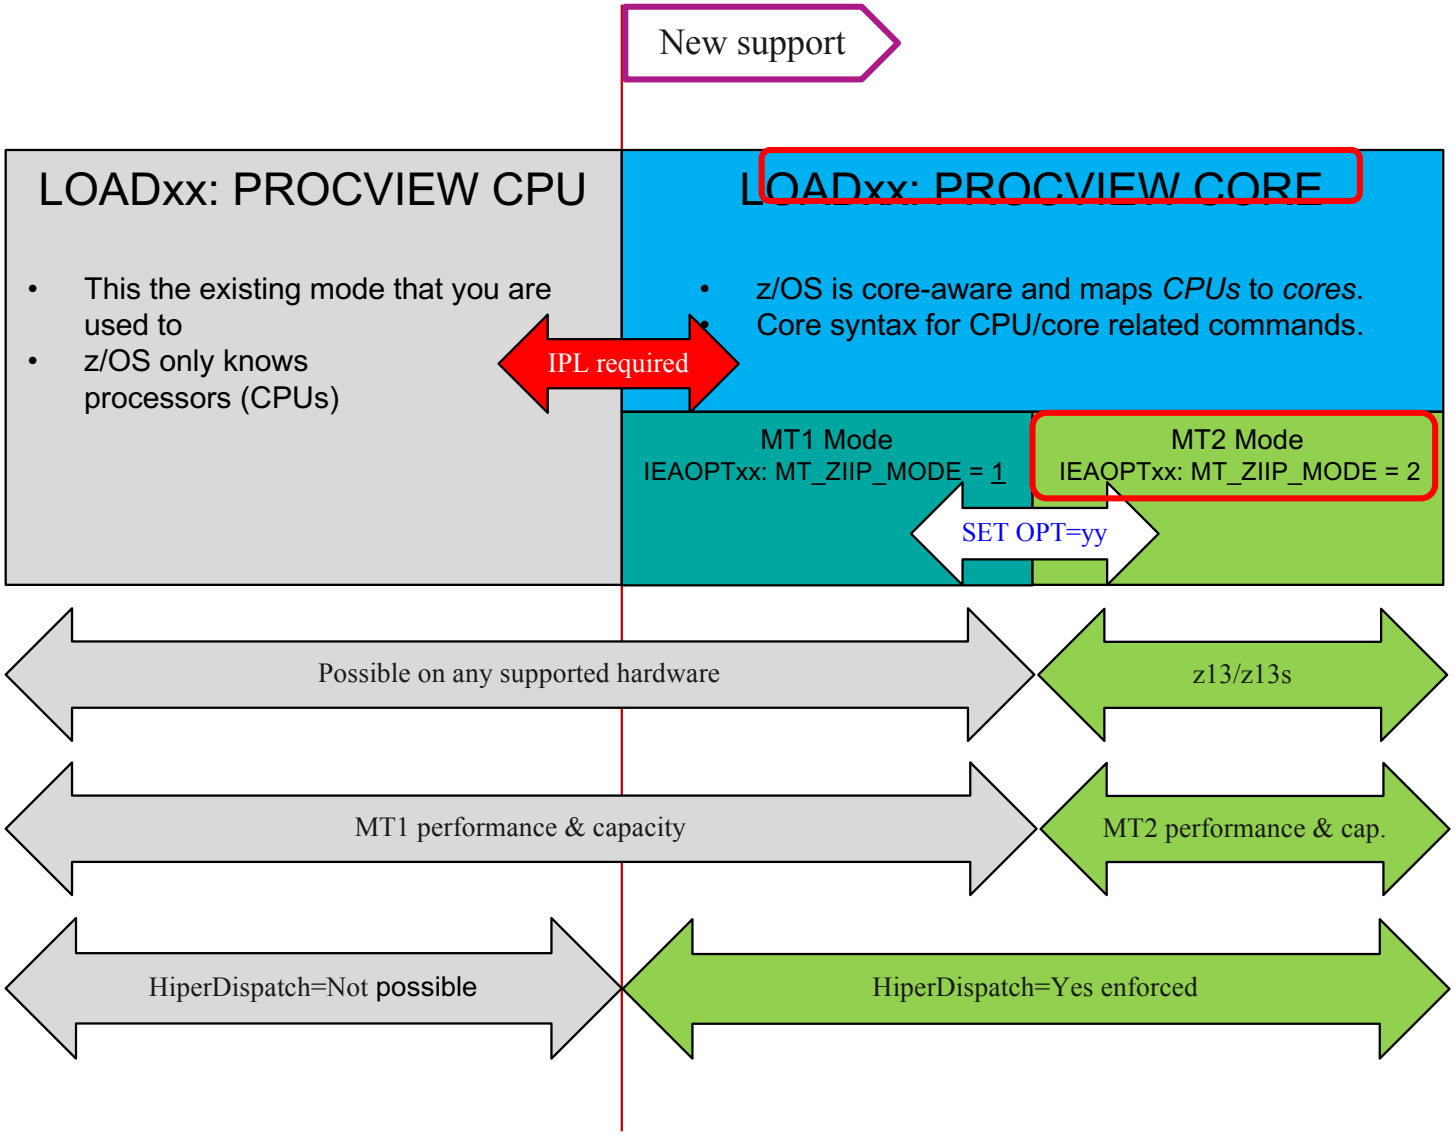

z/VM SMT implementation

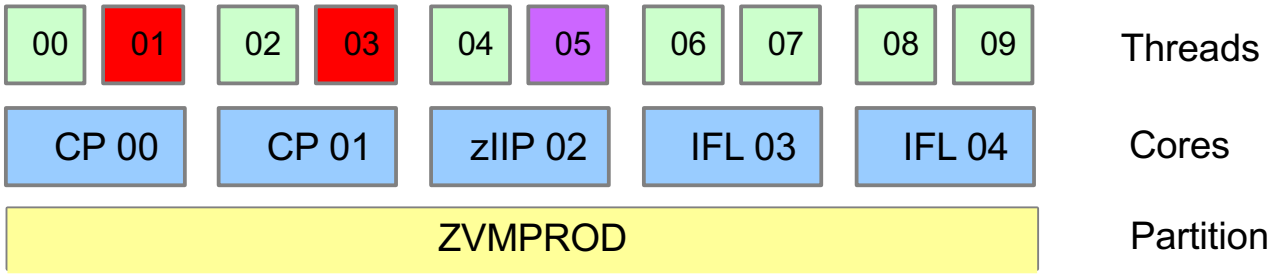

- **CPU address expansion**
  - Without SMT
    - CPU x0014 = 0000 0000 0001 0100
  - With SMT
    - Core x0014 thread 0 = 0000 0000 0010 1000 (CPU x0028)
    - Core x0014 thread 1 = 0000 0000 0010 1001 (CPU x0029)
  - Non-IFL processor odd address unavailable or unused

Table of Content

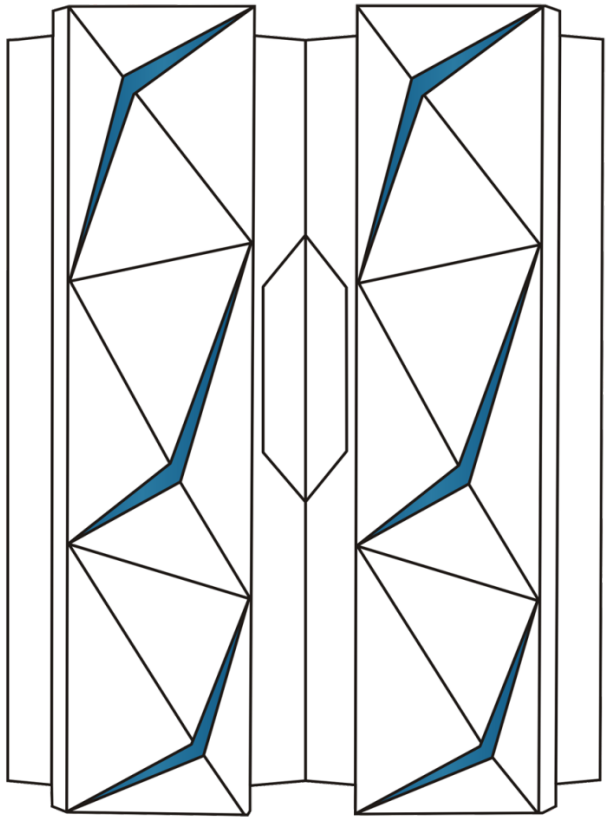

Why Single Instruction Multiple Data (SIMD) on z Systems

- **Background**
  - The amount of data is increasing exponentially
    IT shops need to respond to the diversity and volume of data
  - Enterprises use traditional integer, floating point, string,
    and XML character-based data
  - It's becoming more important for customers to do computations,
    analytics closer to the data

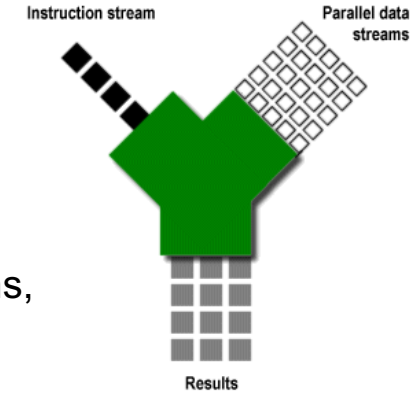

- **Customer *perception* of Analytics and z Systems**
  - z Systems handle OLTP and Batch jobs types of workload
  - Mathematical and data intensive operations can lead to unaffordable MIPS usage

- ***Reality* of Analytics and z Systems**
  - For the past 2-3 generations, z Systems processor has changed its capabilities in compute-intensive processing (analytics)
  - SIMD provides next phase of enhancements for analytics and compute-intensive competitiveness on z Systems

SIMD (Single Instruction Multiple Data) Processing

Increased parallelism to enable analytics processing

- Smaller amount of code helps improve execution efficiency
- Process elements in parallel enabling more iterations
- Supports analytics, compression, cryptography, video/imaging processing

Value

- ✓ Enable new applications
- ✓ Offload CPU
- ✓ Simplify coding

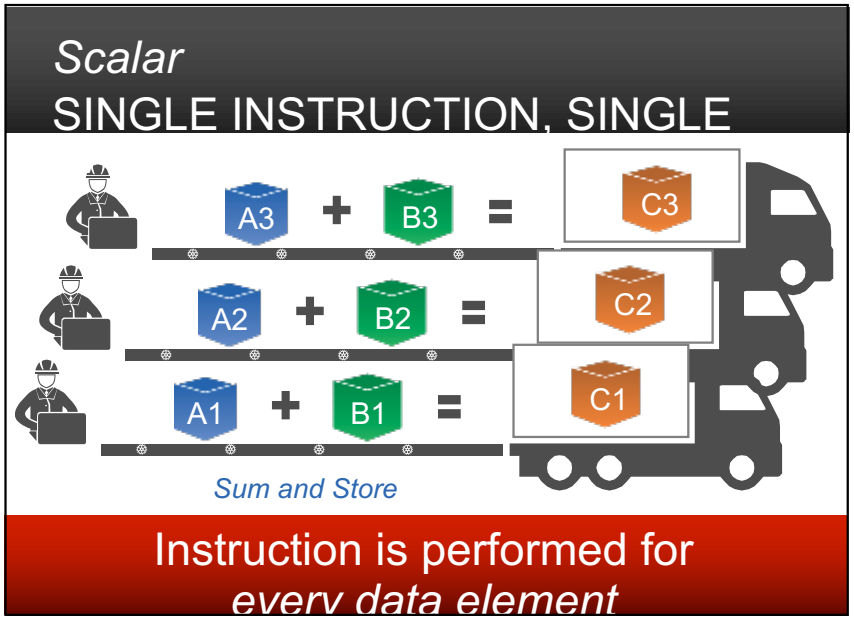

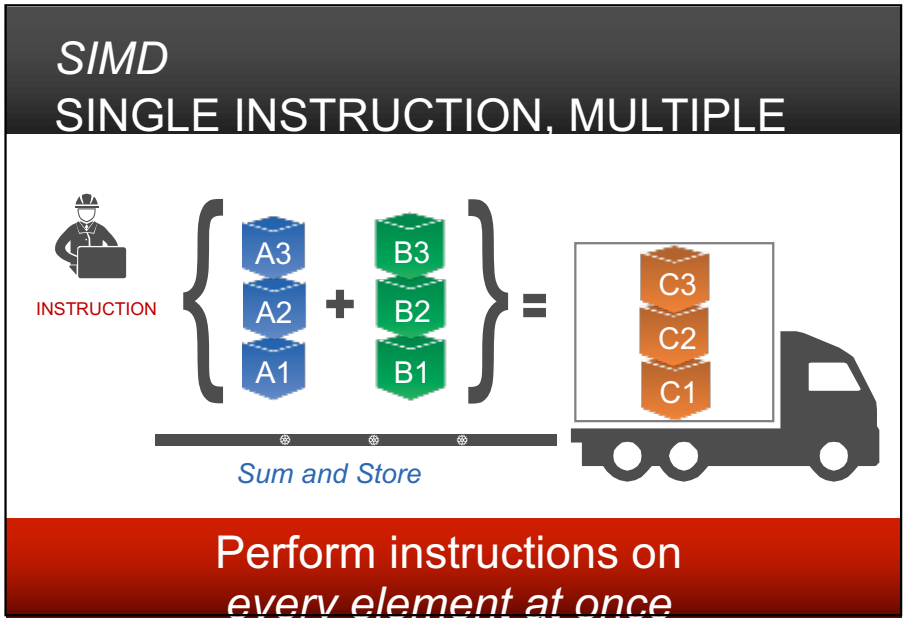

SIMD – Single Instruction Multiple Data

- **Old: Single Instruction Single Data (64b)**

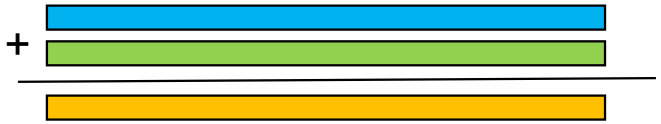

- **New: Single instruction operates on multiple data in parallel**

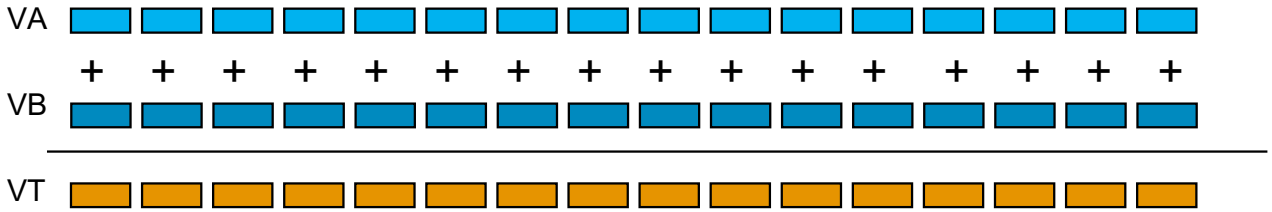

- Each register contains multiple data elements of a fixed size
  - Byte, Halfword, Word, Doubleword, Quadword
  - The collection of elements in a register is also called a **vector**
  - Instructions have a non-destructive operand encoding (T=A+B vs A=A+B)
  - For most operations the CC is not set, but a few instructions a summary condition code is used
  - Field in the instruction word specifies data format type

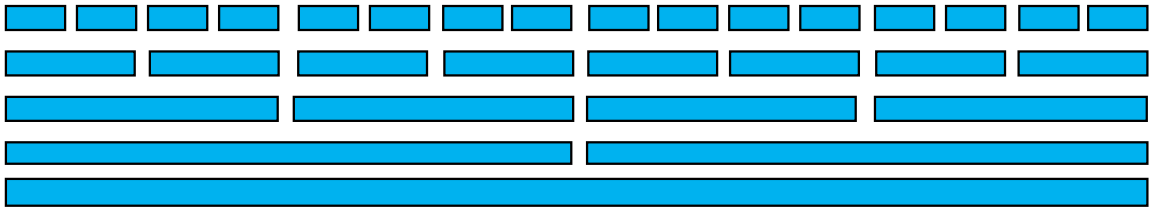

128b wide vector:
16xB, 8xHW, 4xW,
2xDW, 1xQW

SIMD Vector Processing Support on z Systems

- OS/Hypervisor & Software Support for SIMD:
  - z/OS V2.2, or z/OS: 2.1 + PTFs
  - Linux: IBM is working with its Linux Distribution partners to support new functions/features
  - z/VM V6.3 SIMD support will be delivered in 1Q2016 via PTFs for APAR VM65733; z/VM 6.4 (4Q2016)
  - KVM for IBM z V1.1.1 enables guests to exploit SIMD (1Q2016)
  - Compiler exploitation
    - IBM Java 8
    - z/OS XL C/C++ V2R1M1, V2.2 (added AUTOSIMD option)
    - XL C/C++ for Linux on z Systems, V1.1
    - Enterprise COBOL for z/OS v5.2
    - Enterprise PL/I for z/OS v4.5

Increase throughput with IBM CPLEX Optimizer V12.6.1

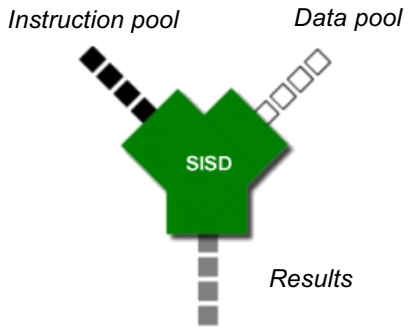

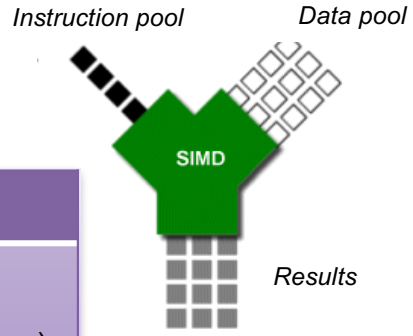

Workloads		
Java 8	C/C++Compiler built-ins for SIMD operations (z/OS and Linux on z Systems)	MASS & ATLAS Math Libraries (z/OS and Linux on z Systems)
SIMD Registers and Instruction Set		

MASS - Mathematical Acceleration Sub-System / ATLAS - Automatically Tuned Linear Algebra Software

SIMD Exploitation and Enablement – Things IBM is doing for you

- SIMD on z Systems Differentiation
  - z Systems brings analytics processing to the operational data – z System, data co-exist in the same environment
    - Enables new workload growth and development on z
    - Port analytics workloads from the distributed/LOB analytics shops; avoid ETL
  - z Systems is building a rich SIMD ecosystem spanning HW, OS, SW/Middleware, and ISV SW

Area	Product	Description*
SIMD Optimized Workloads	z/OS XMLSS	XML Parsing
	ILOG-CPLEX	Mathematical optimization solver
	Java	Workloads with string character or floating point data types
Enabling Libraries	Rational Compiler Suite	MASS Library on z/OS, Linux on z Systems
		ATLAS Library on z/OS, Linux on z Systems
Enabling Compilers / Built-in Functions (String, Integer, Floating Point Processing)	SIMD XLC for z/OS	SIMD XLC Intrinsic and vector data types
	GCC Compiler, Linux Kernel /Runtimes	Default Linux C Compiler; SIMD context save/restore support, binutils, glibc
	Enterprise COBOL for z/OS	COBOL intrinsics (INSPECT), string processing facilities
	Java8 Compiler	Java string character conversions, auto-vectorization
	PL/I	Optimizer and checkout compiler
Tools	Linux gdb	Debugger for Linux OS Programs
	PD Tools (Fault Analyzer, Debug Tool, Application Performance Analyzer)	Source level Debugger for z/OS C, C++ Programs

* See Speaker Notes

- SIMD Exploitation and Enablement
  - Exploitation: Workloads with targeted usage of SIMD based on known execution characteristics (XMLSS, Java string)
  - Enablement: Allows workloads to be independently targeted by developers for exploitation of instructions and register
  - Enablement Stack: Runtimes (Java), Tools (XL C/C++ compiler), Library (MASS, ATLAS), Firmware (String Millicode Instructions); for developers wanting to SIMDize their own workload
  - IBM is building a robust ecosystem that is capable of driving the growth of workloads for analytics and those with compute- and data-intensive properties

Table of Content

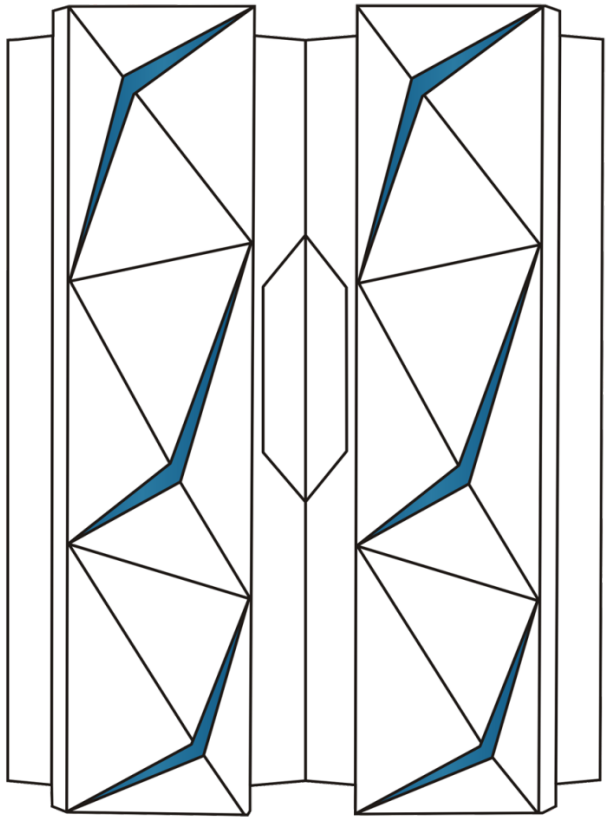

5-Channel RAIM Memory Controller Overview

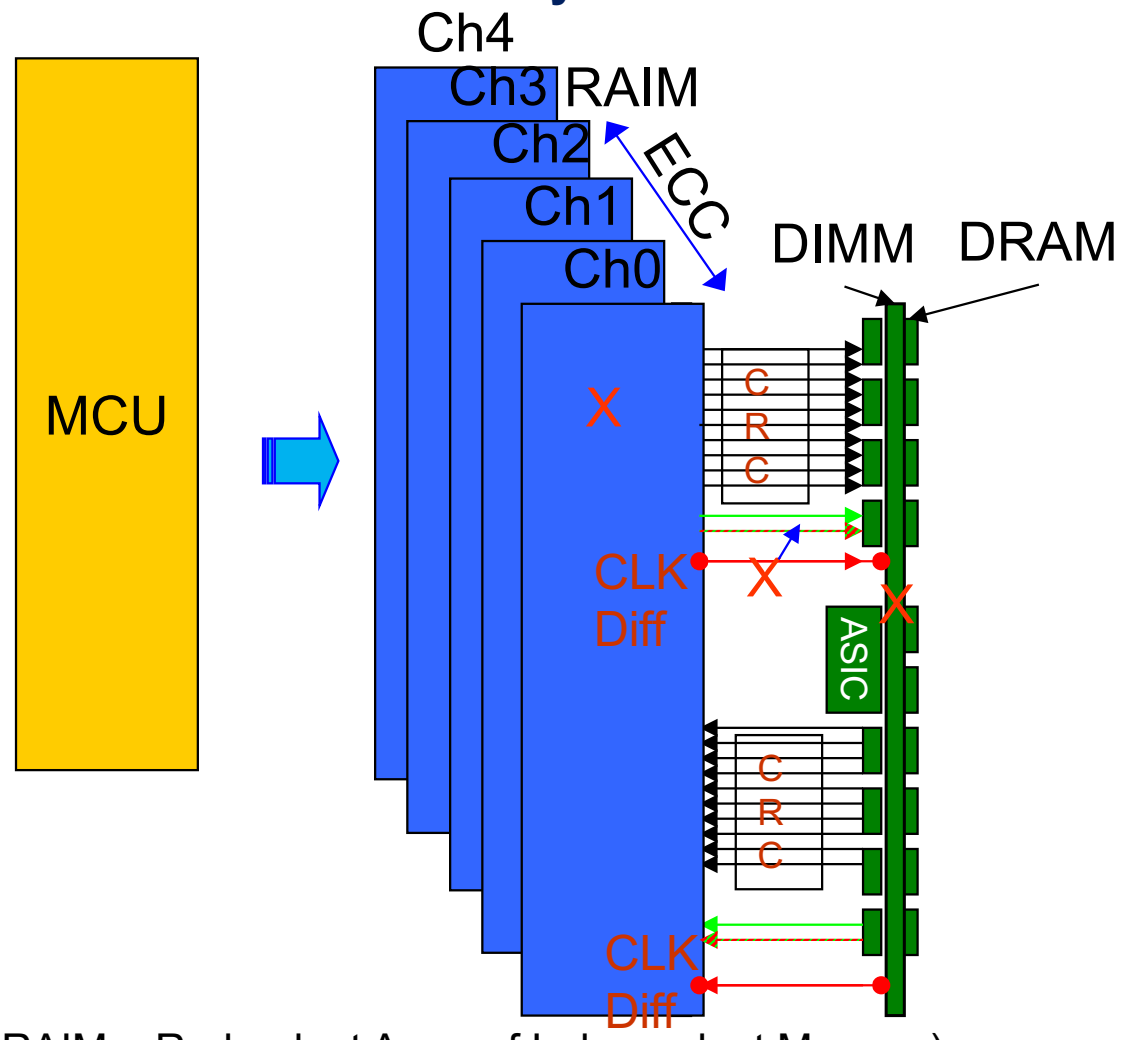

Layers of Memory Recovery

ECC
- Powerful 90B/64B Reed Solomon code

DRAM Failure
- Marking technology; no half sparing needed
- 2 DRAM can be marked
- Call for replacement on third DRAM

Lane Failure
- CRC with Retry
- Data – lane sparing
- CLK – RAIM with lane sparing

DIMM Failure (discrete components, VTT Reg.)
- CRC with Retry
- Data – lane sparing
- CLK – RAIM with lane sparing

DIMM Controller ASIC Failure
- RAIM Recovery

Channel Failure
- RAIM Recovery

RAIM = Redundant Array of Independent Memory)

Each memory channel supports only one DIMM (no DIMM cascading)

z14 Memory Usage and Allocation

- Installed Physical Memory (DIMM capacity) in configuration reports is RAIM Array size. Addressable Memory for customer partitions and HSA is 20 percent smaller.
- Servers are configured with the most efficient configuration of memory DIMMs that can support Addressable Memory required for Customer Ordered Memory plus HSA. In some cases, there will be Available Addressable Memory that might support one or more concurrent LIC CC Customer Memory upgrades with no DIMM changes.

  Note: DIMM changes require a disruptive POR on z14 Model M01. They are always done without a POR on models with multiple drawers using Enhanced Drawer Availability (EDA). On those models, some or all LPARs can continue to run with one drawer out of service to have DIMMs changed or added. Probably all LPARs, if Flexible Memory is selected.

*HSA size is 192 GB on z14

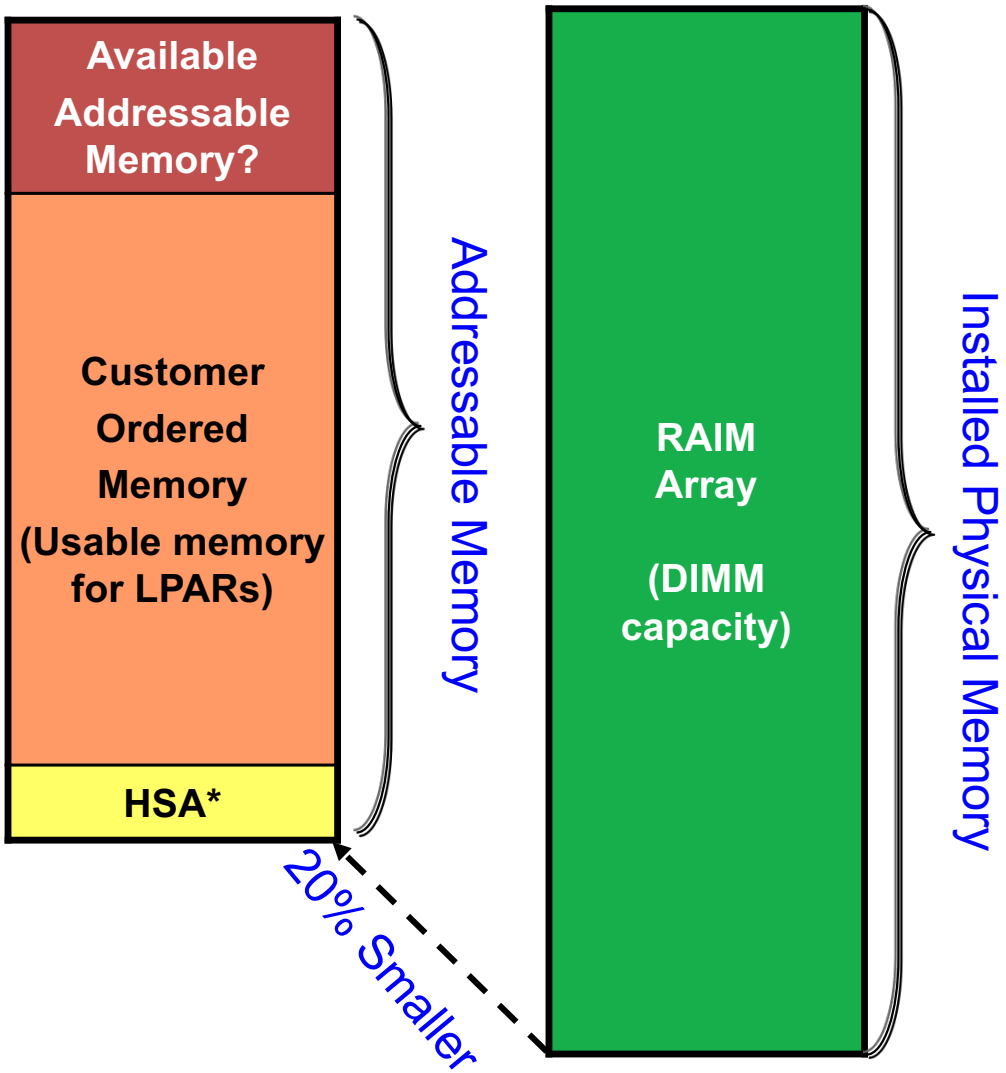

IBM z14 Large Memory

- Improving response time
  - Consistent fast transactional response time can result in an improved customer experience
  - Near immediate response time can drive productivity accelerating the velocity of development
  - Caching and other memory related techniques can help increase service levels to new heights
- Availability
  - Organizations trained to conserve memory can now relax restrictions to “enable the possible”
  - Tuning knobs can be adjusted to their max to further exploit memory
  - Examples:
    - Increased ability to handle workload spikes
    - Faster workload startup
    - Improved performance even given I/O disruptions
- Economics
  - Incentive pricing encourages customers to experiment with more memory and surface new use cases.
- Innovation
  - With mega memory, organizations can rethink and simplify application design for new business advantages
  - Example: collocate analytics and in memory data stores for high performance data mining solutions
    - IBM z/OS Platform for Apache Spark
    - IBM Machine Learning for z/OS

Quick summary of Physical Planning changes

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

Table of Content

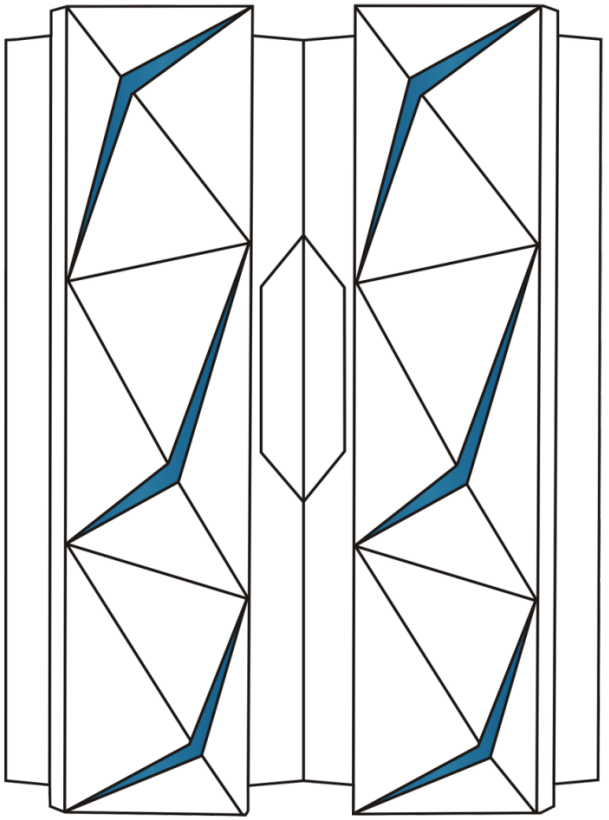

z14 “New Build” I/O and MES Features Supported

New Build Features

- Features – PCIe I/O drawer
  - FICON Express16S+ (SX and LX, 2 SFPs, 2 CHPIDs)
  - OSA-Express6S
    - 10 GbE LR and SR (1 SFP, 1 CHPID)
    - GbE SX, LX, and 1000BASE-T (2 SFPs, 1 CHPID)
  - 10GbE RoCE Express2 (new 10GbE feature code, FC0412)
  - zEDC Express
  - Crypto Express6S
  - zHyperLink Express (up to 16 features, 32 ports)
  - Coupling Express LR

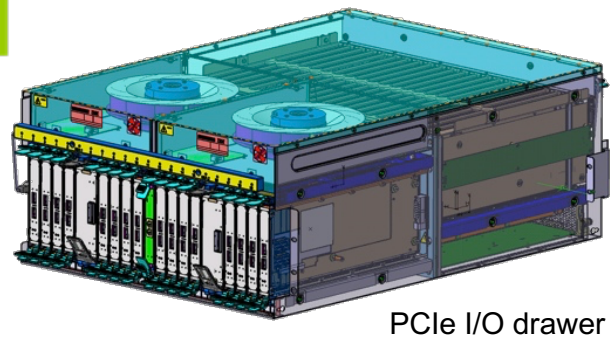

PCIe I/O drawer
32 I/O Slots

- PCIe Coupling Link Feature (Fanout)
  - ICA SR two 8GBps PCIe Gen3 Coupling Links

- InfiniBand Coupling Features (Fanouts)
  - HCA3-O two 12x 6GBps InfiniBand DDR Coupling Links
  - HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links

Note - “Plan Ahead” for I/O drawers is not offered on z14

Last machine to support InfiniBand features!!

z14 “Carry Forward” Features Supported

Carry Forward Features

- Features – PCIe I/O drawer
  - FICON Express16S (SX and LX, 2 SFPs, 2 CHPIDs)
  - FICON Express8S (SX and LX, 2 SFPs, 2 CHPIDs)
  - OSA-Express5S (All)
  - OSA-Express4S 1000BASE-T (from zEC12 only)
  - 10GbE RoCE Express
  - zEDC Express
  - Crypto Express5S
  - Coupling Express LR ← available on z13/z13S the same date as z14 GA
  - TKEs - FC0842, FC0847

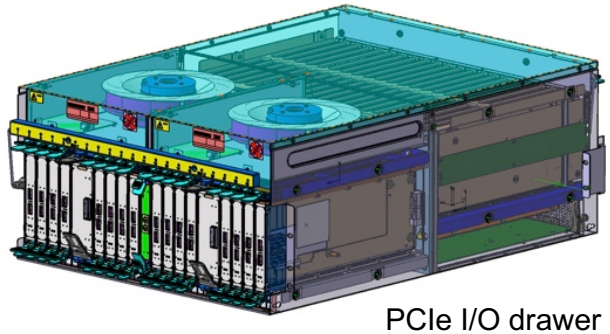

PCIe I/O drawer
32 I/O Slots

- PCIe Coupling Link Feature (Fanout)
  - ICA SR two 8GBps PCIe Gen3 Coupling Links

- InfiniBand Coupling Features (Fanouts)
  - HCA3-O two 12x 6GBps InfiniBand DDR Coupling Links
  - HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links

Last machine to support InfiniBand features!!

PCIe 32 I/O slot drawer

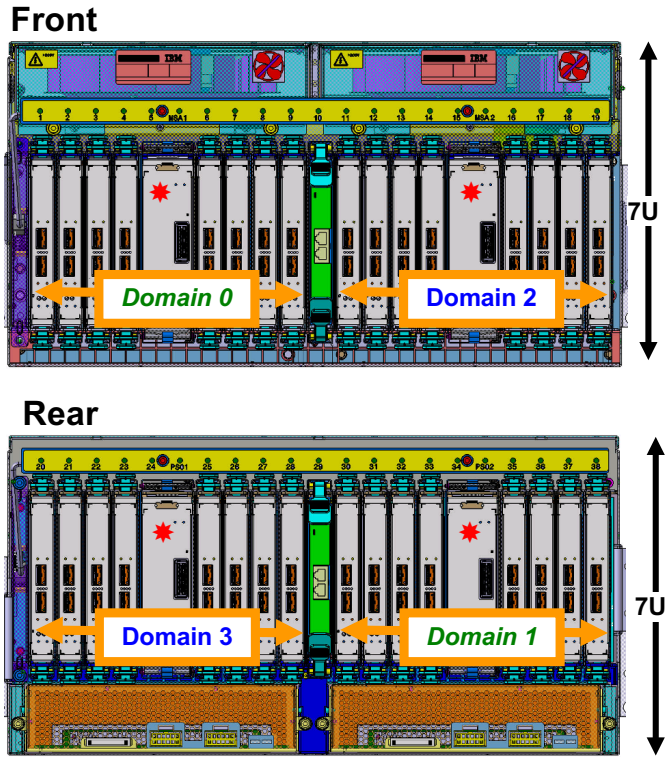

- Supports only PCIe I/O cards
  - –z14: Up to five drawers
- Supports 32 PCIe I/O cards, 16 front and 16 rear, vertical orientation, in four 8-card domains (shown as 0 to 3).
- Requires four 16 GBps PCIe switch cards (✸), each connected to a 16 GBps PCIe I/O interconnect to activate all four domains.

- To support Redundant I/O Interconnect (RII) between front to back domain pairs *0-1* and 2-3 the two interconnects to each pair will be from 4 different PCIe fanouts. (All four domains in one of these cages can be activated with two fanouts.)
- Concurrent field install and repair.
- Requires 7 EIA Units of space (12.25 inches ≈ 311 mm)

Logical channel subsystems (LCSS), subchannel sets (SS), Function Definitions, and Logical Partitions on z14

- Six Logical Channel Subsystems (LCSS) each with four subchannel sets (SS) and up to 256 channels
  - Maximum channel count includes channels spanned to more than one LCSS
  - Total physical channels depend on I/O features configured
  - Up to 63.75k base IODEVICEs in SS 0 and 64 k alias IODEVICEs each in SS 1 to SS 3 per LCSS
- FUNCTION definition support for virtualized RoCE, zEDC and zHyperLink independent of LCSS
- Up to 85 Logical Partitions: 15 each in LCSS 0 – 4, 10 in LCSS 5
  - Only channels and IODEVICEs defined in its LCSS can be assigned to an LPAR
  - Any defined FUNCTION can be assigned to any LPAR

z14					
Function Definitions for up to16 RoCE features (31 LPARs each) and 8 zEDC features (15 LPARs each)					
LCSS 0	LCSS 1	LCSS 2	LCSS 3	LCSS 4	LCSS 5
Up to 15 Logical Partitions	Up to 15 Logical Partitions	Up to 15 Logical Partitions	Up to 15 Logical Partitions	Up to 15 Logical Partitions	Up to 10 Logical Partitions
Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k SS 3 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k SS 3 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k SS 3 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k SS 3 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k SS 3 – 64 k	Subchannel Sets: SS 0 – 63.75 k SS 1 – 64 k SS 2 – 64 k SS 3 – 64 k
Up to 256 Channels	Up to 256 Channels	Up to 256 Channels	Up to 256 Channels	Up to 256 Channels	Up to 256 Channels

LPARS B-F Reserved

z14 I/O Connectivity Summary

Features	Offered As	Maximum # of features	Maximum channels	Increments per feature	Purchase increments
Storage					
FICON Express16S+	NB	160	320 channels	2 channels	2 channels
FICON Express16S	CF	160	320 channels	2 channels	2 channels
FICON Express8S	CF	160	320 channels	2 channels	2 channels
Networking					
OSA-Express6S	NB	48	96 (48 for 10 GbE) ports	1 (10 GbE) / 2	1 feature
OSA-Express5S	CF	48	96 (48 for 10 GbE) ports	1 (10 GbE) / 2	1 feature
OSA-Express4S 1000BASE-T Only	CF	48	96 ports	2	1 feature
Crypto					
Crypto Express6S	NB	16 PCIe adapters	16	1 PCIe adapter	2 features**
Crypto Express5S	CF	16 PCIe adapters	16	1 PCIe adapter	2 features**
Special purpose					
10GbE RoCE Express 2	**NB**	**8 PCIe adapters**	**16 ports**	**2 ports per adapter**	**1 feature**
10GbE RoCE Express	**CF**			**2 ports per adapter**	**1 feature**
zHyperLink Express	NB	16 PCIe adapters	32 ports	2 ports per adapter	1 feature
zEDC Express	NB/CF	16 PCIe adapters	16	1 PCIe adapter	1 feature
Coupling Links					
ICA SR	NB/CF	40 PCIe adapters	64 ports	2 ports per adapter	1 feature
Coupling Express LR (PCIe Adapter)	NB/CF	32 PCIe adapters	64 ports	2 ports per adapter	1 feature
HCA3-O (12x)	NB/CF	16 HCA3 adapters	32 ports	2 ports per adapter	1 feature
HCA3-O LR (1x)	NB/CF	16 HCA3 adapters	64 ports	4 ports per adapter	1 feature

* 2 features initially, one thereafter
CF = Carry Forward , NB = New build, Migration Offering, IBM Z Exchange Program, and if previously offered carried forward

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)**

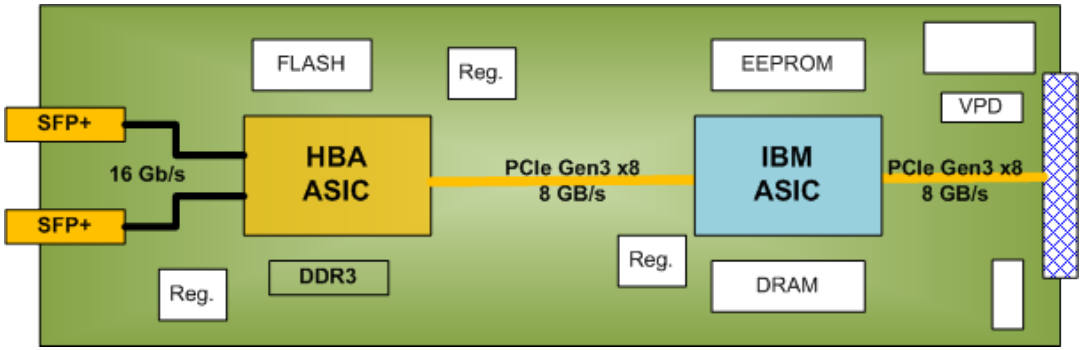

FC #0427 – 10KM LX, FC #0428 – SX

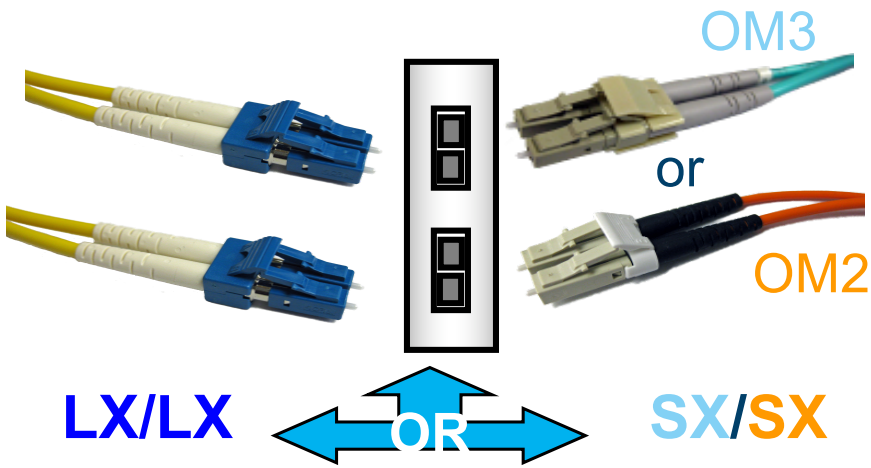

FICON Express16S+ and IOCP Rules

- Both ports must be same CHPID type (either FICON or FCP).
  - z14 only.
  - z14 can also have older FICON adapters (FICON Express16S/FICON Express8S) from an upgrade which does not have this restriction.

- To permit the mix for an older card, the user will need to specify a new keyword on at least one channel for an adapter where a mix is desired.
  - The new keyword is: ***MIXTYPE***.
    The keyword only needs to be on one of the PCHIDs for the card.

- The IOCP PCHID Summary and Channel Path Reports will be updated to indicate this. The current design is to add a suffix of '(M)' after the PCHID number.

- IOCP will ignore MIXTYPE keyword for processors prior to z14 – warning message

FICON Express16S+ and IOCP Rules (cont.)

Both: FC or FCP

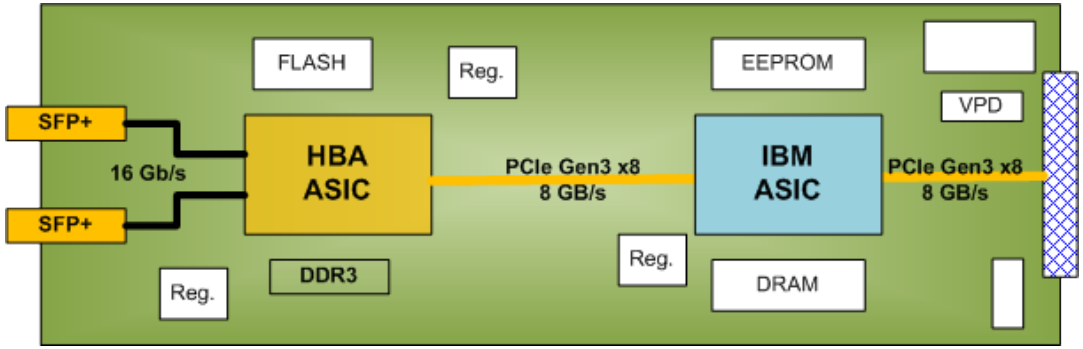

FICON Express16S+

Valid syntax examples

```
FICON EXPRESS16S          CHPID PCHID=100,PATH=(CSS(0),80),SHARED,TYPE=FC,MIXTYPE
FICON EXPRESS16S+         CHPID PCHID=101,PATH=(CSS(0),81),SHARED,TYPE=FCP
FICON EXPRESS16S+         CHPID PCHID=104,PATH=(CSS(0),82),SHARED,TYPE=FC
FICON EXPRESS8S           CHPID PCHID=105,PATH=(CSS(0),83),SHARED,TYPE=FCP,MIXTYPE
FICON EXPRESS8S           CHPID PCHID=108,PATH=(CSS(0),84),SHARED,TYPE=FCP
FICON EXPRESS8S           CHPID PCHID=109,PATH=(CSS(0),85),SHARED,TYPE=FC,MIXTYPE
FICON EXPRESS8S           CHPID PCHID=10C,PATH=(CSS(0),86),SHARED,TYPE=FCP,MIXTYPE
FICON EXPRESS16S          CHPID PCHID=10F,PATH=(CSS(0),87),SHARED,TYPE=FC,MIXTYPE
FICON EXPRESS16S+         CHPID PCHID=110,PATH=(CSS(0),88),SHARED,TYPE=FCP
FICON EXPRESS16S+         CHPID PCHID=114,PATH=(CSS(0),89),SHARED,TYPE=FC
```

109 -- FC

10C -- FCP

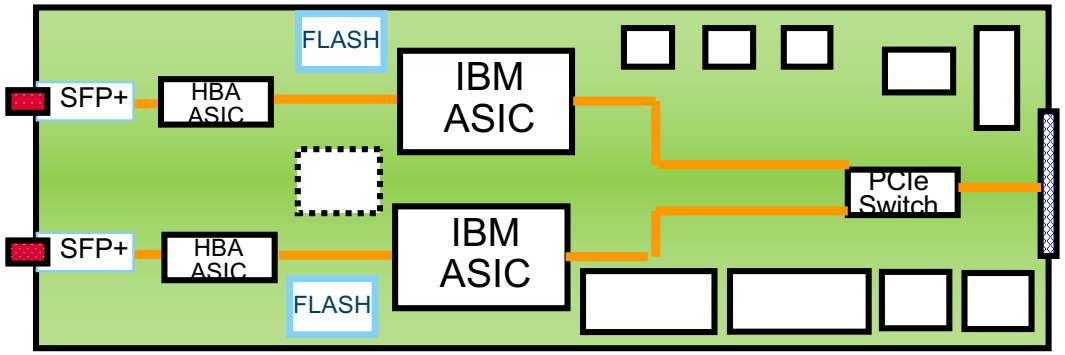

FICON Express16S or FICON Express8S

zHPF and FICON Performance* z14

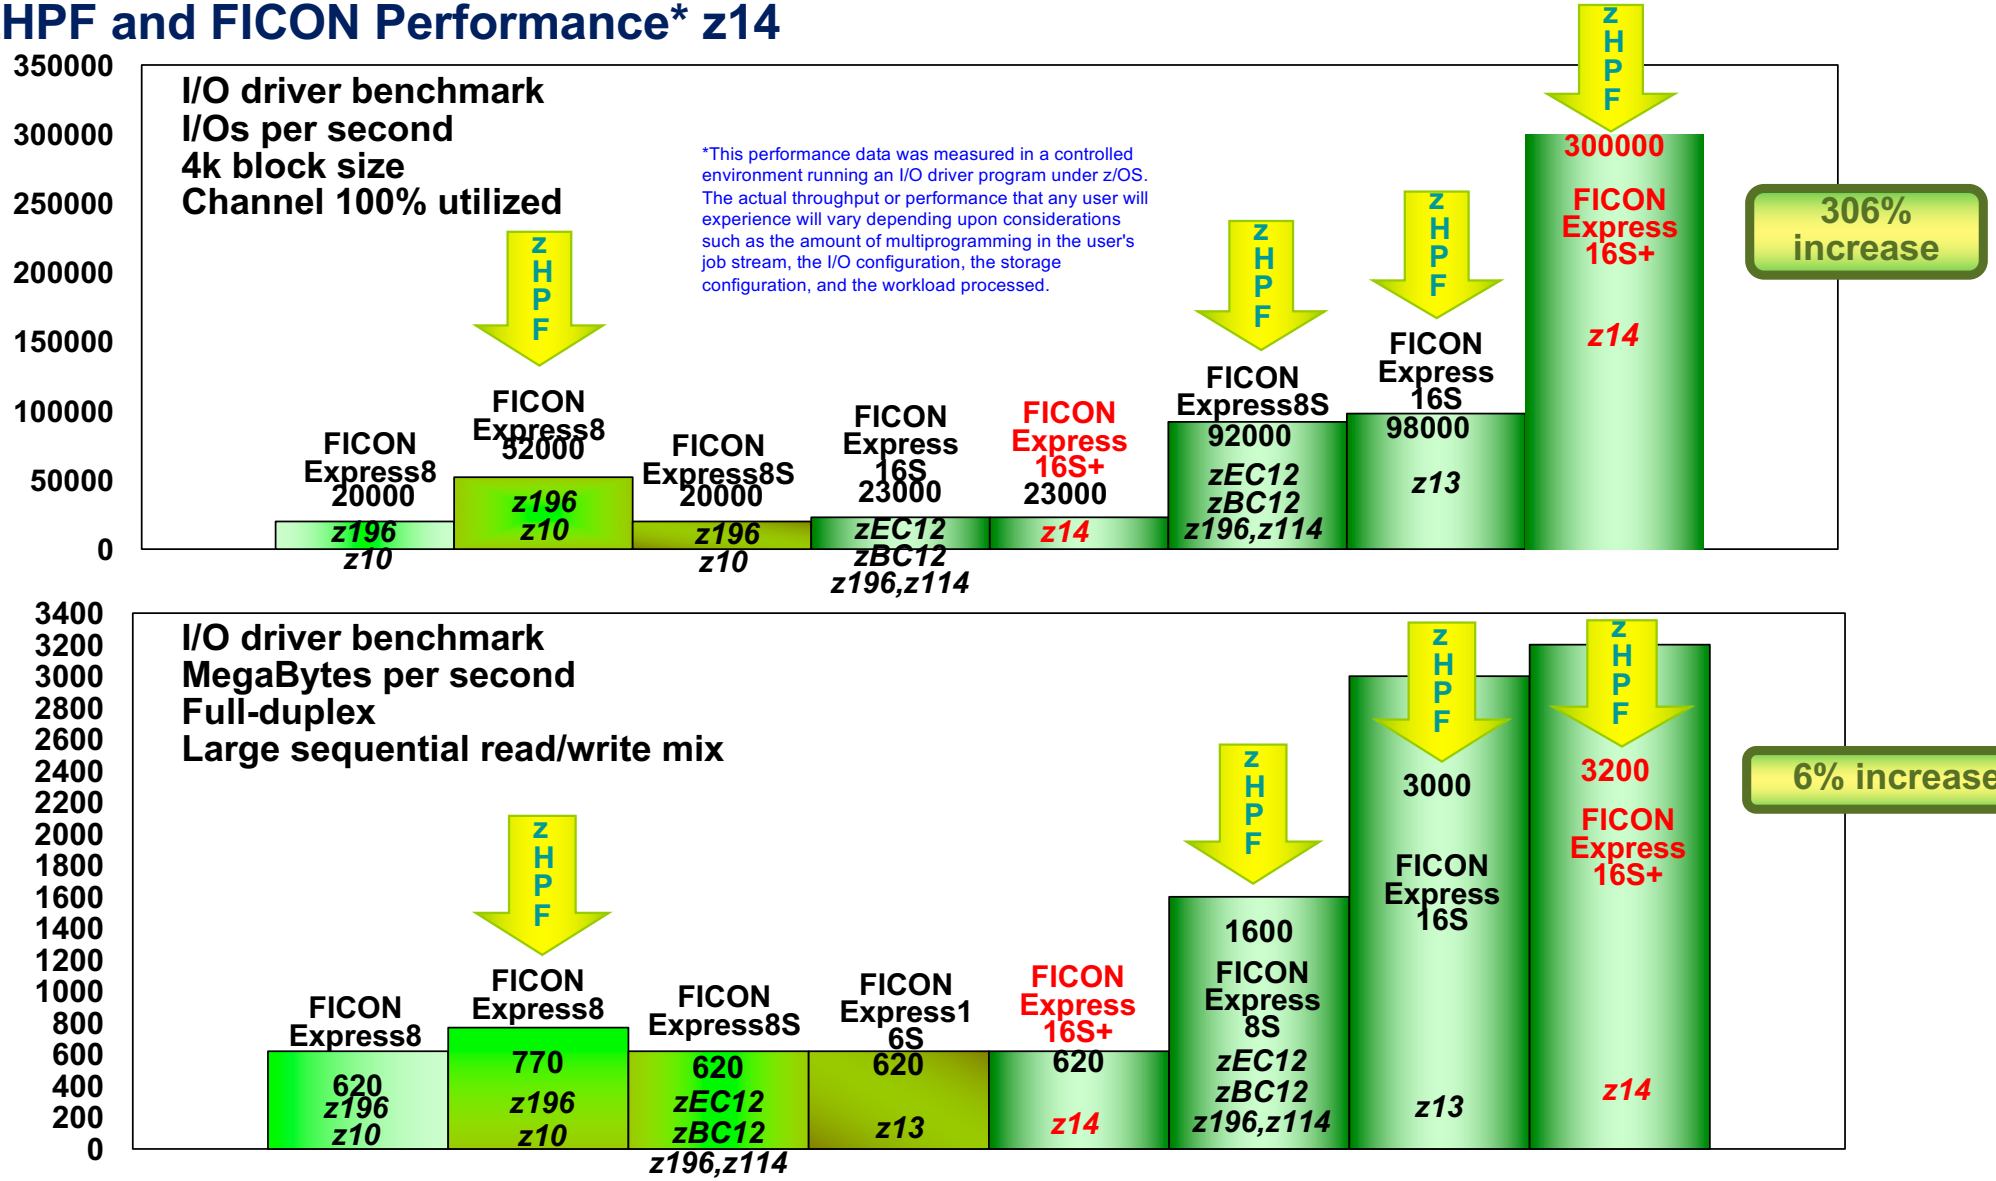

FCP Performance* for z14

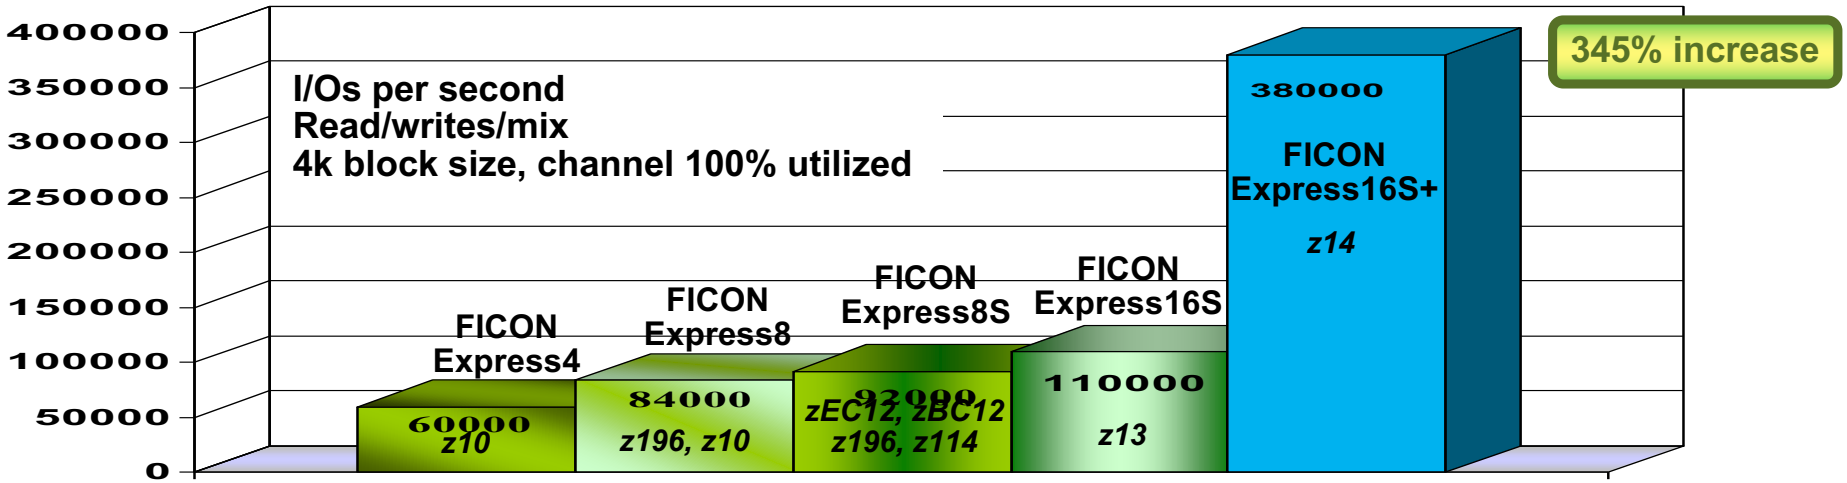

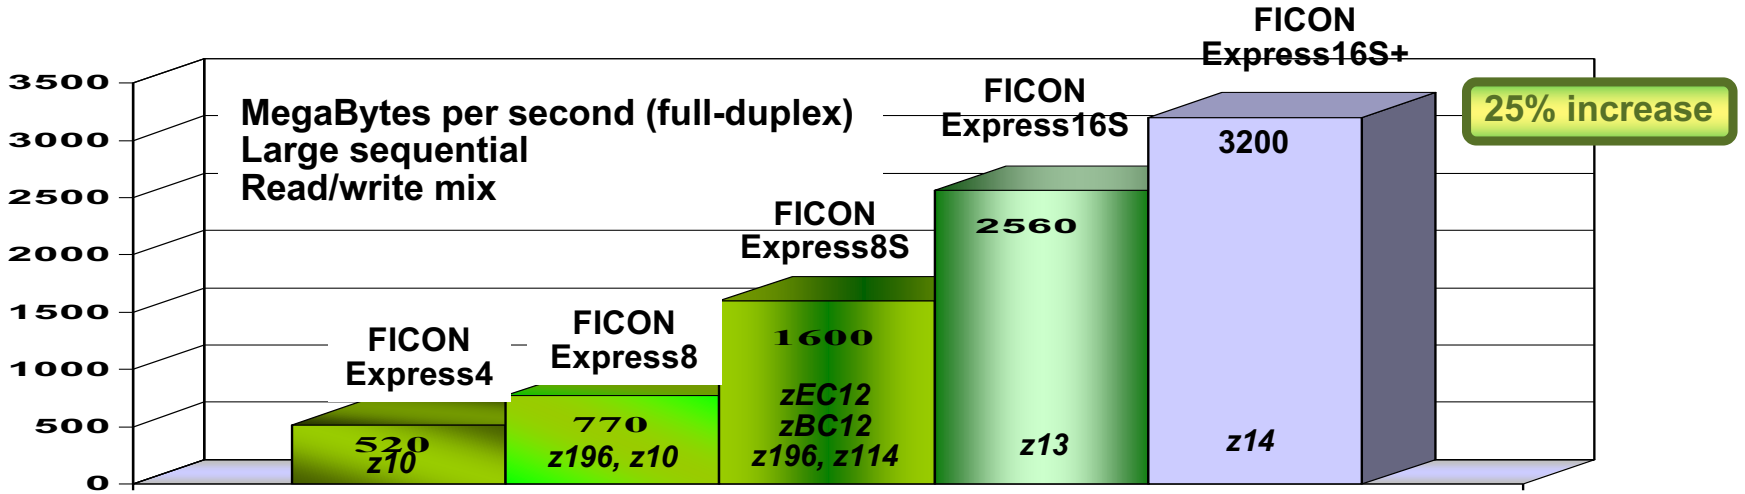

*This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration,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

Reminder - New FICON Function Introduced with IBM z13

- 16 Gbps Link Speeds (March 9, 2015)
  - Designed to reduce I/O latency to improve response time for performance-critical middleware and to shrink the batch window required to accommodate I/O bound batch work
- 6th Logical Channel Subsystem (March 9, 2015)
  - Up to 85 Logical Partitions: More flexibility for server consolidation
- 4th Subchannel Set (March 9, 2015)
  - Simplifies I/O configurations for a 2nd synchronous copy of data
  - With multi-target PPRC, can do HyperSwap and still maintain synchronous copy for 2nd HyperSwap
- Preserve Virtual WWPNs for NPIV configured FCP channels
  - Designed to simplify migration to a new-build z13 (March 9, 2015)
- Support for a maximum of 32K devices per FICON channel (March 9, 2015)
  - Up to 85 Logical Partitions: More flexibility for server consolidation
- zHPF Extended I/O execution at Distance (June 26, 2015)
  - Up to 50% I/O service time improvement for remote write
  - Designed to help GDPS HyperSwap configurations with secondary DASD in remote site
- FICON Dynamic Routing (September 25, 2015)
  - Designed to allow ISL sharing by FC and FCP traffic to optimize use of ISL bandwidth in the SAN fabric for both types of traffic
- Forward Error Correction Codes (September 25, 2015)
  - Designed to address high bit-error rate on high frequency (>= 8Gb/s) links
  - Estimated equivalence to doubling optical signal power

Fiber Channel Physical Interface Standard

- Applies to FICON Express16S (4, 8, 16 Gps) FICON Express8S and FICON Express8 (2, 4, 8 Gbps)
- CHPID types FC (FICON, zHPF, CTC) and FCP (Fibre Channel Protocol)
- Unrepeated distances in kilometers (km), meters (m), and feet (ft)

*	2 Gbps		4 Gbps		8 Gbps		16 Gbps		10 Gbps ISLs	
Fiber Core (µ) Light source	**Distance meters feet**	*** Link loss budget**	**Distance meters feet**	*** Link loss budget**	**Distance meters feet**	*** Link loss budget**	**Distance meters feet**	*** Link loss budget**	**Distance meters feet**	*** Link loss budget**
9µ SM LX laser	10 km 6.2 miles	7.8 dB	10 km 6.2 miles	7.8 dB	10 km 6.2 miles	6.4 dB	10 km 6.2 miles	6.4 dB	10 km 6.2 miles	6.0 dB
9µ SM LX laser	4 km # 2.5 miles	4.8 dB #	4 km # 2.5 miles	4.8 dB #	N/A	N/A	N/A	N/A	N/A	N/A
50µ MM OM4 4700 MHz-km SX laser	500 m 1640 ft	3.31 dB	400 m 1312 ft	3.0 dB	190 m 623 ft	2.2 dB	125 m 410 ft	1.9 dB	500 m 1804 ft	3.1 dB
50µ MM OM3 2000 MHz-km SX laser	500 m 1640 ft	3.31 dB	380 m 1247 ft	2.88 dB	150 m 492 ft	2.04 dB	100 m 328 feet	1.86 dB	300 m 984 ft	2.6 dB
50µ MM OM2 500 MHz-km SX laser	300 m 984 ft	2.62 dB	150 m 492 ft	2.06dB	50 m 164 ft	1.68 dB	35 m 115 feet	1.63 dB	82 m 269 ft	2.3 dB
62.5µ MM OM1 200 MHz-km SX laser	150 m 492 ft	2.10 dB	70 m 230 ft	1.78 dB	21 m 69 ft	1.58 dB	15 m 49 feet	1.56 dB	33 m 108 ft	2.4 dB

Inter-Switch Links (ISLs) is the link between two FICON directors; FICON features do not operate at 10 Gbps

* The link loss budget is the channel insertion loss as defined by the standard.

\# This distance and dB budget applies to FICON Express4 4KM LX features

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

FICON Optical Lengths

There are several factors that limit distance for a FICON connection, DB loss, number of connections, cable type and speed all impact distance. Customers that plan on ordering and using FICON Express16S SX connections should plan out their installations carefully !

Table list various speeds, FICON Express16 only supports 4, 8 and 16 Gbps

LX = Long wavelength 1310 nm transceiver SX - Short wavelength 850 nm

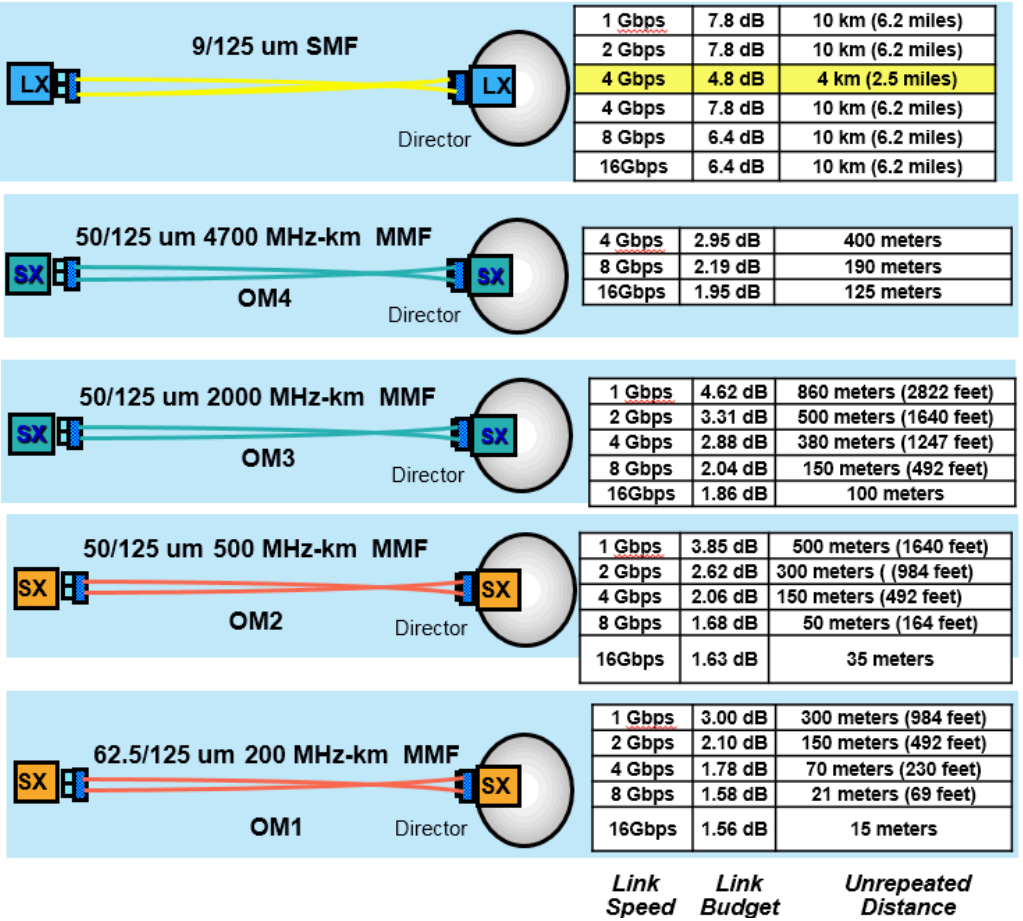

Link Speed	Link Budget	Unrepeated Distance
1 Gbps	7.8 dB	10 km (6.2 miles)
2 Gbps	7.8 dB	10 km (6.2 miles)
4 Gbps	4.8 dB	4 km (2.5 miles)
4 Gbps	7.8 dB	10 km (6.2 miles)
8 Gbps	6.4 dB	10 km (6.2 miles)
16Gbps	6.4 dB	10 km (6.2 miles)
4 Gbps	2.95 dB	400 meters
8 Gbps	2.19 dB	190 meters
16Gbps	1.95 dB	125 meters
1 Gbps	4.62 dB	860 meters (2822 feet)
2 Gbps	3.31 dB	500 meters (1640 feet)
4 Gbps	2.88 dB	380 meters (1247 feet)
8 Gbps	2.04 dB	150 meters (492 feet)
16Gbps	1.86 dB	100 meters
1 Gbps	3.85 dB	500 meters (1640 feet)
2 Gbps	2.62 dB	300 meters ((984 feet)
4 Gbps	2.06 dB	150 meters (492 feet)
8 Gbps	1.68 dB	50 meters (164 feet)
16Gbps	1.63 dB	35 meters
1 Gbps	3.00 dB	300 meters (984 feet)
2 Gbps	2.10 dB	150 meters (492 feet)
4 Gbps	1.78 dB	70 meters (230 feet)
8 Gbps	1.58 dB	21 meters (69 feet)
16Gbps	1.56 dB	15 meters

OSA-Express6S Fiber Optic Features – PCIe Drawer

- 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

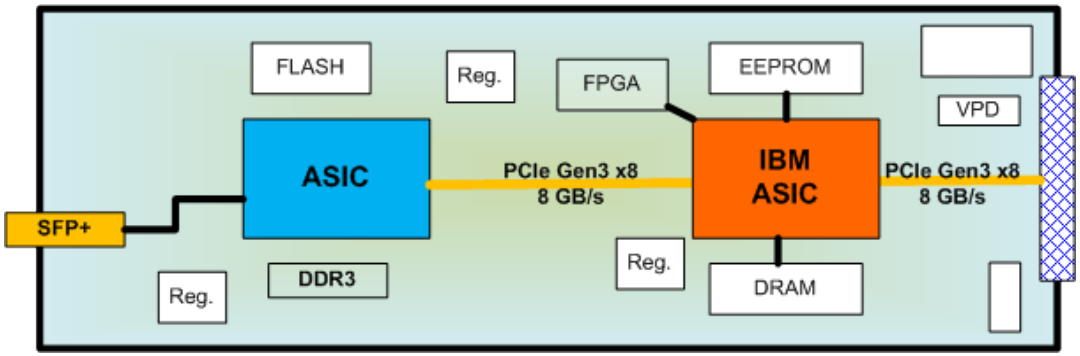

FC #0424 – 10 GbE LR, FC #0425 – 10 GbE SR

- Gigabit Ethernet (1 GbE)
  - CHPID types: OSD (OSN not supported)
  - 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

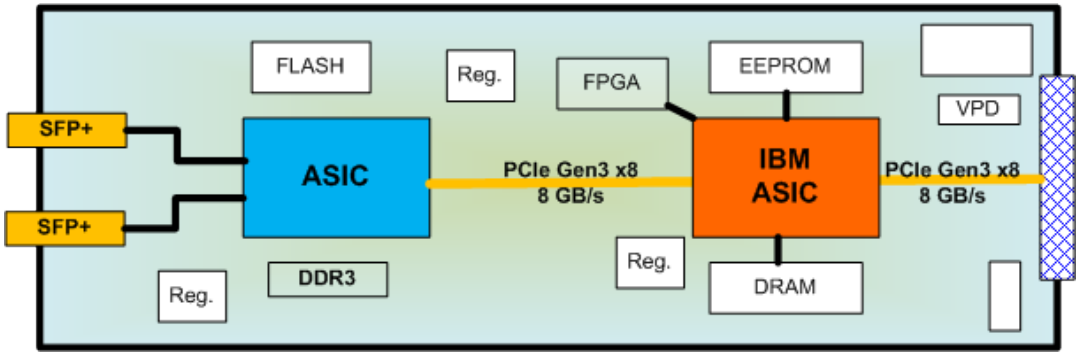

FC #0422 – GbE LX, FC #0423 – GbE SX

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* or 1000 Mbps and full duplex only on Category 5 or better copper
  - No 10Mbps
  - RJ-45 connector
  - Operates at “line speed”

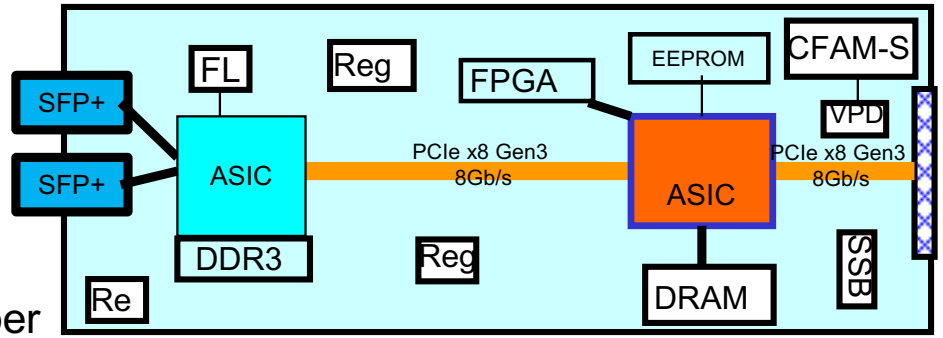

FC #0426

Connector = RJ-45

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)
~~OSA for NCP (LP-to-LP)~~	~~OSN~~	~~NCPs running under IBM Communication Controller for Linux (CCL)~~

* OSA-Express6S 1000BASE-T adapters (#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.

Networking Enhancements

- Allow or disallow the takeover of MAC address by another OS (or hypervisor).
  - Transparent to existing operating systems
  - KVM hypervisor only today
  - Ensures that only one OS can takeover MAC that belongs to another OS
- Separate MCL streams for OSA3270 (OSA-ICC) and OSA3215 (TPF)
  - Availability focused
- Global OSA Display via OSA/SF
- OSA/SF on HMC: Add APIs to existing User Interface capabilities

New Networking Functionality on z13 GA2

- OSA-ICC (OSC Channel) Secure Sockets Layer Support (z13s and z13 GA2)
  - Designed to improve security of console operations
  - Up to 48 secure sessions per CHPID (the overall maximum of 120 connections unchanged).

- Shared Memory Communication – Direct Memory Access (SMC-D) (z13s and z13 GA2)
  - High bandwidth, low latency LPAR-to-LPAR TCP/IP traffic using the direct memory access software protocols over virtual Internal Shared Memory (ISM) devices. Designed to provide application-transparent RDMA communications to TCP endpoints for sockets-based connections with reduced latency, improved throughput, and reduced CPU cost compared to HiperSockets, OSA, or SMC-R (RoCE).
  - Streaming workload test case results*:
    Up to 89% reduction in latency, 9 times the throughput, and 87% reduction in CPU cost compared to HiperSockets*
    Up to 95% reduction in latency, 20 times the throughput, and 83% reduction in CPU cost compared to OSA*
    Up to 94% reduction in latency, 16 times the throughput, and 58% reduction in CPU cost compared to SMC-R (RoCE)*
  - See the z/OS Communications Server web page for a link to information on detailed SMC-D performance test results:

- OSA OSD Channel Multiple VSWITCH Link Aggregation (LAG) Support (z13s from z13 GA)
  - Designed to improve z/VM V6.3 virtual networking capabilities and to permit sharing of supporting OSD channels among multiple z/VM V6.3 images

*This performance data was measured in a controlled environment under z/OS. The actual latency, throughput, and CPU cost that any client will experience will vary depending upon considerations such the I/O configuration, the storage configuration, and the characteristics of the communications workload.

Table of Content

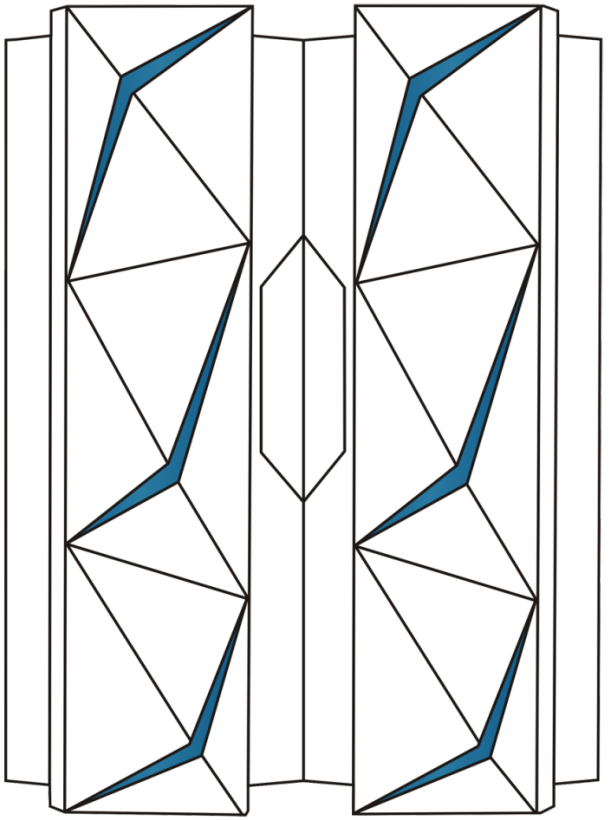

What is SMC?

- SMC – Shared Memory Communications
- Two forms of implementations:
  - SMC-R – Shared Memory Communication – Remote Direct Memory Access using 10GbE RoCE Express adapter – Intra- and inter- CPC communications.
    SMC-R is an *open* sockets over RDMA protocol that provides transparent exploitation of RDMA for TCP based applications, while preserving key functions and qualities of service from the TCP/IP ecosystem that enterprise level servers/network depend on
  - SMC-D – Shared Memory Communications – Direct Memory Access over Internal Shared Memory (ISM) – Intra CPC communications for TCP based applications.
    When ISM is exploited by z/OS using SMC-D, the combined solutions provide improved transaction rates for interactive (request/response) workloads due to reduced network latency and lower CPU cost for workloads with larger payloads (i.e. analytics, streaming, big data, or web services).
- Current implementation supports z/OS only
  - Any z/OS TCP sockets-based workload can seamlessly use SMC-R or SMC-D without application changes
  - SMC Applicability Tool (SMCAT) helps assess benefit of SMC-R and SMC-D for your environment
  - Connection level security is preserved with SMC-R and SMC-D
  - Both SMC-R and SMC-D require z/VM 6.3 + PTFs for (z/OS) guest exploitation

z14: Shared Memory Communications – (SMC-R and SMC-D)

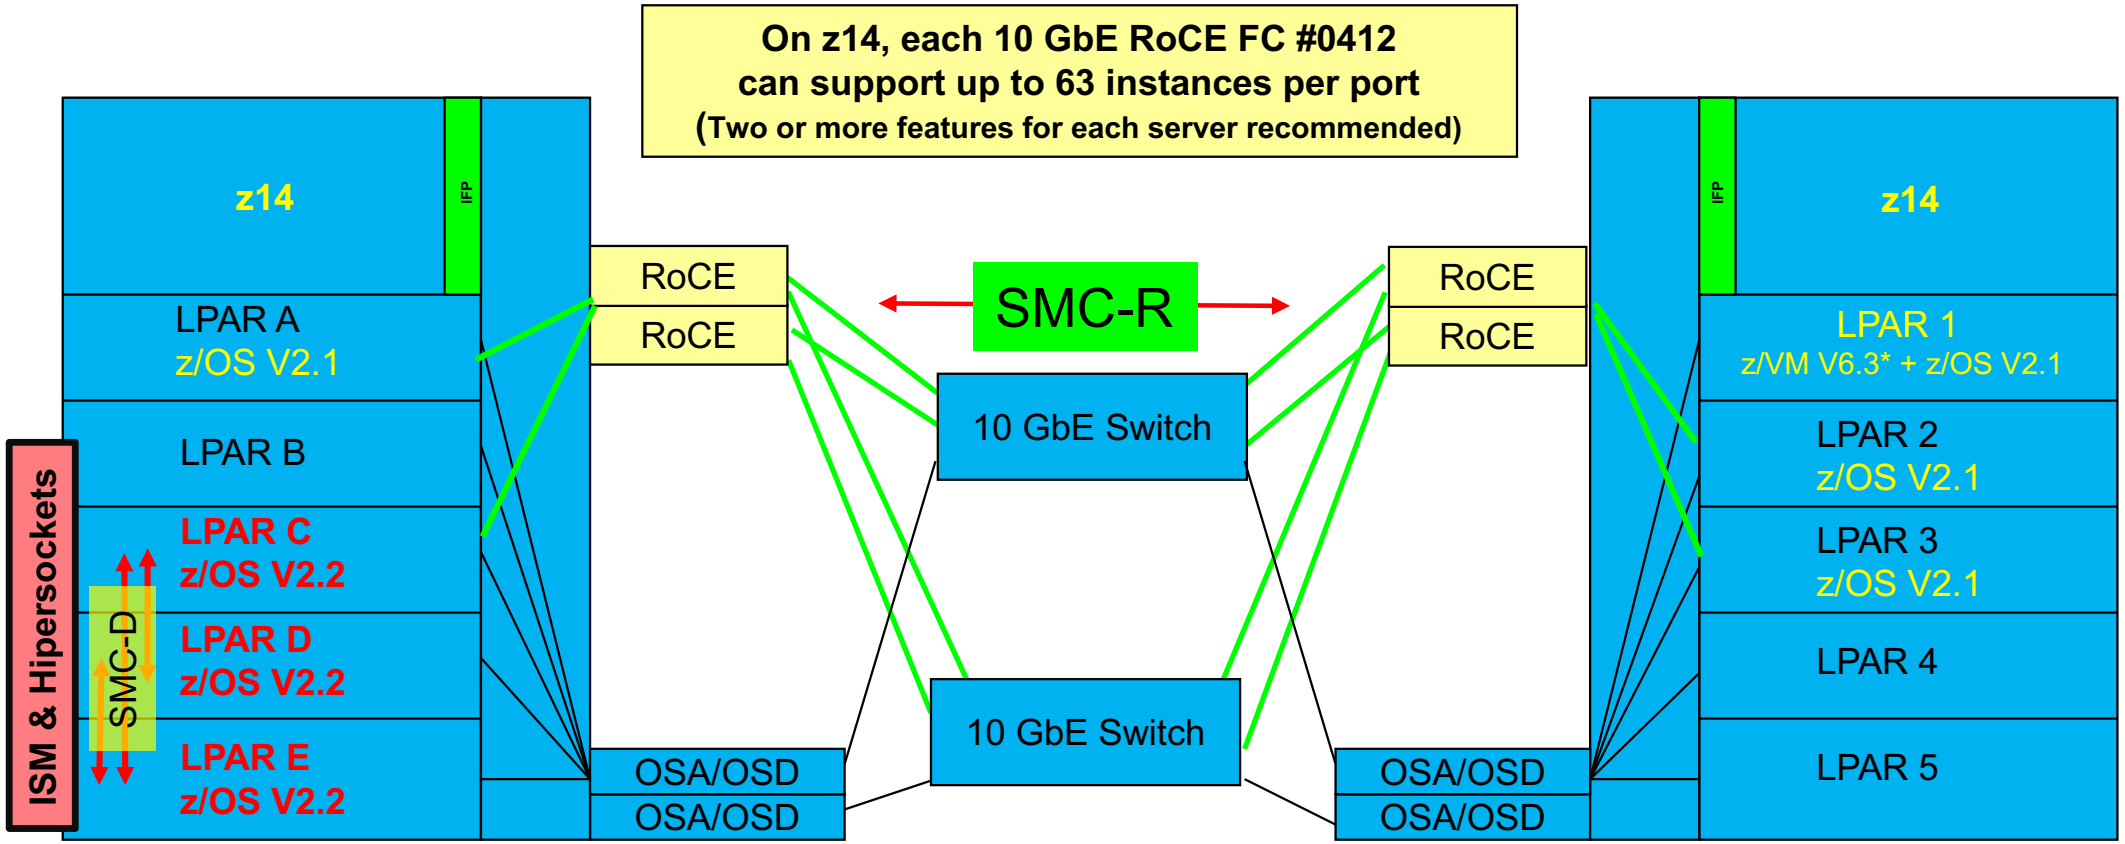

- This configuration allows SMC-D connectivity among LPAR C, LPAR D, and LPAR E.
- SMC-D within one machine is better than using HiperSockets alone.
- For LPAR to LPAR, HiperSockets or OSD connections are required to establish the SMC-D communication.
- ISM = Internal Shared Memory
- No additional hardware purchase required.
- z/VM Guest support

"Shared Memory Communications over RDMA" (SMC-R) concepts

Clustered Systems

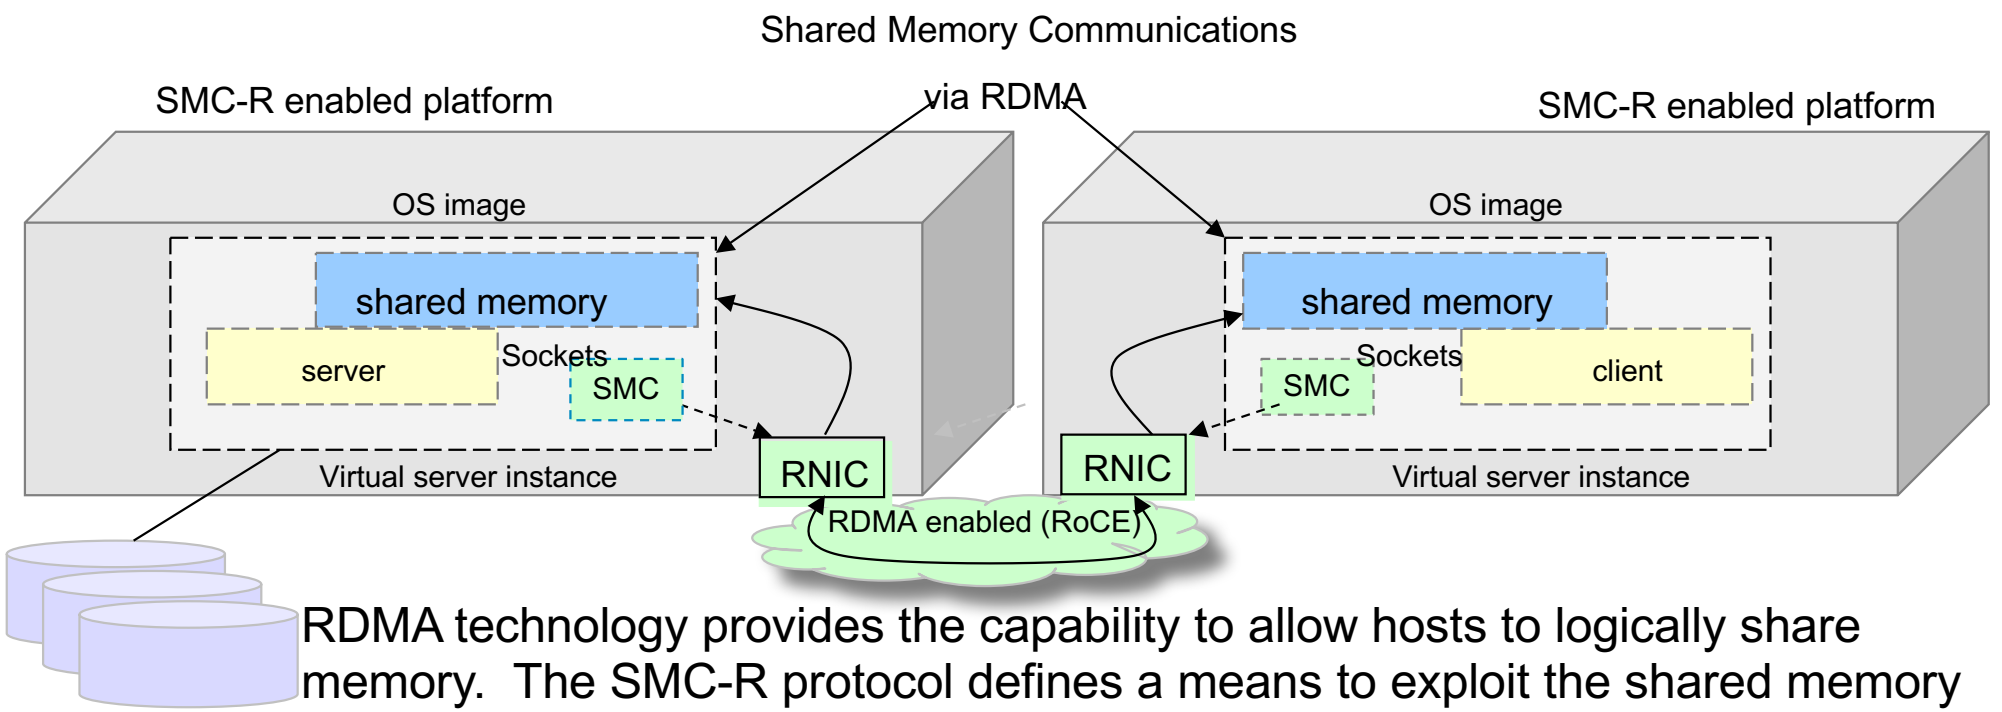

RDMA technology provides the capability to allow hosts to logically share memory. The SMC-R protocol defines a means to exploit the shared memory for communications - *transparent* to the applications! SMC-R RFC: https://datatracker.ietf.org/doc/rfc7609/

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.

Dynamic Transition from TCP to SMC-R

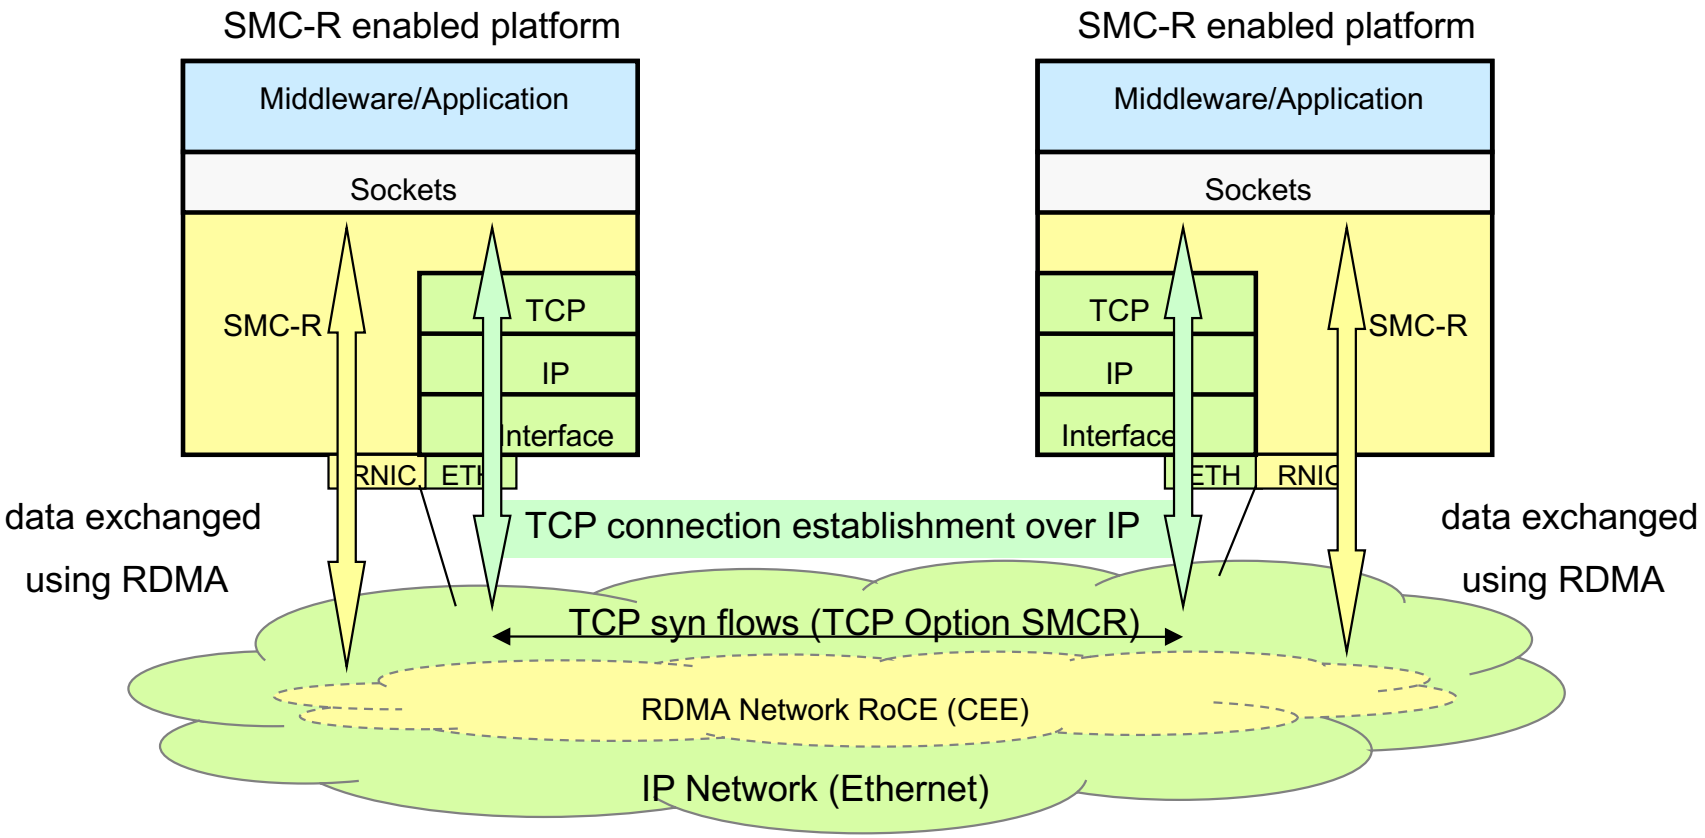

Dynamic (in-line) negotiation for SMC-R is initiated by presence of TCP Option (SMCR)

TCP connection transitions to SMC-R allowing application data to be exchanged using RDMA

Exploitation Considerations

SMC-R Review:

- 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 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 by field personnel or the customer**
  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). This aspect will benefit the Linux shared RoCE environment.

- **z/VM guest support for both SMC-R and SMC-D**

SMC-D Overview (introduced with z13 GA2)

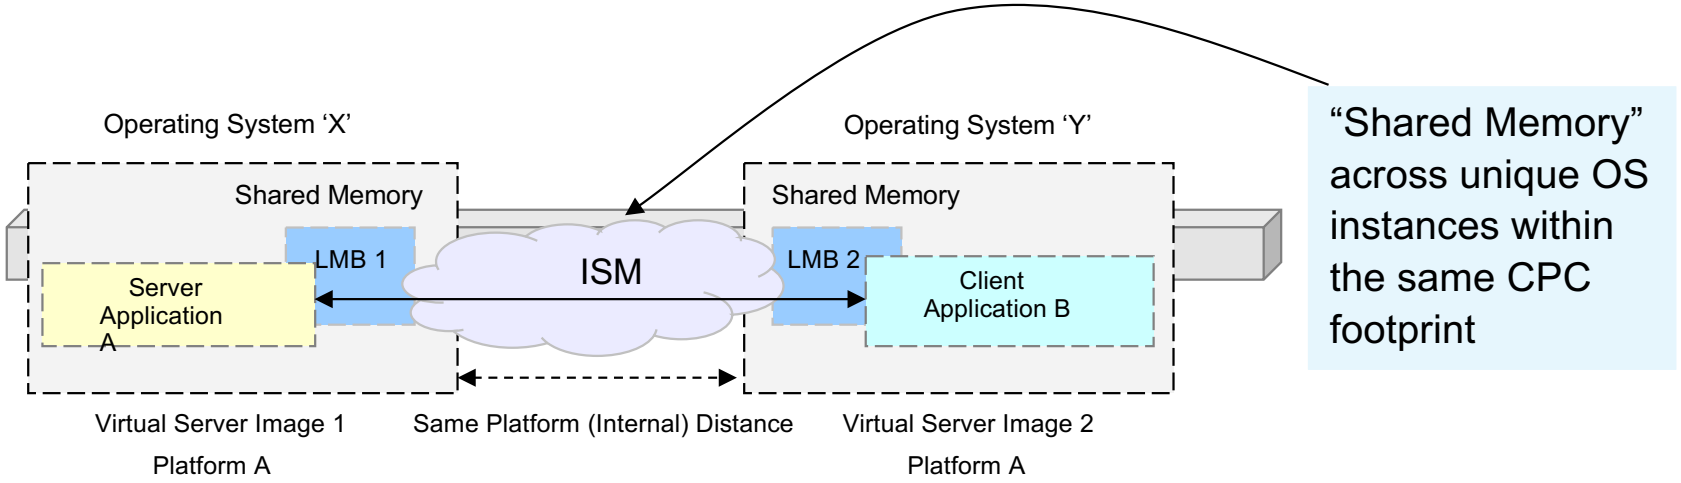

SMC-D over ISM is very similar to SMC-R over RoCE. SMC-D extends the benefits of SMC-R to same CPC operating system instances without requiring physical resources (RoCE adapters, PCI bandwidth, ports, I/O slots, network resources, 10GbE switches.)

Note: Reference performance information:

http://www-01.ibm.com/software/network/commserver/SMCR/

z14: Shared Memory Communication – Direct Memory Access (SMC-D)

- Local Shared Memory Communications for LPAR-to-LPAR links within a machine.
- Like HiperSockets but faster, less overhead and lower latency
- Performance Improvements
  - Up to 61% CPU savings for FTP file transfers across z/OS systems versus HiperSockets
  - Up to 9x improvement in throughput with more than a 88% decrease in CPU consumption and a 90% improvement in response time for streaming workloads versus using HiperSockets
  - Up to 91% improvement in throughput and up to 48% improvement in response time for request/response workloads versus using HiperSockets
- RMF is updated for SMC-D channel path activity and PCI activity

Shared Memory Communications Within the Enterprise Data Center (RoCE) and Within z Systems CPC (ISM)

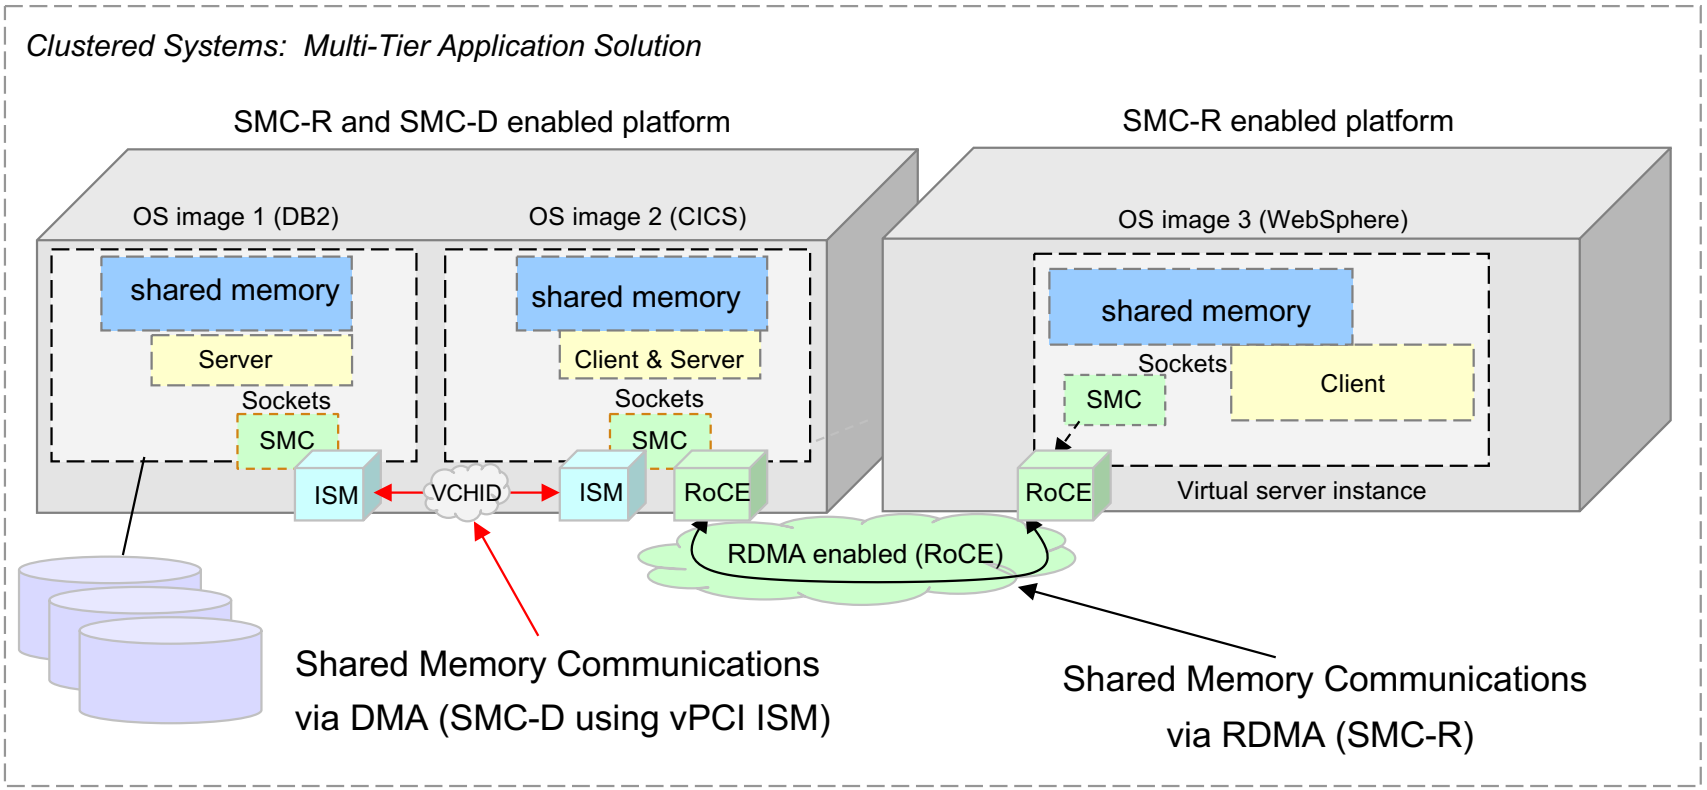

Both forms of SMC combine to provide a highly optimized solution. Shared Memory Communications: via z Systems PCI architecture:

1. RDMA (SMC-R for cross platforms via RoCE)
2. DMA (SMC-D for same CPC via ISM)

Table of Content

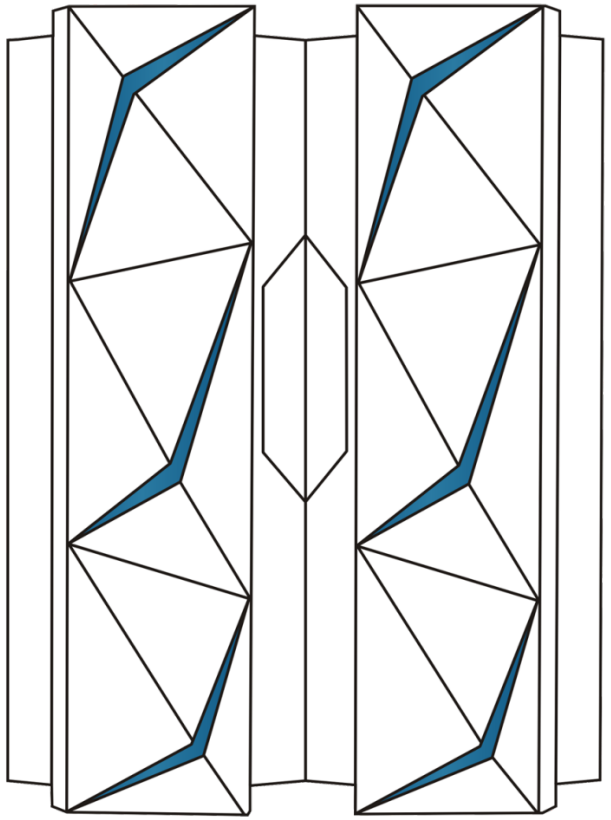

zHyperLink Express – A better Data Server !

- **Coupling links attain 8 microsecond response time.**
  - **FICON/zHPF attains approximately 145-170 microsecond read/write.**
- **Point-to-point connection using PCIe Gen3.**
- **The optical MPO cables and transceivers are the same as those defined for CS5 sysplex coupling links (ICA SR), with a maximum cable length of 150 meters (492 feet).**

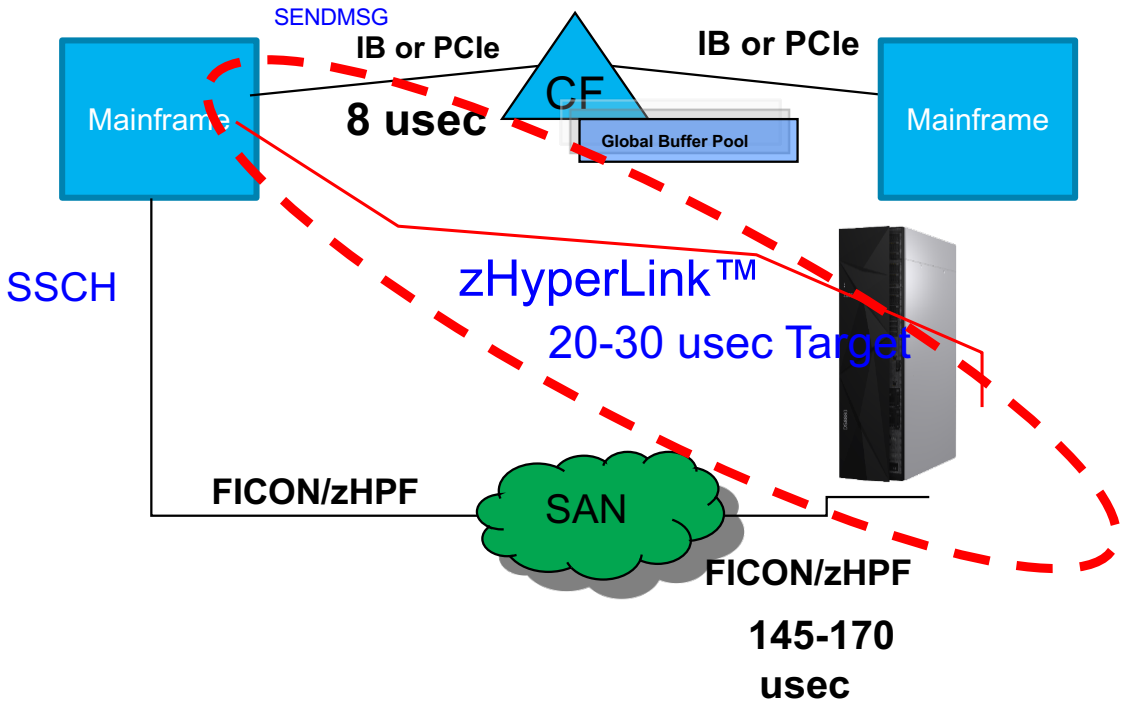

Size	I/O Operation	Average RT
4K	zHPF Read	126.9 usec
	zHyperLink Read	16.3 usec
4K	zHPF Write	TBD
	zHyperLink Write	TBD

Speed matters: Performance innovation for IBM Z

- Feature Code 0431
  - Min/Max, 0-16. Increment = 1 (2 ports)
- zHyperLink improves application response time, cutting I/O sensitive workload response time by up to 50% without requiring application changes.
- Designed to accelerate Db2 for z/OS Log write by up to 5x.
  - Significantly improve Db2 transaction rate.
- Reads: fast enough to serve data to Db2 without an undispatch during IOP
  - Improved Application and Middleware scaling

- Future exploitation ?
  - VSAM
  - MQ
  - IMS WADS/OLDS
  - CICSVR
  - ISV Enablement
  - z/OS System Logger

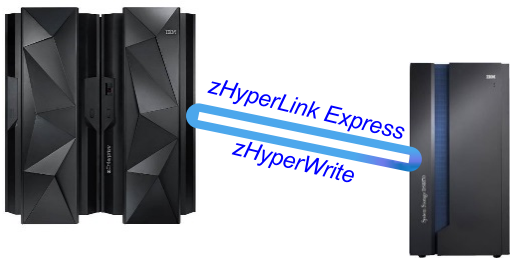

First Exploiter: Db2 V11 or V12

Stretch Sysplex with Asymmetric Performance

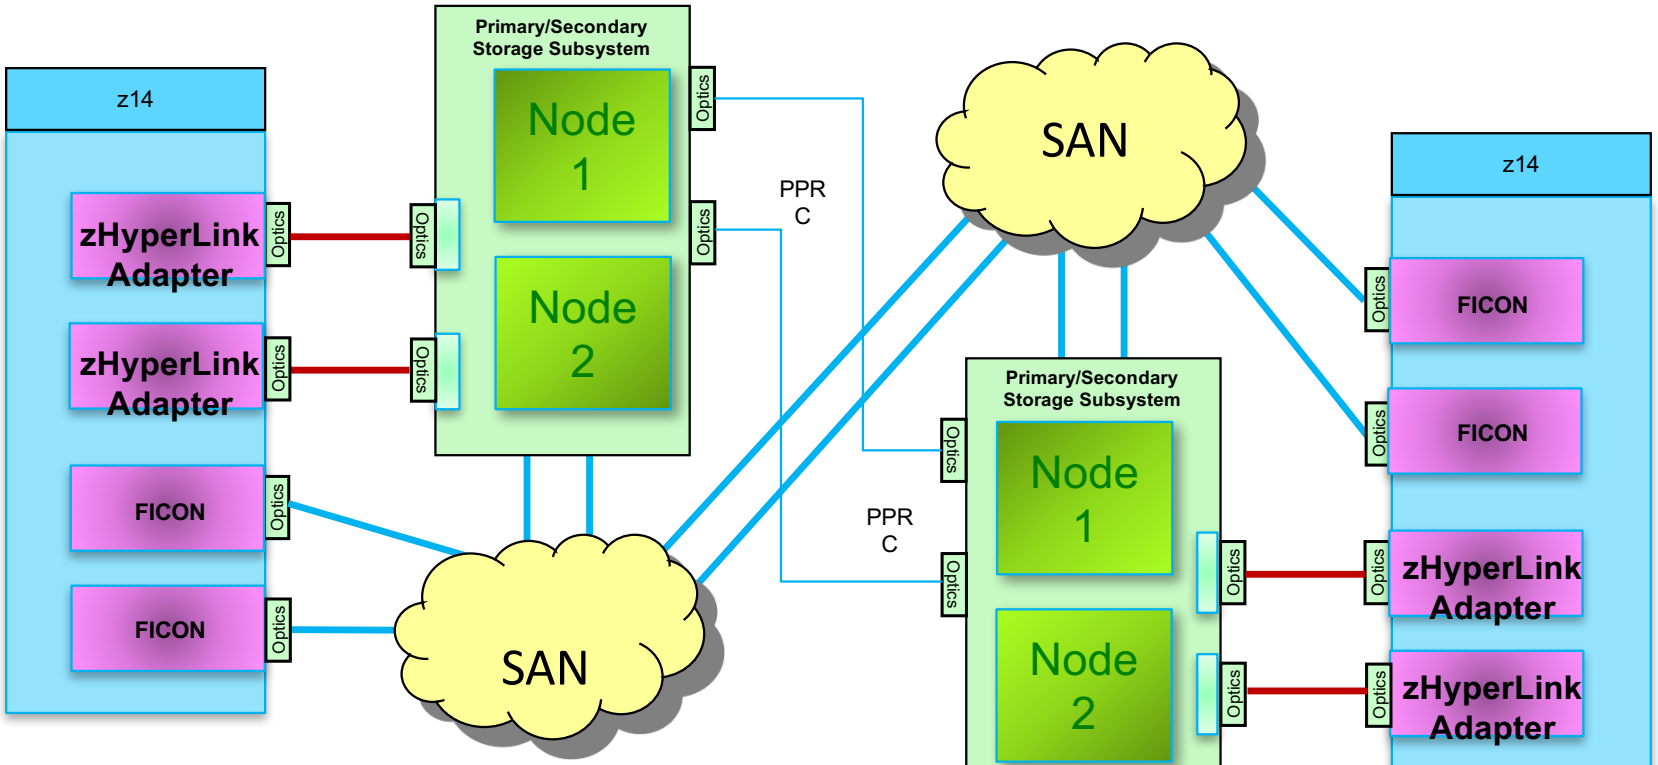

Transactions that can run against local DASD can run faster with zHyperLinks

zHyperLink Express – Current Restrictions

- Only ECKD supported
  - Fixed Block/SCSI to be considered for future release
- FICON & zHPF paths required in addition to zHyperLink Express
- Only native LPAR supported
  - z/VM and KVM guest support to be considered for a future release

zHyperLink Express & IOCP

- Function ID Type = HYL
- PCHID keyword
- Maximum of 16 features (32 ports)
  - Port keyword in the Function Statement (value = port 1 & port 2)
- Up to 127 Virtual Functions (VFs) per PCHID
  - Split VFs between both ports
- PNETID not allowed

FUNCTION PCHID=100,PORT=2,FID=1000,VF=16,TYPE=HYL,PART=((LP1),(...))

DS8880 zHyperLink™ Ports

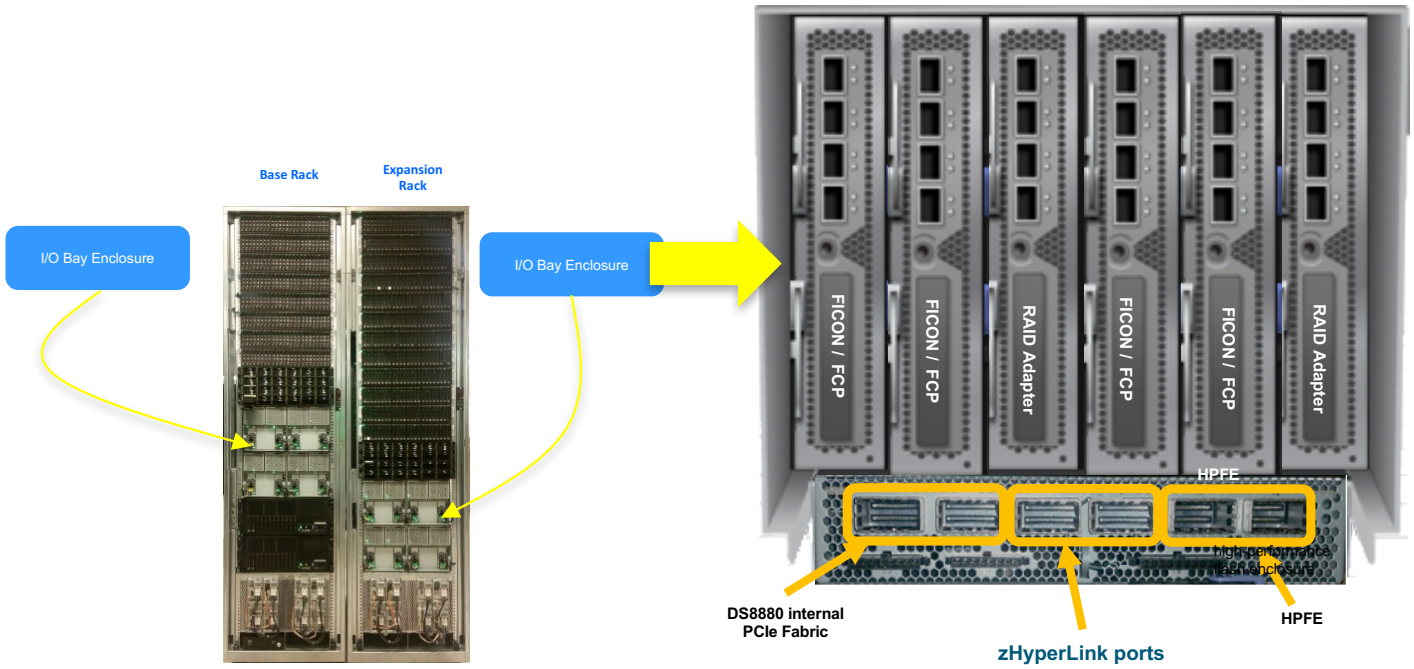

Investment Protection – DS8880 hardware shipping 4Q2016 (models 984, 985, 986 and 988), older DS8880's will be field upgradeable at z14 GA

Components of zHyperLinks

- DS8880 -- Designed for Extreme Low Latency Access to Data and Continuous Availability
  - New zHyperLink is an order of magnitude faster for simple read and write of data
  - zHyperWrite protocols built into zHyperLink protocols for acceleration of database logging with continuous availability
  - Investment protection for clients that already purchased the DS8880
  - New zHyperLinks compliment, do not replace, FICON channels: *Note: A standard FICON channel (CHPID type FC) is required for exploiting the zHyperLink Express feature
- z14 – Designed from the Casters Up for High Availability, Low Latency I/O Processing
  - New I/O paradigm transparent to client applications for extreme low latency I/O processing
  - End-to-end data integrity policed by IBM Z CPU cores in cooperation with DS8880 storage system
- z/OS, DB2 --New approach to I/O Processing
  - New I/O paradigm for the CPU synchronous execution of I/O operations to SAN attached storage. Allows reduction of I/O interrupts, context switching, L1/L2 cache disruption and reduced lock hold times typical in transaction processing work loads.

- IBM intends to deliver VSAM exploitation of z14 and DS8880 zHyperLink Express. zHyperLink Express is a short distance mainframe attach link designed for up to 10x lower latency than High Performance FICON.

Table of Content

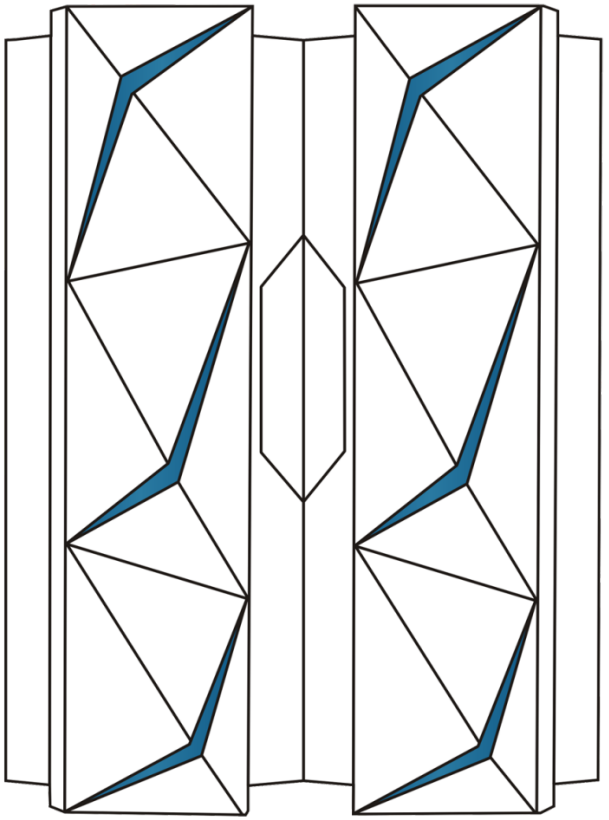

z14 Coupling Connectivity

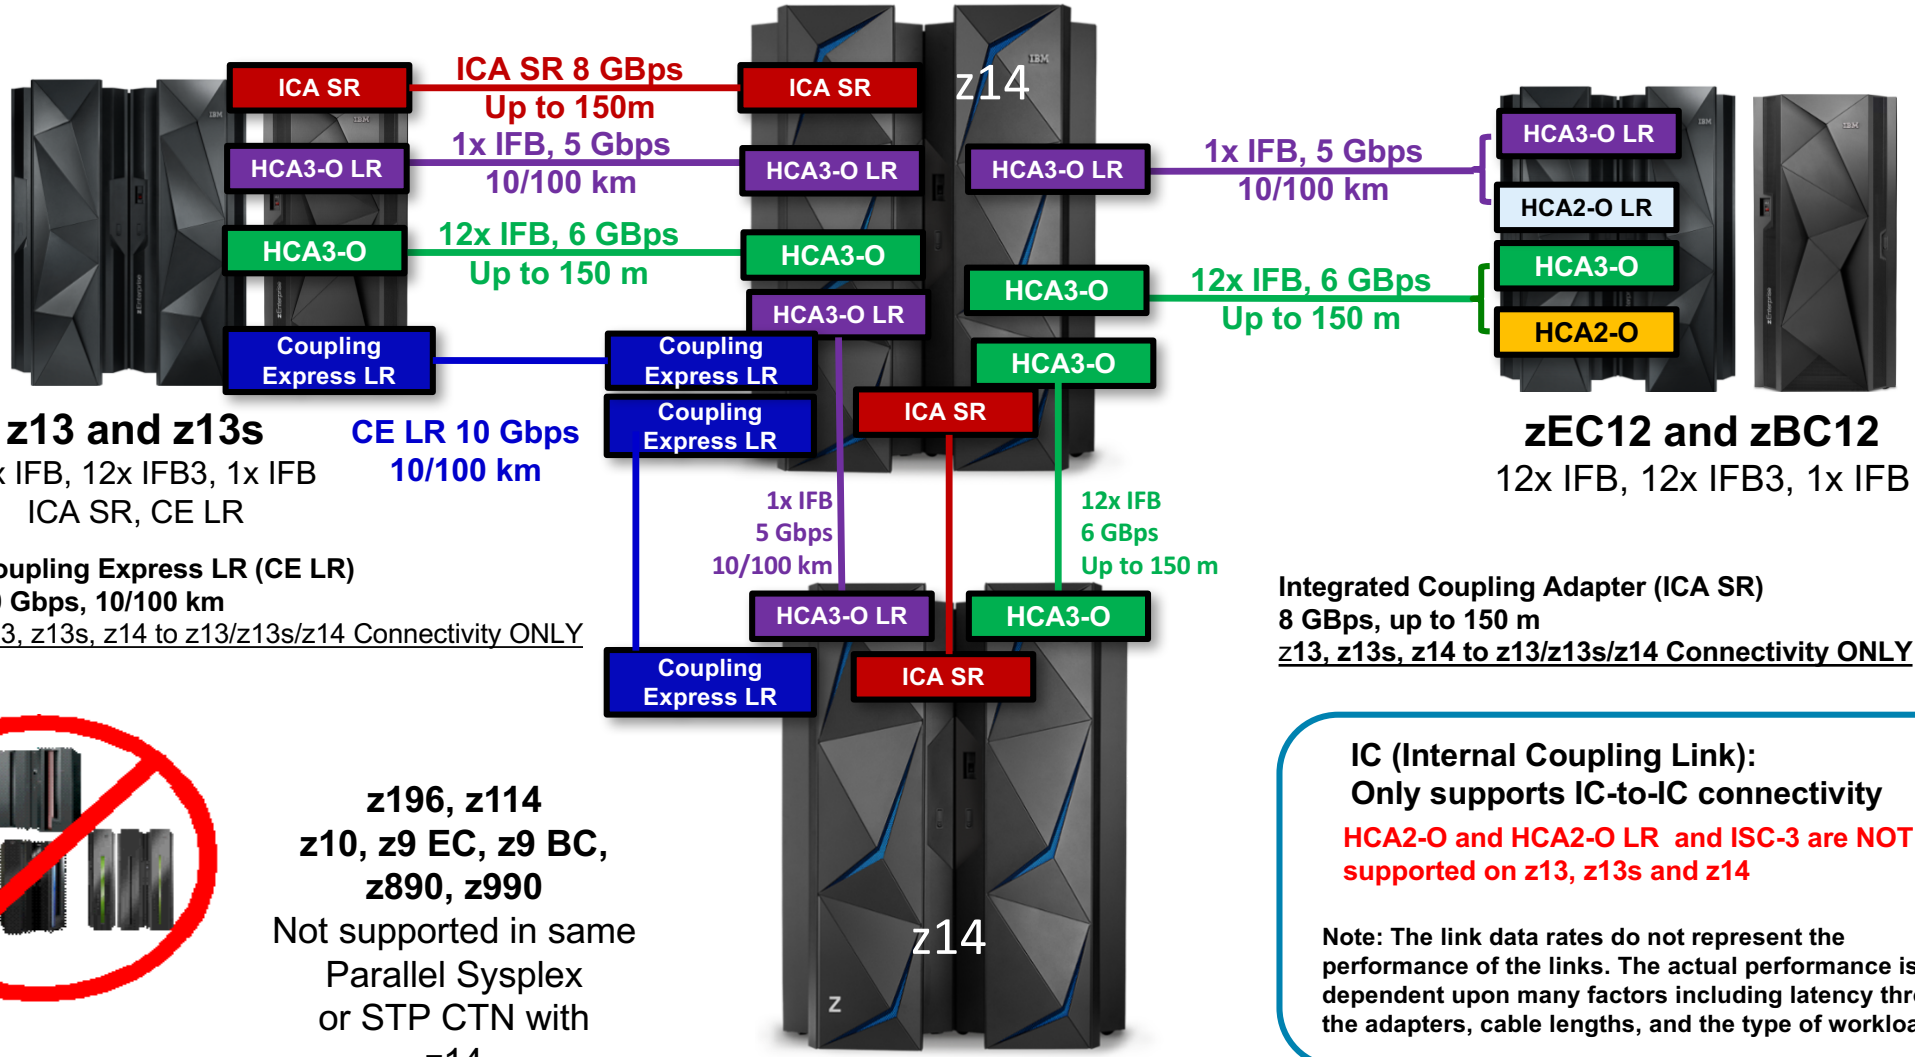

z13 and z13s
12x IFB, 12x IFB3, 1x IFB
ICA SR, CE LR

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

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

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

z196, z114
z10, z9 EC, z9 BC,
z890, z990
Not supported in same Parallel Sysplex or STP CTN with z14

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 and 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.

z14 Processor Drawer Connectivity for I/O and Coupling

Last machine to support HCA3 features

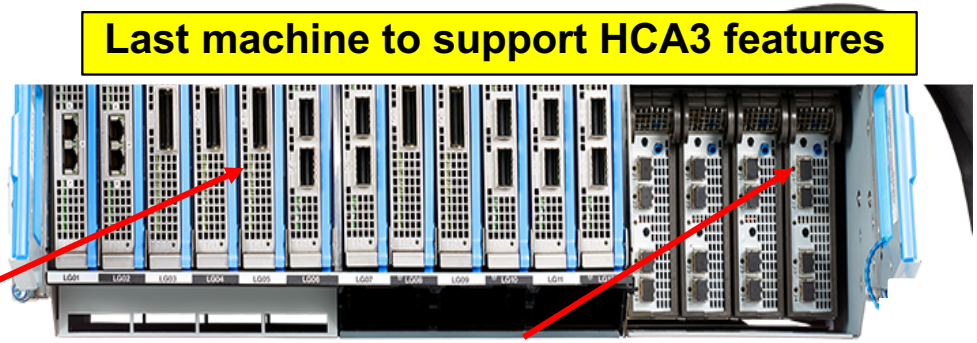

- 10 PCIe fanout slots per CPC drawer

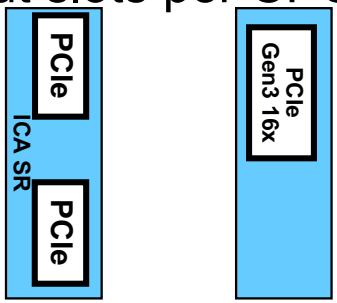

- 4 IFB HCA3-O (GX++) fanout slots per drawer

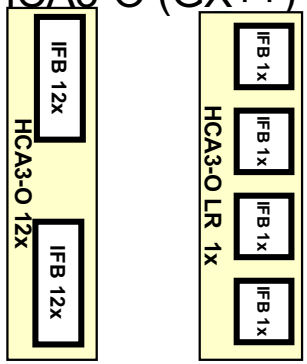

(quantities per CPC Drawer)

Coupling Express LR – FC 0433

- Coupling Connectivity: z14/z13/z13s -to- z14/z13/z13s
- Same adapter as 10 GbE RoCE Express2 but with Coupling Optics and firmware
- 10 Gbps - Actual performance depends upon many factors including latency through the adapters, cable lengths, and the type of workload
- Distance: 10 km unrepeated; up to 100 km with a qualified DWDM
- Cabling: Utilizes same 9u, Single Mode fiber type as 1X IFB and ISC-3
  - 9/125 micrometer SM optic with LC Duplex Connector

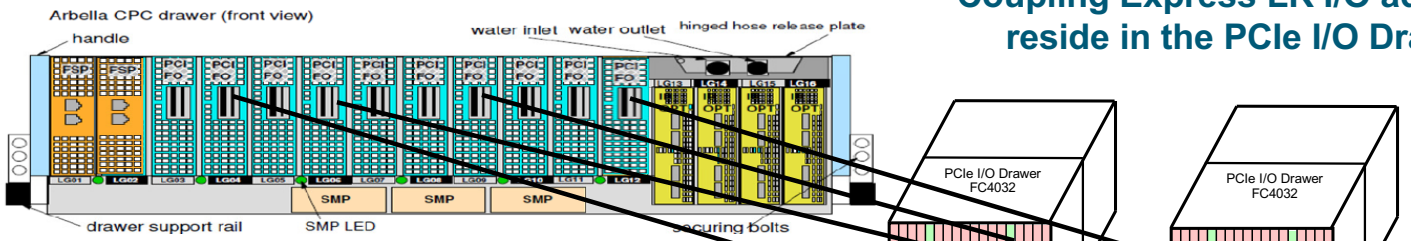

Coupling Express LR I/O adapters reside in the PCIe I/O Drawer

PCIe Gen3 fanouts drive to domains in PCIe I/O drawer to enable PCIe I/O features.

A minimum of two PCIe Fanouts required for the PCIe I/O Drawer

Parallel Sysplex Coupling Links – 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 - CHPID is identified by VCHIDs
- 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*, Up to 4 CHPIDs per port, 32 buffers (i.e. 32 subchannels) per CHPID
- Maximum configuration supported:
  - 64 ports (32 features) for z14 and z13 , 32 ports (16 features) for z13s

- Card is Identified by PCHID in IOCDS (card/slot -- unlike AID which is used for PSIFB)
  - CHPID PATH=(CSS(0),34),TYPE=CL5,CPATH=35,PCHID=234,PORT=1,PART=LP1,CSYSTEM=SYST
  - Point to Point (no switching)
- Cabling: Utilizes same 9u, Single Mode fiber type as 1X IFB and ISC-3
- z13 GA2+, z13s GA+ and z14 availability
- No Going Away Signal for STP

*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 Express LR - Migration Considerations

- Hardware Infrastructure requirements
  - PCIe I/O Drawer required; may drive additional infrastructure (ex: stand alone CF users).
  - On a Stand Alone CF, adding features will generate a PCIe I/O Drawer requirement.
  - A minimum of two PCIe Gen3 Fanouts are required for the PCIe I/O drawer, using RII.
- Ports/optics are not FRUs (field replaceable units).
- No Going Away Signal (GAS) in Server Time Protocol (STP) with CE LR
  - Only a concern if no Arbiter.
  - Console Assisted Recovery (CAR) still works as designed.
- Defined in Resource Groups (4), like 10GbE RoCE Express2 (SMC-R).
  - Separate EC stream from 10 GbE RoCE Express, and global pause not required.
    - Switch setting
  - Identified by PCHID instead of AID
  - Reduced SAP utilization

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)*.
  - 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

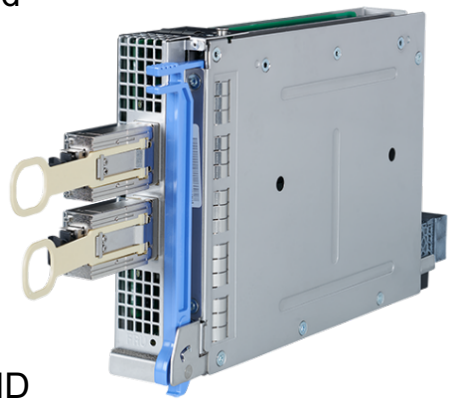

*Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

CFCC Enhancements

- **Enhanced Security and Data Protection through CF Structure Encryption**
  - Today, customer data that flows through the Coupling Facility (CF) and the CF Link infrastructure is vulnerable to potential for exposure because the data is not encrypted
  - With z/OS 2.3 support for CF Encryption:
    - Via CFRM policy, you can direct that CF structures be transparently encrypted, with no middleware or application changes needed
    - For those structures, host-based CPACF protected key AES-256-CBC encryption will be used to encrypt (decrypt) data and adjunct data that is written to (read from) the designated structure
  - Encryption provides improved protection against potential security breaches that could otherwise expose sensitive data

- **Asynchronous duplexing of CF lock structures provides significant value:**
  - The availability advantages of having a duplexed lock structure to provide quick failover and operational redundancy of sysplex lock information, in either a standalone CF or internal CF configuration, avoiding problems with
    - Rebuild recovery delays, and
    - Group-wide outages due to “lost locks” when rebuild is not possible
  - Simplex-like lock structure CF response times and throughput/performance
  - Efficient lock structure duplexing, even at extended GDPS distances – enabling a “continuous availability” multi-site GDPS DR environment in which all critical lock, list, and cache data for a datasharing workload is maintained at both sites, and can fail over in a site outage scenario
  - Available today – supported on z13 and z13s servers at CFCC service level 02.16 and higher, as well as on z14

z/OS Host-Based Data Encryption

New SMF metrics will let you
Estimate (or measure) cost of buffer copies and encryption (via zBNA tooling)

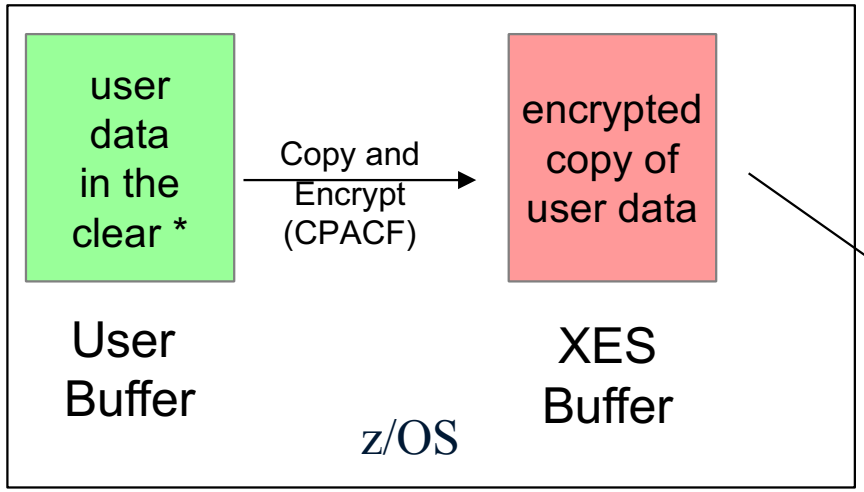

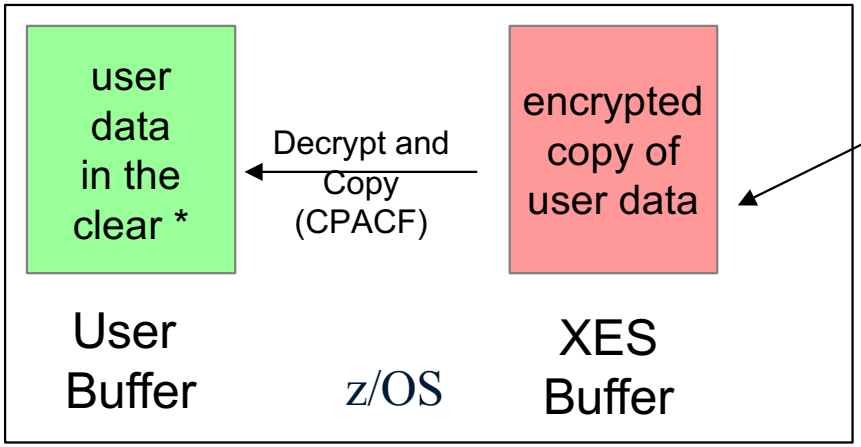

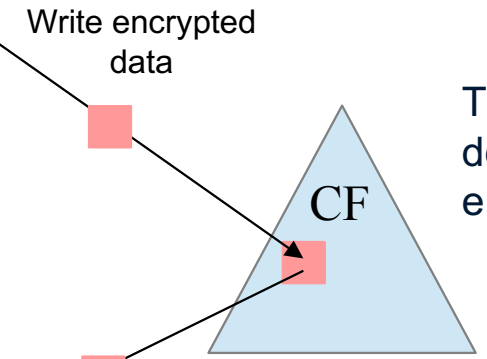

** User buffer assumed to be in the clear…. it gets encrypted even if it's not*

The Coupling Facility does NOT do any encryption or decryption

The CF Links do NOT do any encryption or decryption

Overview of CF Structure Encryption

- Encryption of each CF structure is controlled through a new z/OS Coupling Facility Resource Management (CFRM) policy keyword:
  - ENCRYPT(NO): structure should NOT be encrypted (the default)
  - ENCRYPT(YES): structure should be encrypted
- Dynamic, non-disruptive changes to the encryption state for an allocated structure occur through the existing CF structure rebuild mechanism
  - Clear to encrypted, encrypted to clear, encrypted to encrypted (key changes)
- CF Structure encryption is completely transparent to CF-exploiting middleware and applications
  - CF exploiters are neither aware of, nor involved in, the encryption
- All data and adjunct data flowing between z/OS and the CF is encrypted
  - Intent is to encrypt all data that might be sensitive
  - Internal control information and related request metadata is not encrypted, including locks / lock structures
- z/OS generates the required structure-related encryption keys and does much of the key management automatically via CFRM using secure, protected keys (never clear keys)
  - Secure keys maintained in CFRM couple dataset
- z/OS supports dynamic structure-level key changes and master key changes

Must rebuild structure for dynamic changes to CF structure encryption

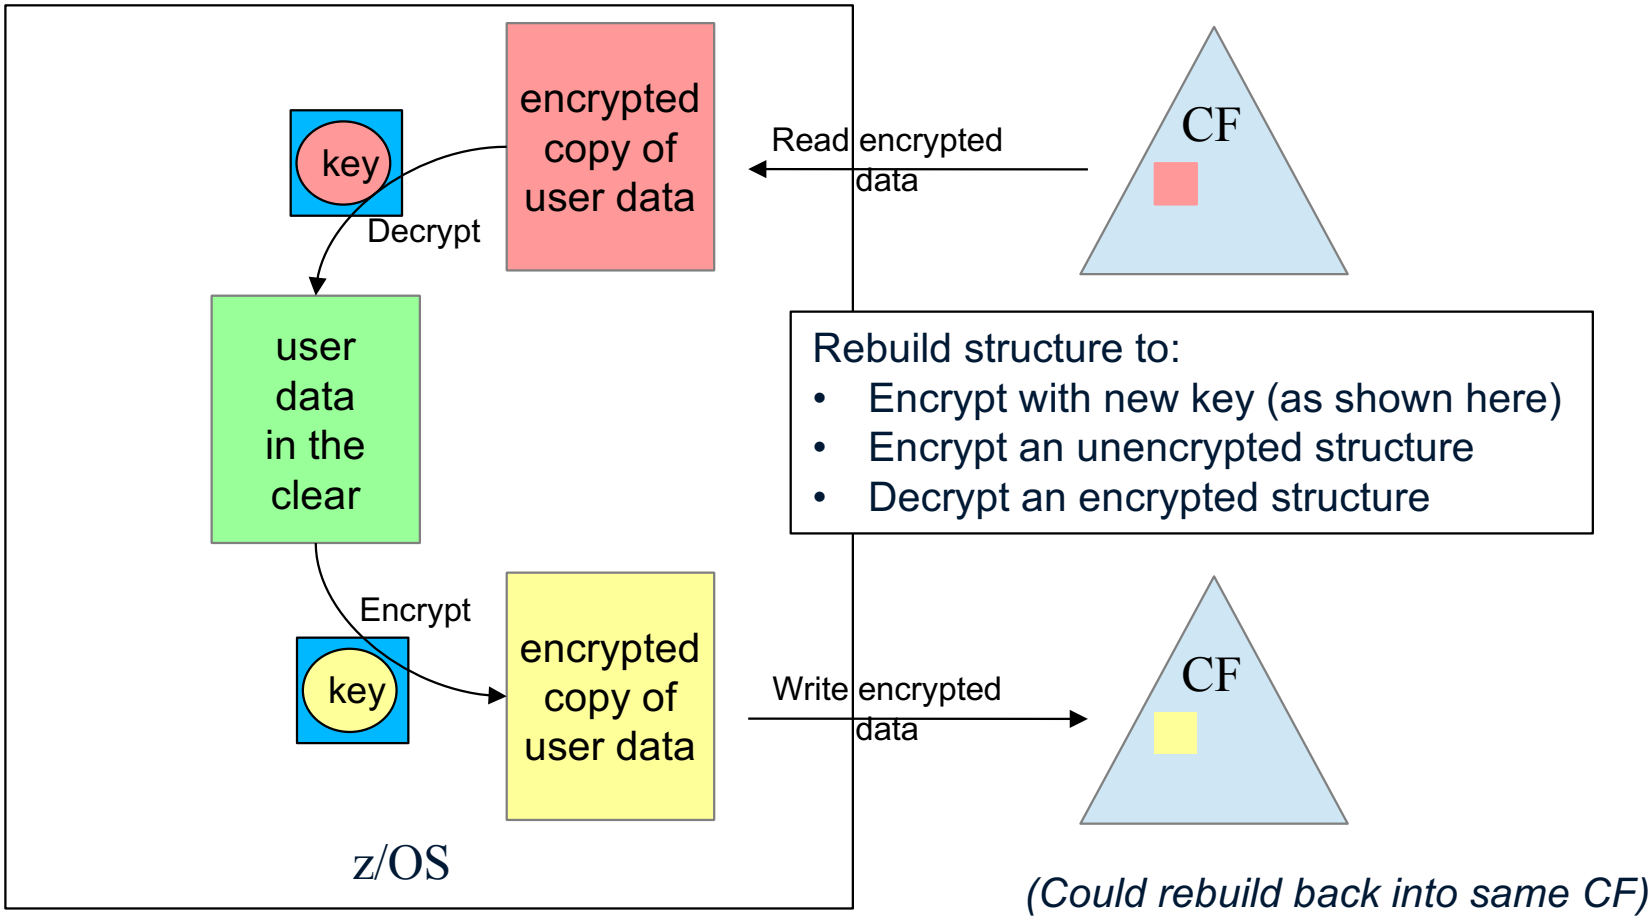

CF Structure Encryption Requirements Summary

z/OS V2.3:

Product/Feature	Required Level	Description
Hardware		
z/OS: Minimum HW	zEC12	Minimum supported for z/OS 2.3
	Crypto Express3	Required for Protected-key CPACF
z/OS: Recommended HW	z14	AES-CBC CPACF encrypt/decrypt performance improvements
CF: Recommended HW	z14	Simplified recovery for sysplex-wide CF reconciliation scenarios when using a new/changed CFRM couple dataset
Operating System – Base Support		
z/OS	z/OS 2.3	z/OS XCF/XES support for CF encryption
Additional Support		
zSecure	zSecure 2.3	zSecure Audit support for CF encryption
zBNA	zBNA	zBatch Network Analyzer support for CF encryption

Summary of CF Structure Encryption

Enhanced Security and Data Protection for CF Structures with z/OS 2.3

- Through a Coupling Facility Resource Manager (CFRM) policy specification, one can request that user data in designated Coupling Facility (CF) structures be encrypted
- Changes to encryption state of CF structures are accomplished via structure rebuild
- Central Processor Assist for Cryptographic Functions (CPACF) is used to encrypt and decrypt the data
  - Provides high performance encryption/decryption with low latency
- Data will be in the encrypted state:
  - While being transmitted between z/OS and the CF (in either direction)
  - When at rest in the CF
- Data is in the clear only when it resides in the z/OS host (e.g. local bufferpools)

Coupling Facility Asynchronous Duplexing for Lock Structures

- New asynchronous duplexing protocol for lock structures.
  - z/OS sends command only to primary CF
  - Primary CF processes command and returns result
  - Primary CF forwards description of required updates to secondary CF
  - Secondary CF updates secondary structure instance asynchronously
- Provided only for lock structures
  - Most performance-sensitive
  - DB2 V12 is first exploiter
- Benefit / Value
  - DB2 locking receives performance similar to simplex operations
    - Reduced CPU and CF link overhead
  - Duplexing failover much faster than log-based recovery
  - Targeted at multi-site clients who run split workloads at distance to make duplexing lock structures at distance practical

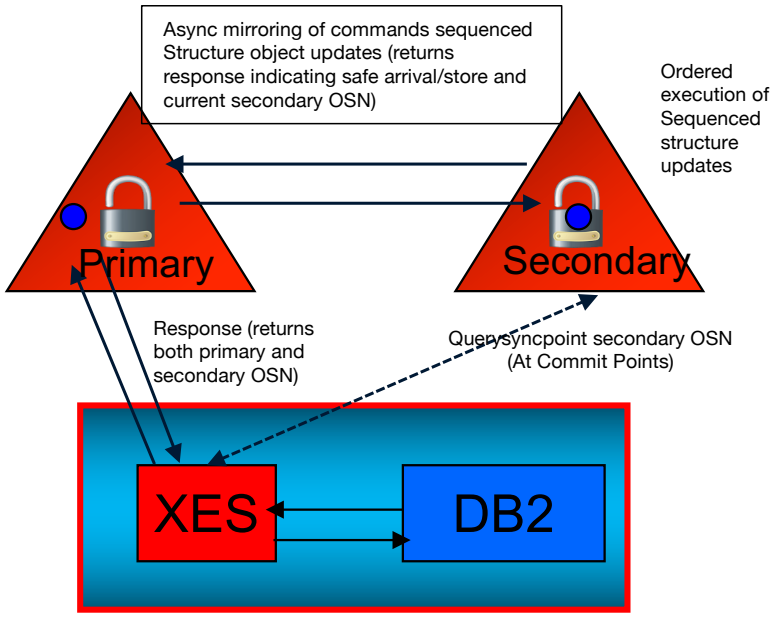

Duplexing of lock structures now practical for customers that could not afford the overhead

Table of Content

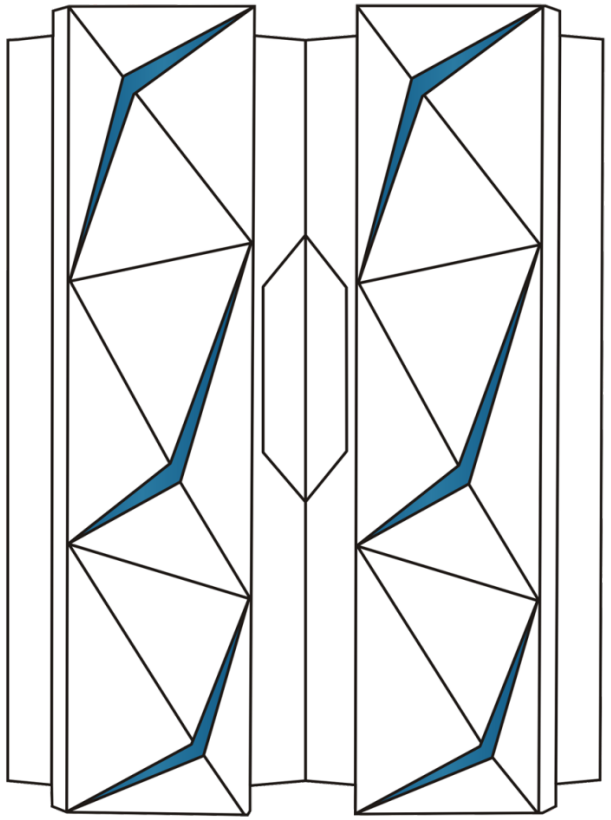

IBM Virtual Flash Memory (VFM) – Feature Code 0604

- IBM Virtual Flash Memory
  - Replacement for IBM Flash Express I/O features – same use cases.
  - Up to 10% end to end performance improvement - estimation
    - 1000X improvement in Read/Write latency - estimation
  - Saves at least two PCIe I/O Drawer Slots
  - Less power consumption
  - During upgrade, Feature Conversion for IBM Flash Express
  - VFM Concurrent Add if physical memory is available or when there are multiple CPC Drawers

- Increment Size
  - Up to four features/increments
    - 1.5 TB, 3.0 TB, 4.5 TB, 6.0 TB
  - Customer ordered memory + HSA (192GB) + VFM (n x 1.5 TB) + RAIM = physical memory installed.
  - If flexible memory is ordered, VFM flexible memory will be included by a feature code.
  - The memory associated with the VFM can be evacuated to another processor drawer along with the “regular” customer memory.

Activation Profiles/Activation Profile Wizards

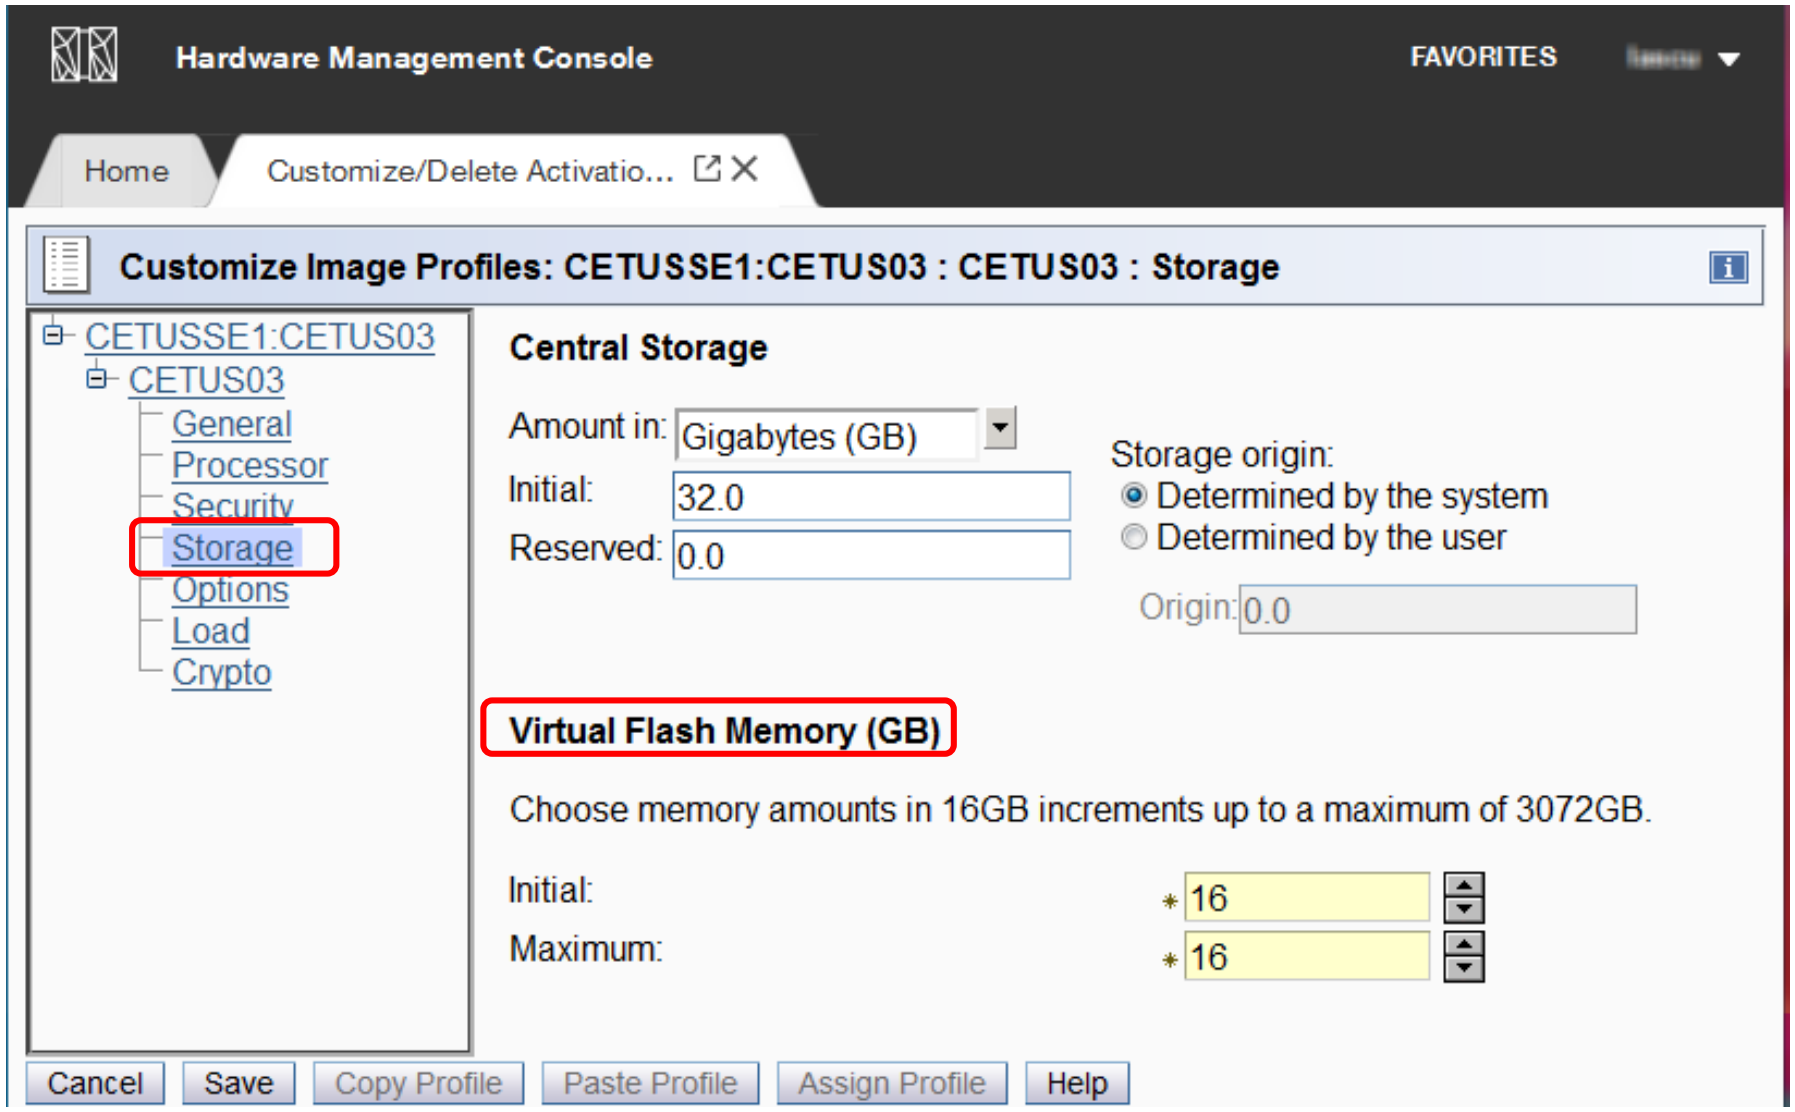

Migration considerations

- VFM is a REPLACEMENT for Flash Express and has slightly different behavior characteristics.
- There are some situations that cannot be handled non-disruptively.
- Pre-planning is therefore critical !
  - Question: “Hot plug” no longer available to increase amount of installed Flash Memory (VFM). So how does the customer increase the amount of VFM for a partition?
  - Answer: “Over-commitment” of VFM at partition activation time is possible. The maximum amount of VFM can be set to the LICCC value in each partition.
- Allows more storage to be added to partitions subject to amount currently not assigned (same rule applied to Flash Express implementation)
- If total amount of VFM storage allocated to all active partitions is equal to the LICCC value, but the sum of active partition maximums is larger than the installed amount, customer may be able to concurrently add VFM and increase allocations without re-activating partitions.

Table of Content

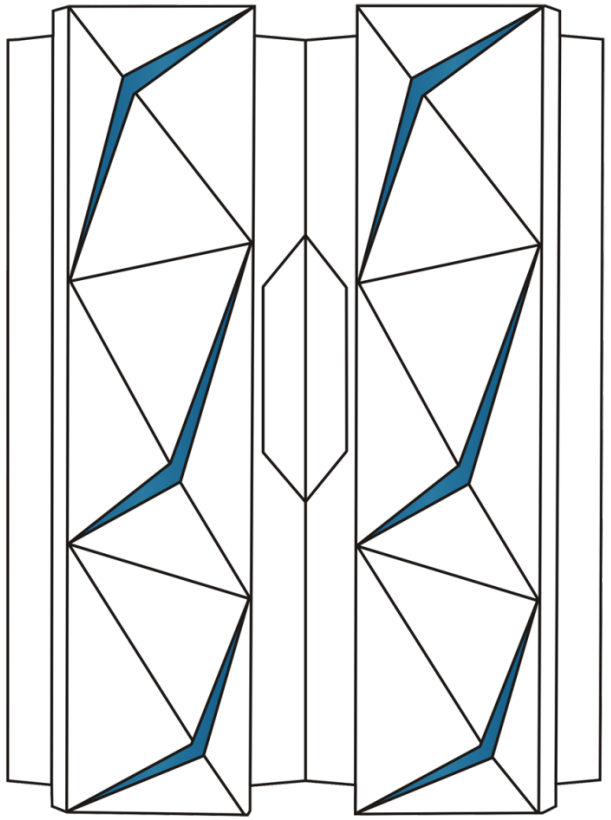

Deploying zEDC (zEnterprise Data Compression) – Review*

- Operating system requirements
  - Requires z/OS 2.1 (with PTFs) and later with the zEDC Express for z/OS feature
  - z/OS V1.13 and V1.12 offer software decompression support only
  - z/VM V6.3 and later support for z/OS V2.1 and later guest exploitation
- Server requirements
  - Available on zEC12, zBC12, z13, z13s and z14
  - zEDC Express feature for PCIe I/O drawer (FC#0420)
    - Each feature can be shared across up to 15 LPARs
    - Up to 8 features available on zEC12/zBC12/z13/z13s/z14
  - Recommended high availability configuration per server is four features**
    - This provides up to 4 GB/s of compression/decompression
    - Provides high availability during concurrent update (half devices unavailable during update)
    - Recommended minimum configuration per server is two features
  - Steps for installing zEDC Express in an existing zEC12/zBC12/z13/z13s/z14
    - Apply z/OS Service; Hot plug zEDC Express; update your IODF, and Dynamic Activate

* For the full zEDC benefit, zEDC should be active on ALL systems that might access or share compressed format data sets. This eliminates instances where software inflation would be used when zEDC is not available.

** Availability rules relaxed for zEDC installed on a system with 4 RGs (z13 D27, z13s, with MCL, z14)

QSAM/BSAM Data Set Compression with zEDC

- Reduce the cost of keeping your sequential data online
- zEDC compresses data up to 4X, saving up to 75% of your sequential data disk space
- Capture new business opportunities due to lower cost of keeping data online
- Better I/O elapsed time for sequential access
- Potentially run batch workloads faster than either uncompressed or QSAM/BSAM current compression
- Sharply lower CPU cost over existing compression
- Enables more pervasive use of compression
- Up to 80% reduced CPU cost compared to tailored and generic compression options
- Simple Enablement
- Use a policy to enable the zEDC

Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Example Use Cases

SMF Archived Data can be stored compressed to increase the amount of data kept online up to 4X

zSecure output size of Access Monitor and UNLOAD files reduced up to 10X and CKFREEZE files reduced by up to 4X

Up to 5X more *XML* data can be stored in sequential files

The IBM Employee Directory was stored in up to 3X less space

z/OS SVC and Stand Alone DUMPs can be stored in up to 5X less space

QSAM/BSAM Data Set Compression with zEDC

- Setup is similar to setup for existing types of compression (generic and tailored)
  - It can be selected at either or both the data class level or system level.
    - Data class level
      In addition to existing tailored (T) and generic (G) values, new zEDC Required (ZR) and zEDC Preferred (ZP) values are available on the COMPACTION option in data class.
      When COMPACTION=Y in data class, the system level is used
    - System level
      In addition to existing TAILORED and GENERIC values, new zEDC Required (ZEDC_R) and zEDC Preferred (ZEDC_P) values are available on the COMPRESS parameter found in IGDSMSxx member of SYS1.PARMLIB.
    - Activated using SET SMS=xx or at IPL
      Data class continues to take precedence over system level. The default continues to be GENERIC.
- zEDC compression for new extended format data sets is Optional
  - All previous compression options are still supported
  - For the full zEDC benefit, zEDC should be active on ALL systems that might access or share compressed format data sets. This eliminates instances where software inflation would be used when zEDC is not available

BSAM/QSAM zEDC Compression Results

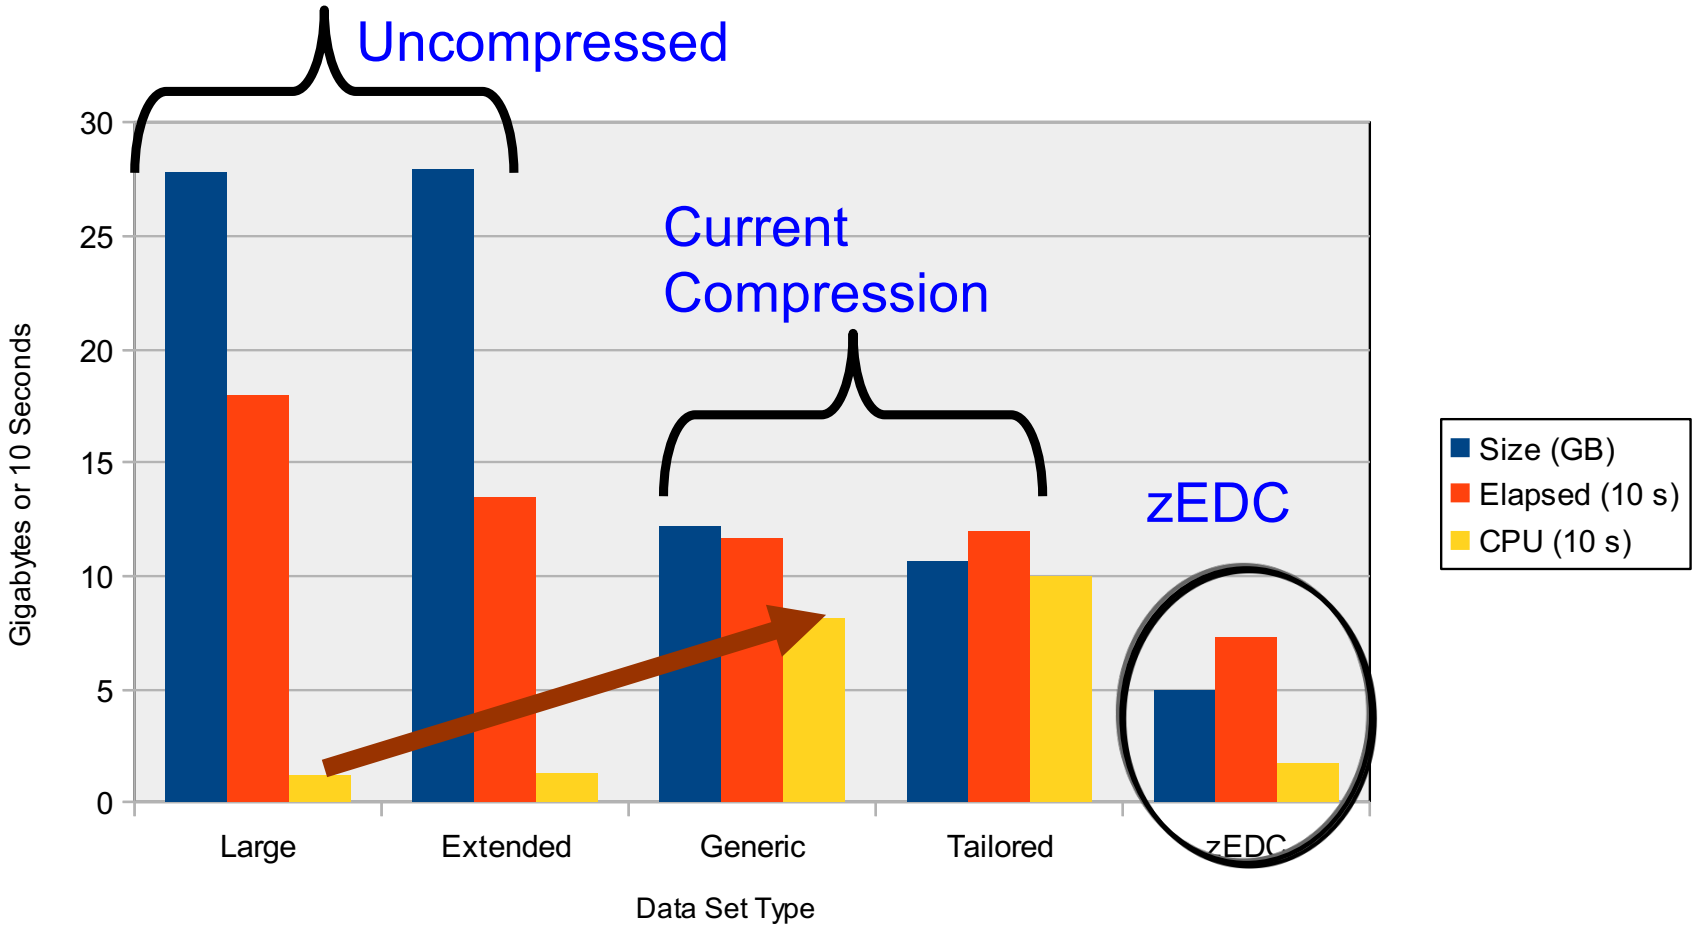

*Measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Compression Coprocessor (CMPSC) vs. zEDC

Using the right hardware compression acceleration for each of your workloads

	Compression Coprocessor	z Enterprise Data Compression
	On Chip In every IBM Z server Mature: Decades of use by Access Methods and DB2® Work is performed jointly by CPU and Coprocessor Proprietary Compression Format	PCIe Adapter Introduced with IBM zEnterprise® EC12 GA2 and IBM zEnterprise BC12 Mature: Industry Standard with decades of software support Work is performed by the PCIe Adapter Standards Compliant (RFC1951)
Use Cases		
	Small object compression ▪ Rows in a database	**Large Sequential Data** ▪ QSAM/BSAM Online Sequential Data ▪ Objects stored in a data base	**Industry Standard Data** ▪ Cross Platform Data Exchange
	Users ▪ VSAM for better disk utilization ▪ DB2 for lower memory usage ▪ The majority of customers are currently compressing their DB2 rows	**Users** ▪ QSAM/BSAM for better disk utilization and batch elapsed time improvements ▪ SMF for increased availability and online storage reduction	**Users** ▪ Java for high throughput standard compression via java.util.zip ▪ Encryption Facility for z/OS for better industry data exchange ▪ IBM Sterling Connect: Direct® for z/OS for better throughput and link utilization ▪ ISV support for increased client value

Table of Content

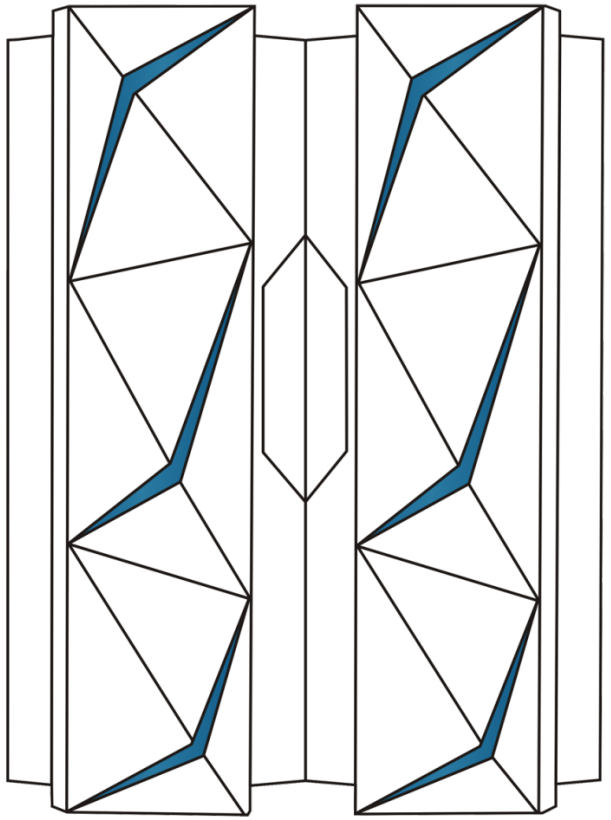

IBM Secure Service Container Overview - What is it?

- IBM Secure Service Container (SSC*)
  - provides the base infrastructure to create an **IBM Z Appliance** including operating system, middleware, SDK and firmware support
  - Deploy an appliance (that provides a function or a service) in minutes instead of days
  - Protect the workload from being accessed by a system administrator or an external attacker

- Objective
  - To streamline the IBM Z application experience so it is comparable to installing an app on a mobile device.

- Terminology
  - **IBM Z Appliance** is an integration of operating system, middleware and Software components that work autonomously and provides core services and infrastructure focusing on consumability and security
  - **SSC Partition** is a LPAR type running an **IBM Z Appliance** based on **Secure Service Container**. The Partition type provides main aspects of the security features.

*Formerly IBM Z Appliance Container Infrastructure - zACI)

IBM Z and Appliances Integration

Multiple virtual appliances integrated into IBM Secure Service Container LPARs on IBM Z servers

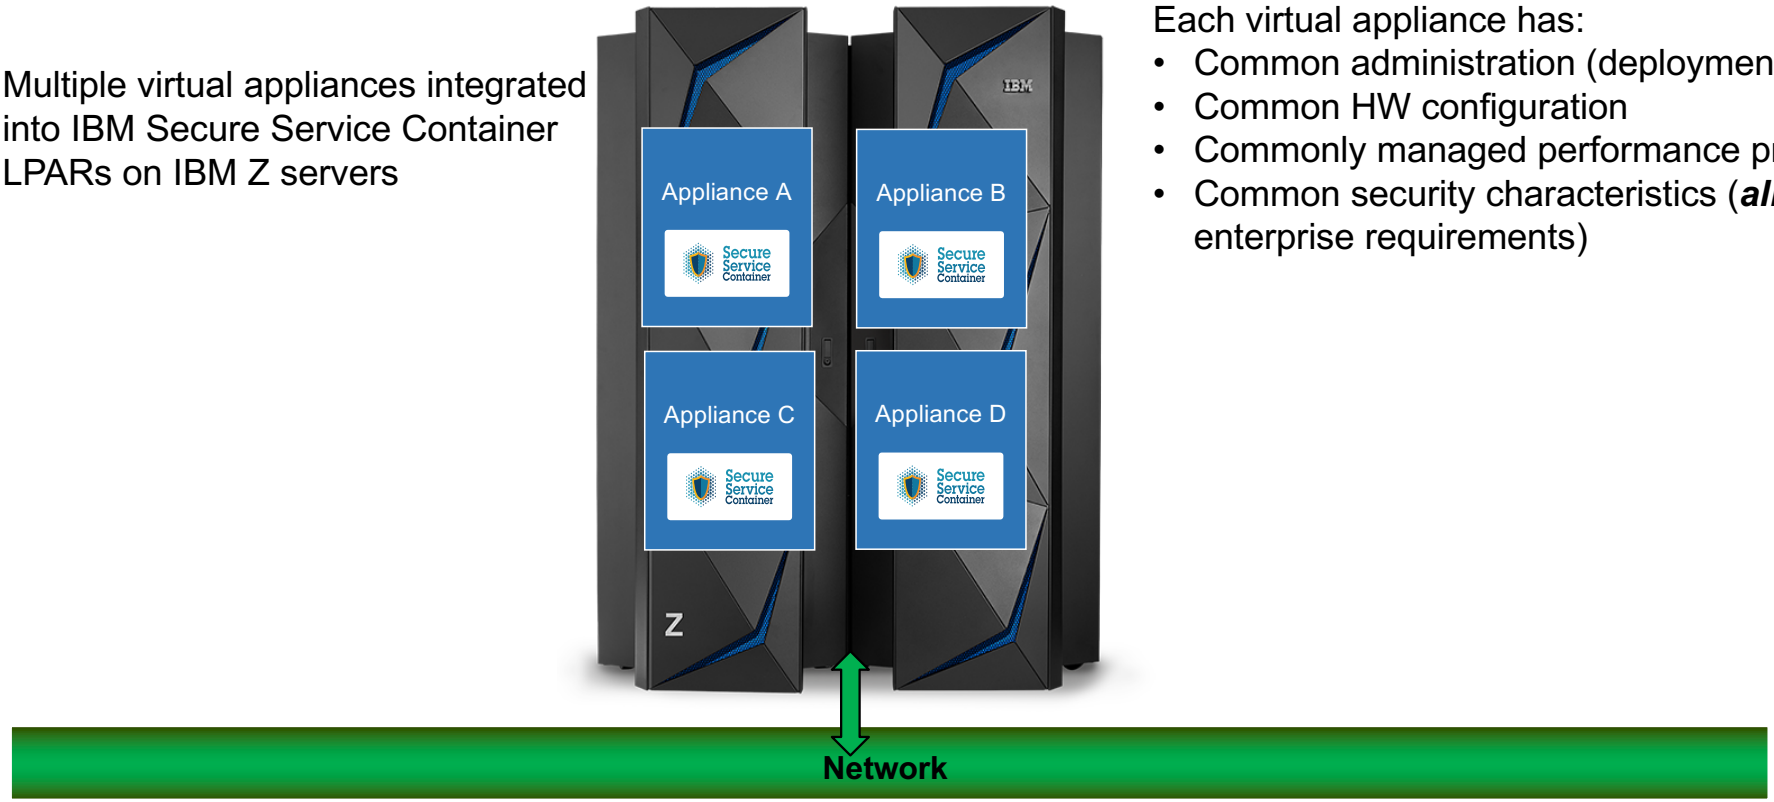

Each virtual appliance has:

- Common administration (deployment)
- Common HW configuration
- Commonly managed performance profiles
- Common security characteristics (***aligned*** with enterprise requirements)

IBM Secure Service Container Deployment Overview

Deployment in 5 steps

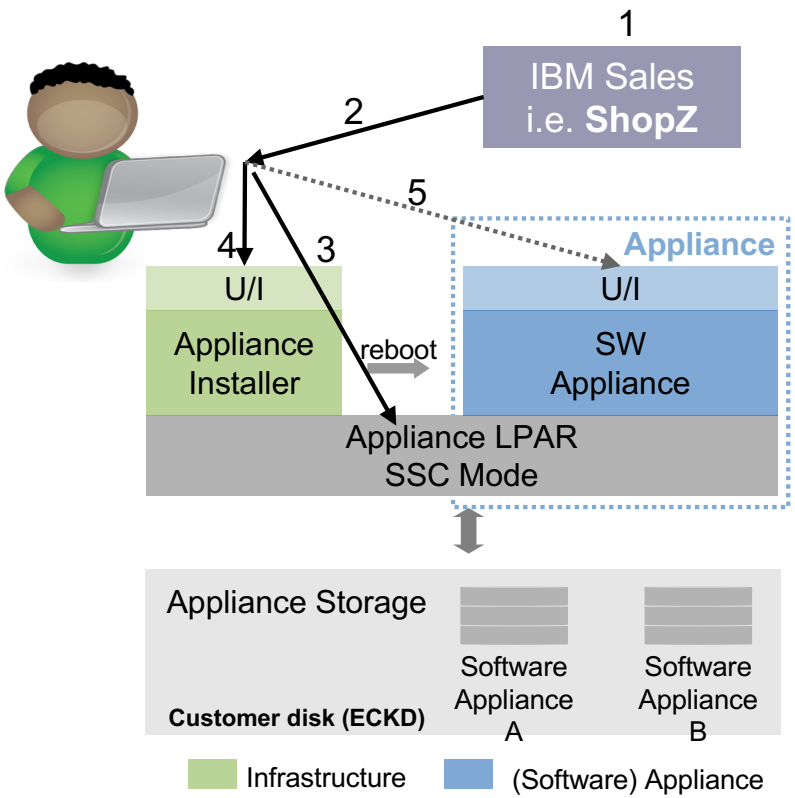

1) Buy a Software Appliance (e.g. IOAz/***zAware***)

2) Download the Appliance image

3) Create and activate an appliance (SSC) LPAR

4) Deploy Appliance using Appliance Installer

5) Configure and use Appliance through REST API or web U/I

SSC Supported Appliances: z/VSE Network Appliance

- Linux Fast Path - z/VSE IP Assist
- Customer has to
  - download the LFP image based on Secure Service Container SDK
  - install / configure the image
- Linux acts as a router for z/VSE
- No TCP/IP stack required on z/VSE
- No Linux distribution required

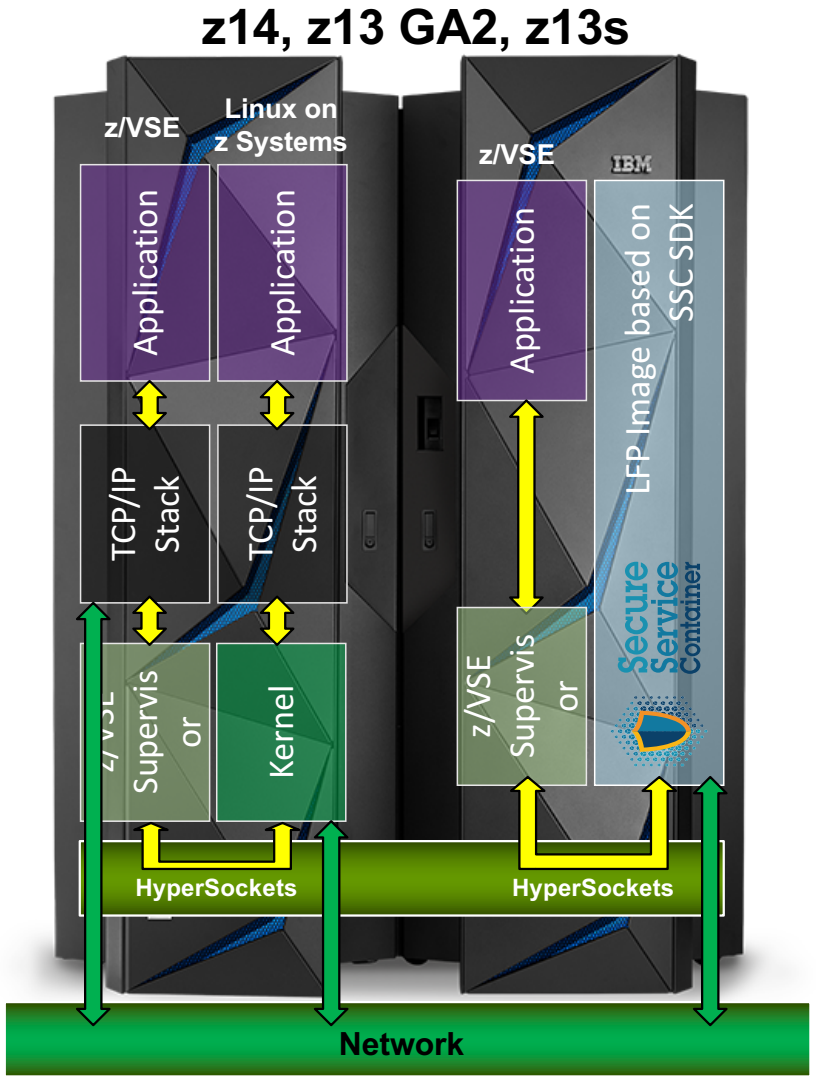

Supported Appliances (z13 GA2, z13s): IBM zAware* based on SSC

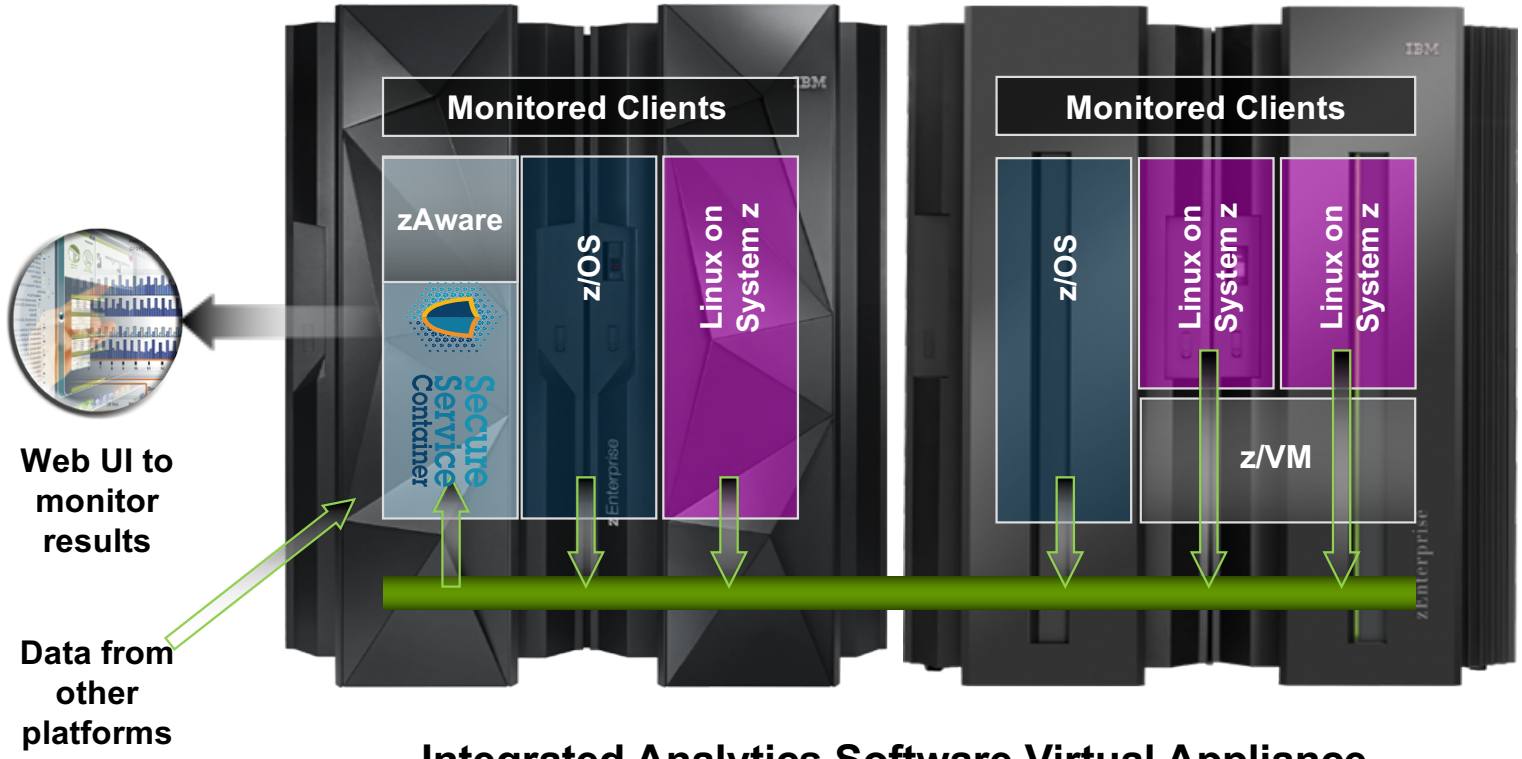

- Provides anomaly detection, search and reporting
- Identify unusual system behavior of z/OS and Linux on System z images
- Analyze message logs for all servers in the enterprise from single browser in real-time

***Note: IBM System z® Advanced Workload Analysis Reporter (IBM zAware), is now deployed as software and integrated with IBM Operations Analytics for IBM Z (IOAz)**

Supported Appliances: Blockchain

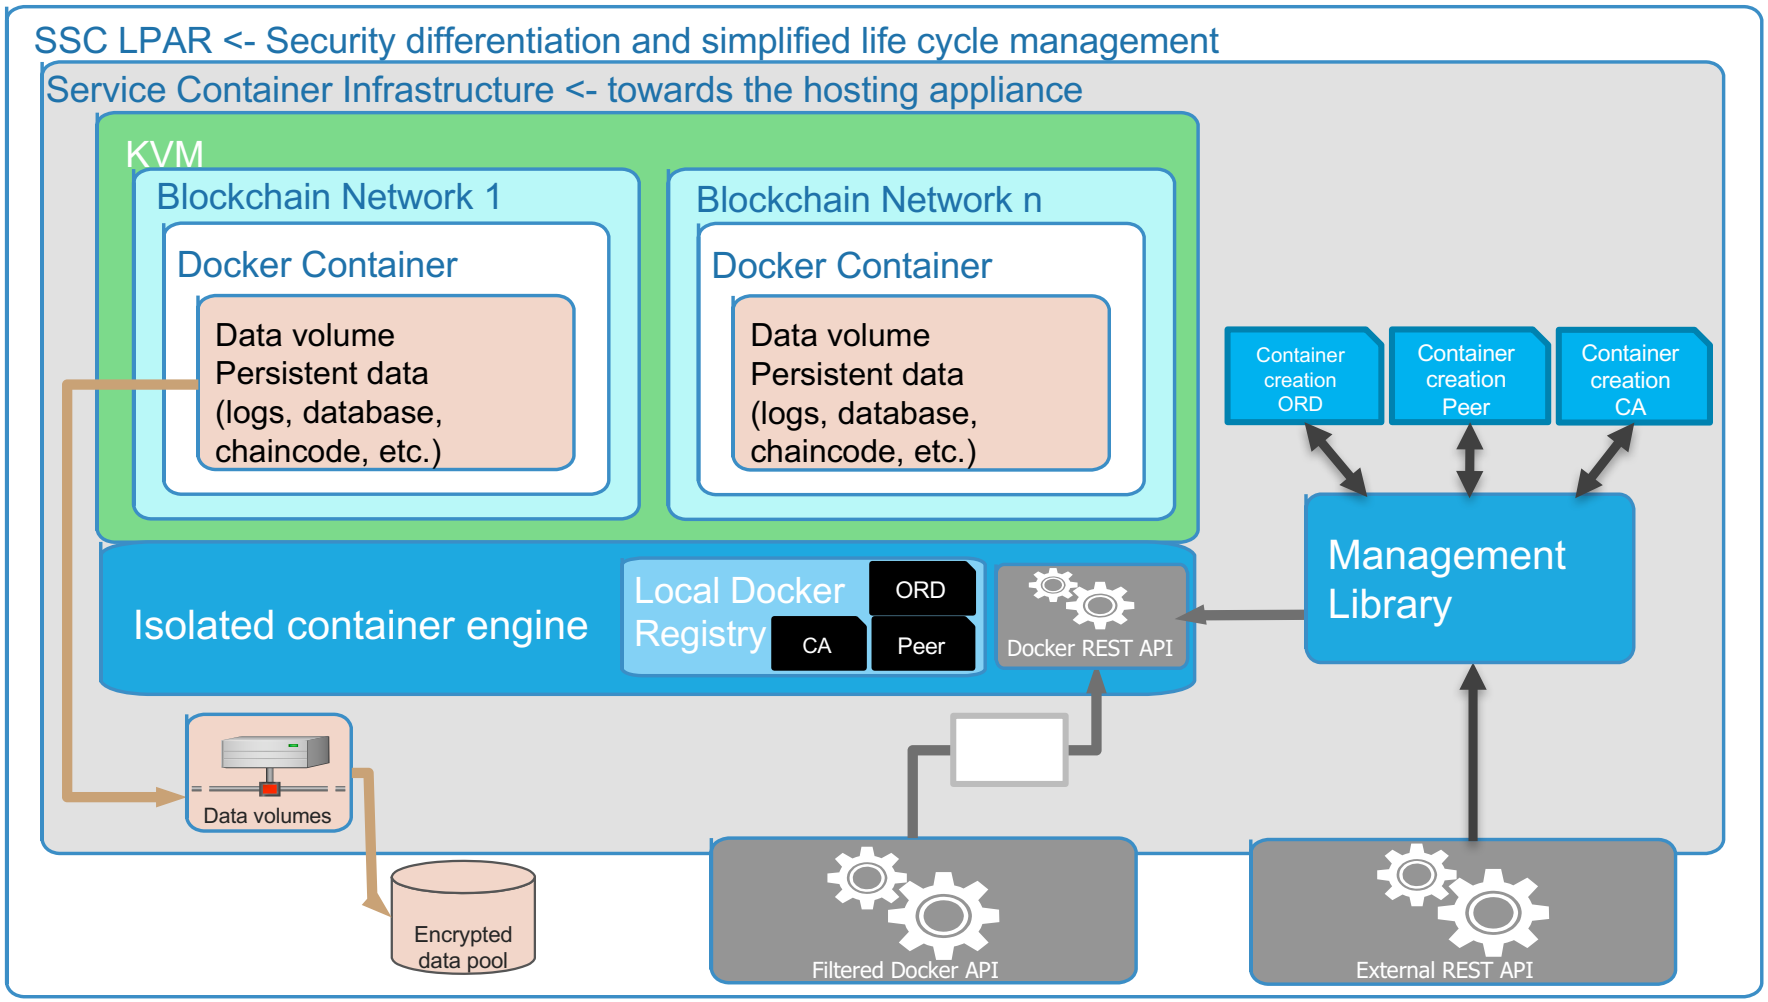

Table of Content

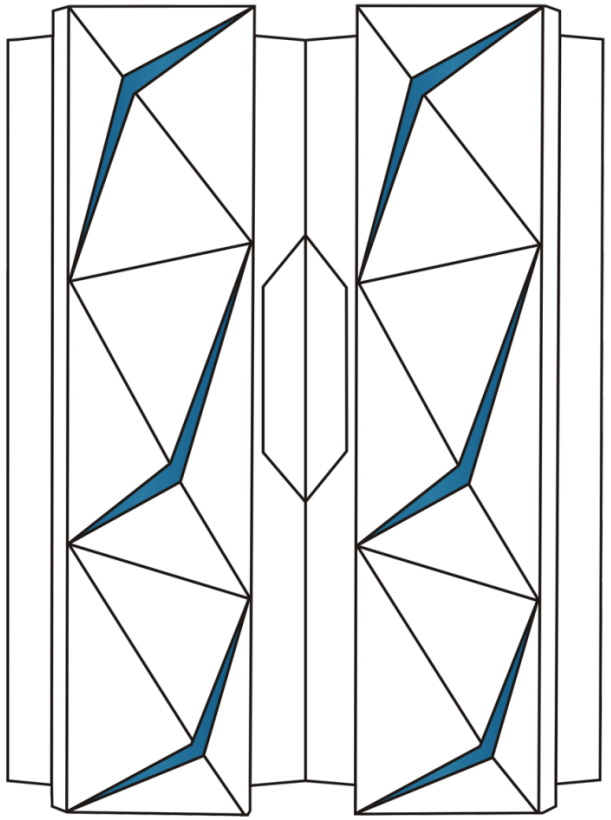

The IBM z14 Makes Pervasive Encryption Achievable

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

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

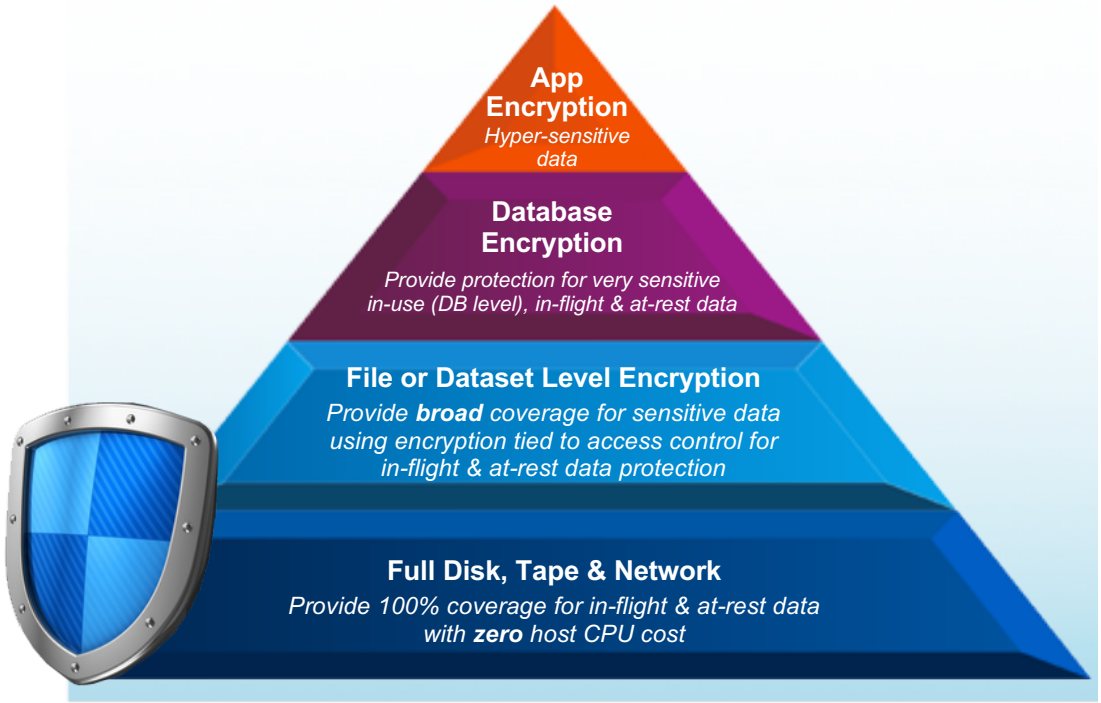

Hardware Crypto support in IBM Z

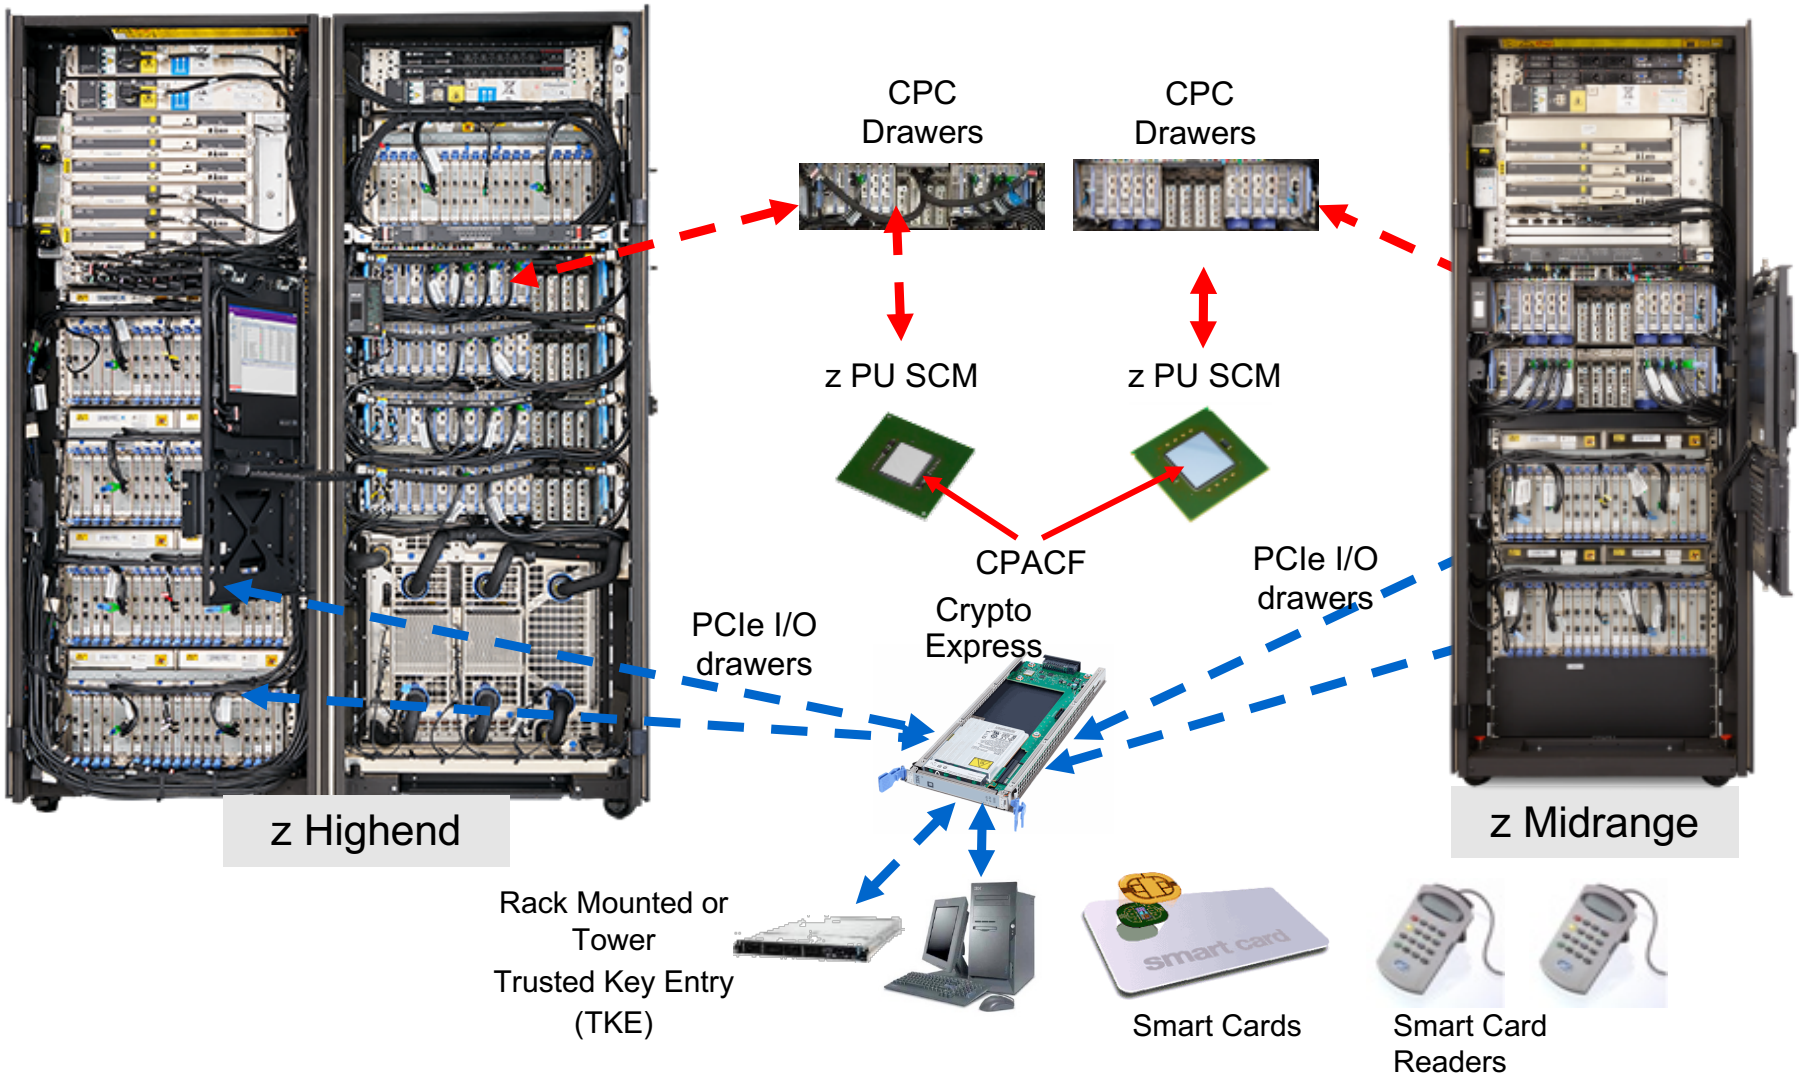

Central Processor Assist for Cryptographic Function (CPACF)

Making Pervasive Encryption Affordable

- Feature Code 3863, No Charge
- Hardware accelerated encryption on every core with the Central Processor Assist for Cryptographic Function (CPACF) which is designed to provide faster encryption and decryption than previous servers.
- CPACF – 2-6X faster encryption than z13 of data in-flight and at-rest.
- Key Management requires Crypto Express5S/6S.

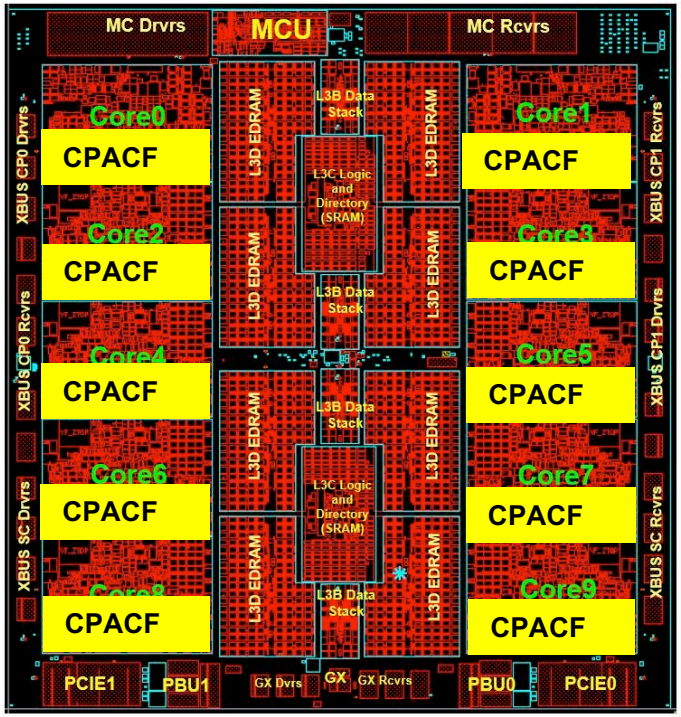

Cryptography Express6S

- One coprocessor per feature
  - Initial order – two features
- Up to 16 features per server
- Designed to be FIPS 140-2 Level 4 compliant
- Support for SHA-3
- Average, 1.5X to 2X performance increase over Crypto Express5S

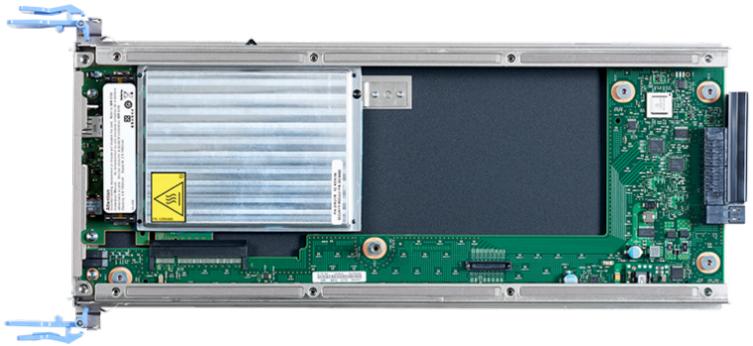

Feature Code 0893

Three Crypto Express6S configuration options

- Only one configuration option can be chosen at any given time
- Switching between configuration modes will erase all card secrets
  - Exception: Switching from CCA to accelerator or vice versa

Accelerator		CCA Coprocessor		EP11 Coprocessor	
TKE	N/A	TKE	OPTIONAL	TKE	REQUIRED
CPACF	NO	CPACF	REQUIRED	CPACF	REQUIRED
UDX	N/A	UDX	YES	UDX	NO
CDU	N/A	CDU	YES(SEG3)	CDU	NO
Clear Key RSA operations and SSL acceleration		*Secure Key crypto operations*		*Secure Key crypto operations*	

"MiniBoot" is the secure code used to initialize the crypto card at power on. Provides additional trust there is nothing subversive buried in the card. Code moved from a firmware approach to a more secure hardware based method.

Trusted Key Entry (TKE)

Description	FC	z14
TKE 8.1 LIC	0878	N/A
TKE 9.0 LIC	0879	Yes
Workstation	0096 Tower	Yes
Workstation	0085 Rack	Yes
Workstation	0841	No
Workstation	0842 Tower	Convert*
Workstation	0847 Tower	Convert*
Workstation	0097 Rack	Convert*
Workstation	0098 Tower	Convert*
Workstation	0080 Rack	Yes w/4768
Workstation	0081 Tower	Yes w/4768

Notes:
Convert*: This is accomplished by replacing and returning the Crypto adapter card (level 4765 or 4767) installed in the TKE workstation with the newer level 4768 TKE Crypto Adapter card. Level 4768 is FC 0844.

Table of Content

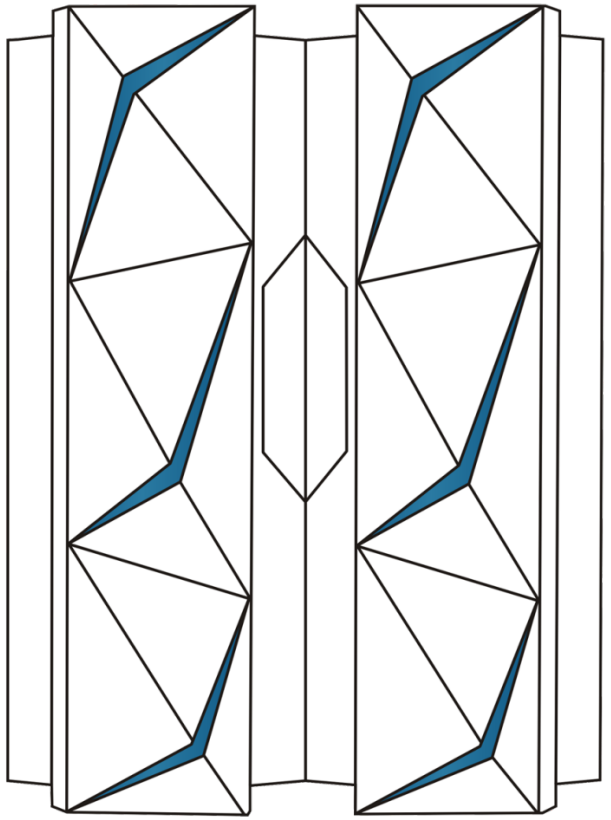

Available HMC configurations

FC 0083 (1U HMC)
- z14 only - Includes chassis/server/keyboard/display, maximum of 10
- Customers responsibility to order rack for mounting

NEW

FC 0082 (MiniTower)
- Can be ordered as new for z14, maximum of 10

NEW

FC 0092/ FC 0095 (MiniTower)
- Carry forward only* from z13, maximum of 10
- Will be checked for proper driver and memory

FC 0094 / FC 0096 (1U HMC)
- Carry Forward only* - Includes chassis/server/keyboard/display
- Will be checked for proper driver and memory maximum of 10 allowed

**** Carry Forward (Not applicable for LinuxONE)***

HMC Enhancements 2.14.0

- Workspace Direction/Enhancements – No more Classic UI
- Enhanced security for z14 HMC and SEs
- Java Applet Removal for z14
  - With HMC/SE 2.14.0 can now edit IOCDS source directly on HMC console
- Firmware management enhancements
- IBM Enhanced Support Facility – Changes to Call-Home infrastructure
- Crypto PCI-HSM Compliance Level
- Global OSA/SF
- LPAR Resource Assignment
- New characters allowed for load parameters
- FTP Through HMC
- Secure Console-to-Console Communications
- SNMP/BCPii API Enhancements
- Remote Browser IP Address Limiting
- Manage System Time (STP interface enhancements)
- Multi-factor Authentication (MFA for HMC authentication)

HMC 2.14.0 System support

- The new HMC Version 2.14.0 (GA1) will support the systems shown in the table.
- z9 EC/BC (Driver 67, SE version 2.9.2) systems are no longer supported.

Machine Family	Machine Type	Firmware Driver	SE Version	Ensemble Node Potential
z14 HMC	3906	32	2.14.0	Yes
z13s	2965	27	2.13.1	Yes
z13	2964	27	2.13.0 2.13.1	Yes
zBX Node	2458-004	22	2.13.0	Required
zBC12	2828	15	2.12.1	Yes
zEC12	2827	15	2.12.1	Yes
z114	2818	93	2.11.1	Yes
z196	2817	93	2.11.1	Yes
z10 BC	2098	79	2.10.2	No
z10 EC	2097	79	2.10.2	No

Workspace Enhancements - Masthead

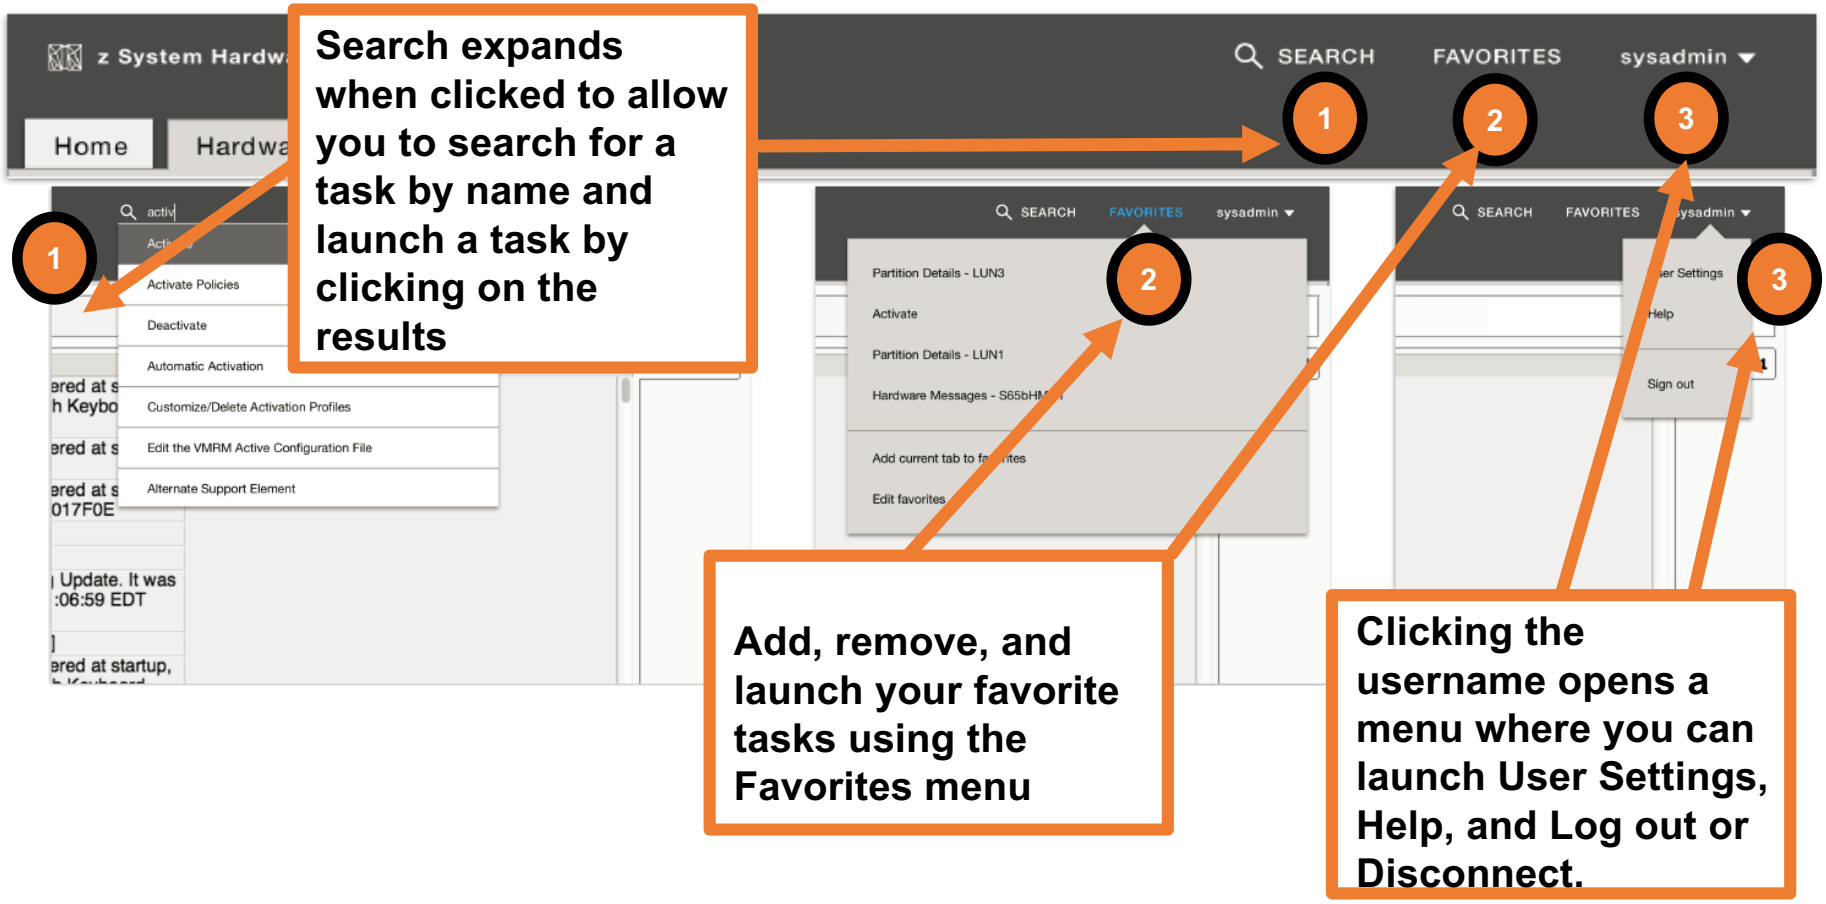

Workspace Enhancements - Tabbed Tasks

- Home tab pinned to the left, New tasks opens a new tab,
- Close tabs, Pop-out tabs, Pop-into tabs, Group related tabs

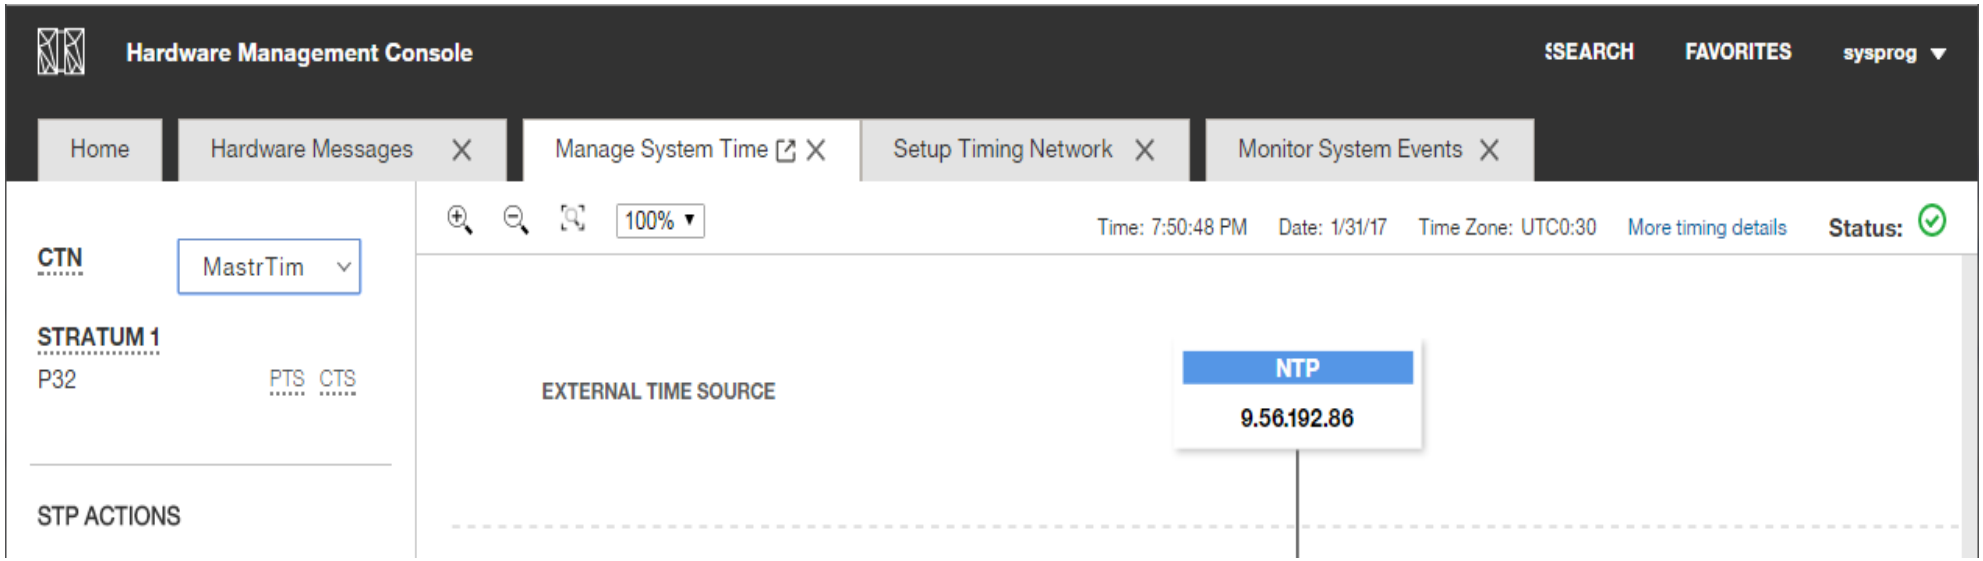

IBM HMC Mobile for z Systems and LinuxONE

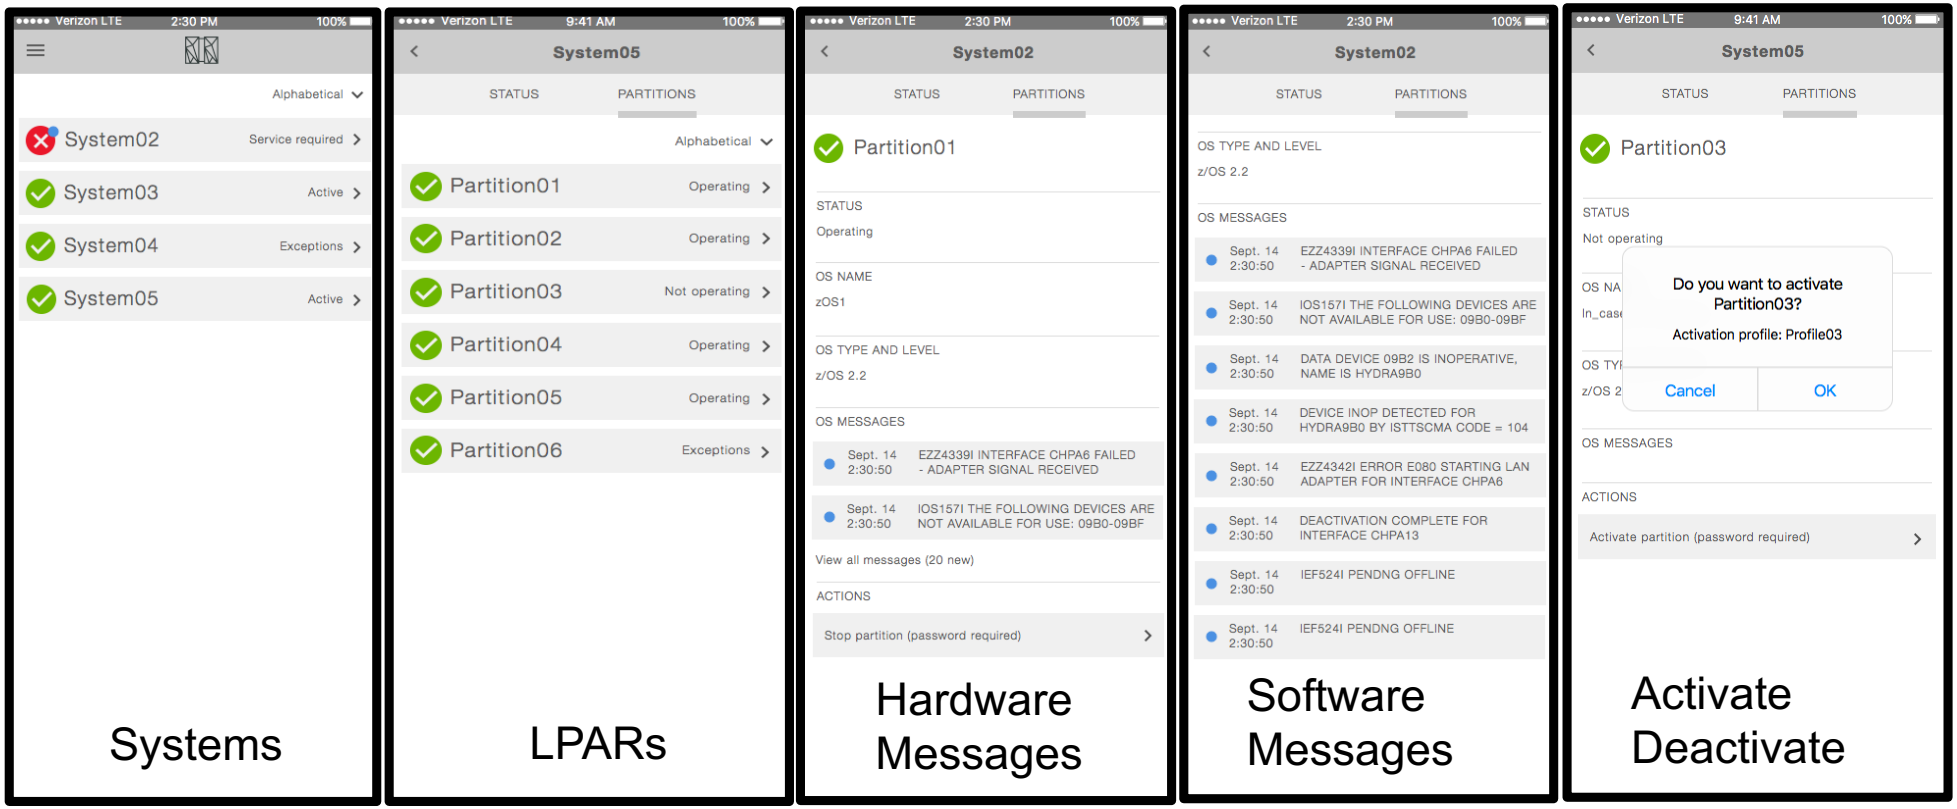

(Available December 2017)

Multi-factor Authentication (MFA) – SE, HMC, TKE – Two Factors

- 1st factor: HMC/SE/TKE userid and password
- 2nd factor: smartphone that provides a “multi-factor authentication code”
  - TOTP – Time-based One-Time Password
    - Defined by RFC 6238
    - A cryptographic function that combines a secret key with the current time to generate a one-time-use password
    - Can be used only once
    - Changes every 30 seconds
      - HMC/SE/TKE accepts current, previous and next TOTP
  - Compatible with existing freely-available smartphone apps
    - e.g., Google Authenticator and FreeOTP

IP addresses for call home

- **IBM Enhanced support facility**
  - IPV4
    - 129.42.56.189
    - 129.42.60.189
    - 129.42.54.189
  - IPV6
    - 2620:0:6c0:200:129:42:56:189
    - 2620:0:6c2:200:129:42:60:189
    - 2620:0:6c4:200:129:42:54:189

Open the firewall for these addresses.

z14 will not call home without the Enhanced Support.

- **Legacy IBM support facility**
  - IPV4
    - 129.42.26.224
    - 129.42.42.224
    - 129.42.50.224
  - IPV6
    - 2620:0:6c0:1::1000
    - 2620:0:6c2:1::1000
    - 2620:0:6c4:1::1000

These addresses are not used on the z14. Legacy machines can still use these addresses, but recommend using the Enhanced Support Facility.

Graphical Representation of an STP Network

- Simplify understanding of STP network and evaluating changes to network
- Easier problem isolation and determination (i.e. broken or inoperative links)
- Show STP timing network NTP connections and IP Addresses
- Validate intended configuration

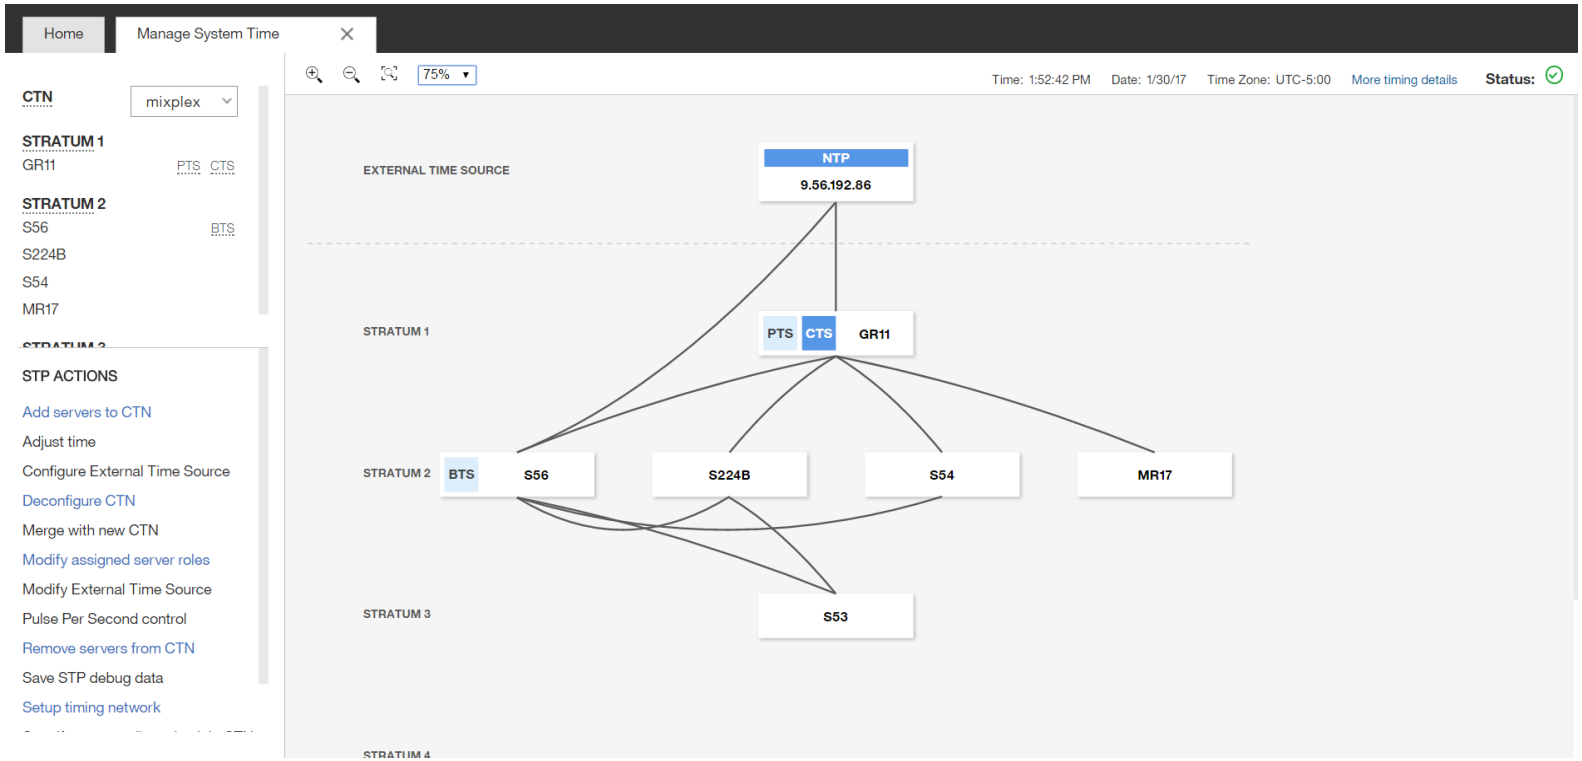

Additional STP Stratum Level

Provide additional flexibility/safety and cost savings when making configuration changes.

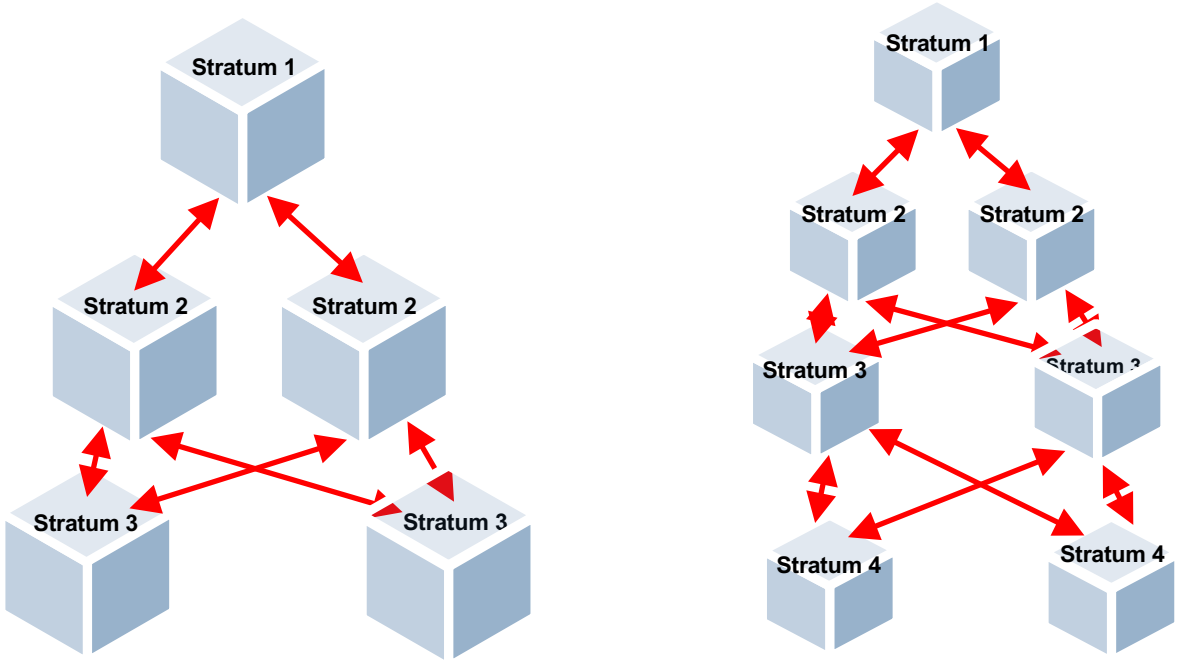

Temporary placement in S4.
Avoid Stratum 0 (unsynchronized) during migrations.

Table of Content

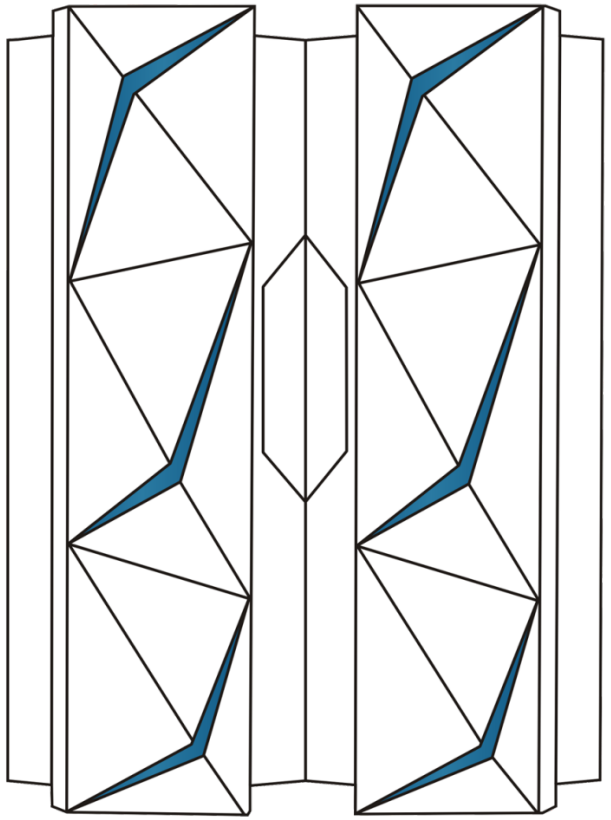

Operating Systems – zArchitecture and ESA/390 Architecture Modes

- z14 is the first IBM Z server to support only z/Architecture mode
  - This applies to all operating systems running native on PR/SM as well as second level guests

- IBM operating systems that run in ESA/390 mode are either no longer in service or only currently available with extended service contract, and they will not be usable on IBM Z beyond IBM z13.
  - All 24-bit and 31-bit problem state application programs originally written to run on the ESA/390 Architecture will be unaffected by this change

z/Architecture & Compilers

- If you don't keep your compilers up to date, you may not get the benefit of the architecture. Exploit the new instructions.

- In most cases, no source changes required.
  Just re-compile with new compiler.
  Compute intensive workloads can see up to a 5-10% performance improvement.
  Cobol V6, PL/1, C/C++

- Automatic Binary Optimizer V1.2 (ABO)
  - Requires no recompilation, or source migration
  - Increases application performance; Reduces CPU usage and operating costs
  - Targets latest IBM Z (e.g. zEC12, zBC12. z13, & z13s, z14)
  - Performance and Scalability Improvements

IBM z14 Operating Systems Support

z/OS

- z/OS 2.3 Sept. 29, 2017 GA
- z/OS 2.2 with PTFs
- z/OS 2.1 with PTFs
- z/OS 1.13 (compatibility only)
  - IBM Software Support Services purchase
  - September 2016, EoS

z/VM

- z/VM 6.4 with PTFs
- z/VM 6.3 with PTFs

Linux on IBM Z

- Minimum Distributions
- SLES 12 SP2
- SLES 11 SP4
- RHEL 7.3
- RHEL 6.8
- Ubuntu 16.04

IBM cannot legally discuss z14 exploitation prior to GA from distributors.

Officially Tested:
www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

z/VSE

- z/VSE 6.2 Preview 4/11/17
- z/VSE 6.1 with PTFs
- z/VSE 5.2 with PTFs (October 31, 2018 = EoS)
- z/VSE 5.1
  - June 30, 2016 = EoS, limited toleration
- Earlier releases can not IPL

z/TPF

- z/TPF 1.1 with PTFs (Compatibility, including Crypto Express5S support)
- HyperDispatch Support

Summary – 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.
P PTF is required, use SMP/E FIXCAT for identification
Y Support is in the base
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
AO Requires the ICSF web deliverable for FMID HCR77A0 minimally, with PTF.
W A web deliverable is required, available at http://www-03.ibm.com/systems/z/os/zos/downloads/

z/OS Support Summary

Release	z9 EC z9 BC WdfM OoS	z10 EC z10 BC WdfM	z196 z114 WdfM	zEC12 zBC12	z13 z13s	z14	End of Service	Extended Defect Support[1]
z/OS V1.13	X	X	X	X	X	X[1]	9/16	9/19*
z/OS V2.1	X	X	X	X	X	X	9/18*	9/21*
z/OS V2.2		X	X	X	X	X	9/20*	9/23*
z/OS V2.3[2]				X	X	X	9/22*	9/25

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

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

WdfM – Server has been withdrawn from Marketing
OoS -- Out of support

Legend

IBM Software Support Services required for z/OS support
Generally supported

z/VSE Support for IBM z14

- 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)
    - CHPID-type OSC, OSE, OSD and OSX
    - Layer 2 and layer 3 support
  - Crypto Express6S toleration
    - Support of greater than 16 domains
    - CCA Coprocessor and Accelerator mode, EP11 not supported
    - Exploited by OpenSSL

z/VSE Support for IBM z14 – Migration Considerations

- Customers should collect reference information before migration
  - Performance data, CPU utilization of reference workload, 1/0 activity, elapsed times
    - Required to size z14 and compare workload characteristics after migration
  - See "z/VSE Release and Hardware Upgrade" document for more details ftp://public.dhe.ibm.com/eserver/zseries/zos/vse/pdf3/zVSE_Release_and_Hardware_Upgrade.pdf
- Apply PTFs for z14
  - PSP (Preventive Service Planning) bucket for z14, z/VSE and z/VM PSPs list PTFs https://www-304.ibm.com/webapp/set2/psearch/tips?domain=psp
  - Apply required PTFs for IBM and ISV software
    - z/VSE 5.1 (end of service since June 2016):
      - APAR DY47654 (PTF UD54170) required to IPL on z14 http://www.ibm.com/support/docview.wss?uid=isg1DY47654
      - Crypto Express6S not supported
    - z/VSE 5.2, z/VSE 6.1: PTFs are required for Crypto Express6S (PTFs not yet available). Additional PTFs may be required (see PSP)
- Compatibility Issues / Requirements
  - Only z/VSE 5.1 (unsupported), z/VSE 5.2 and z/VSE 6.1 can IPL on z14
  - **z14 IPLs in z/Architecture mode only**

z/VM Support for IBM z14 - Compatibility Support

- Compatibility → z/VM support for host / guests on z14 at the z13 functional level with limited exploitation of new functions (some transparent)
  - Support available as PTFs concurrently at announce
  - Includes PTFs for EREP, IOCP, HCD, and HLASM
  - z/VM 6.3 and 6.4 will not IPL on z14 without compatibility support
  - z14 EC will be the last server supported by z/VM 6.3 (z14 MR will not be supported)
- Compatibility Support for several new architectural facilities (see next chart)
- z/VM 6.3, and 6.4 will be supported on z14 EC
  - Supported releases:
    - z/VM 6.3 GA 7/23/2013, EoS 12/31/2017, requires a z10 ALS
    - z/VM 6.4 GA 11/11/2016, No announced EoS, requires a z196/z114 ALS
  - z/VM 5.4 will not support, nor IPL on z14, EoS 12/31/2017 and no longer tied to z9
  - z/VM 6.2 will not support, nor IPL on z14, EoS 06/30/2017
- Guest Exploitation support for the new CEX6S crypto adapter
- Stack and Netstat support for the new OSA-Express6S adapter

z/VM Support for IBM z14 - Compatibility Support

- z/VM 6.3 and 6.4 support:
  - Crypto – CEX6S, CPACF enhancements (MSA6/SHA3, MSA7/True RNG, MSA8/GCM)
    - Crypto Express6S dedicated to or virtualized for guest exploitation; Both interfaces will include clear key ECC support; dedicated also allows secure key
  - SIMD and compression enhancements
  - Miscellaneous-Instruction-Extensions Facility 2
  - Vector Enhancements Facility 1 and Vector Packed Decimal Facility
  - TLB Purge, HPMA2 IBR Revalidation (6.4 only), CMM2
  - TCP/IP Stack and NETSTAT OSAINFO support for OSA-Express6S
  - RoCE Express2 guest attachment (guest exploitation)
  - Dynamic I/O for Sync I/O (Hyper I/O), RoCE Express2, OSA-Express6S, FICON Express 16S/16S+, Regional Crypto, Coupling Express LR
  - ESA/390 compatibility Mode support
  - CPUMF Counter enhancements “Diag x’204’ Busy indication facility”
- ESA/390 removal roll-back from z/VM 6.4 to 6.3; The CPLOAD nucleus, Stand-Alone Program Loader (SAPL), DASD Dump Restore (DDR), Stand Alone Dump, ICKDSF run entirely in z/Architecture mode.
- No z/VM 6.3 support for defining Z mode virtual machines.

z/VM Pre-requisite service

- The PTF for APAR VM65207 must be applied to z/VM V6.2 or z/VM V6.3 when z/VM V6.4 is a guest of either release and will host a MACHINE Z guest.
- If you intend to exploit HyperPAV paging from z/VM V6.4 running as a guest of z/VM V6.3, you need to apply the PTF for APAR VM65748 to the z/VM V6.3 first level system.
- The PTF for APAR VM65976 provides infrastructure support in z/VM V6.3 and V6.4 for the z14 server, satisfying the statement of direction for the removal of ESA/390 architecture support, and must be installed on all the members of an SSI cluster before any member will be running on the new server.
- The PTF for APAR VM65867 is required on all z/VM V6.2 and V6.3 systems in an SSI cluster that includes a z/VM V6.4 system.
- z/VM 6.3 support for ESA/390 removal requires:
  - APARs VM65856 (CP), VM65921 (CP SAD), VM65922 (CMS SAD)

- **z/VM 6.3 cannot be installed on a z14 EC**

- **Install of z/VM 6.4 on a z14 EC requires updated install media, availability August 25, 2017**

SSI Cluster and z14 Machine

APAR VM65976 is needed on **ALL** members of an SSI cluster prior to bringing up **ANY** member on a z14 machine.

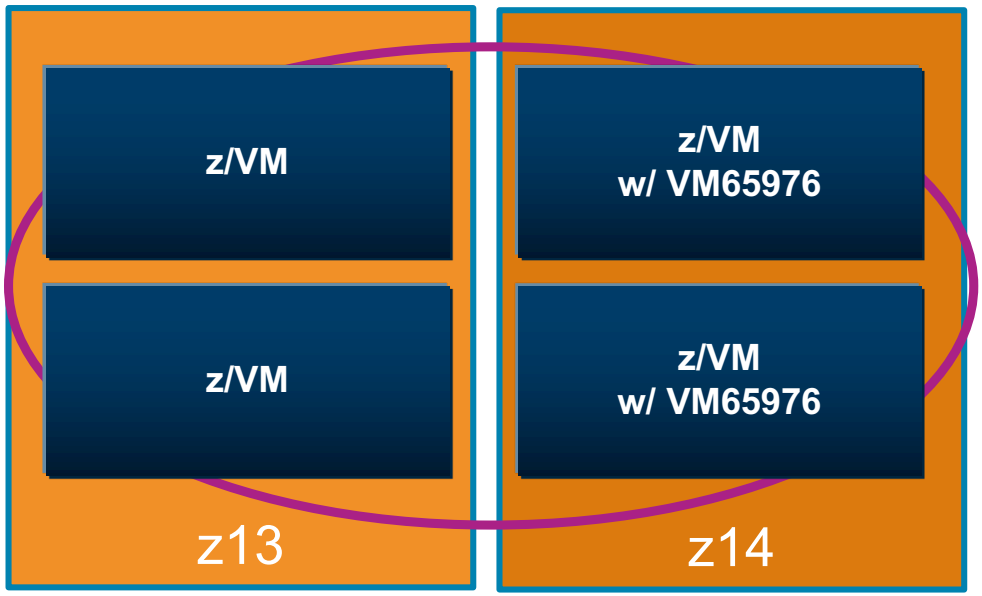

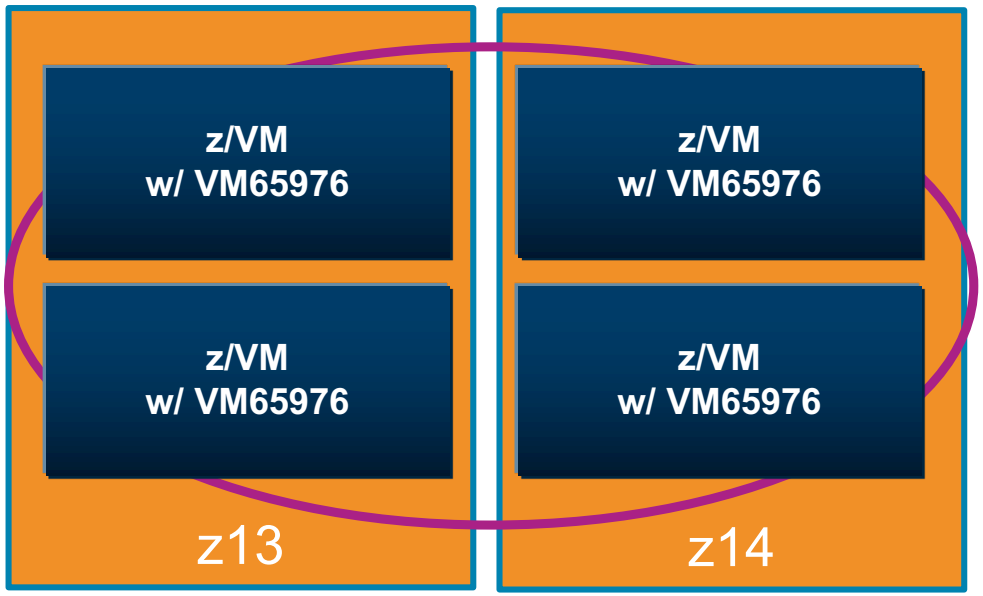

z14 GA1 items z/VM will not support, including Guest Exploitation

- z/VM does not currently have plans to support the following at z14 GA1
  - Time Slice Instrumentation
  - TOD Multiple Epochs
  - zHyperLink I/O (Dynamic I/O only)
  - Virtual Flash Memory (VFM) (Flash Express is not supported by z/VM)
  - SCSI-IPL Normal Support; no perceived guest exploitation benefit

- Still no z/VM support for the following:
  - Runtime Instrumentation (RI) support
  - zHPF Extended distance II
  - Enhanced DAT Facility 2
  - CHSC & I/O: Guest Support for Multiple-Subchannel-Sets
  - Flash Express Adapter

z/VM Compatibility Support Schedule

- z14 compatibility APARs (z/VM 6.3 and 6.4 at GA):
  - VM65942: GA1 processor compatibility and CEX6S crypto Adapter support
  - VM65942: GA1 Memory Management support
  - VM65942: GA1 I/O compatibility (dynamic I/O) support
  - VM65865: Coupling Express LR dynamic I/O support
  - PI73016: TCP/IP Stack and NETSTAT support for OSA-Express6S
  - PI46151: ICKDSF support (6.3 only, part of 6.4 installation media)
  - VM65952: EREP support
  - PI62275 and PI65715: HLASM support
  - VMHCD: VM65843 GA1; VM66009 LinuxONE; VM65849 Coupling Express
  - VMHCM: VM65266 GA1; VM65208 Coupling Express LR
  - IOCP: VM65939 GA1; VM65880 Coupling Express LR

- Updates will be available on the z/VM web page after announce: http:/www.ibm.com/vm

- Review hardware PSP bucket 3906DEVICE z/VM subset

z/VM Support for IBM z14 – Exploitation Support

- z/VM 6.4 post z14 GA1 (late 4Q2017) support:
  - APAR VM65986: Guest exploitation support for the Instruction Execution Protection Facility; z14 functionality to improve the security of programs running on IBM Z by allowing virtual memory elements to be identified as containing only data. If an attempt is made to fetch an instruction from an address in such an element or if an address in such an element is the target of an execute-type instruction a Protection Exception will occur.
  - APAR VM65987: Guest exploitation support for Pause-Less Garbage Collection; support of z14 functionality designed to improve the performance of garbage-collection processing by various languages, in particular Java.
  - APAR VM65993: Encrypted paging, in support of pervasive encryption; Ciphering will occur as data moves to and from active memory onto a paging volume owned by CP
    - Dynamic encryption support will be available as part of the SPE

Not related to z14 Exploitation:

- APAR VM65918: Multi-VSwitch Link Aggregation Enhancements focusing on improving the area of load balancing to leverage both horizontal and vertical growth in both single and cross virtual switch networking configurations. (December).

KVM New Directions

- IBM Linux on z distribution partners will announce and provide KVM.
  - Announcement March 7, 2017
  - Last Date to Order August 28, 2017
- IBM is withdrawing KVM for IBM Z
  - V1.1.2 will be the last release delivered by IBM
  - V1.1.2 Support Lifecycle through March 31, 2018
  - No new functional additions.
- Todays options from Linux distributors
  - SuSE Linux Enterprise Server (SLES) 12 SP2 and later
  - Canonical Ubuntu 16.04 LTS and later
  - Migration considerations – KVM for IBM Z users have been contacted by IBM
- IBM continues to load z related open source code that the Linux distributors can add to their products and exploit.

Table of Content

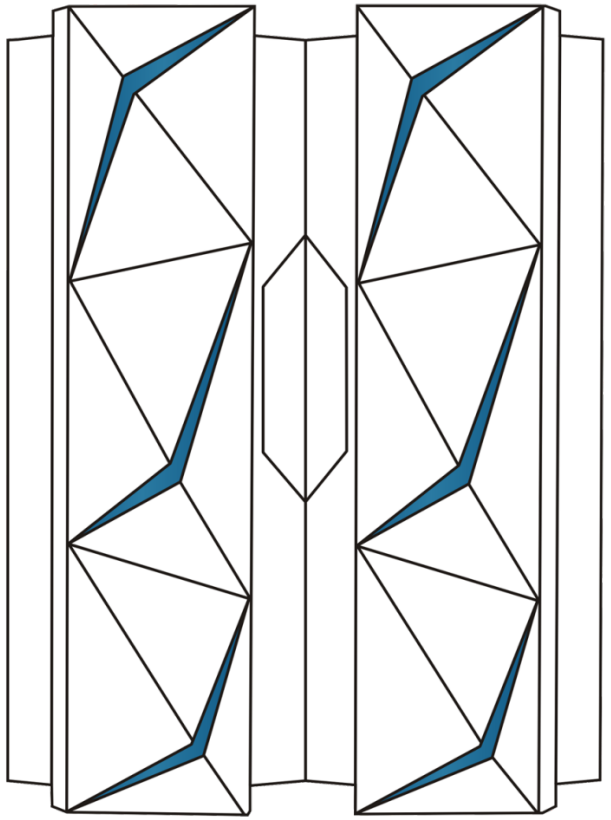

IBM z14 at a glance

System, Processor, Memory
Five hardware models: M01, M02, M03, M04, M05
10 core 5.2GHz 14nm PU SCM
1 - 170 PUs configurable as CPs, zIIPs, IFLs, ICFs, up to 196 Pus
Increased Uniprocessor capacity
Up to 33 sub capacity CPs at capacity settings 4, 5, or 6
CPC Drawers and backplane Oscillator
Enhanced SMT and new instructions for SIMD
Enhanced processor/cache design with 1.5x more on-chip cache sizes
Up to 32 TB DRAM, protected by Redundant Array of Independent Memory (RAIM)
Virtual Flash Memory (VFM)
192 GB HSA
Improved pipeline design and cache management

Announce: July 17, 2017

I/O Subsystem, Parallel Sysplex, STP, Security
PCIe Gen3 I/O fanouts with 16 GBps Buses
6 CSS, 4 Subchannel sets per CSS
0 – 5 PCIe I/O Drawer Gen3 (no I/O Drawer)
Next generation FICON Express16S+
10 GbE RoCE Express2
Integrated Coupling Adapter (ICA SR) and Coupling express LR for coupling links
Support for up to 256 coupling CHPIDs per CPC
CFCC Level 22
Crypto Express6S and CMPSC compression and Huffman Coding compression
STP configuration and usability enhancements (GUI)
IBM zHyperLink Express
OSA-Express6S
Secure Service Container

RAS, simplification and others	
L3 Cache Symbol ECC	**Acoustic and thin covers (space saving)**
N+1 radiator design for Air Cooled System	**Drop "Classic" HMC UI**
ASHRAE Class A3 design	**Enhanced SE and HMC Hardware (security)**
Support for ASHRAE Class A3 datacenter	TKE 9.0 LICC
Largesum TCP/IP hardware Checksum (OSA-Express6S)	Pause-less garbage collection
Universal Spare SCM s (CP and SC)	Simplified and enhanced functionality for STP configuration
Enhanced Dynamic Memory Relocation for EDA and CDR	Virtual Flash Memory (replaces IBM zFlash Express)

PR/SM
Up to 170 CPUs per partition
IBM Dynamic Partition Manager updates
Up to 85 LPARs
16 TB Memory per partition

Table of Content

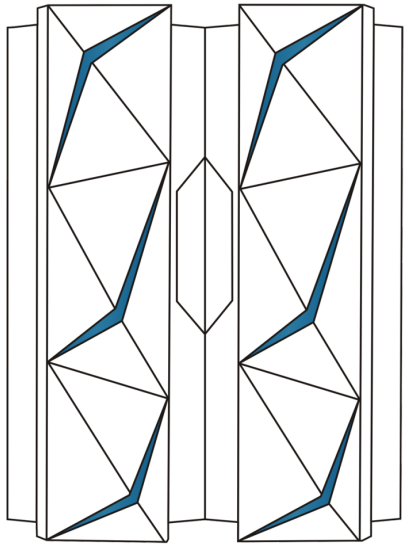

Statements of General Direction

- **Stabilization of z/VM V6.3 support**: IBM z14 is planned to be the last z Systems and IBM Z high-end server and z13s is planned to be the last mid-range z Systems server supported by z/VM V6.3 and the last z Systems and IBM Z servers that will be supported when z/VM V6.3 is running as a guest (second level). z/VM V6.3 will continue to be supported until December 31, 2017, as announced in announcement letter # 915-025.
- **Future z/VM release guest support**: z/VM V6.4 will be the last z/VM release supported as a guest of z/VM V6.2 or older releases.
- **Disk-only support for z/VM dumps**: z/VM V6.4 will be the last z/VM release to support tape as a media option for stand-alone, hard abend, and snap dumps. Subsequent releases will support dumps to ECKD DASD or FCP SCSI disks only.
- **IBM intends to deliver IMS exploitation of IBM z14 and DS8880 zHyperLink™ WRITE operations.** zHyperLink Express is a direct connect short distance IBM Z I/O adapter designed to work in conjunction with a FICON or High Performance FICON SAN infrastructure.
- **IBM z14 will be the last z Systems and IBM Z high-end server to support FICON Express8S:** z14 will be last z Systems and IBM Z high-end server to support FICON Express8S (#0409 and #0410) channels. Enterprises should begin migrating from FICON Express8S channels to FICON Express16S+ channels (#0427 and #0428). FICON Express8S will not be supported on future high-end IBM Z servers as carry forward on an upgrade.

Statements of General Direction (Cont.)

- **IBM z14 will be the last z Systems and IBM Z server to support HCA3-O and HCA3-O LR adapters:** z14 will be last z Systems and IBM Z server to support HCA3-O fanout for 12x IFB (#0171) and HCA3-O LR fanout for 1x IFB (#0170). As announced previously, z13s is the last mid-range z Systems server to support these adapters. Enterprises should begin migrating from HCA3-O and HCA3-O LR adapters to ICA SR and/or Coupling Express Long Reach (CE LR) adapters on z14, z13, and z13s. For high-speed short-range coupling connectivity, enterprises should migrate to the Integrated Coupling Adapter (ICA-SR). For long-range coupling connectivity, enterprises should migrate to the new Coupling Express LR coupling adapter. For long-range coupling connectivity requiring a DWDM, enterprises will need to determine their desired DWDM vendor's plan to qualify the planned replacement long-range coupling link.
- **OSA-Express6S 1000BASE-T adapters:** OSA-Express6S 1000BASE-T adapters (#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.
- **Dynamic Partition Manager support of ECKD:** IBM intends to deliver support for adding and configuring ECKD FICON disks to partitions created in Dynamic Partition Manager (DPM) mode for Linux running in LPAR, under KVM on z, and under z/VM 6.4.

Statements of General Direction (Cont.)

- **New DB2 Analytics Accelerator deployment option on the IBM Z infrastructure**: The DB2 Analytics Accelerator for z/OS extends IBM Z and DB2 for z/OS to form a hybrid environment, which is capable of running both transactional and analytical SQL query workload. This extension is currently delivered via two deployment options: either as an appliance, based on IBM PureData System for Analytics, or as a hosted cloud environment. In addition to these two form factors, IBM intends to deliver a new DB2 Analytics Accelerator deployment option on the IBM Z infrastructure. This would further extend deployment options available to DB2 Analytics Accelerator clients. This new deployment option would allow for deeper integration with the IBM Z infrastructure. Clients would benefit with the flexibility to deploy the form factor that best suits their requirements to enable unified homogeneity of service, support and operations and deeper integration with their processes, e.g. for their disaster recovery.
- **IBM intends to deliver VSAM exploitation of z14 and DS8880 zHyperLink Express**. zHyperLink Express is a short distance mainframe attach link designed for up to 10x lower latency than High Performance FICON.

Table of Content

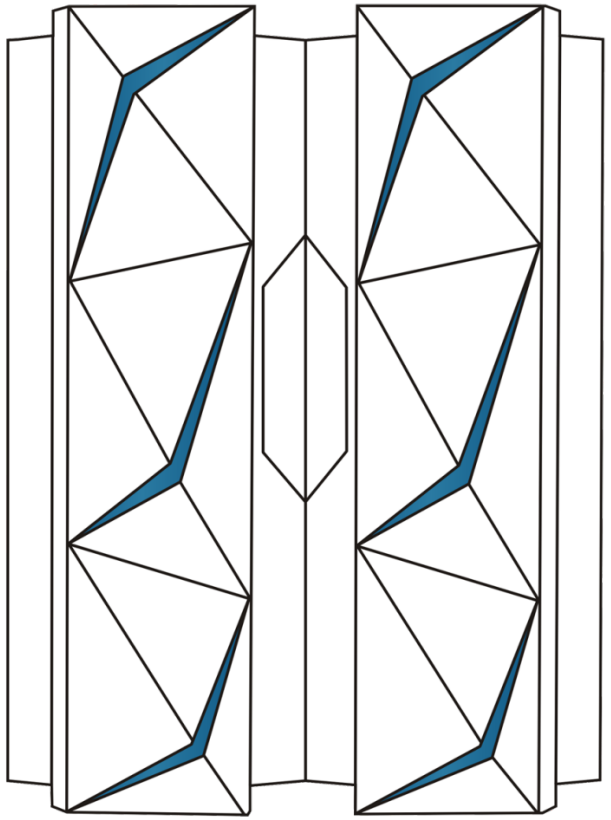

IBM Dynamic Partition Manager (DPM)

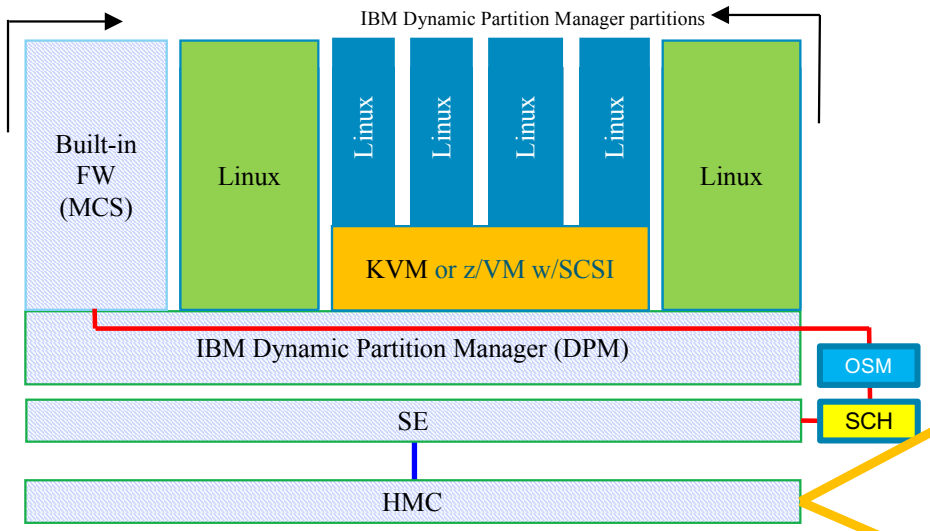

Provides simplified, consumable, enhanced IBM Z experience for reducing the barriers of adoption for new and existing clients.

New Auto-Config
New IPV6 Support

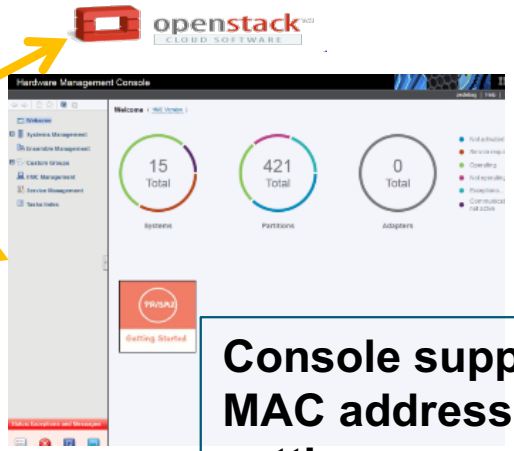

Deliver dynamic IBM Z hardware and virtual infrastructure management including integrated dynamic I/O management with general virtualization knowledge while requiring minimal mainframe knowledge.

Supports IBMs Cloud Strategy & IAAS.

Console support.
MAC address setting.
VLAN enforcement.

IBM z14 Functional Comparison to IBM z13

Performance and Scale	▪ Uniprocessor Performance	▪ New up to 10% performance improvement over IBM z13 (z13)[1]
	▪ System Capacity	▪ New up to 35% system total z/OS capacity performance improvement over z13[1]
	▪ SMT	▪ New 2nd generation SMT delivers virtualization benefits for Linux and up to 25% performance improvement for Linux on z Systems® and zIIP workloads vs non-SMT on z14
	▪ SIMD	▪ New instructions for perform boost to traditional workloads and new analytics workloads versus SIMD on z13
	▪ Cache	▪ New z14 has 1.5x more on-chip cache per core versus z13
	▪ Models	▪ Five models with up to four CPC drawers (z13 also has five CPC models and four drawers)
	▪ Processing cores	▪ New up to 170 cores to configure, up to 141 on z13
	▪ Granular Capacity	▪ New up to 269 capacity settings versus 231 on the z13
	▪ Memory	▪ New up to 32 TB RAIM memory versus 10 TB RAIM memory on z13
	▪ Compression	▪ New improvement in CMPSC compression and Huffman Coding compression ratio using zEDC Express versus z13
Virtualization	▪ LPAR virtualization	▪ 85 partitions – same as z13
	▪ RoCE adapter	▪ New 10 GbE RoCE Express2 with additional virtual functions (VFs) per physical port (10GbE Exoress on z13)
	▪ Simplified LPAR management	▪ Enhanced IBM Dynamic Partition Manager allows for config and mgmt of system resources– new support for z/VM and for ECKD disk
Infrastructure Efficiency	▪ Networking	▪ New OSA-Express6S with improvements over z13 using OSA-Express5S
	▪ HiperSockets and SMC-D	▪ Up to 32 HiperSockets (same as z13) **and** memory-to-memory communications with SMC-D offers within-the-box communications for z/OS
	▪ FICON	▪ New FICON Express16S+ will provide an increase in I/O rates over FICON Expresss16S
	▪ zHPF	▪ 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
	▪ IBM zHyperLink	▪ New IBM zHyperLink - New short distance z14 channel that can be installed on IBM DS8880 System Storage® for lower latency not on z13
	▪ Forward Error Correction	▪ Industry standard FEC for optical connections for substantially reduced I/O link errors same as z13
	▪ FICON dynamic routing	▪ Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches same on z13
	▪ LCSS/Subchannel sets	▪ Up to six LCSS and 4 Subchannel sets – same as z13
	▪ WWPN	▪ I/O serial number migration allows keeping same serial number on replacement server same as z13
	▪ HMC	▪ 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)
	▪ Pause-less garbage collection	▪ New enterprise scale JAVA applications to run without periodic pause for garbage collection on larger & larger heaps
	▪ Upgradeability	▪ Upgradeable from z13 and zEnterprise EC12 (zEC12)
	▪ IBM Virtual Flash Memory	▪ New memory replacement for Flash Express helping improve availability – available only on z14
Resiliency and Availability	▪ Coupling – HCA-3	▪ Coupling with HCA-3 InfiniBand Coupling Links – long and short distance – same as z13
	▪ Coupling – ICA SR	▪ Short distance coupling with PCIe-based links (ICA SR) – same as z13
	▪ Coupling Express LR	▪ New Coupling Express LR – Coupling Express LR will be available on z13
	▪ STP	▪ New Simplified STP management with HMC enhancements not available on z13
	▪ Sparing	▪ Enhanced integrated sparing on z14 and z13 reducing the number of on site service and maintenance events
	▪ Rack Mounted Accessories	▪ Rack-mounted HMC and TKE options to save space in the data center
	▪ Environmentals	▪ Optional non raised floor, overhead cabling, water cooling and DC power plus New ASHRAE A3 ration – (not ASHRAE rating on z13)
Security	▪ Cryptographic Coprocessor	▪ CPACF for improved performance and new true Random Number Generator versus z13
	▪ Crypto Express	▪ Nes Crypto Express6S with performance increase plus new algorithms for elliptic curve, SHA, VISA FPE versus z13 Crypto Express5S
	▪ IBM Secure Service Container	▪ Secure deployment of software virtual appliances
	▪ Secure Console Access	▪ Protection of sensitive data by using Transport Layer Security (TLS) support in the Open Systems Adapter-Integrated Console Controller (OSA-ICC)

https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocumen

IBM z14 Functional Comparison to IBM zEC12

Performance and Scale	▪ Uniprocessor Performance	▪ New up to **22%** performance improvement over IBM zEnterprise EC12[1]
	▪ System Capacity	▪ New up to **89%** system total z/OS capacity performance improvement over zEC12[1]
	▪ SMT	▪ 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
	▪ SIMD	▪ Vector processing (SIMD) model provides construction of richer, complex analytics models, increased programmer productivity, faster analytics to traditional workloads versus no SIMD on zEC12
	▪ Cache	▪ New z14 has up to 88% more on-chip cache per core versus zEC12
	▪ Models	▪ Five models with up to four CPC drawers (zEC12 also has five CPC models and four books)
	▪ Processing cores	▪ New up to 170 cores to configure, up to 101 on zEC12
	▪ Granular Capacity	▪ New up to 269 capacity settings versus 161 on the zEC12
	▪ Memory	▪ New up to 32 TB RAIM memory versus 3 TB RAIM memory on zEC12
	▪ Compression	▪ New improvement in CMPSC compression and Huffman Coding compression ratio using zEDC Express versus zEC12
Virtualization	▪ LPAR virtualization	▪ 85 partitions versus 60 on zEC12
	▪ RoCE adapter	▪ New 10 GbE RoCE Express2 with additional virtual functions (VFs) per physical port (versus dedicated 10GbE Express on zEC12)
	▪ Simplified LPAR management	▪ 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
Infrastructure Efficiency	▪ Networking	▪ New OSA-Express6S with improvements over zEC12 using OSA-Express5S
	▪ HiperSockets and SMC-D	▪ 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)
	▪ FICON	▪ New FICON Express16S+ and increased FICON subchannels to 32K versus FICON Express 8S and 24K on zEC12
	▪ zHPF	▪ 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
	▪ IBM zHyperLink	▪ New IBM zHyperLink - New short distance z14 channel that can be installed on IBM DS8880 System Storage® for lower latency not on zEC12
	▪ Forward Error Correction	▪ Industry standard FEC for optical connections for substantially reduced I/O link errors not available on zEC12
	▪ FICON dynamic routing	▪ Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches not on zEC12
	▪ LCSS/Subchannel sets	▪ Up to six LCSS and 4 Subchannel sets – versus four LCSS and 3 subchannel sets on zEC12
	▪ WWPN	▪ I/O serial number migration allows keeping same serial number on replacement server not on zEC12
	▪ HMC	▪ 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)
	▪ Pause-less garbage collection	▪ New enterprise scale JAVA applications to run without periodic pause for garbage collection on larger & larger heaps not on zEC12
	▪ IBM Virtual Flash Memory	▪ New memory replacement for Flash Express helping improve availability – available only on z14
Resiliency and Availability	▪ Coupling – HCA-3	▪ Coupling with HCA-3 InfiniBand Coupling Links – long and short distance – same as zEC12
	▪ Coupling – ICA SR	▪ Short distance coupling with PCIe-based links (ICA SR) – not available on zEC12
	▪ Coupling Express LR	▪ New Coupling Express LR – Coupling Express LR – not available on zEC12
	▪ STP	▪ New Simplified STP management with HMC enhancements not available on zEC12
	▪ Sparing	▪ Enhanced integrated sparing on z14 reducing the number of on site service and maintenance events not on zEC12
	▪ Rack Mounted Accessories	▪ Rack-mounted HMC and TKE options to save space in the data center not available on zEC12
	▪ Environmentals	▪ Optional non raised floor, overhead cabling, water cooling and DC power plus New ASHRAE A3 ration – (not ASHRAE rating of zEC12)
Security	▪ Cryptographic Coprocessor	▪ CPACF for improved performance and new true Random Number Generator not on zEC12
	▪ Crypto Express	▪ Nes Crypto Express6S with performance increase plus new algorithms for elliptic curve, SHA, VISA FPE versus zEC12 Crypto Express4S
	▪ IBM Secure Service Container	▪ Secure deployment of software virtual appliances not on zEC12
	▪ Secure Console Access	▪ Protection of sensitive data by using Transport Layer Security (TLS) support in the Open Systems Adapter-Integrated Console Controller (OSA-ICC) not on zEC12

https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocumen