Dr. Manfred Gnirss

IBM Client Center, Böblingen

G02: Neues von IBM z Systems - z13s und Rockhopper

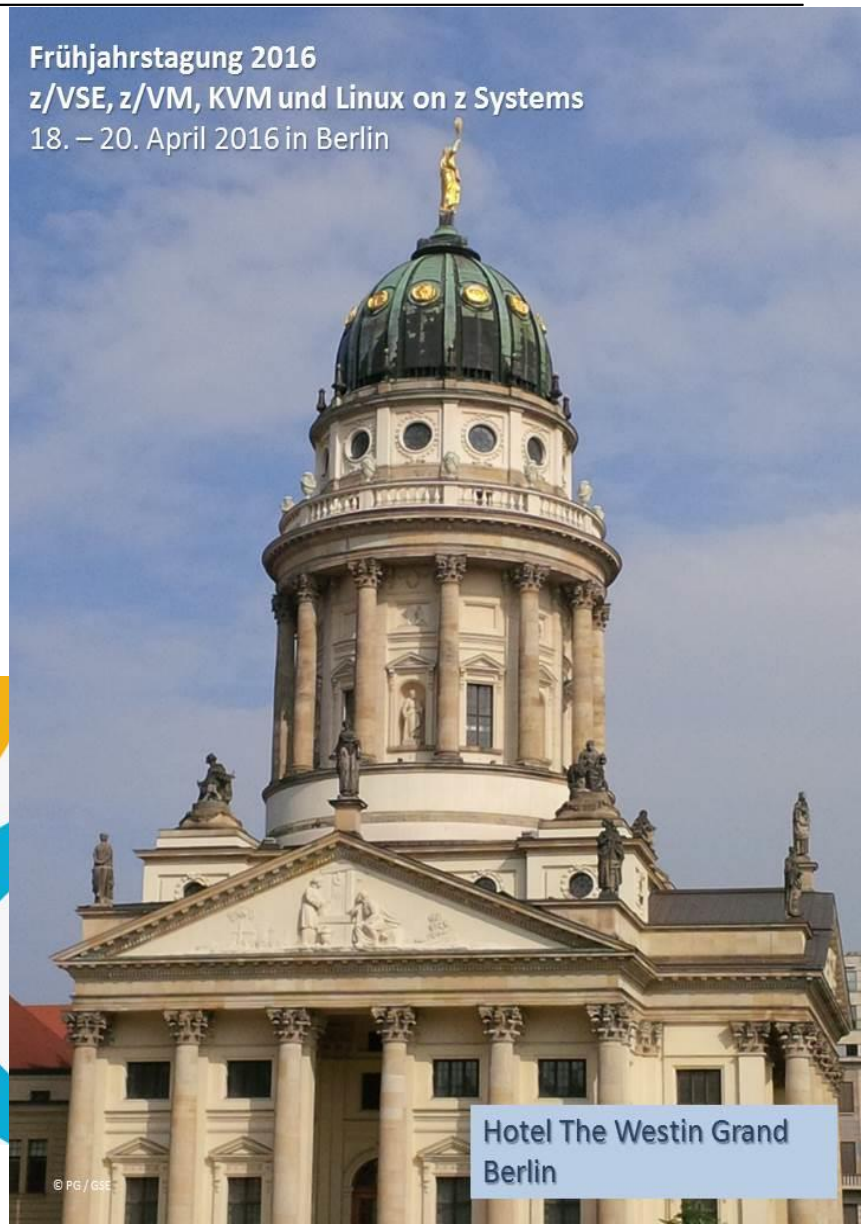

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.
Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a more complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*BladeCenter®, CICS®, DataPower®, DB2®, e business(logo)®, ESCON, eServer, FICON®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, PureFlex, PureSystems, S/390®, ServerProven®, Sysplex Timer®, System p®, System p5, System x®, z Systems®, System z9®, System z10®, WebSphere®, X-Architecture®, z13™, z Systems™, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zEnterprise®, zSeries®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

*** All other products may be trademarks or registered trademarks of their respective companies.**

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured 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.

Notice Regarding Specialty Engines

Any information contained in this document regarding Specialty Engines (“SEs”) and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs).
IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”).
No other workload processing is authorized for execution on an SE.
IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

IBM z Systems Evolution

Jan 12, 2015 – z13
"A new generation of z Systems built for digital business and the new mobile app economy"

Feb 16, 2016 – z13s
"IBM Combats Cyber Threats with Mainframe built for Hybrid Cloud"

Digital Disruption
Hyper transactions
Hybrid Cloud
In-Transaction Analytics
Trusted Security

Growth in Hybrid Cloud and Mobile
Rising Security Threats
Security from infrastructure to endpoint
Enhanced, optimized and integrated capabilities
IBM Ready for Security Intelligence Program

Our Servers

IBM z13

IBM z13s

IBM Emperor

IBM Rockhopper

IBM z Systems Generations

N-4

z9 Enterprise Class
- Announced 7/2005
- 1.7 GHz
- Up to 54 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 512 GB Memory

z9 Business Class
- Announced 4/2006
- 1.4 GHz
- Up to 7 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 64 GB Memory

N-3

z10 Enterprise Class
- Announced 2/2008
- 4.4 GHz
- Up to 64 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 1.5 TB Memory

z10 Business Class
- Announced 10/2008
- 3.5 GHz
- Up to 10 cfg cores (5 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 248 GB Memory

N-2

zEnterprise 196
- Announced 7/22/2010
- 5.2 GHz
- Up to 80 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 3 TB Memory

zEnterprise 114
- Announced 7/12/2011
- 3.8 GHz
- Up to 10 cfg cores (5 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 248 GB Memory

N-1

zEnterprise EC12
- Announced 8/28/2012
- 5.5 GHz
- Up to 101 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 3 TB Memory

zEnterprise BC12
- Announced 7/23/2013
- 4.2 GHz
- Up to 13 cfg cores (6 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 496 GB Memory

N

IBM z13
- Announced 1/14/2015
- 5.0 GHz
- Up to 141 cfg cores
- CP, IFL, ICF, zIIP
- Up to 10 TB Memory

IBM z13s
- Announced 2/16/2016
- 4.3 GHz
- Up to 20 cfg cores (6 CP)
- CP, IFL, ICF, zIIP
- Up to 4 TB Memory

IBM z13 (Driver 27) GA2 and z13s – March 2016

IBM z13 (2964)

- **Announce – 02/16, GA1 – since 01/15**
- **5 models – NE1, NC9, N96, N63, N30**
  - Up to 141 customer configurable engines
- **Sub-capacity Offerings for up to 30 CPs**
- **PU (Engine) Characterization**
  - CP, IFL, ICF, zIIP, SAP, IFP (No zAAPs)
- **SIMD instructions, SMT for IFL and zIIP**
- **On Demand Capabilities**
  - CoD: CIU, CBU, On/Off CoD, CPE
- **Memory – up to 10 TB**
  - Up to 10 TB per LPAR (if no FICON Express8)
  - 96 GB Fixed HSA
- **Channels**
  - PCIe Gen3 16 GBps channel buses
  - Six LCSSs, up to 85 LPARs
  - Four Subchannel Sets per LCSS
  - FICON Express16S or 8S (8 Carry forward)
  - OSA Express5S (4S carry forward)
  - HiperSockets™ – up to 32
  - Flash Express
  - zEnterprise Data Compression (zEDC)
  - RDMA over CE (RoCE) with SR-IOV Support
  - Shared Memory Communications - Direct Memory Access (SMC-D)
- **Crypto Express5S**
- **Parallel Sysplex clustering, PCIe Coupling, InfiniBand Coupling**
- **IBM zAware: z/OS and Linux on z Systems**
- **Operating Systems**
  - z/OS®, z/VM®, z/VSE, z/TPF, Linux on z Systems, KVM for IBM z
- **IBM Dynamic Partition Manager (DPM)**
- **z Appliance Container Infrastructure (zACI)**

IBM zBX Model 004 (2458)

- **Announce – 01/15**
- **Upgrade ONLY stand alone Ensemble node converted from an installed zBX Model 2 or 3**
- **Doesn't require a 'owning' CPC**
- **Management – Unified Resource Manager**
- **zBX Racks (up to 4) with:**
  - Dual 1U Support Elements, Dual INMN and IEDN TOR switches in the 1st rack
  - HMC LAN attached (no CPC BPH attachment)
  - 2 or 4 PDUs per rack
- **Up to 8 BladeCenter H Chassis**
  - Space for 14 blades each
  - 10 GbE and 8 Gbps FC connectivity
  - Advanced Management Modules
  - Redundant connectivity, power, and cooling
- **Up to 112 single wide IBM blades**
  - IBM BladeCenter PS701 Express
  - IBM BladeCenter HX5 7873
  - IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise (M/T 2462-4BX)
  - IBM WebSphere DataPower® Integration Appliance XI52 Virtual Edition on System x
- **Operating Systems**
  - AIX 5.3 and higher
  - Linux on IBM System x®
  - Microsoft Windows on System x
- **Hypervisors**
  - KVM Hypervisor on System x
  - PowerVM™ Enterprise Edition

IBM z13s (2965)

- **Announced 02/16**
- **2 models – N10, N20**
  - Up to 20 customer configurable engines. Max 6 CPs
- **Higher levels of granularity available**
  - 156 Capacity Indicators
- **PU (Engine) Characterization**
  - CP, IFL, ICF, zIIP, SAP, IFP (No zAAPs)
- **SIMD instructions, SMT for IFL and zIIP**
- **On Demand Capabilities**
  - CoD, CIU, CBU, On/Off CoD. CPE
- **Memory – up to 4 TB for Server**
  - 40 GB Fixed HSA
- **Channels**
  - PCIe Gen3 16 GBps channel bus
  - Three LCSSs, up to 40 LPARs
  - Three Subchannel Sets per LCSS
  - FICON Express16S or 8S (8 Carry forward)
  - OSA Express5S (4S carry forward)
  - HiperSockets – up to 32
  - Flash Express
  - zEnterprise Data Compression (zEDC)
  - RDMA over CE (RoCE) with SR-IOV Support
  - Shared Memory Communications - Direct Memory Access (SMC-D)
- **Crypto Express5S**
- **Parallel Sysplex clustering, PCIe Coupling, InfiniBand Coupling**
- **IBM zAware: z/OS and Linux on z Systems**
- **Operating Systems**
  - z/OS, z/VM, z/VSE, z/TPF, Linux on z Systems, KVM for IBM z,
- **IBM Dynamic Partition Manager (DPM)**
- **z Appliance Container Infrastructure (zACI)**

IBM z Systems – Reliable, Scalable, Secure and Virtualized

An integrated, highly scalable computer system that allows many different pieces of work to be handled at the same time, sharing the same information as needed with protection, handling very large amounts of information for many users with security, without users experiencing any failures in service

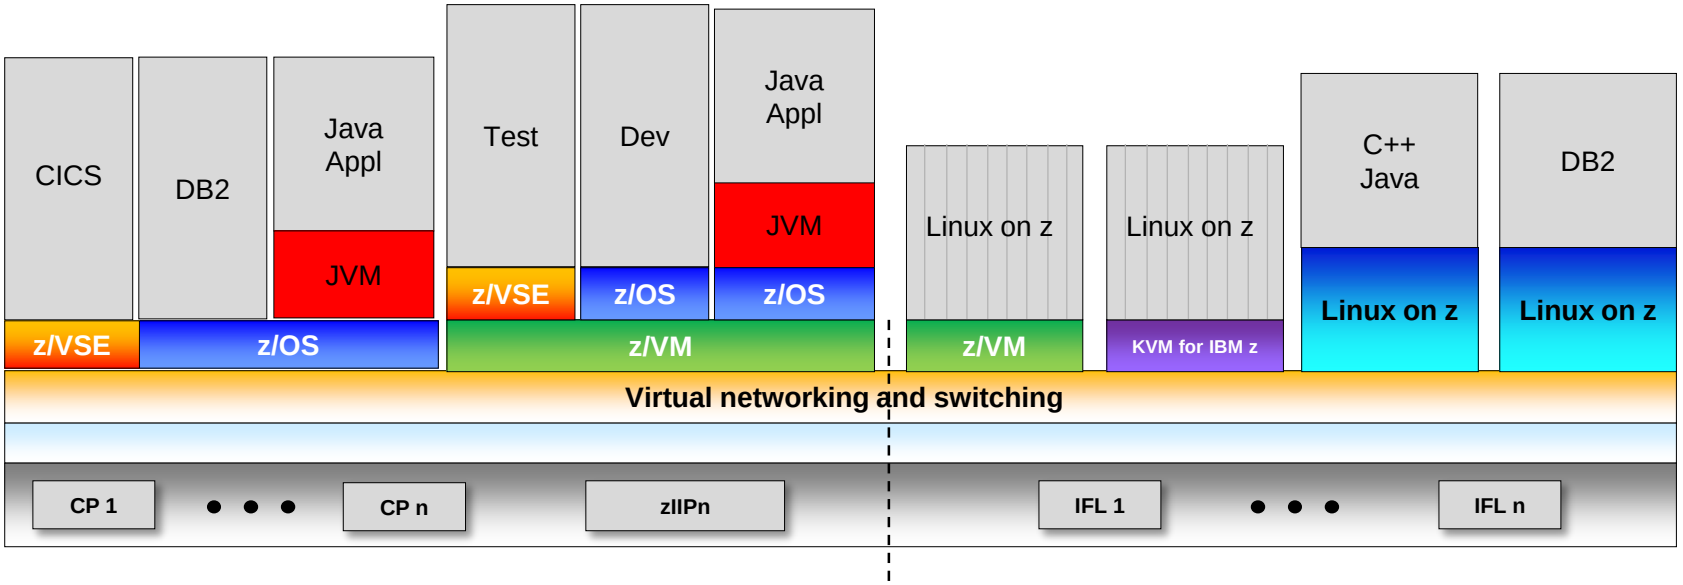

- Large scale, robust consolidation platform
- Built-in Virtualization
- 100s to 1000s of virtual servers on z/VM
- Intelligent and autonomic management of diverse workloads and system resources

Accelerate Key Workloads with Special-Purpose Hardware

- **On-processor**
  - Crypto (CPACF), Compression, SIMD, SMT
  - Tight, synchronous integration with instruction stream
- **PCIe Gen3**
  - Accessible and sharable by all processors
  - Faster time to market for new functions
  - Compression (zEDC), Crypto, Flash Express
- **Network Acceleration**
  - Shared Memory Communications over RDMA -- SMC-R (RDMA over Converged Ethernet -- RoCE)
  - Shared Memory Communication – Direct Memory Access (SMC-D)
- **Integrated External Accelerators**
  - Integrated by software
  - IBM DB2 Analytics Accelerator for DB2 Query Acceleration
- **Specialty Engines and Firmware Partitions**
  - Leverage flat SMP design, enable price flexibility
  - zIIP for DB2 and Java, IFL for Linux on z Systems
  - IBM zAware
  - IBM z Appliance Container Infrastructure (zACI)

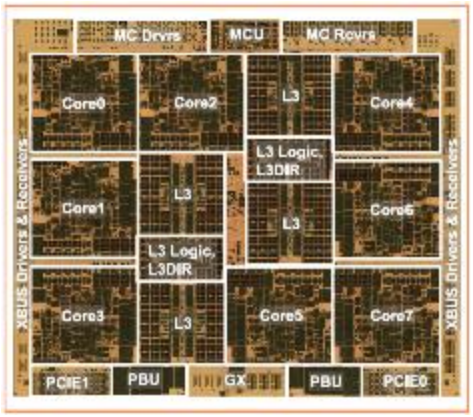

IBM DB2 Analytics Accelerator built on Netezza Technology

z13 Continues the CMOS Mainframe Heritage Begun in 1994

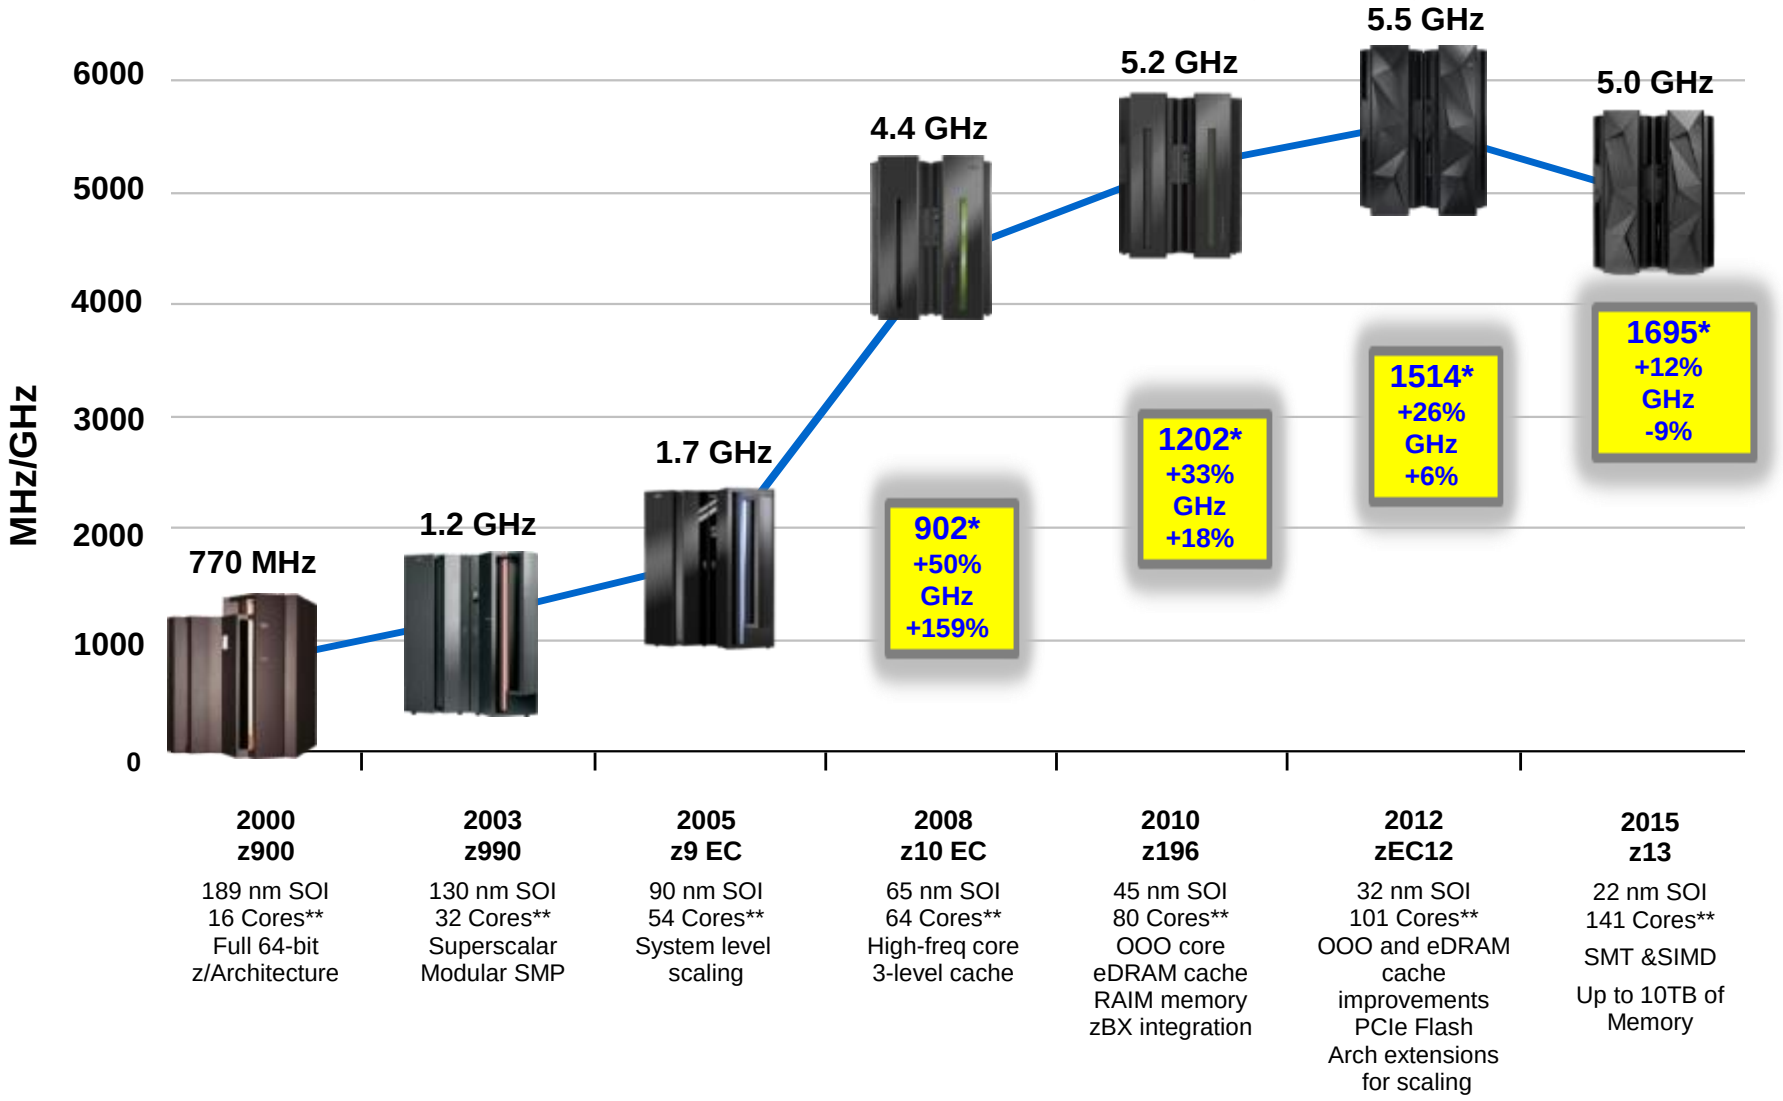

* MIPS Tables are NOT adequate for making comparisons of z Systems processors. Additional capacity planning required

** Number of PU cores for customer use

z13s continues the CMOS Mainframe Heritage

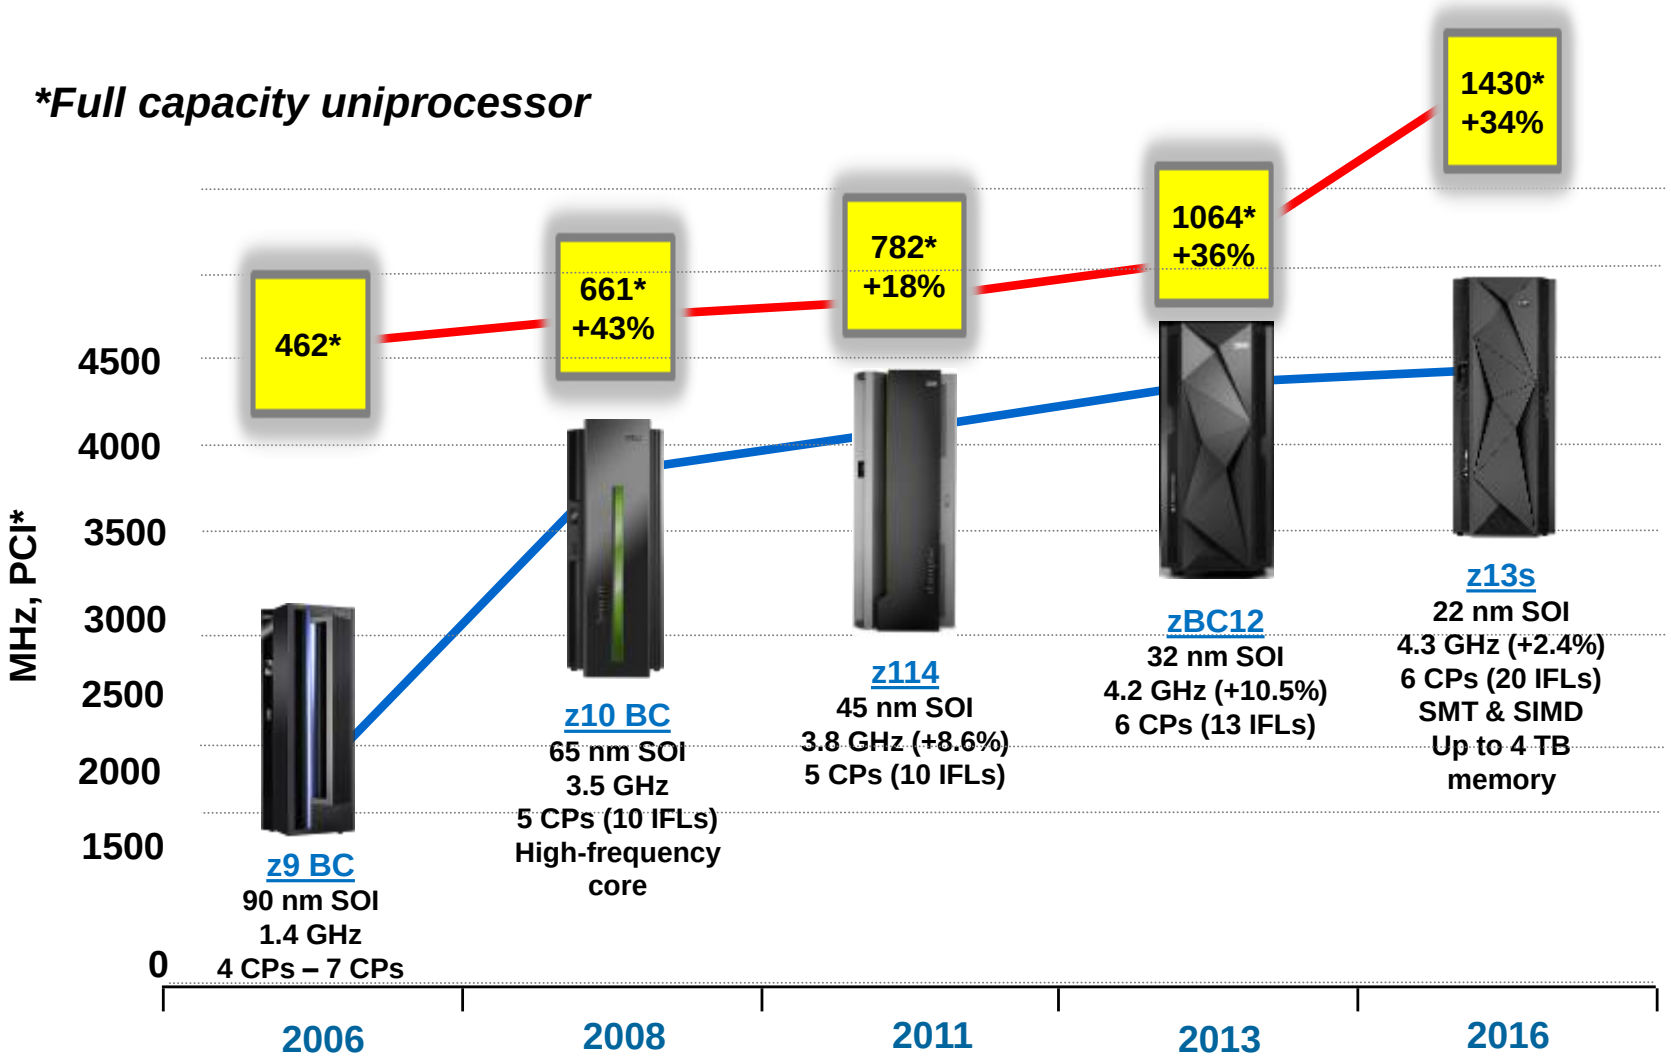

*NOTE: MIPS Tables are NOT adequate for making comparisons of z Systems processors in proposals

Advanced System Design Optimized for Digital Business

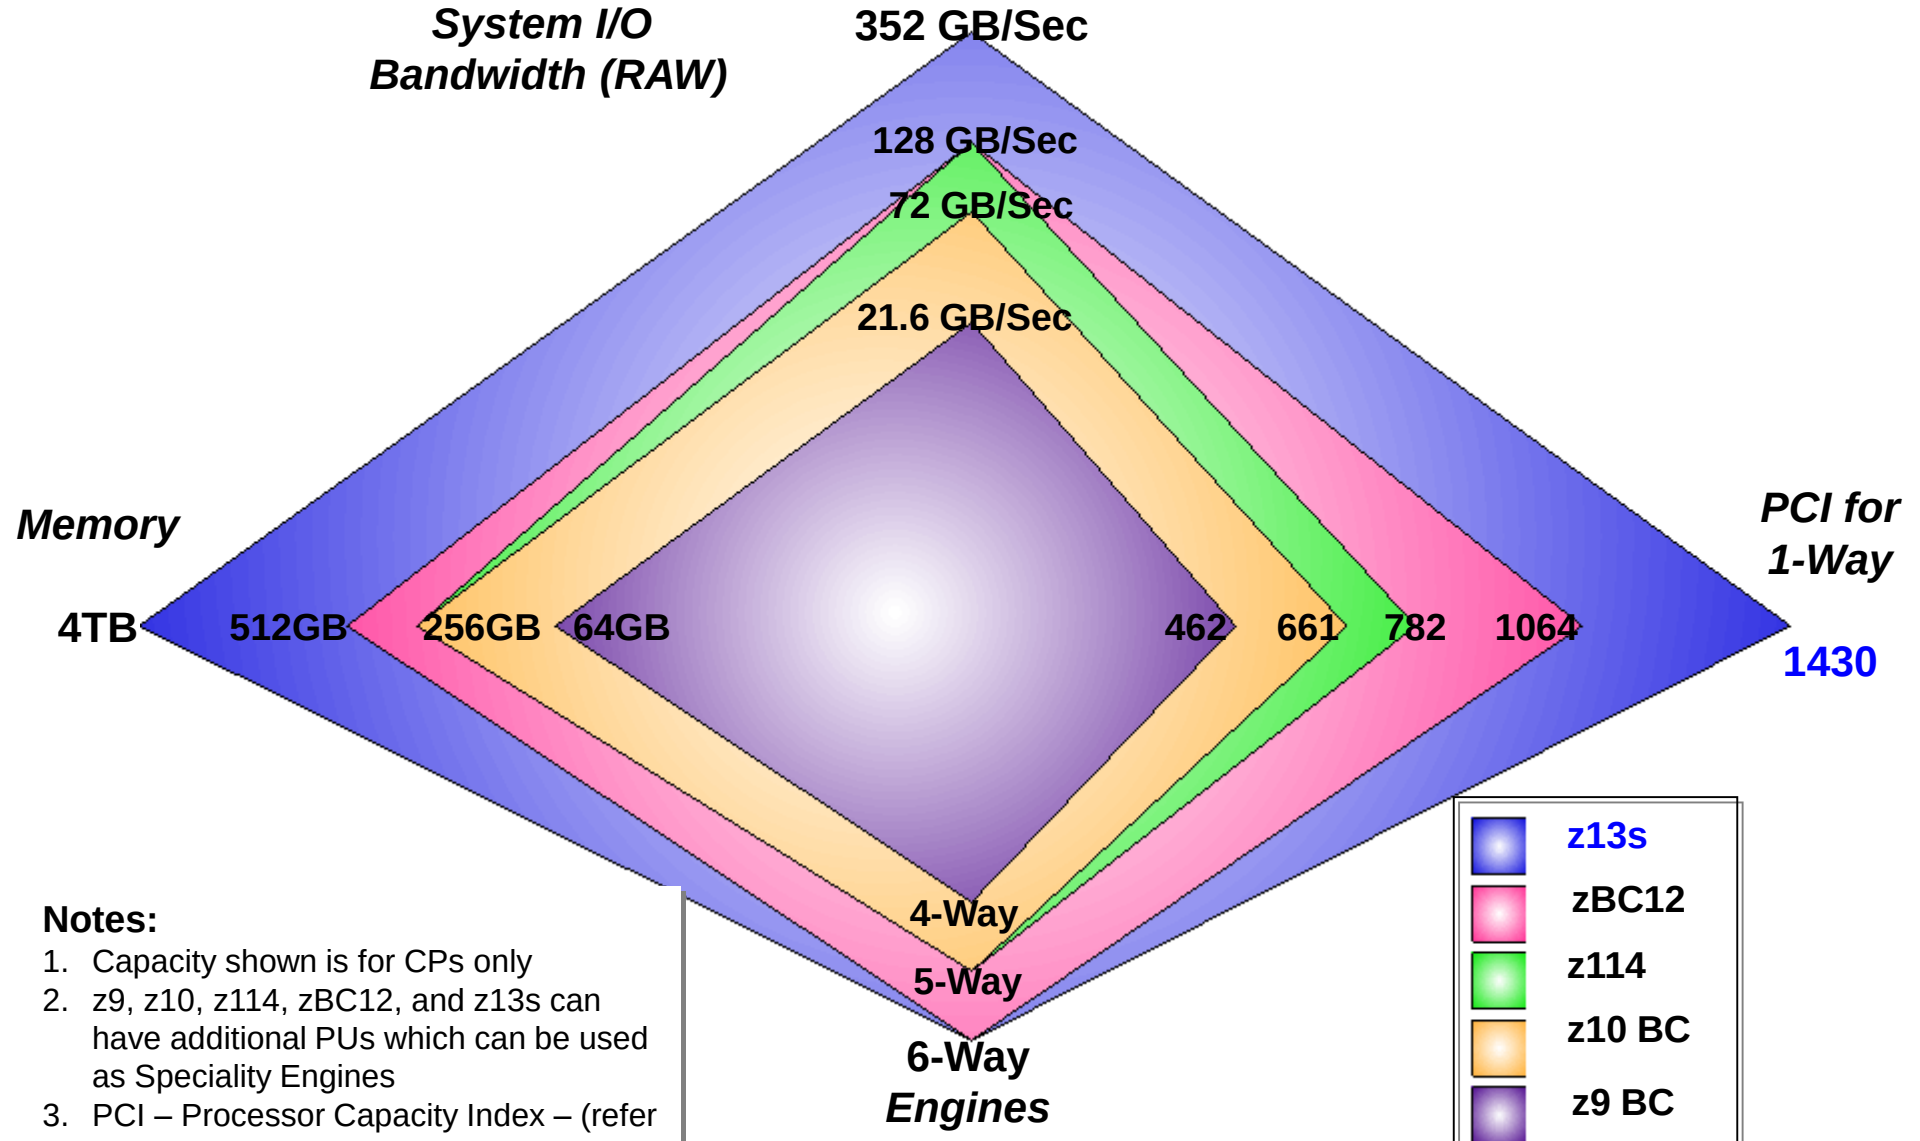

Notes:

1. Capacity shown is for CPs only
2. z9, z10, z114, zBC12, and z13s can have additional PUs which can be used as Speciality Engines
3. PCI – Processor Capacity Index – (refer to LSPR)

z13s at a Glance

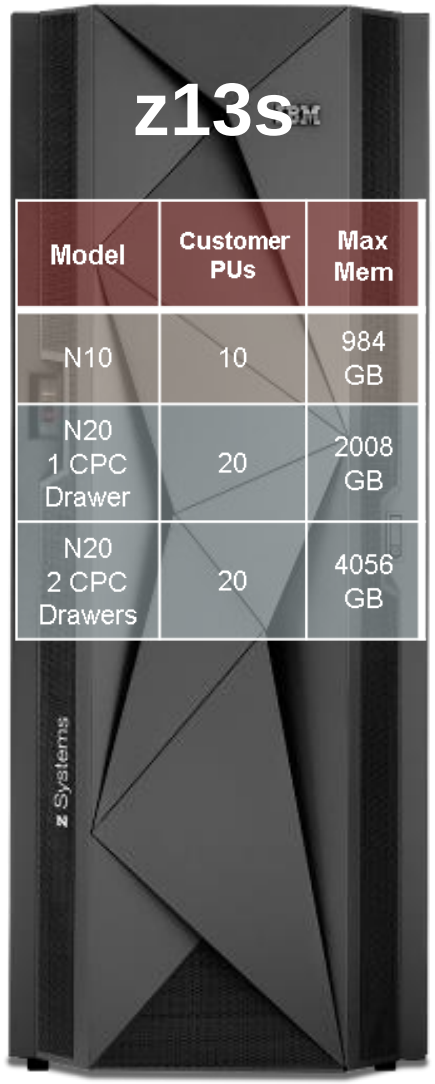

- **Machine Type**
  - –2965
- **2 Models**
  - –N10 and N20
    - N20 available as one- or two- processor drawer model
    - The 2nd drawer in the N20 is driven by I/O and/or memory requirements
  - –Single frame, air cooled
  - –Non-raised floor option available
  - –Overhead Cabling and DC Power Options
- **Processor Units (PUs)**
  - –13 PU active cores (model N10 – 10 client configurable) or 26 PU active cores (model N20 – 20 client configurable)
  - –Up to 3 standard SAPs per system (2 for model N10, 3 for model N20)
  - –2 spares designated for Model N20
  - –1 Integrated firmware processor (IFP)
  - –Dependent on the H/W model - up to 10 (N10) or 20 (N20) PU cores available for characterization:
    - •Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), IBM z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs), Integrated firmware processor (IFP)
  - •156 capacity settings
- **Memory**
  - –Up to 4 TB including:
    - •System minimum = 64 GB
    - •40 GB fixed HSA separately managed
    - •RAIM standard
    - •Maximum for customer use 4056 GB (Model N20-2 drawer)
    - •Increments of 128 to 1024 GB
    - •Flash Express Read/Write Cache in HSA (0.5 GB)
- **I/O**
  - –PCIe Gen3 channel subsystem
    - Up to 64 PCIe Channel features
  - Support for non-PCIe Channel features (max one I/O drawer via carry forward)
  - –Up to 3 Logical Channel Subsystems (LCSSs)
- **STP - optional (No ETR)**

IBM z Systems Single Frame Comparison

	z114 M05	z114 M10	zBC12 H06	zBC12 H13	z13s N10	z13s N20(1)
Uniprocessor Performance	782 MIPS		1064 MIPS		1430 MIPS (+34%)	
Frequency	3.8 GHz		4.2 GHz		4.3 GHz	
z/OS Capacity	26 - 3139 MIPS		50 – 4958 MIPS		80 to 7123 MIPS	
Total System Memory	120 GB	248 GB	240 GB	496 GB	1TB	4TB
Configurable Specialty Engines	5	10	6	13	10	20
Configurable CPs	0-5		0 – 6		0 to 6	
LPARS/LCSS	30/2		30/2		40/3	
HiperSockets	32		32		32	
PCIe I/O drawer, I/O drawer, Max	2, 2, 3	2, 2, 3	2, 1, 3	2, 1, 3	1, 1, 2 (2)	2, 1, 3 (2)
I/O slots per I/O drawers/ PCIe I/O drawers	8/32		8/32		8/32	
FICON® Channels	128		128		64(3)	128(3)
OSA Ports	96		96		64(3)	96(3)
ESCON® Channels	240		0(4)		0(4)	
IFB host bus Bandwidth, PCIe Bandwidth	6.0 GB/sec (IFB), 8.0 GB/sec (PCIe Gen2)		6.0 GB/sec (IFB), 8.0 GB/sec (PCIe Gen2)		6.0GB/sec(IFB), 16.0 GB/sec (PCIe Gen3)	
ISC-3, PSIFB, PCIe	48, 8 -16, 0	48, 16 – 32, 0	32, 8 -16 0	32, 16- 32 0	0(5), 4-8(6), 8	0(5), 16-32(7), 16
zIIP/zAAP Maximum Qty	2	5	Up to 4 / 3	Up to 8 / 6	Up to 6(8) / 0	Up to 12(8) / 0
IFL Maximum Qty	5 (3139 MIPS)	10 (5390 MIPS)	6 (4958 MIPS)	13 (8733 MIPS)	10 (10767 MIPS)	20 (18335 MIPS)
ICF Maximum Qty	5	10	6	13	10	20
Capacity Settings	130	130	156	156	156	156
Upgradeable	M05 to M10, to zBC12 or to z13s		H06 to H13, H13 to zEC12 H20 (Radiator-based air cooled only) or to z13s		Within z13s (N10 to N20), N20 to z13 N30 (Radiator-based /air cooled only)	

Note

(*)See next page for Notes

IBM z Systems Comparison, z13s vs. z13 Model N30

	z13s N10	z13s N20[1]	z13 Model N30 *(new Build)*
Uniprocessor Performance	1430 MIPS		1695 MIPS
z/OS Capacity	80 – 7123 MIPS		250 – 32671 MIPS
Total System Memory	1 TB	4 TB	2.5 TB
Configurable Engines	10	20	30
Configurable CPs	0 – 6		0 – 30
LPARS/CSS	40/3		85/6
HiperSockets	32		32
I/O drawers / PCIe I/O drawers	1/1 [2]	1/2 [2]	0/3
I/O slots per I/O drawer/ PCIe I/O drawer	8/32		8/32
FICON® Channels	64[3]	128[3]	160
OSA Ports (10GbE/1GbE/1000BASE-T)	32/64/64[3]	48/96/96[3]	48/96/96
IFB host bus Bandwidth PCIe Gen3 Bandwidth	6.0 GB/sec (IFB) 16.0 GB/sec (PCIe)		6.0 GB/sec (IFB) 16.0 GB/sec (PCIe)
ISC-3 / PSIFB / ICA-SR	0[5]/4 -8[6]/8	0[5]/16 - 32[7]/16	0[5]/8 – 16/20
zIIP Maximum Qty	Up to 6[8]	Up to 12[8]	Up to 20[8]
IFL Maximum Qty	10 (10767 MIPS)	20 (18335 MIPS)	30 (32671 MIPS)
ICF Maximum Qty	10	20	30
Capacity Settings	156	156	120
Upgradeable	Within z13s (N10 to N20), N20 to z13 N30 (Radiator-based air cooled only)		z13 N63, N96, NC9, NE1 (Only radiator to radiator or water to water cooled as a MES)

See previous chart for footnotes

z13s Model N20 (One CPC Drawer) – Under the Covers

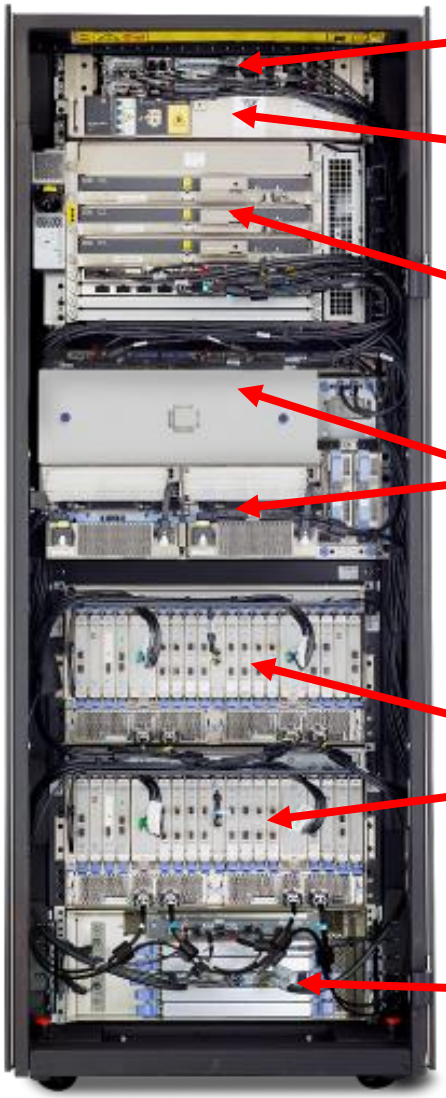

Rear View

Support Elements

Internal Batteries (optional)

Power Supplies

1 x CPC Drawer + slot for 2nd CPC Drawer Memory & Fanouts

PCIe I/O Drawers

I/O Drawer

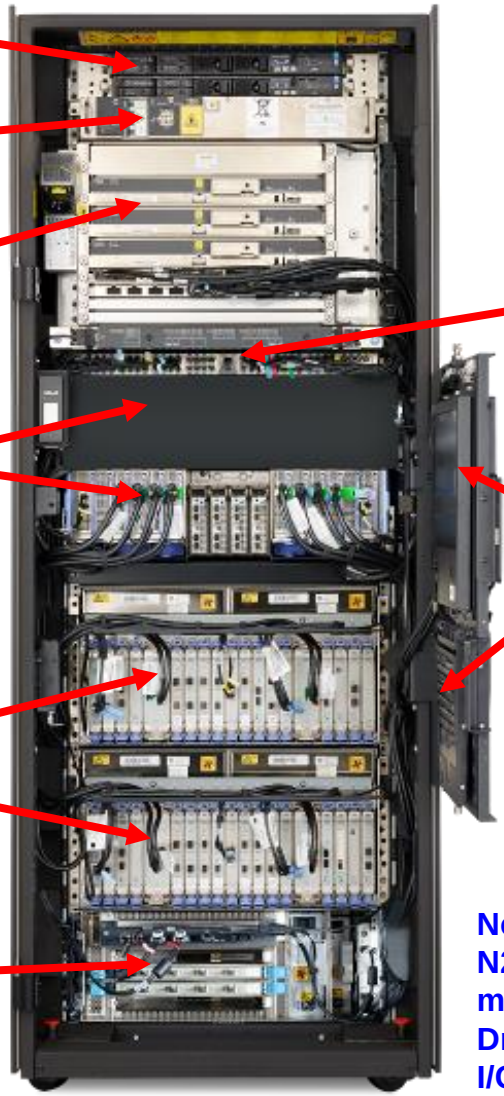

Front View

Ethernet cables for internal System LAN connecting Flexible Service Processor (FSP) cage controller cards (not shown)

SCH (x2)

Monitors + keyboards for SEs

Note: A 2 CPC drawer N20 can only support a maximum of 2 PCIe I/O Drawers OR one PCIe I/O Drawer and one I/O Drawer

Processor design and structure: What Is New?

- The IBM z13s symmetric multiprocessor (SMP) system is the next step in an evolutionary trajectory that began with the introduction of the IBM System/360 in1964. Over time, the design was adapted to the changing requirements that were dictated by the shift toward new types of applications that clients depend on.
- The z13s has ultra-high frequency, large high-speed buffers (caches) and memory, superscalar processor design, out-of-order core execution, simultaneous multithreading (SMT), single-instruction multiple-data (SIMD) execution, and flexible configuration options. It is the next implementation of z Systems to address the ever-changing IT environment.
- The microprocessor of z13s has the same design as z13. The difference is the frequency - 4.3GHz vs. 5GHz for z13.
- The drawer packaging is also derived from z13. The z13s PU chips are air cooled (vs. water cooled for z13) as the lower frequency generates less heat.
- z13s has several microprocessor improvements upon zBC12:
  - higher instruction execution parallelism (improved OOO instruction handling)
  - two vector execution units (SIMD)
  - simultaneous multithreaded (SMT) architecture that supports concurrent execution of two threads
  - eight-core processor chip (6 or 7 cores per chip will be used with z13s)
  - a robust cache hierarchy
- Benefits:
  - performance gains in legacy online transaction processing and business analytics workloads.

z13s Model N10 CPC Drawer (Top View)

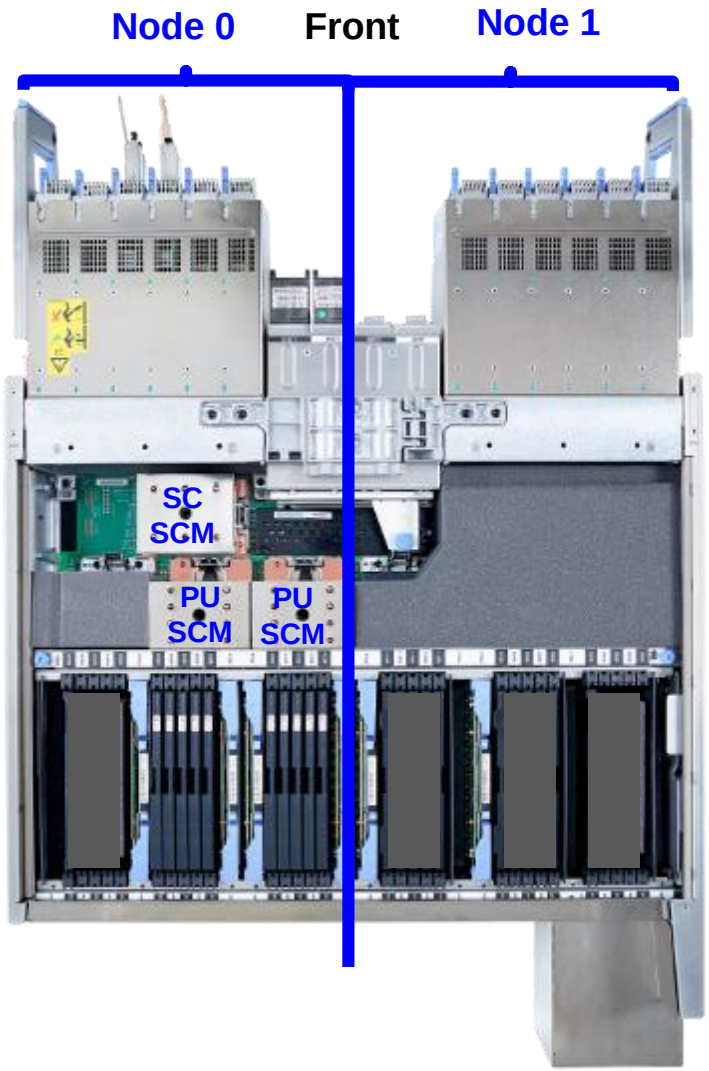

- **Model N10 supports a single Node (Node 0)**
- **Each Node Supports:**
  - One System Control (SC) chip (480 MB L4 cache)
  - Two Processing Units (PU) chips running at 4.3GHz
    - Eight-core per PU chip design
    - Six or seven active cores per PU chip
  - One memory controller per PU chip (two per node)
  - Five DDR3 DIMM slots per memory controller: 10 total per node (up to 1024GB per node)
  - Two Flexible Service Processors
  - Four PCIe fanout slots
  - Two slots for IFB fanouts or PSIFB coupling link fanouts

z13s Model N20 CPC Drawer (Top View)

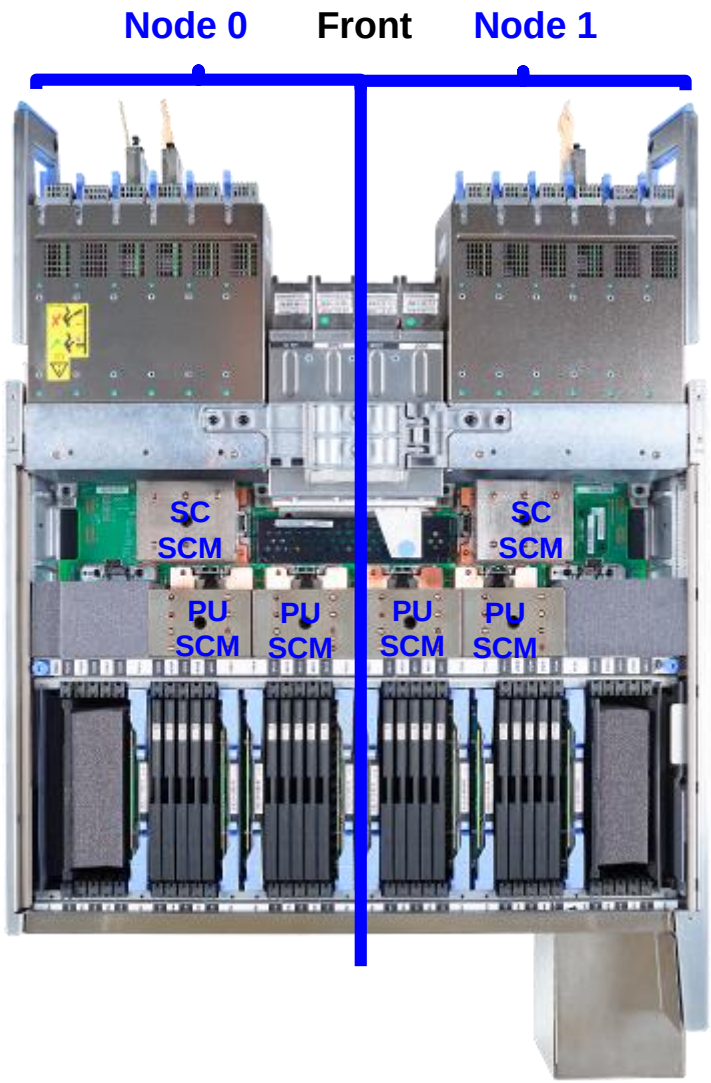

- **Model N20 CPC drawer has two nodes: Node 0 and Node 1**
- **Each node:**
  - One System Control (SC) chip (480 MB L4 cache)
  - Two Processing Units (PU) chips running at 4.3GHz
    - Eight-core per PU chip design
    - Up to seven active cores per PU chip
  - One memory controller per PU chip (two per node)
  - Five DDR3 DIMM slots per memory controller: 10 total per node (up to 1024GB per node)
  - Two Flexible Service Processors
  - Four PCIe fanout slots
  - Two slots for IFB fanouts or PSIFB coupling link fanouts

z13s Model N20 Single CPC Drawer Layout Details

Front

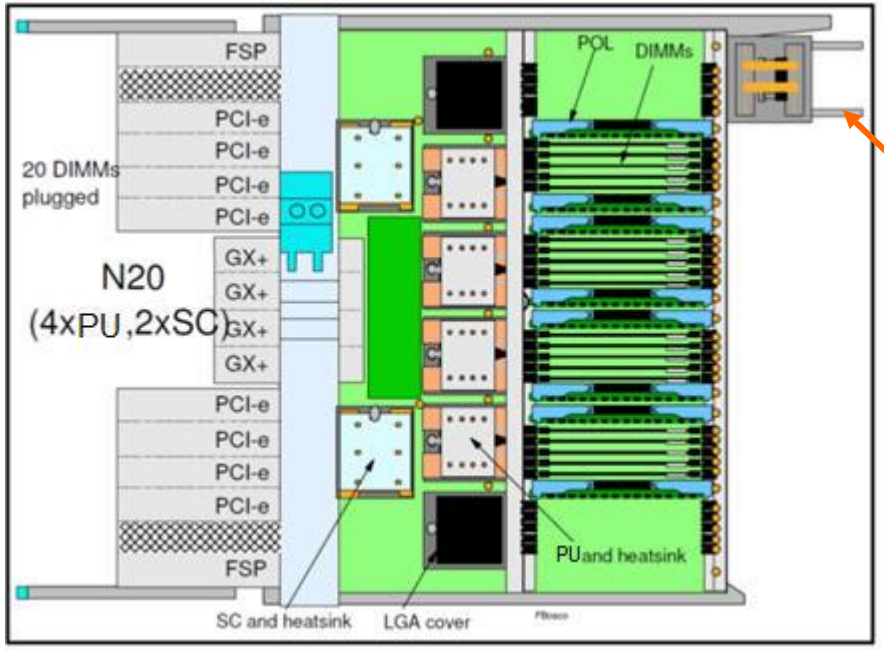

Rear

OSC connection to the OSC back plane

SMP connectors for inter-Drawer connections – one on each side of the CPC drawer.

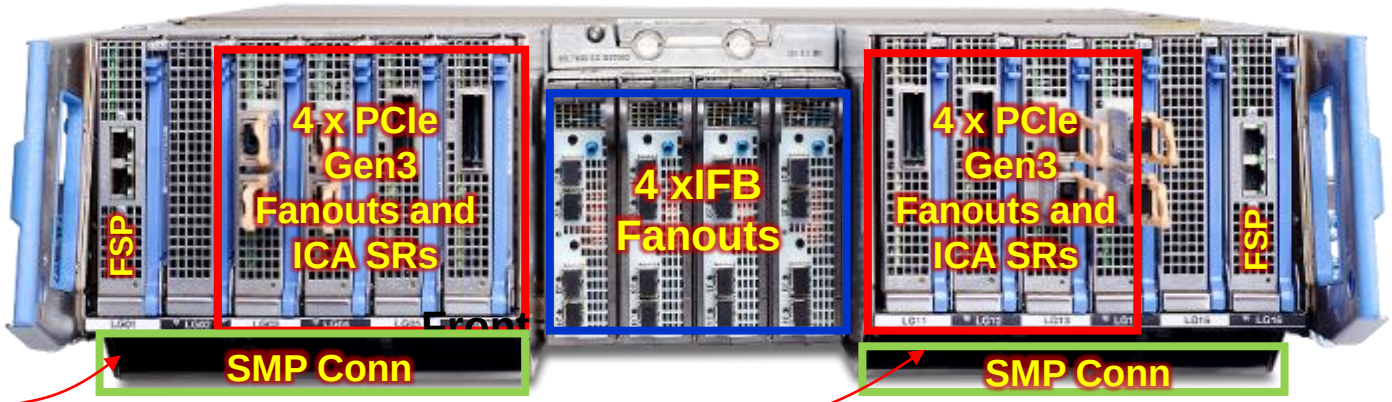

z13s N20 – Two CPC Drawers (Front View)

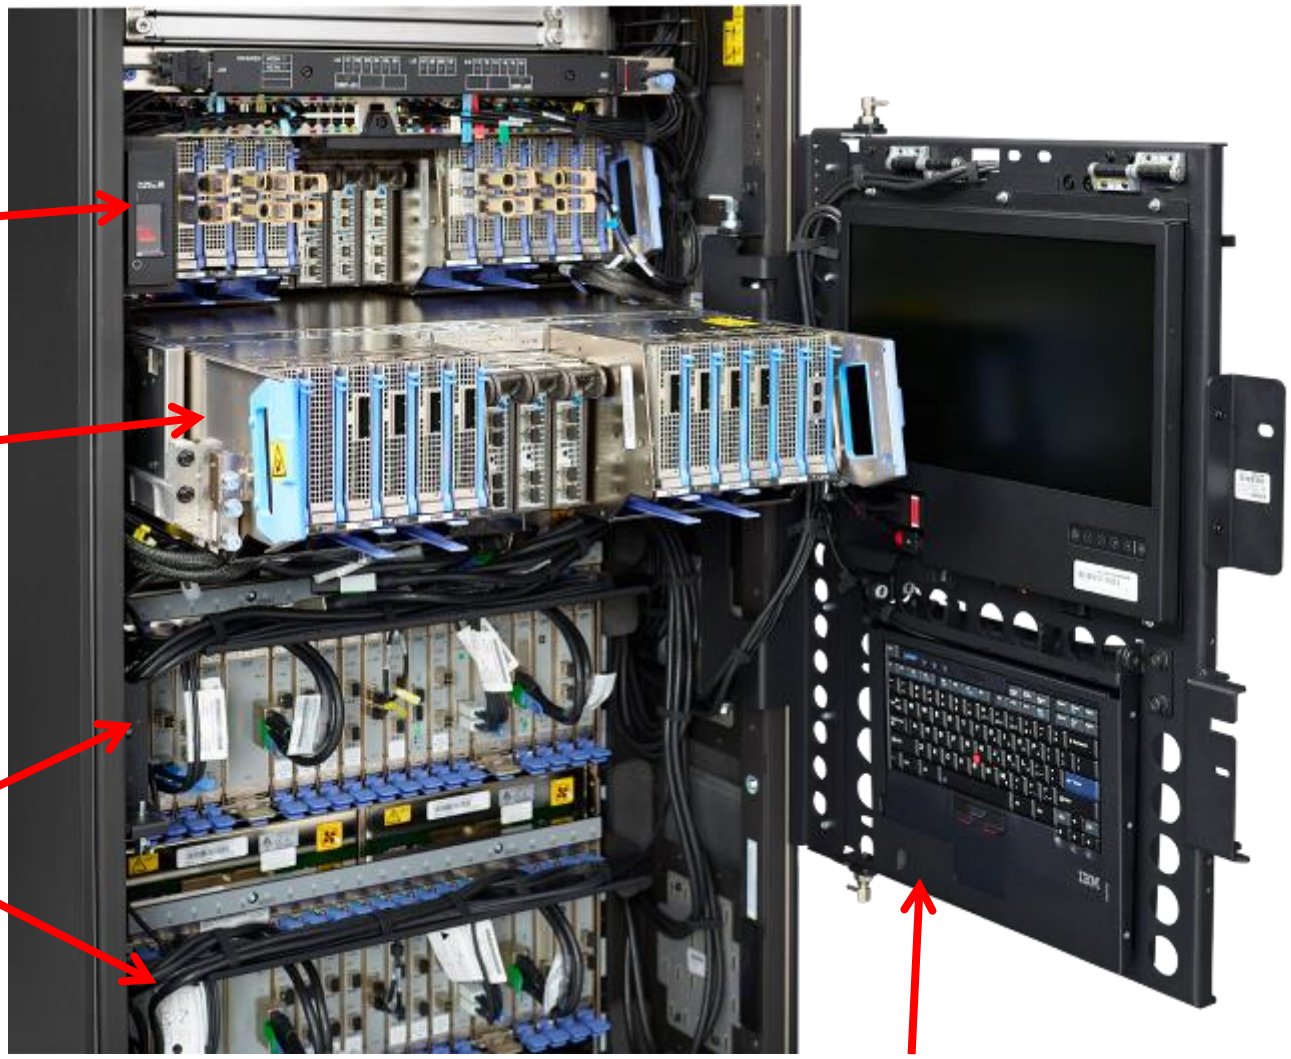

z13s Processor Unit (PU) Chip Overview

- Double width instruction pipeline
  - Improves IPC for all modes
  - Symmetry simplifies dispatch/issue rules
  - Required for effective SMT
- Added additional FXU and BFU
  - 4 FXUs
  - 2 BFUs, DFUs
  - 2 new SIMD units
- SIMD unit plus additional registers
- Pipe depth re-optimized for power/performance
  - Processor performance increased
- SMT support
  - Wide, symmetric pipeline
  - Full architected state per thread
  - SMT-adjusted CPU usage metering

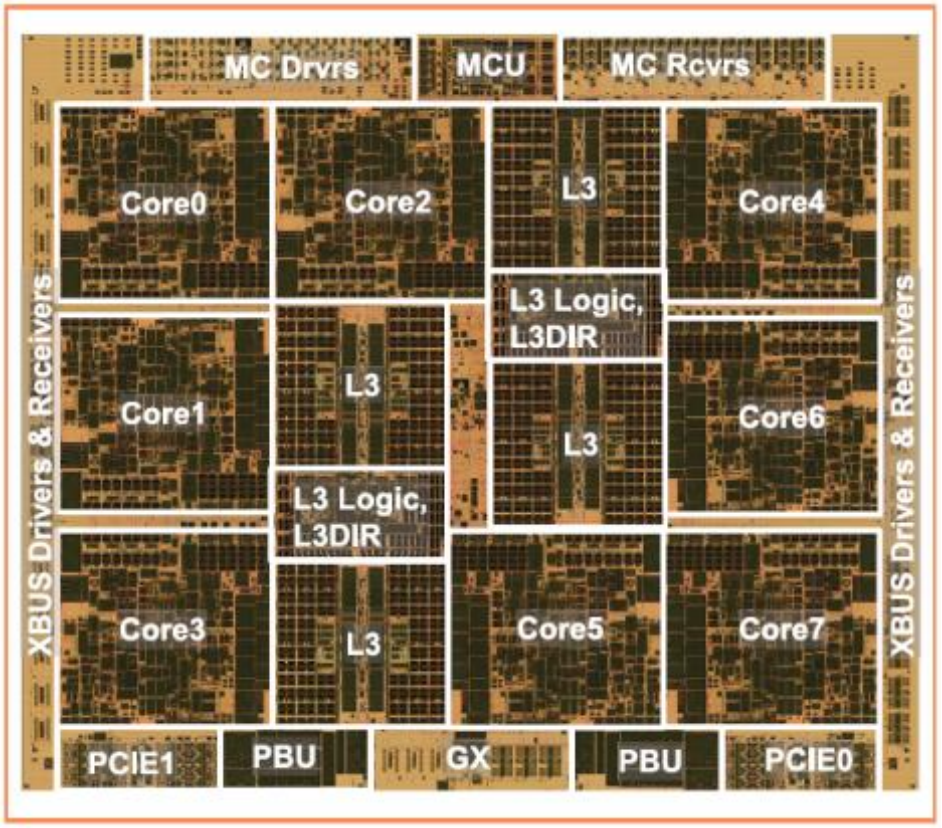

z13s Processor Unit (PU) Chip Details

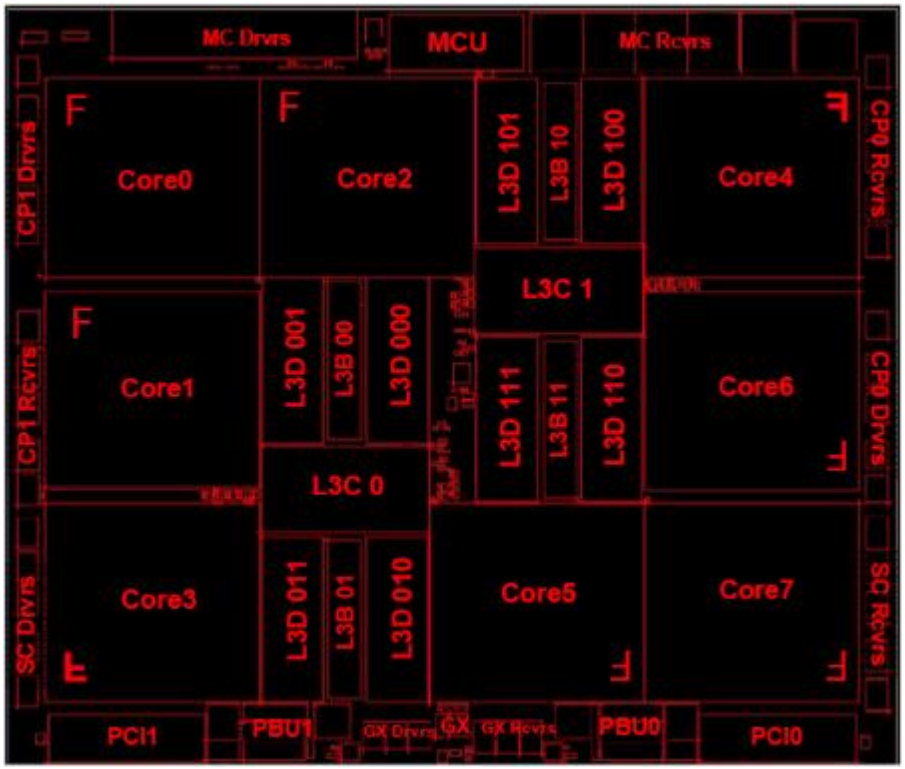

- **14S0 22nm SOI Technology**
  - 17 layers of metal
  - 3.99 Billion Transistors
  - 13.7 miles of copper wire
- **Chip Area**
  - 678.8 mm^2
  - 28.4 x 23.9 mm
  - 17,773 power pins
  - 1,603 signal I/Os

- **6 or 7 *active* cores (PUs) per chip on z13s (designed with 8 cores total)**
  - 4.3 GHz (vs. 4.2GHz for zBC12)
  - L1 cache/ core
    - 96 KB I-cache
    - 128 KB D-cache
  - L2 cache/ core
    - 2M i+2M d Byte eDRAM split private L2 cache
- **Single Instruction/Multiple Data (SIMD)**
- **Single thread or 2-way simultaneous multithreading (SMT) operation**
- **Improved instruction execution bandwidth:**
  - Greatly improved branch prediction and instruction fetch to support SMT
  - Instruction decode, dispatch, complete increased to 6 instructions per cycle
  - Issue up to 10 instructions per cycle
  - Integer and floating point execution units
- **On chip 64 MB eDRAM L3 Cache**
  - Shared by all cores
- **I/O buses**
  - One InfiniBand I/O bus
  - Two PCIe I/O buses
- **Memory Controller (MCU)**
  - Interface to controller on memory DIMMs
  - Supports RAIM design

z13s Storage Controller (SC) Chip Detail

- CMOS 14S0 22nm SOI technology
  - 15 Layers of metal
  - 7.1 Billion transistors
  - 12.4 Miles of copper wire
- Chip area
  - 28.4 x 23.9 mm
  - 678 mm2
  - 11,950 power pins
  - 1,707 Signal Connectors
- eDRAM shared L4 cache
  - 480 MB per SC chip (Non-inclusive of L3)
    - 960 MB on a **two node** CPC drawer model (N20)
  - 224 MB L3 NIC directory
    - 448 MB L3 NIC on a **two node** CPC drawer model (N20)
- Interconnects (L4 – L4)
  - Two buses to PU chips intra node
  - One bus to second SC in drawer (inter-node)
  - Two busses to SCs in remote drawer (1 per node)
- Six clock domains

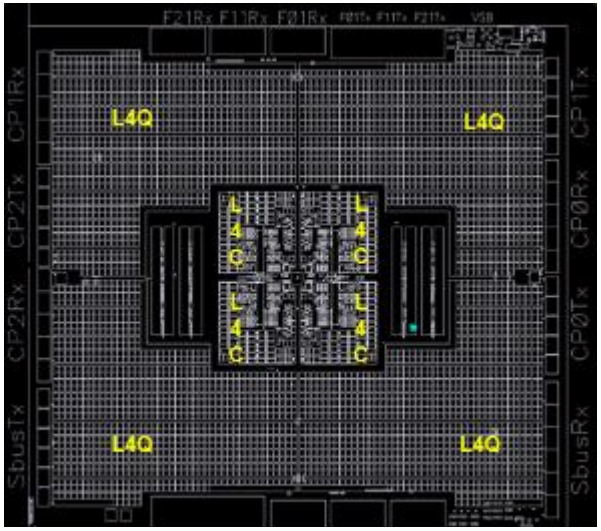

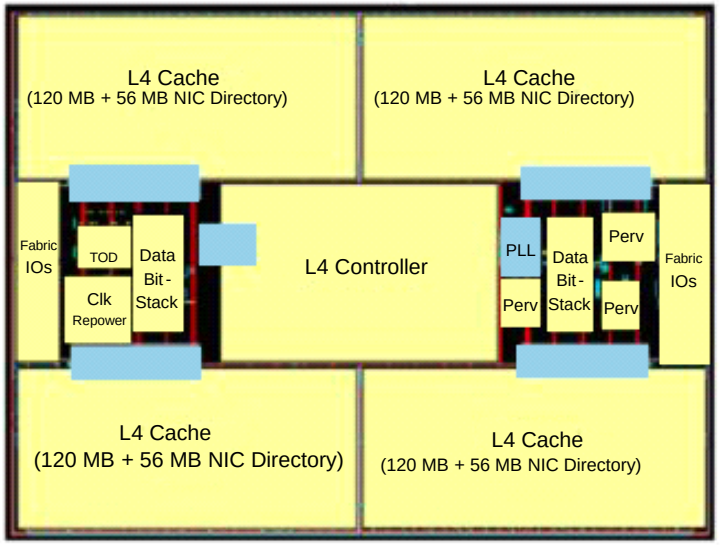

(480 MB + 224 MB NIC Directory)

PU and SC SCMs and CPC Drawer Structure Details

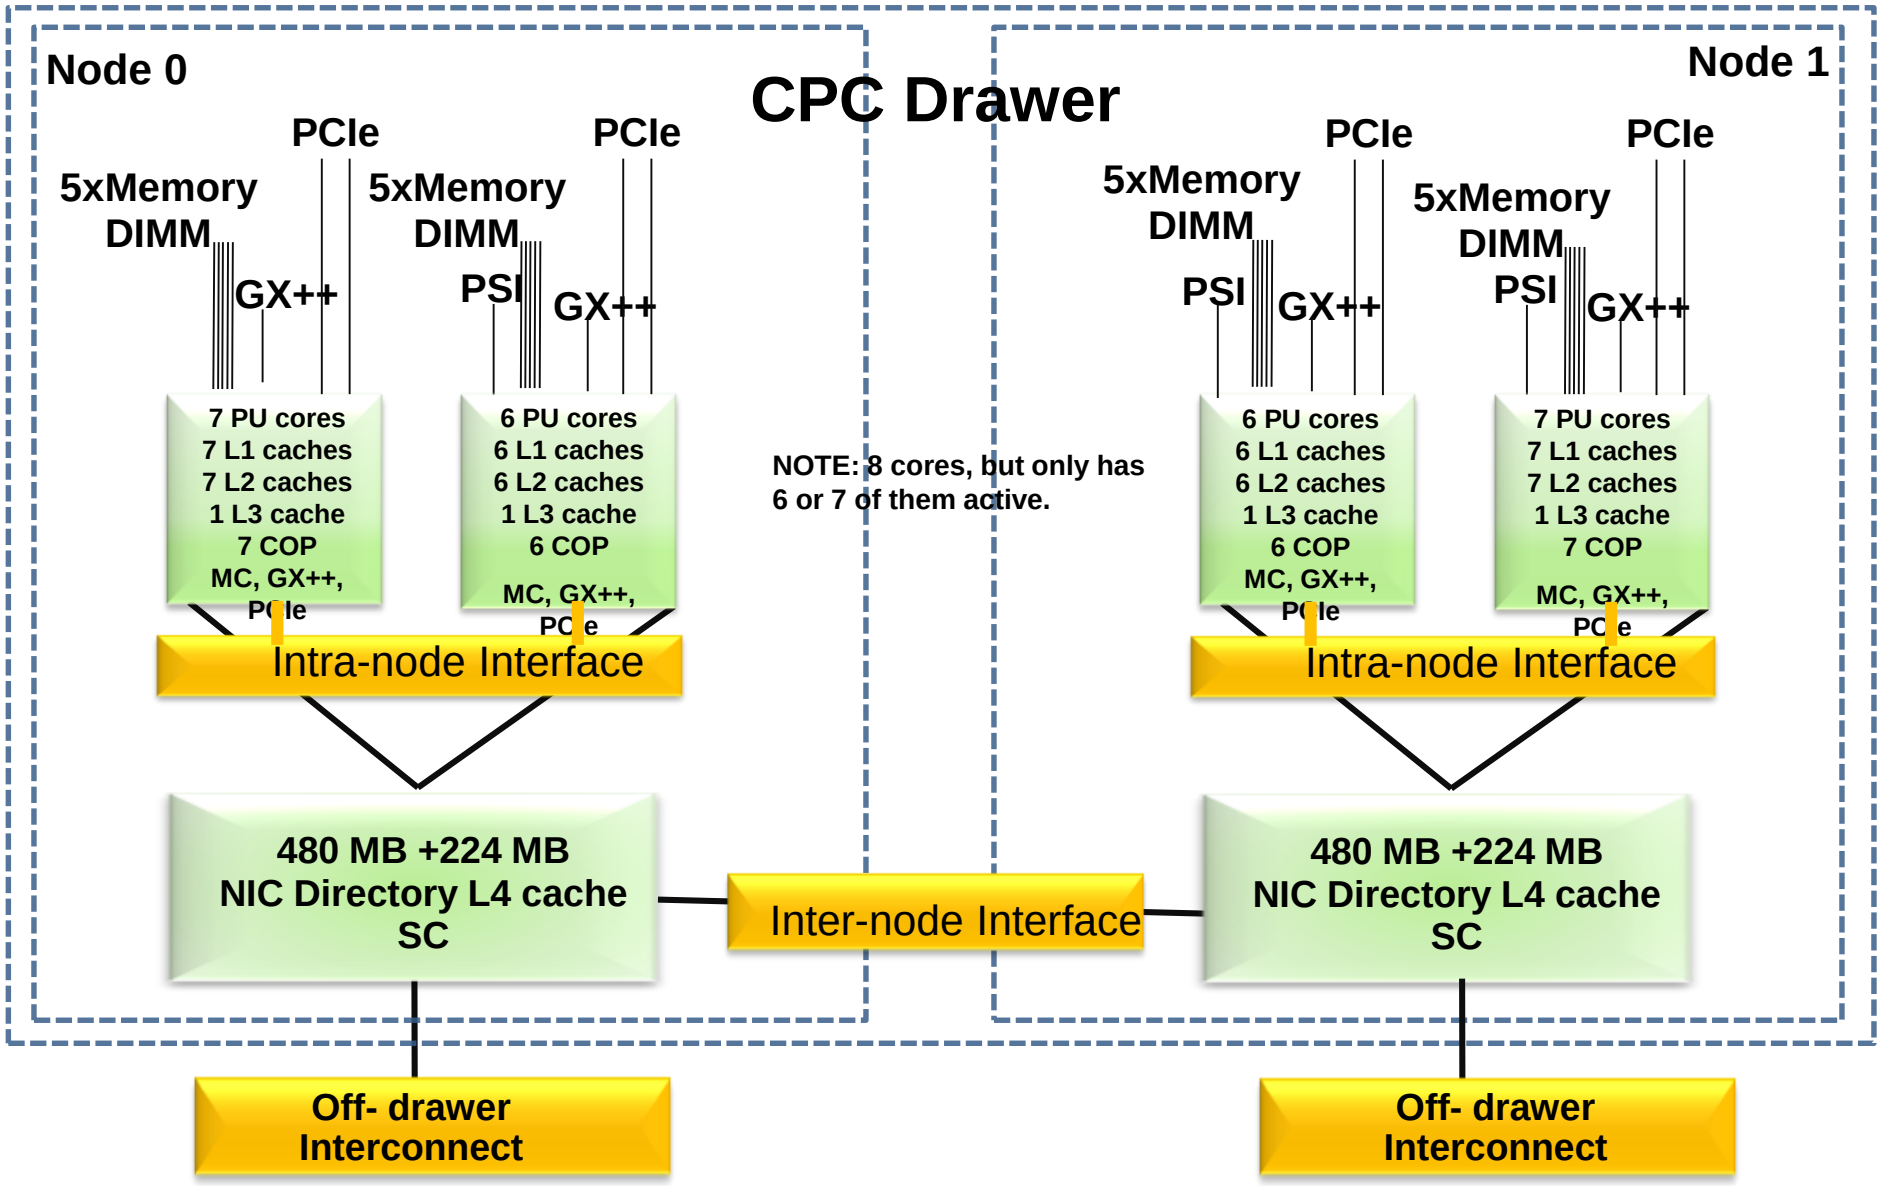

z13s Model N10 Drawer Structure and Interconnect

Partially Populated Drawer

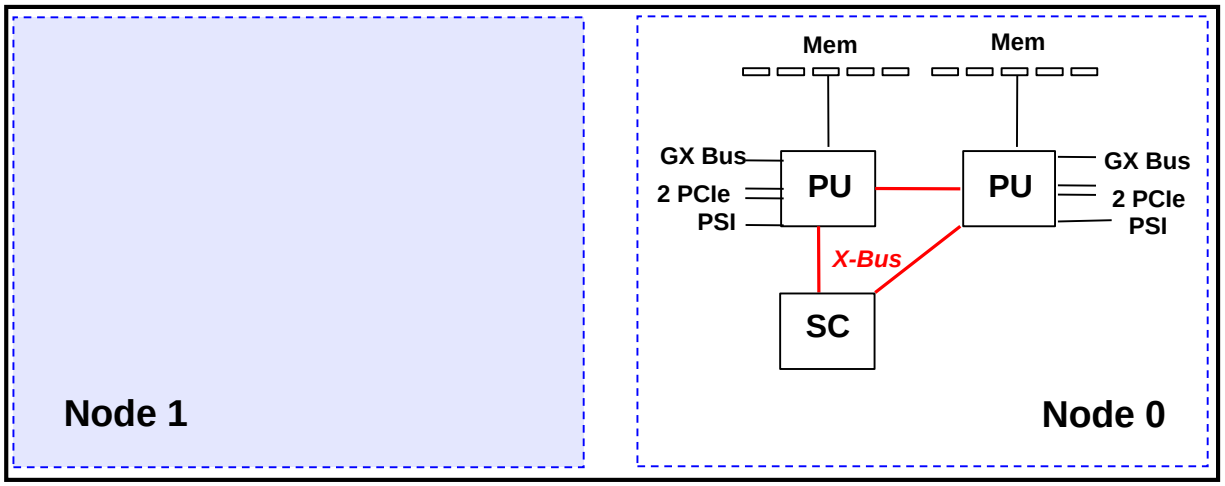

1 CPC Drawer
1 Node (N10)

One Physical node

- **Chips**
  - Two PU chips
  - One SC chip
- **Redundant array of independent memory (RAIM) Memory**
  - Two Memory Controllers: One per **PU**
  - Five DDR3 DIMM slots per PU Chip Memory Controller
  - Populated DIMM slots: 10
- **SC and PU Chip Interconnects**
  - SC and PU Chip Interconnects
  - X-bus: SC L4 and CP L3s to each other (same node)

Note: Both z114 and zBC12 have one node per CPC drawer with two PU SCMs, five DIMMs per PU SCM and one System Control SCM. The PU SCM L3 Caches are not connected directly. L4 cache design is inclusive of the L3s.

z13s Model N20 with 2 Drawers Structure and Interconnect

Fully Populated Drawer

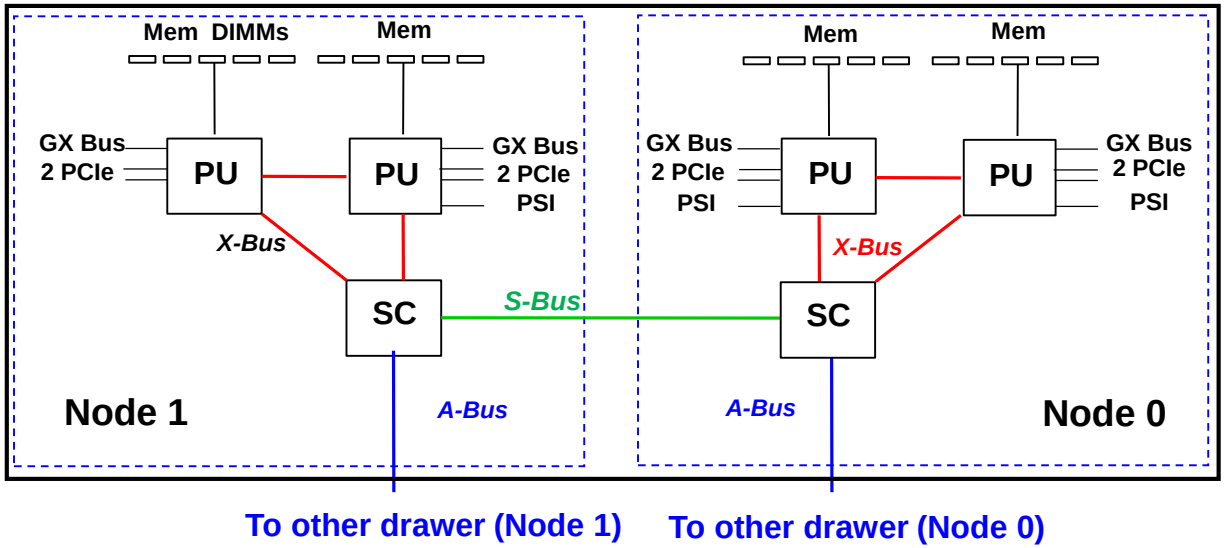

2 Drawers
4 Nodes

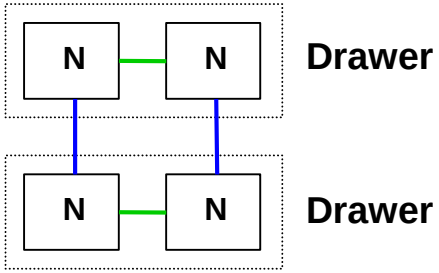

Four Physical Nodes

- **Chips**
  - Two PU chips per node
  - One SC chip per node
- **RAIM Memory**
  - Four Memory Controllers: Two per Drawer
  - Five DDR3 DIMM slots per Controller
  - Populated DIMM slots: 20, 30 or 40
  - SC and CP Chip Interconnects
- **SC and CP Chip Interconnects**
  - X-bus: SC L4 and CP L3s to each other (same node)
  - S-bus: SC L4 to SC L4 in the same drawer
  - A-bus: SC L4 to SC L4 in the remote drawer

Note: Both z114 and zBC12 have one node per CPC drawer with two PU SCMs, five DIMMs per PU SCM and one System Control SCM. The PU SCM L3 Caches are not connected directly. L4 cache design is inclusive of the L3s.

z13s N20 CPC Drawer Cache Hierarchy

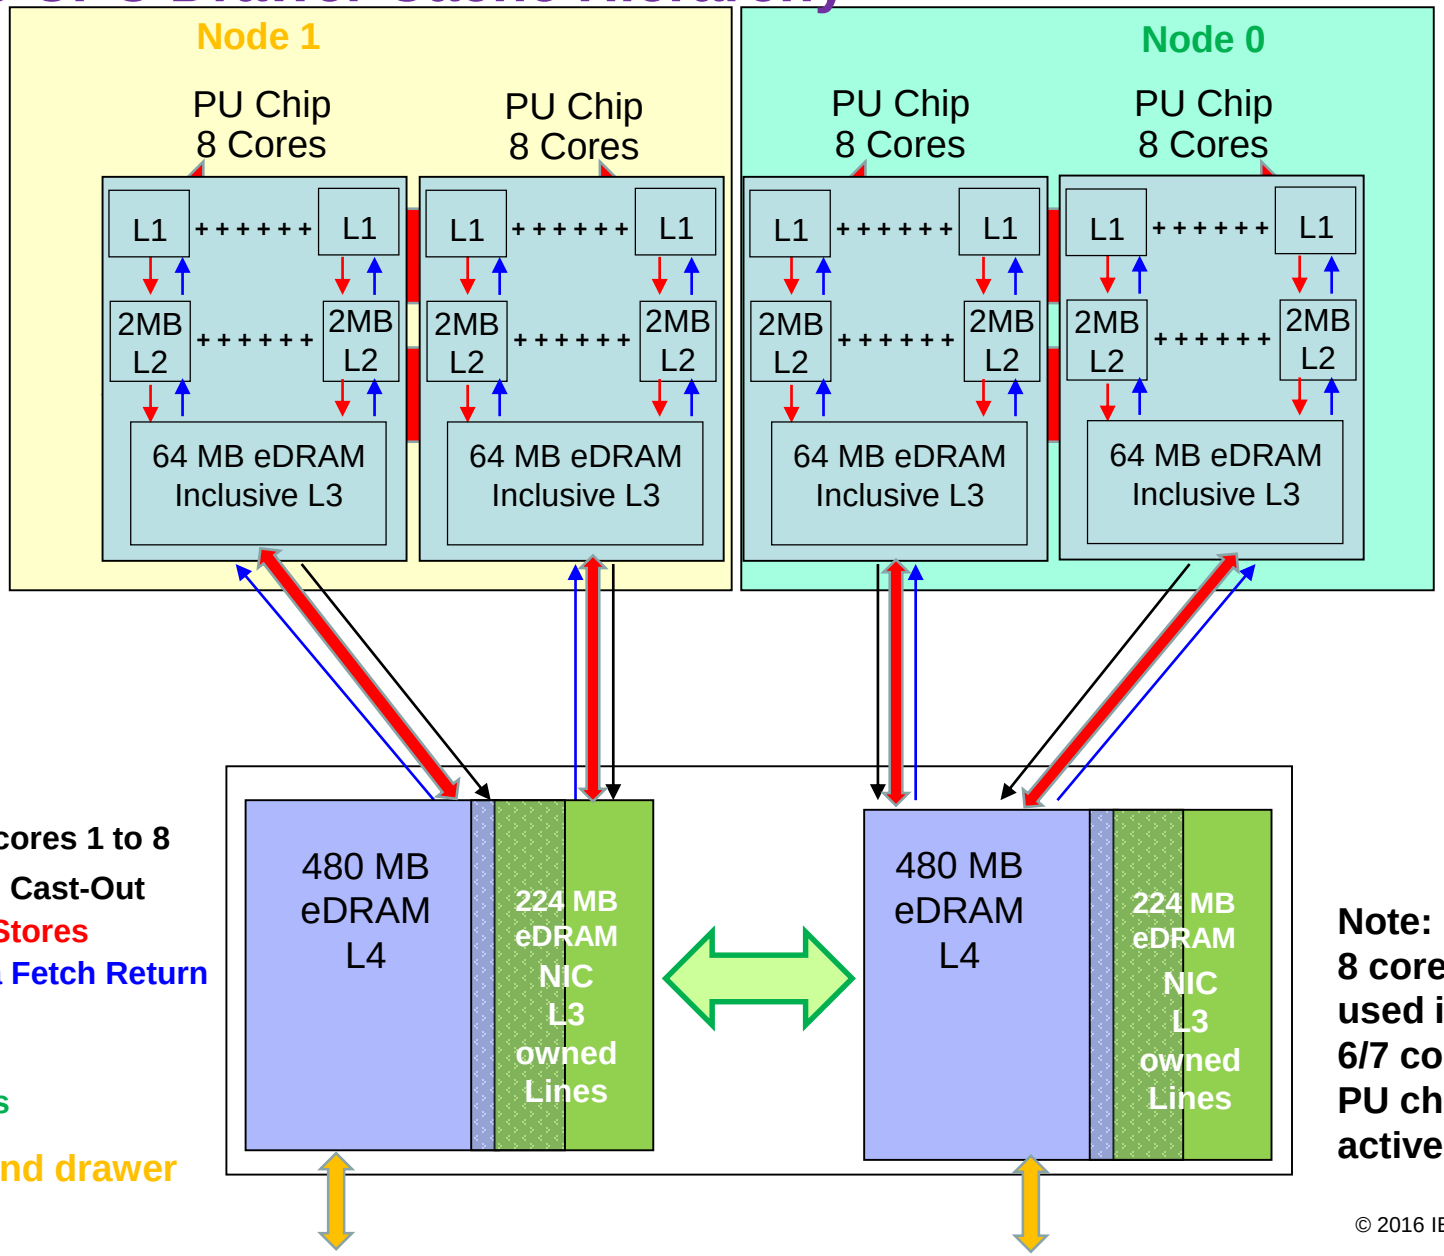

Note: Not all 8 cores are used in z13s. 6/7 cores per PU chip are active.

z13s Processor Unit Allocation and Usage (2:1 zIIP to CP ratio)

Model*	Drawers /PUs	CPs	IFLs uIFLs	zIIPs	ICFs	Std SAPs	Optional SAPs	Std. Spares	IFP
N10	1/13	0-6	0-10	0-6	0-10	2	0-2	0	1
N20	1/26	0-6	0-20	0-12	0-20	3	0-3	2	1
N20	2/26	0-6	0-20	0-12	0-20	3	0-3	2	1

*** See speaker notes**

z13s N20 model is a one- or two- drawer system with same processor feature counts for both configurations.
N20 - second drawer is added when additional fanouts or more than 2TB memory are needed;
The maximum number of logical ICFs or logical CPs supported in a CF logical partition is 16
The integrated firmware processor (IFP) is used for native PCIe I/O support functions
Upgrades from N10 to N20 and N20(1) to N20(2) are disruptive
SMT is supported with processor type IFL, zIIP.

What can you do with 8X the memory?

Improve ease of use and performance!

- **Accommodate growing workloads without changes to applications**
  Helps to reduce CPU usage by avoiding I/O
- **Reduce need for admin to fine tune memory definitions**
  Leverage tuning capabilities in Linux, DB2®, IMS™ and CICS®

Typical Client Use Cases:

- **More memory can help free up cycles to build richer transactions for improved user experience (i.e. fraud detection)**
- **Keeping the entire Cognos® Dynamic Cubes application online in memory allows for faster decision making**
- **Large memory allows taking advantage of larger Java heaps without an increase in paging**
- **Additional memory improves IBM MQ® V8 ability to manage the increasing messaging volumes generated by mobile and cloud apps**

* **Redbook: Benefits of Configuring More Memory in the z/OS Software Stack (REDP-5238-00)**

DB2: Use more memory to cache more data to reduce response time, increase transaction rates, and reduce CPU usage

*Data movement rates at least **20%*** faster for IBM MQ V8*

*Even small amounts of additional memory can have huge benefits improving response time **18%****

z13s Model N10 Memory

N10 Physical Memory RAIM GB Addressable Memory GB	Client GB	Increment GB
32 RAIM (2 x 16GB) - 128 for HSA (40) + Client (88)	64 72 80 88	8
64 RAIM (2 x 32GB) - 256 for HSA (40) + Client (216)	120 152 184 216	32
128 RAIM (2 x 64GB) - 384 for HSA (40) + Client (344)	248 280 312 344	32
128 RAIM (2 x 64GB) - 512 for HSA (40) + Client (472)	408 472	64
256 RAIM (2 x 128GB) - 640 for HSA (40) + Client (600)	536 600	64
256 RAIM (2 x 128GB) - 1024 for HSA (40) + Client (984)	728 856 984	128

- **z13s has the same RAIM Memory infrastructure as zBC12**
- **Minimum client memory is 64 GB**
- **HSA* is 40 GB (vs. 16 GB on zBC12)**

- **Memory upgrades that require DIMM changes are disruptive.**
- **Plan Ahead Memory can be added to eliminate disruption.**
- **On both models, memory upgrades within each row (same color) are concurrent without adding Plan Ahead Memory.**

*** See Speaker Notes**

z13s Model N20 Memory

- z13s has the same RAIM Memory infrastructure as zBC12
- Minimum client memory is 64 GB
- HSA* is 40GB (vs. 16 GB on zBC12)
- Second drawer required for memory configuration >2008GB

N20 two drawers Physical Memory RAIM GB Addressable Memory GB	Client GB	Increment GB
576 RAIM (4 x 128GB + 4 x 16GB) + 2304 for HSA (40) + Client (2264)	2264	256
640 RAIM (4 x 128 GB + 4 x 32GB) + 2560 for HSA (40) + Client (2520)	2520	256
768 RAIM (4 x 128GB + 4 x 64GB) + 3072 for HSA (40) + Client (3032)	2776 3032	256
1024 RAIM (8 x 128GB) + 4096 for HSA (40) + Client (4056)	3288 3544 3800 4056	256

N20 one drawer Physical Memory RAIM GB Addressable Memory GB	Client GB	Increment GB
32 RAIM (2 x 16GB) + 128 for HSA (40) + Client (88)	64 72 80 88	8
64 RAIM (4 x 16GB) + 256 for HSA (40) + Client (216)	120 152 184 216	32
128 RAIM (4 x 32GB) + 384 for HSA (40) + Client (344)	248 280 312 344	32
128 RAIM (4 x 32GB) + 512 for HSA (40) + Client (472)	408 472	64
256 RAIM (4 x 64GB) + 640 for HSA (40) + Client (600)	536 600	64
256 RAIM (4 x 64GB) + 1024 for HSA (40) + Client (984)	728 856 984	128
512 RAIM (4 x 128GB) + 2048 for HSA (40) + Client (2008)	1112 1240 1368 1496 1624 1752 1880 2008	128

Memory upgrades that require DIMM changes are disruptive. Plan Ahead Memory can be added to eliminate disruption. On both models, memory upgrades within each row (same color) are concurrent without adding Plan Ahead Memory.

*** See Speaker Notes**

What is Simultaneous Multithreading (SMT)?

- Prior generations' z Systems CPUs support a single instruction stream
  - z Systems workloads tend to receive a nontrivial number of cache misses
  - CPU generally unproductive while resolving cache miss
- z13/z13s SMT makes PU Core productive during cache misses and fills other pipeline gaps
  - z13/z13s supports two way SMT (two instruction streams [threads]) per core
    - Each thread has its own unique state information (Registers, PSW, etc.)
    - Cannot necessarily execute instructions instantly and must compete and win the use of desired core resources shared between threads
    - z13/z13s insures that one thread can't lock out the other
  - Current z13/z13s implementation allows following engine types to run in SMT mode
    - zIIPs under z/OS
    - IFLs under z/VM
- READY TO RUN Threads share core
  - Threads NOT READY TO RUN still unproductive while resolving cache miss
  - Core resources are productive when either READY TO RUN thread is executing

SMT Value Example

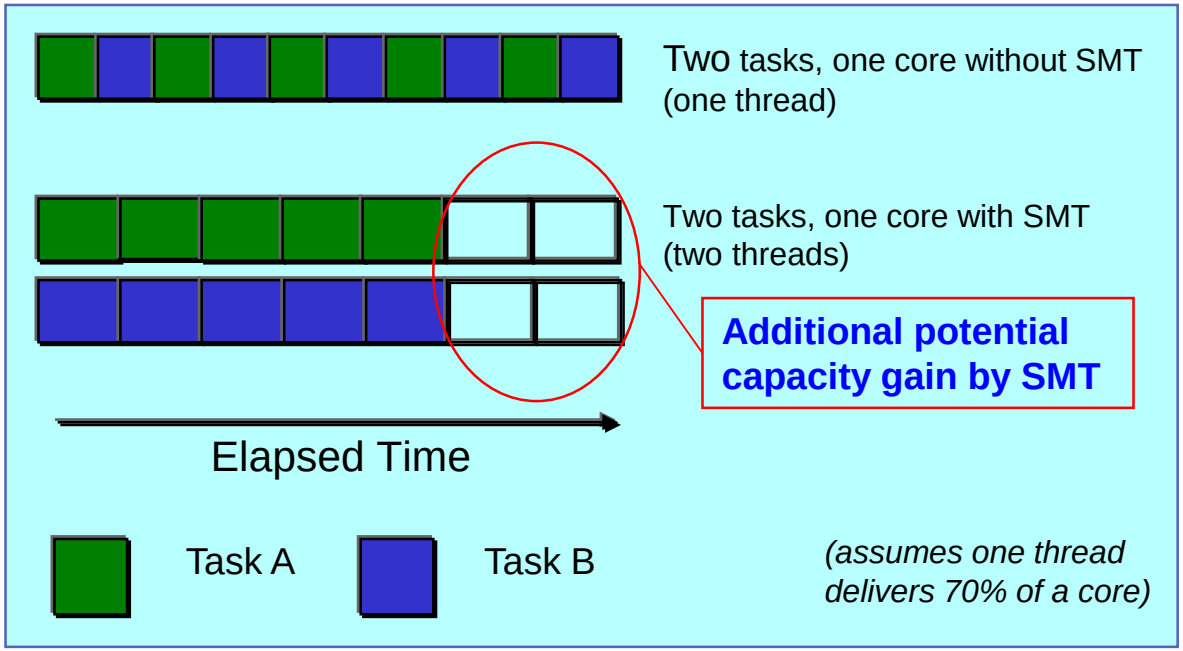

Driver:
The cloud compute capacity required to serve dynamic workloads for hypervisors, operating systems, and applications.

SMT on z13s provides significant throughput improvement with real-time measurements and repeatable metrics for capacity.

Simultaneous Multithreading – The Technology

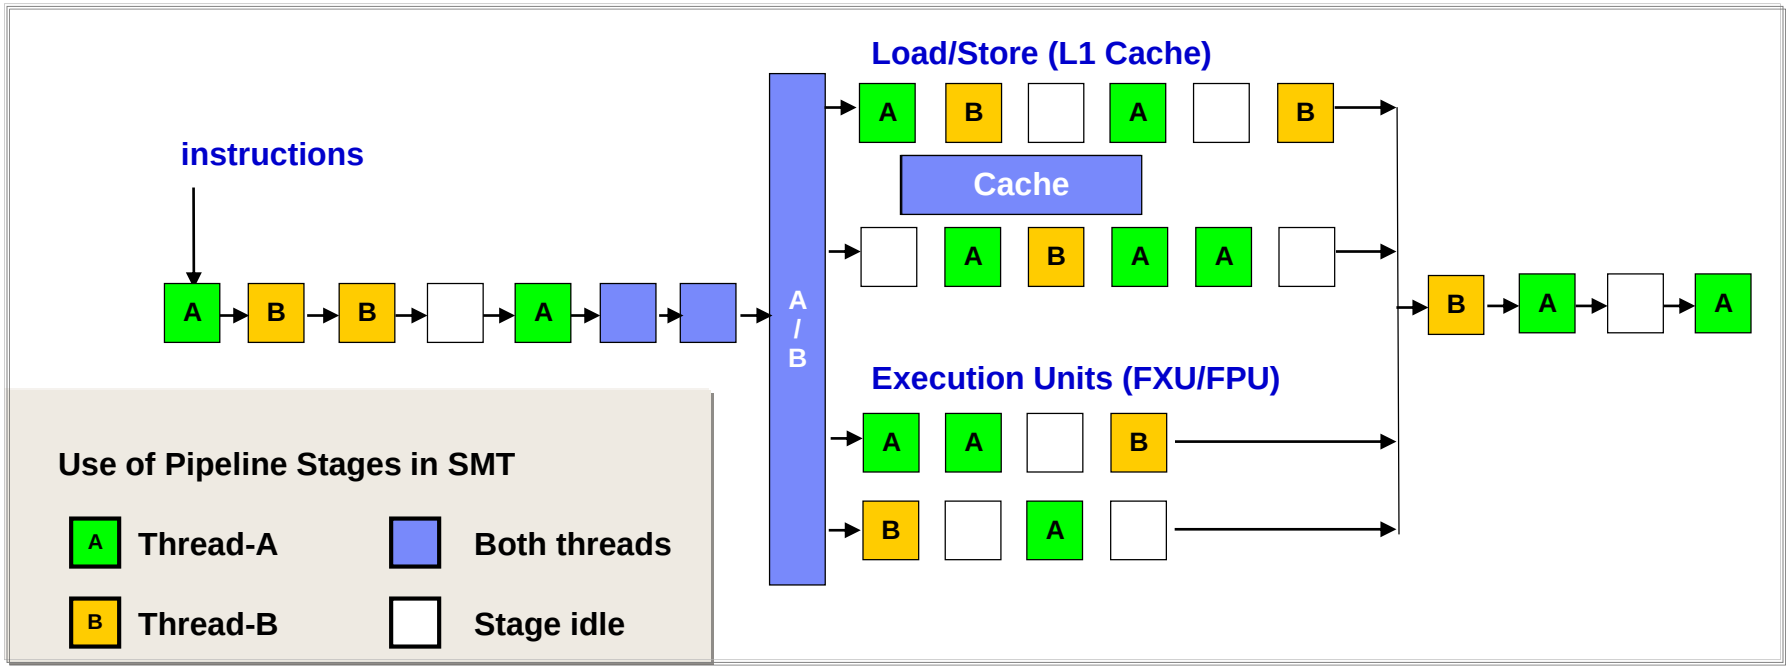

- Simultaneous Multithreading (SMT) technology
  - multiple threads inject instructions into the same core concurrently
  - More efficient use of the core hardware
- Increases overall throughput per core when SMT is active
  - Amount that increase, varies widely with workload – typically 1.x~1.y >1
  - Individual thread runs slower than on a single-threaded core

Active threads share core resources:

In space: data and instruction caches, TLBs, branch history tables, etc.

Cache misses provide opportunities for other thread

Machine ensures fairness between threads

In time: pipeline slots, execution units, address translator, etc.

E.g., take interruptions, start I/O, load wait PSW, signal other threads

Key: Layered SMT Exploitation on z Systems

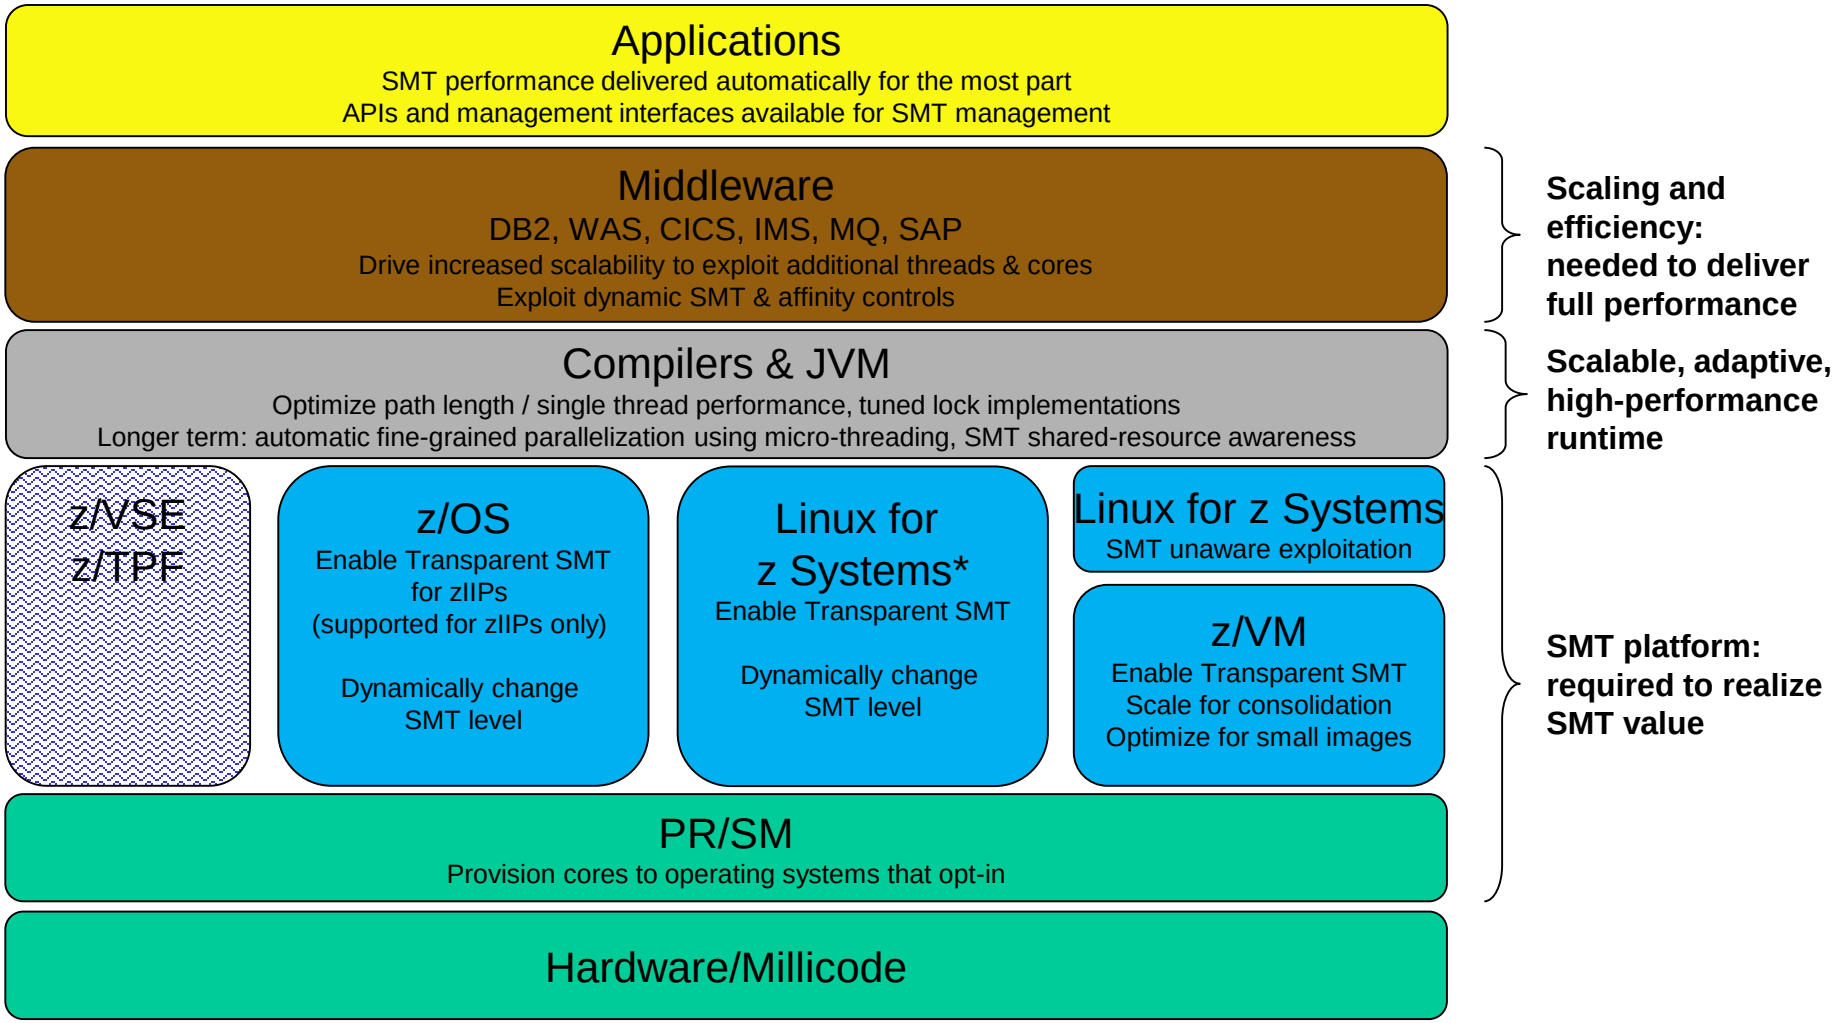

(*) IBM is working with Linux Distribution partners to support new functions and features.

Simultaneous Multithreading (SMT) on z13s

- SMT support one or two threads to execute on a zIIP or IFL.
- Capacity improvement is variable depending on workload. For AVERAGE workloads the estimated capacity of a z13s zIIP/IFL with exploitation of the SMT option is:
  - A z13s zIIP is up to **68%** greater than a zBC12 zIIP
  - A z13s IFL is up to **61%** greater than a zBC12 IFL
  - A z13s zIIP is up to **129%** greater than a z114 zIIP
  - A z13s IFL is up to **119%** greater than a z114 IFL
- SMT exploitation:
  - z/VM V6.3 + PTFs, or z/VM 6.4 for IFLs (GA 4Q2016)
  - z/OS V2.1 + PTFs, or z/OS V2.2 in an LPAR for zIIPs
  - KVM for IBM z V1.1.1 (3/2016)
  - For recommended Linux distribution levels refer to: http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html
- 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 and MT-2 (multi thread) modes with operator commands.
- Notes: 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.

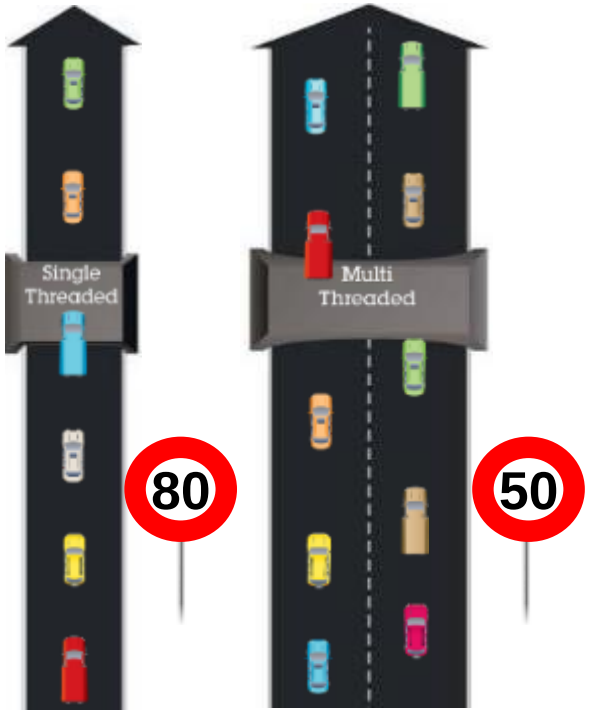

*Which approach is designed for the highest volume*** of traffic? Which road is faster?*

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

(*) 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 .

SIMD on z Systems Overview

- The z13/z13s superscalar processor has 32 vector registers and an instruction set architecture that includes a subset of 139 new instructions.
- This is known as SIMD (Single Instruction Multiple Data), and has been added to improve the efficiency of complex mathematical models and vector processing.
- These new instructions allow a larger number of operands to be processed with a single instruction.
- The SIMD instructions use the superscalar core to process operands in parallel.
- The set of SIMD instructions are a type of data parallel computing and vector processing that can decrease the amount of code and accelerate code that handles integer, string, character, and floating point data types.
- The SIMD instructions improve performance of complex mathematical models and allow integration of business transactions and analytic workloads on z Systems.
- Single Instruction Multiple Data (SIMD) is in fact a type of parallel computing that can accelerate code execution for integer, string, character, and floating point data types

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

Scalar
SINGLE INSTRUCTION, SINGLE DATA

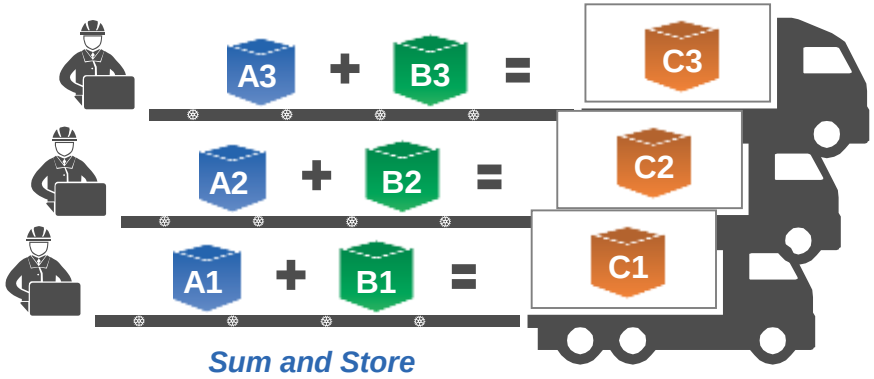

Instruction is performed for *every data element*

SIMD
SINGLE INSTRUCTION, MULTIPLE DATA

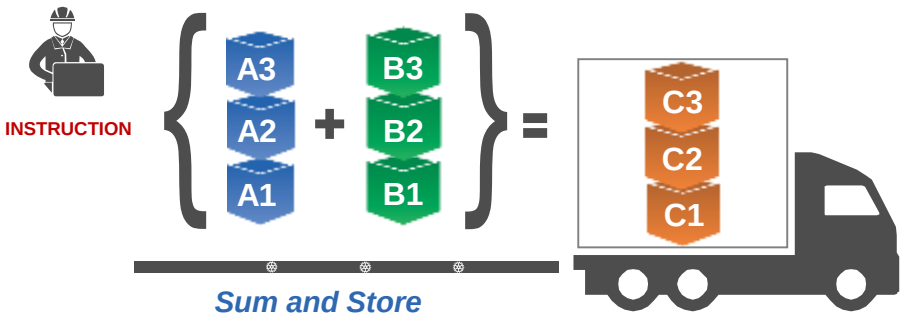

Perform instructions on *every element at once*

SIMD Hardware Accelerator

New SIMD FXU

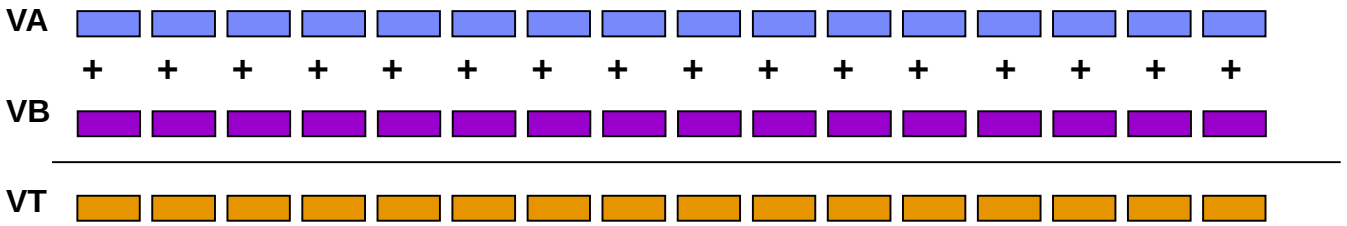

Three distinct data types

Integer	String	Binary Floating-Point
16xB, 8xHW, 4xW, 2xDW, 1xQW		DP only (2xDP)

Implementation

New string engine

Enhancing BFU for SIMD

Exploitation

String processing for Cognos, XMLSS in Cobol, PL1, Java, C/C++

Analytics workloads like ILOG, SPSS, ALGO

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

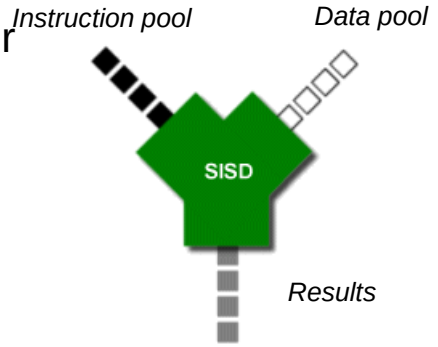

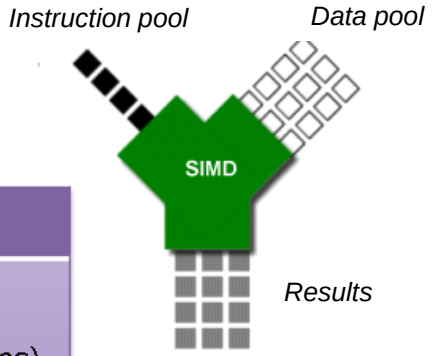

Workloads		
Java 8	C/C++Compiler built-ins for SIMD operations (z/OS and Linux on z Systems)	MASS & ATLAS Math Libraries (z/OS and Linux on z Systems)
SIMD Registers and Instruction Set		

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

z13s Upgrade Paths

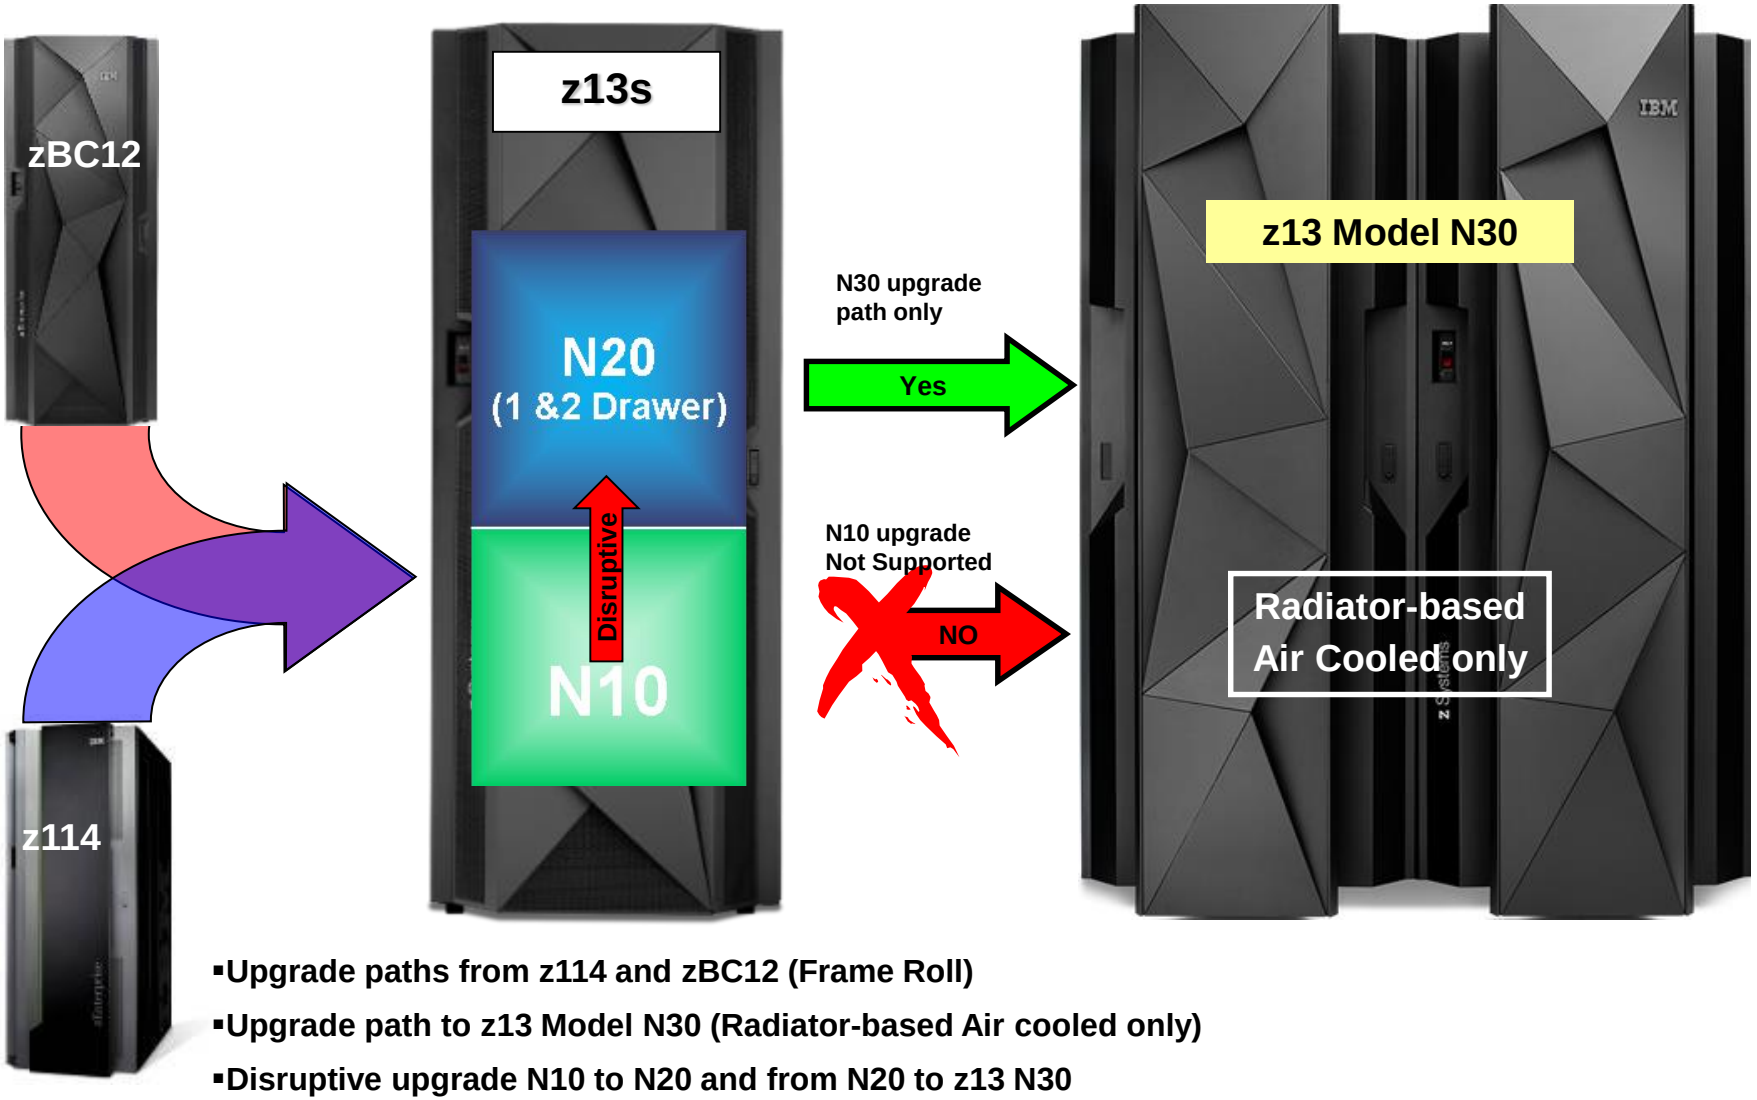

- Upgrade paths from z114 and zBC12 (Frame Roll)
- Upgrade path to z13 Model N30 (Radiator-based Air cooled only)
- Disruptive upgrade N10 to N20 and from N20 to z13 N30
- Disruptive Upgrade N20 (1 drawer) to N20 (2 drawer)

z13s Sub-capacity Processor Granularity

- The z13s has 26 CP capacity levels (26 x 6 = 156)
  - –Up to 6 CPs at any capacity level
    - •All CPs must be the same capacity level
- zAAPs are not available on z13s
- The ratio of zIIPs for each CP purchased is the same for CPs of any speed.
  - –2:1 zIIP to CP ratio – unchanged from zBC12
  - –All specialty engines run at full speed
  - –**Processor Value Unit (PVU) for IFL = 100**

Number of z13s CPs	Base Ratio	Ratio zBC12 to z13s
1 CP	zBC12 Z01	1.34
2 CPs	zBC12 Z02	1.38
3 CPs	zBC12 Z03	1.40
4 CPs	zBC12 Z04	1.42
5 CPs	zBC12 Z05	1.43
6 CPs	zBC12 Z06	1.44

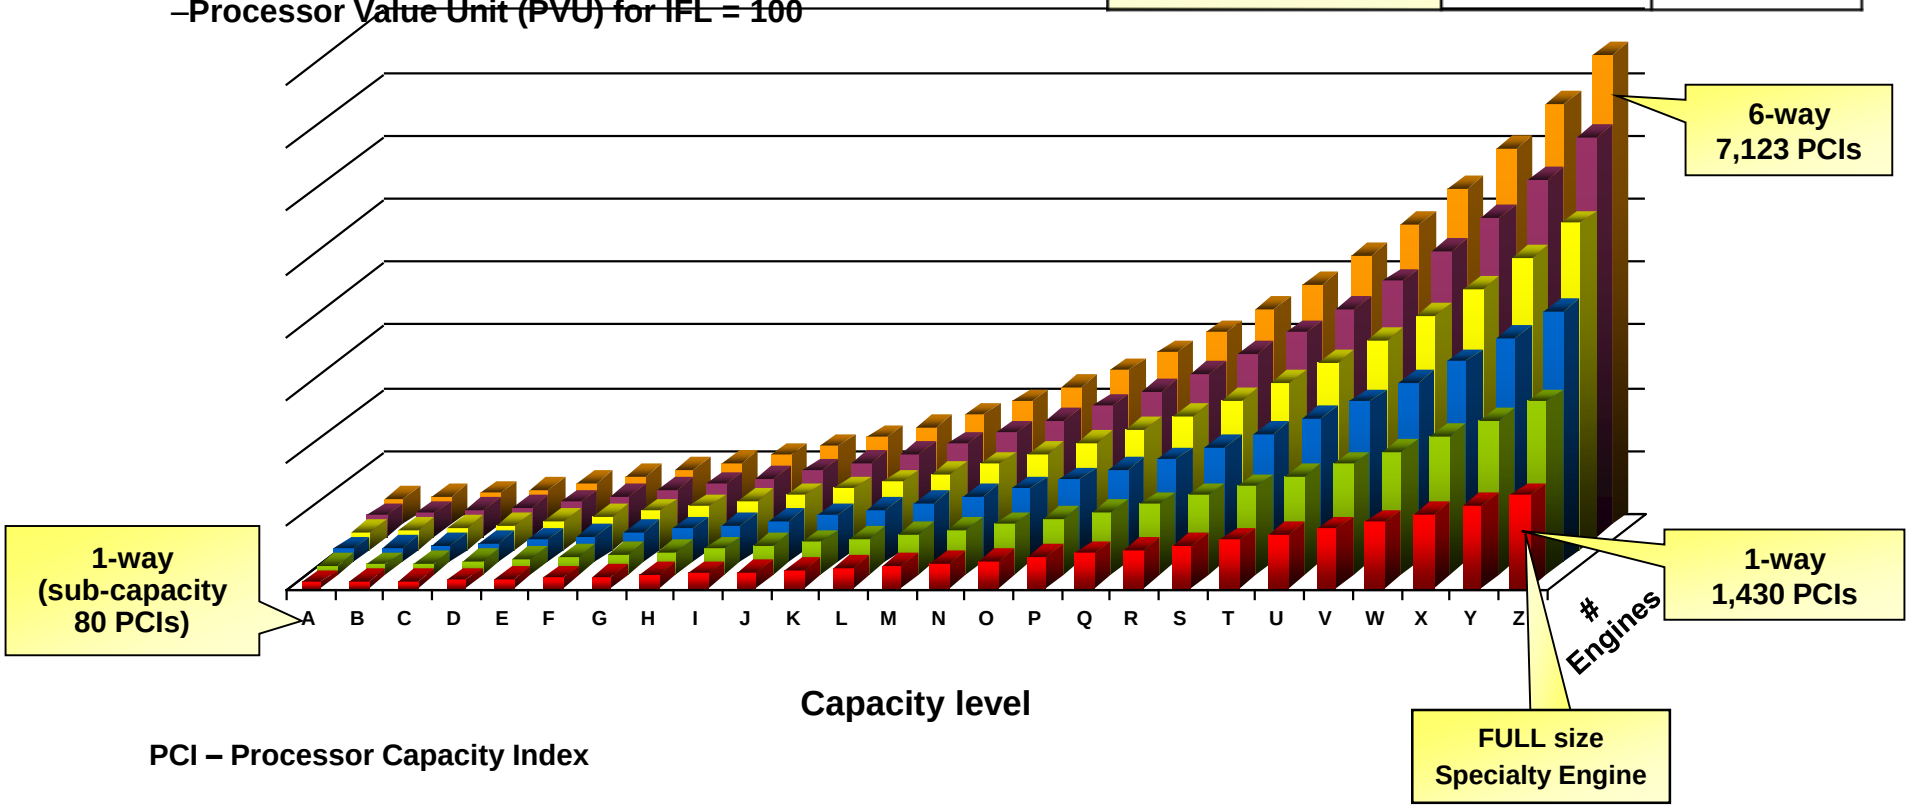

PCI – Processor Capacity Index

z13s LPAR Dynamic PU Reassignment

- PR/SM dynamic relocation of running processor cores to different physical core locations
  - CP, zIIP, IFL and ICF supported
  - Swap an active core to a core in a different PU chip in a different drawer or node
- Designed to optimize physical processor location for the current LPAR's logical processor configuration:
  - Better L3 and L4 cache reuse
  - Move processors to partition memory
- Triggers: Partition activation/deactivation, machine upgrades/downgrades, logical processors configured on/off
- Designed to provide the most benefit for:
  - Multiple drawer machines
  - Dedicated partitions and wide partitions with HiperDispatch active

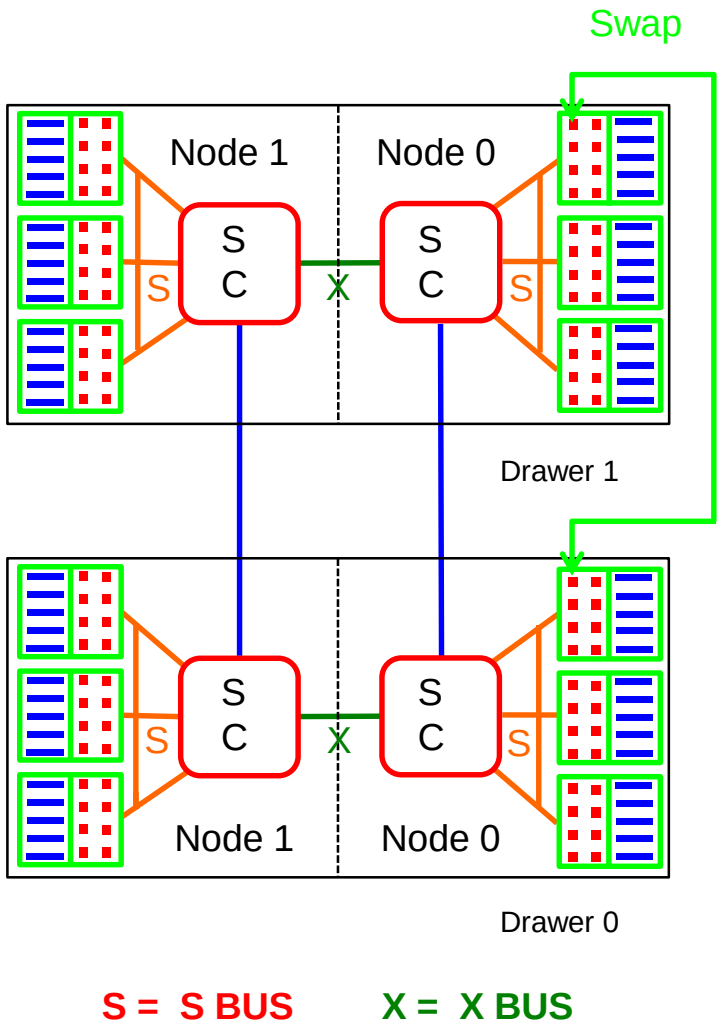

z Systems I/O Subsystem Internal Bus Interconnect Speeds


PCIe Gen3
z13/z13s

16 GBps


PCIe Gen2
zEC12/zBC12
z196/z114

8 GBps


InfiniBand
zEC12/zBC12
z10/z196/z114

6 GBps


eSTI
z9

2.7 GBps


STI
z990/z890

2 GBps

STI: Self-Timed Interconnect

z13s CPC Drawer I/O Fanout and FSP Locations

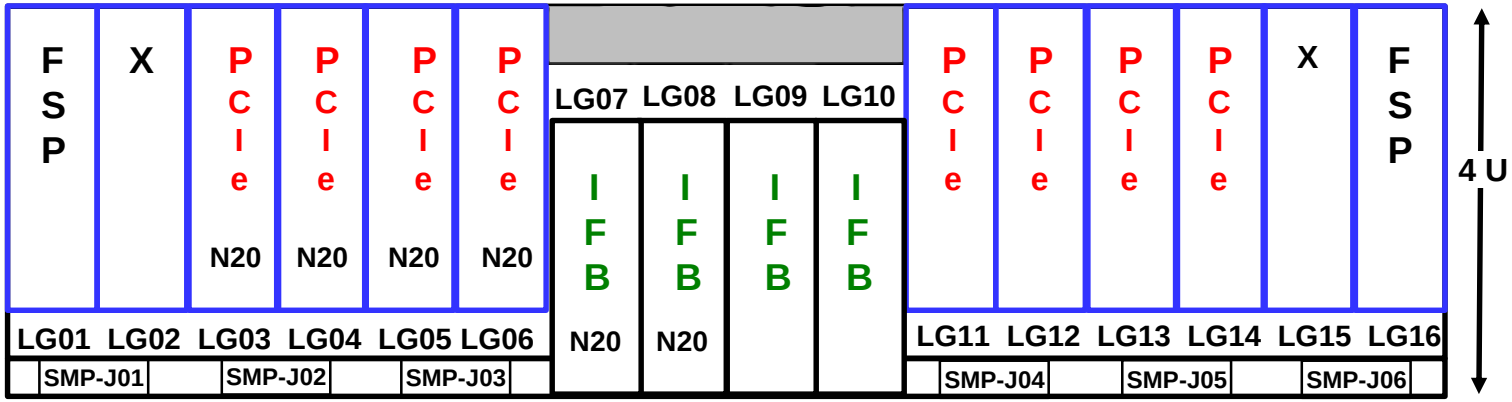

PCIe Fanout Slots: LG03 – LG06 (N20) and LG11 – LG14 (N10 and N20), can support:

- Up to 8 one-port PCIe 16 GBps I/O fanouts to support up to 8 domains in 32-slot PCIe I/O drawers
- Note: A zBC12 Model H06 with four two-port 8 GBps PCIe fanouts supports up to 8 domains in 32-slot PCIe I/O drawers; but **a z13s CPC drawer supports double the bandwidth to each domain**
- Up to 8 ICA (PCIe-SR) two-port coupling fanouts to support up to 16 8 GBps coupling links

IFB Fanout Slots, LG07 – LG08 (N20) and LG09 – LG10 (N10 and N20), can support:

- Up to four HCA3-O 12x InfiniBand coupling fanouts, 8 12x 6 GBps links – Two per fanout
- Up to four HCA3-O LR 1x InfiniBand coupling fanouts 16 1x 5 Gbps links – Four per fanout
- Note: A zBC12 Model H06 with 4 two-port HCA3-O 12x InfiniBand coupling fanouts can support 8 12x links
  A zBC12 Model H06 with 4 four-port HCA3-O LR 1x InfiniBand coupling fanouts can support 16 1x links
- Up to two two-port HCA2-C 6GBps I/O fanouts (2 8-slot I/O drawers) with two slots left

Slots LG01 and LG16 always have Flexible Support Processors (FSPs)

SMP-J01 to J06 connectors are for A-Bus cables to nodes in the other CPC drawer

PCIe I/O Features Support

Traditional z Systems PCIe I/O features:

- **FICON Express16S (shown in picture)**
- **FICON Express8S**
- **OSA-Express5S**
- **OSA-Express4S**

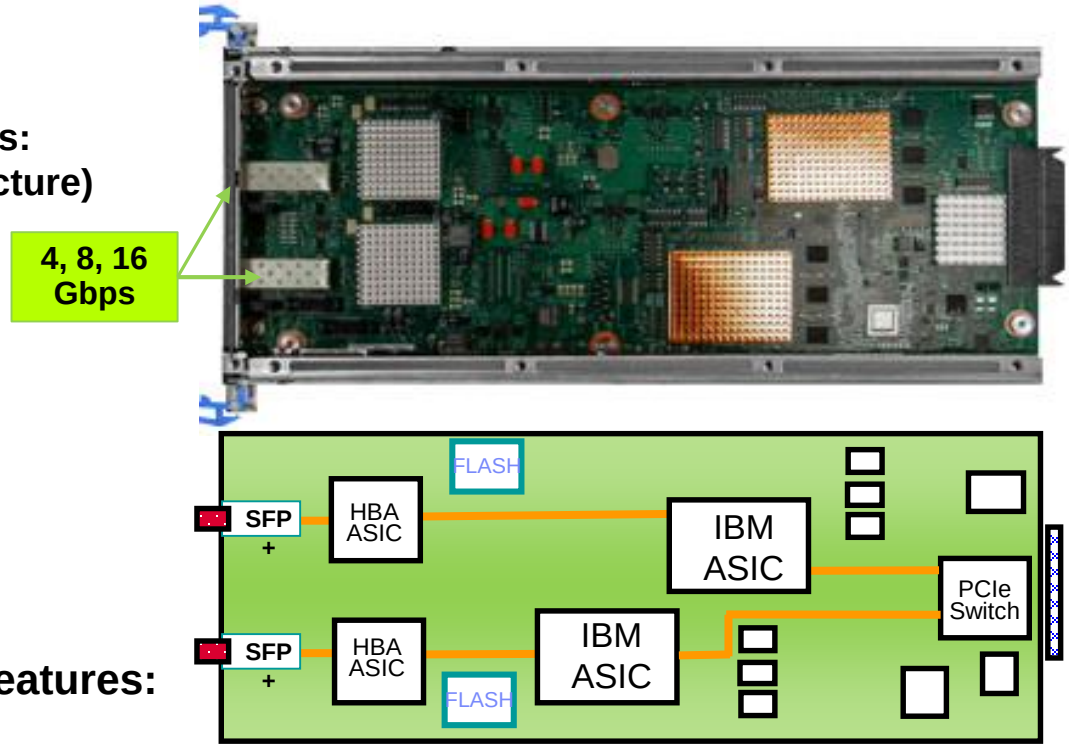

- **Traditional z Systems PCIe I/O features:**
  - One IBM ASIC per Channel/PCHID
  - Definition and LPAR Assignment
    - HCD/IOCP CHPID definition **or**
    - Firmware definition outside HCD/IOCP is possible for some.
      - For example: ***Crypto Express4S*** *is not defined as a CHPID*
  - Virtualization and support by Channel Subsystem LIC on System Assist Processors (SAPs)

Native PCIe I/O Features Support

Native PCIe I/O features supported:
zEDC Express
RoCE Express (shown in picture)

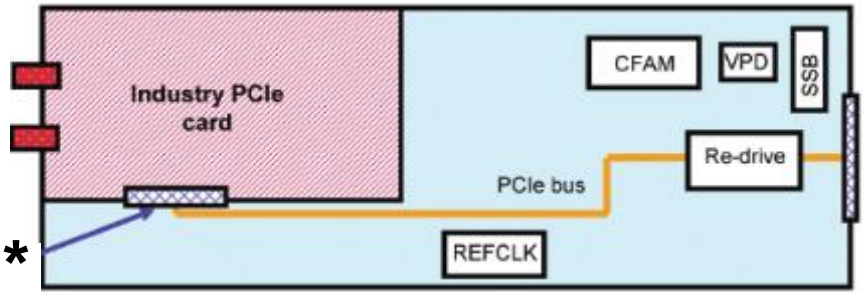

- **Native PCIe I/O features**
  - z System ASIC role moved to the new z System I/O Controller (zIOC) function moved to the z13/z13s PU chip
  - Definition and LPAR Assignment
    - HCD/IOCP FUNCTION definition similar to CHPID definition but with different rules **or**
    - Firmware definition outside HCD/IOCP is possible for some. For example: ***Flash Express*** *and* ***Crypto Express5S*** *are not defined with FUNCTIONs*
  - Virtualization and support by the Redundancy Group LIC running on the Integrated Firmware Processor (IFP)
    (***Note:*** *NOT applicable to Flash Express or Crypto Express5S*)

*PCIe Adapter Connector

z13s "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)
  - FICON Express8S (SX and LX, 2 SFPs, 2 CHPIDs)
  - OSA-Express5S
    - 10 GbE LR and SR (1 SFP, 1 CHPID)
    - GbE SX, LX, and 1000BASE-T (2 SFPs, 1 CHPID)
  - 10 GbE RoCE Express (2 supported SR ports)
  - zEDC Express
  - Crypto Express5S
  - Flash Express (Technology Refresh)

PCIe I/O drawer

32 I/O slots

- **Integrated Coupling Adapter (ICA) Fanout**
  - ICA SR two 8 GBps PCIe Gen3 Coupling Link
- **InfiniBand Coupling Feature Fanouts**
  - HCA3-O two 12x 6GBps InfiniBand DDR Coupling Links
  - HCA3-O LR four 1x 5Gbps InfiniBand DDR or SDR Coupling Links

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.

z13s "Carry Forward" I/O Features Supported

Carry Forward Features

- **Features – PCIe I/O drawer**
  - FICON Express8S (SX and LX, 2 SFPs, 2 CHPIDs)
  - OSA-Express5S (All)
  - OSA-Express4S (All)
  - 10 GbE RoCE Express (Both ports supported on z13s)
  - zEDC Express
  - Flash Express
  - Not Supported: Crypto Express4S

PCIe I/O drawer 32 I/O slots

- **Features – I/O drawer (No MES adds)**
  - FICON Express8 (SX and LX, 4 SFPs, 4 CHPIDs)
  - Maximum of One 8 Slot I/O Drawer
  - Not Supported: ESCON, FICON Express4, OSA-Express3, ISC-3, Crypto Express3,

I/O drawer 8 I/O slots

- **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
  - NOT Supported: HCA2-O 12x, HCA2-O LR 1x InfiniBand Coupling Links

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.

New FICON Function on z13s and z13 GA2

- **Export/Import Physical WWPNs for FCP channels without NPIV support (point to point)**
  - Designed to simplify migration to a new-build z13 or z13s
  - Complements preservation of virtual (NPIV) WWPNs at z13 GA for switched FCP channels
- **Fibre Channel Read Diagnostic Parameters Extended Link Services (ELS)**
  - Supports link failure diagnostics and predictive analysis for optics, cables, ports (except E ports)
- **Dynamic Channel path Management (DCM) support for FICON paths through cascaded FICON Directors**
  - DCM was previously offered only for FICON paths when the channel and the control unit port connected to the same dynamic switch. Adding support for cascaded switches is expected to help improve resiliency for clients using replication technologies such as Global Mirror (XRC).

FICON Express16S – SX and 10KM

For FICON, zHPF, and FCP environments
- CHPID types: FC and FCP
  - Two PCHIDs/CHPIDs

Auto-negotiates to 4, 8, or 16 Gbps
- 2Gbps connectivity NOT supported
- FICON Express8S will be available to order for 2Gbps connectivity

Increased I/O Devices (subchannels) per channel for all FICON features:
- TYPE=FC: Increased from 24k to 32k to support more base and alias devices

Increased bandwidth compared to FICON Express8S

10KM LX - 9 micron single mode fiber
- Unrepeated distance - 10 kilometers (6.2 miles)
- Receiving device must also be LX

SX - 50 or 62.5 micron multimode fiber
- Distance variable with link data rate and fiber type
- Receiving device must also be SX

Two channels of LX or SX (no mix)

Small form factor pluggable (SFP) optics
- Concurrent repair/replace action for each SFP

FC 0409 – 10KM LX, FC 0410 – SX

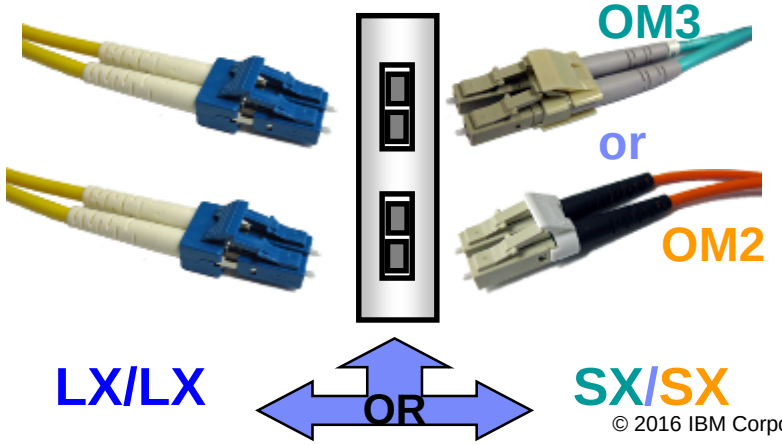

FCP Performance* for z13s

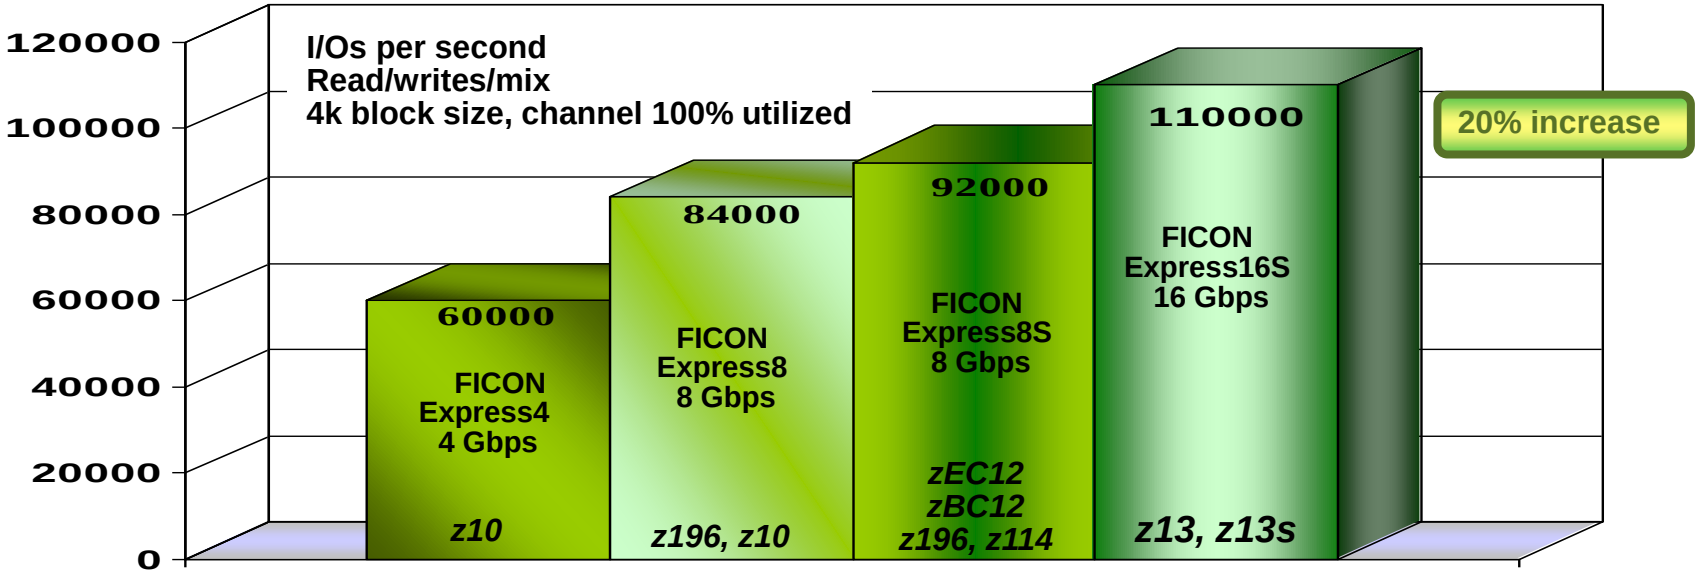

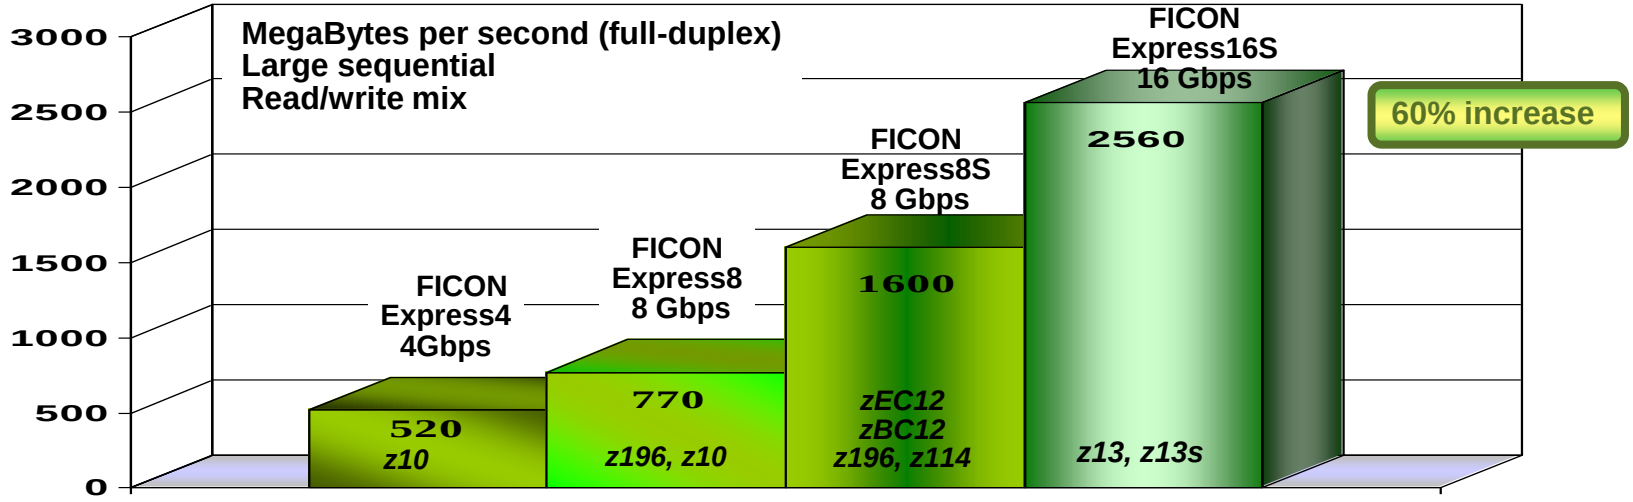

*This performance data was measured in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload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

New Networking Functionality on z13 and z13s

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 supportiing OSD channels among multiple z/VM V6.3 images
- See the “Hypervisors – z/VM” Section for details

10 GbE RoCE Express Virtualization Support (z13s from z13 GA)

- Designed to enable both ports on a RoCE Express feature and to allow sharing of each RoCE Express feature by up to 31 logical partitions

*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.

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

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

Clustered Systems: ***Multi-Tier Application Solution***

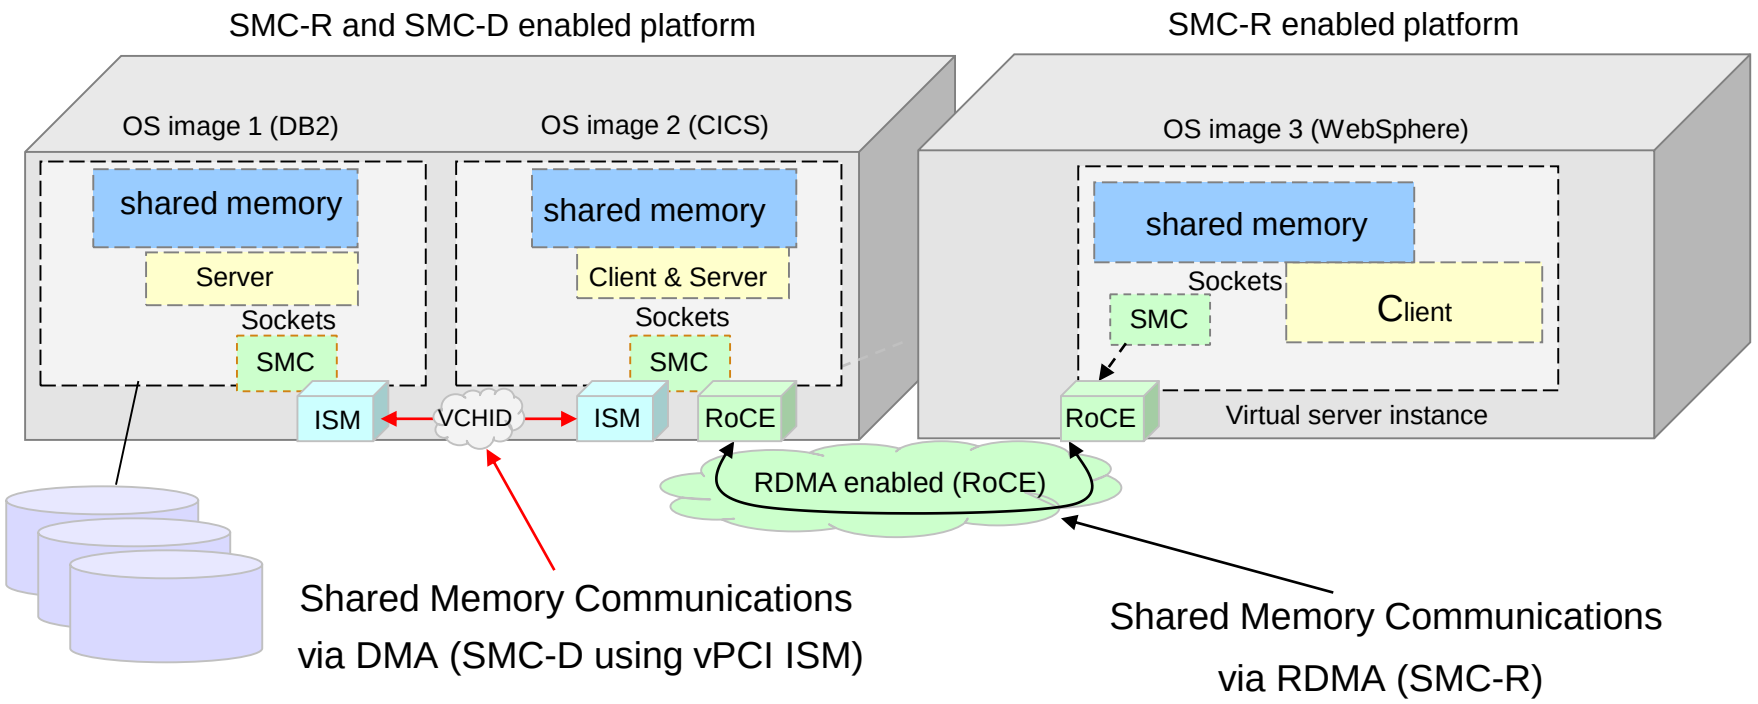

Both forms of SMC combine to provide a highly optimized solution. Shared Memory Communications: via z Systems PCI architecture:

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

IBM z Appliance Container Infrastructure (zACI): What is it ?

- An appliance is a bootable system image that contains all of the layers necessary to provide a specific set of services or functions
  - Operating System, middleware, applications…
- Appliances may be implemented as firmware or software, depending on the environment in which the appliance runs
- Several new appliances are planned for z Systems that are based on a common framework called zAppliance Container Infrastructure (zACI)

Why ?

- The zACI framework simplifies the process a team has to apply to create an appliance, and enforces a common set of behaviors for operations that all appliances have to perform.
- zACI framework provides a set of utilities to implement the common functions that all appliances need (FFDC, network setup, appliance configuration, ...)

z Systems Appliances

» A z Systems Appliance is an integration of operating system, middleware and software components that work autonomously and provide core services and infrastructures focusing on consumability and security «

» The z Appliance Container Infrastructure (zACI) provides the base infrastructure needed to create appliances: Operating System, middleware, SDK and firmware support. «

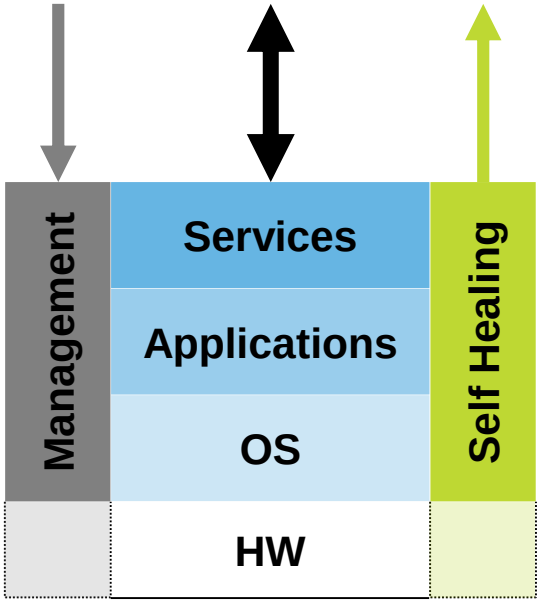

- Encapsulated Operating Systems
- Services provided via Remote APIs and web interfaces
- Embedded monitoring and self-healing
- End-2-End appliance tamper protection
- Protected Intellectual Property of appliance components
- Tested/Qualified by provider for a specific use case

Advantages of this Appliance Approach Based on zACI

Customer

- Receives a complete solution
  - No internals are exposed
  - Clear support structure
- One-stop shop for deployment and management of solutions
- Appliances and their data cannot be altered
  - Protected by firmware
- Appliance data is protected against unauthorized access
- Reliable core infrastructure

IBM

- Re-use of IBM and partner assets
- On-premise protected solutions
  - Encrypted images protect IBM recipes
- Faster time to market
  - Prototypes created in days
  - SW appliances de-coupled from hardware release and service schedules
- Enables additional solutions
  - Supports partners porting solutions to z Systems

zACI Software Appliances

Deployment in 5 Steps

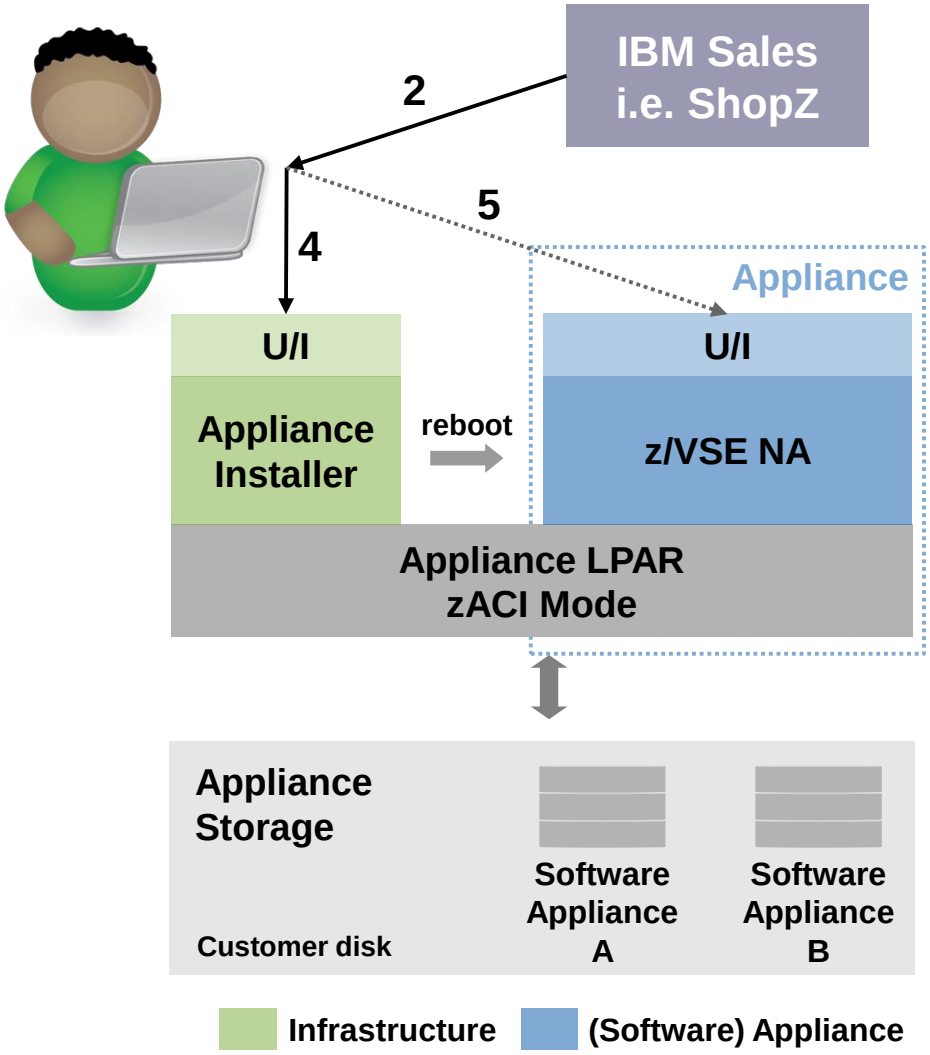

1) Buy a Software Appliance (e.g. z/VSE Network Appliance)
2) Download the installation image
3) Create and activate an appliance (zACI) LPAR
4) Deploy appliance using Appliance Installer
5) Configure and use appliance through REST API or web U/I

Use Cases

- **Firmware IBM zAware will remain available for z13 GA2 and z13s**
  - Will run in zACI LPAR on z13 GA2 and z13s
  - Existing zAware LPAR (on z13) will be converted to zACI
  - No changes required to zAware
  - See the “HMC and SE” and the “zAware” sections of this presentation

- **z/VSE Network Appliance**
  - Planned Availability in June 30, 2016*

- **More software appliances will be available in the future**

*** All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.**

IBM zAware V2 support for z/OS and Linux on z Systems

IBM zAware Partition
Shipped as firmware with z13

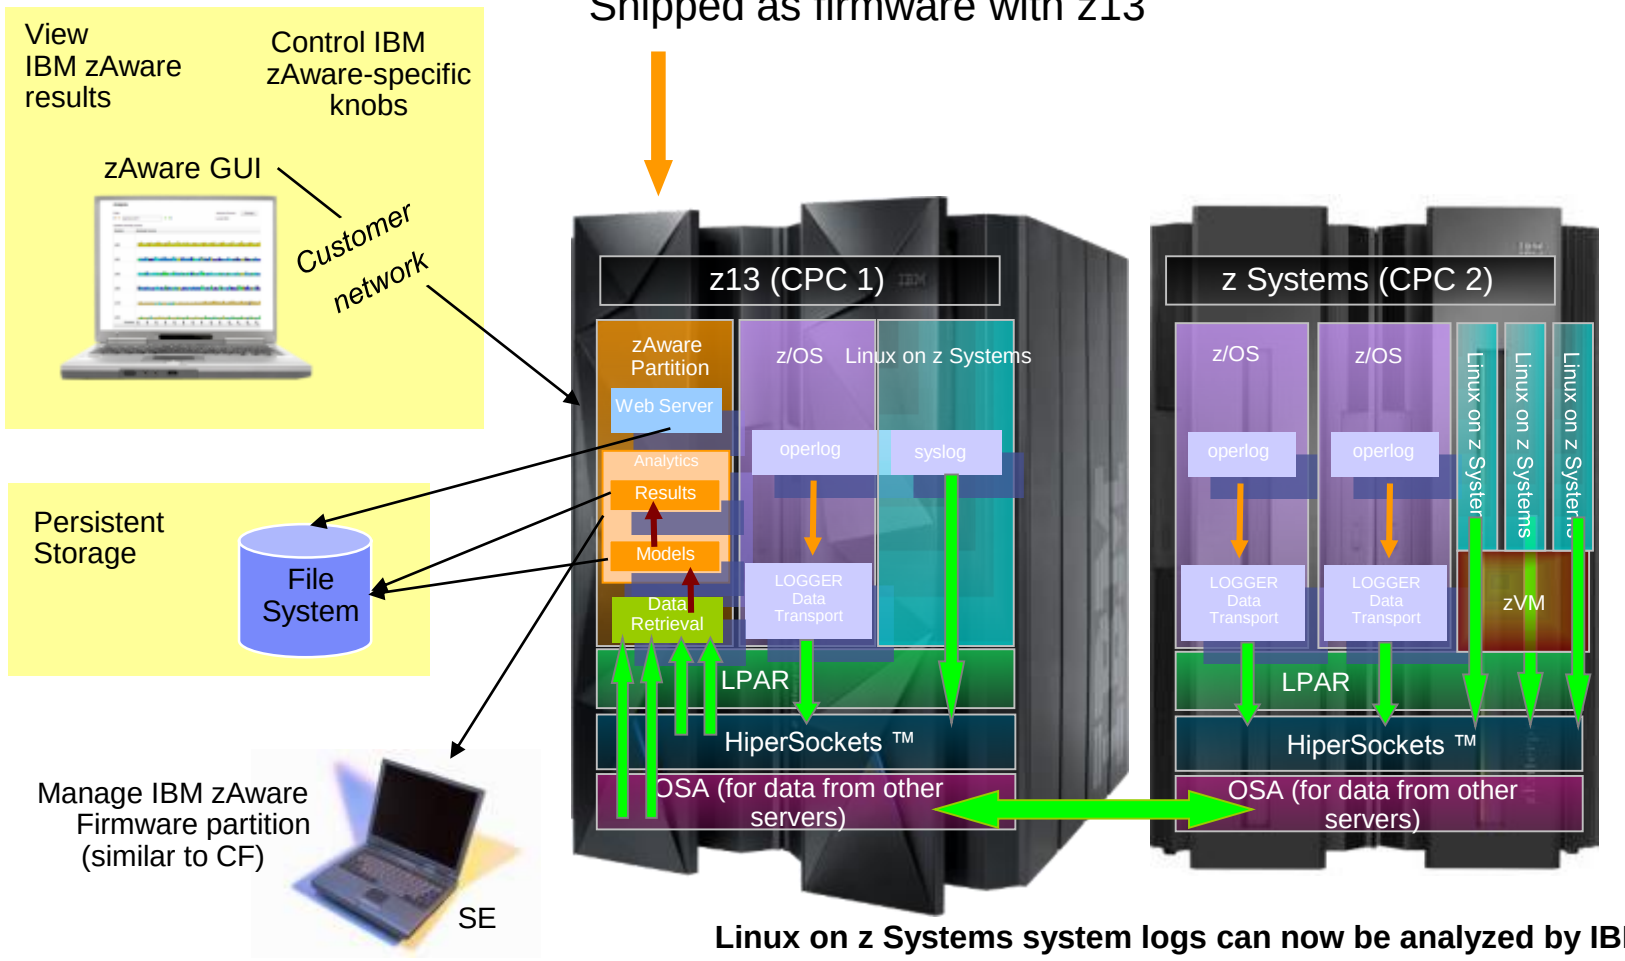

Linux on z Systems system logs can now be analyzed by IBM zAware
Upgraded analytics engine for better results on z/OS analysis
Upgraded internal database for improved RAS
Completely rewritten UI, including heat map views

Hypervisors and Virtualization for z Systems

PR/SM-LPARs

- ***Virtualization capabilities built into the system***
- PR/SM manages and virtualizes all the installed and enabled system resources as a single large SMP system
- ***Full sharing of the installed resources*** with high efficiency and very low overhead
- ***High scalability*** with support for up to 40 (for z13s) or 85 (for z13) logical partitions
- IBM Dynamic Partition Manager ***simplifies management experience***
- Ensured ***workload separation*** based on highest EAL5+ security certification

z/VM v6.4 (preview)

IBM Wave for z/VM

- ***Enables extreme scalability, security and efficiency*** creating cost savings opportunities
- ***Ease Migration*** with upgrade in place infrastructure provides a seamless migration path from previous z/VM releases (z/VM 6.2 and z/VM 6.3) to the latest version
- ***Operational improvements*** by enhancing z/VM to provide ease of use
- ***Improved SCSI support*** for guest attachment of disk and other peripherals, and hypervisor attachment of disk drives
- IBM Wave for z/VM ***simplifies the management*** of virtual Linux servers from a single user interface
- Provides the foundation for cognitive computing on z Systems

KVM on z Systems v1.1

- ***Support new analytics workloads*** with Single Instruction Multiple Data (SIMD) for competitive advantage
- ***Deliver higher compute capacity*** with support for Simultaneous Multithreading (SMT) to meet new business requirements
- ***RAS support enhanced for problem determination and high availability*** setup to reduce down time and quickly react to business needs
- ***Secure and protect*** business data with Crypto exploitation

What is DPM - Product Vision

Dynamic Partition Manager has been designed to deliver dynamic z Systems hardware and virtual infrastructure management including integrated dynamic I/O management with general virtualization knowledge and minimal required mainframe knowledge.

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

Facilitate configuring and operating PR/SM LPARs in a way which is familiar to someone performing these actions on another platform.

Develop towards an adaptive user experience, adjusting to user roles, and reacting to conditions and state

Lay the foundation for a general z Systems user experience overhaul.

IBM Dynamic Partition Manager (DPM)

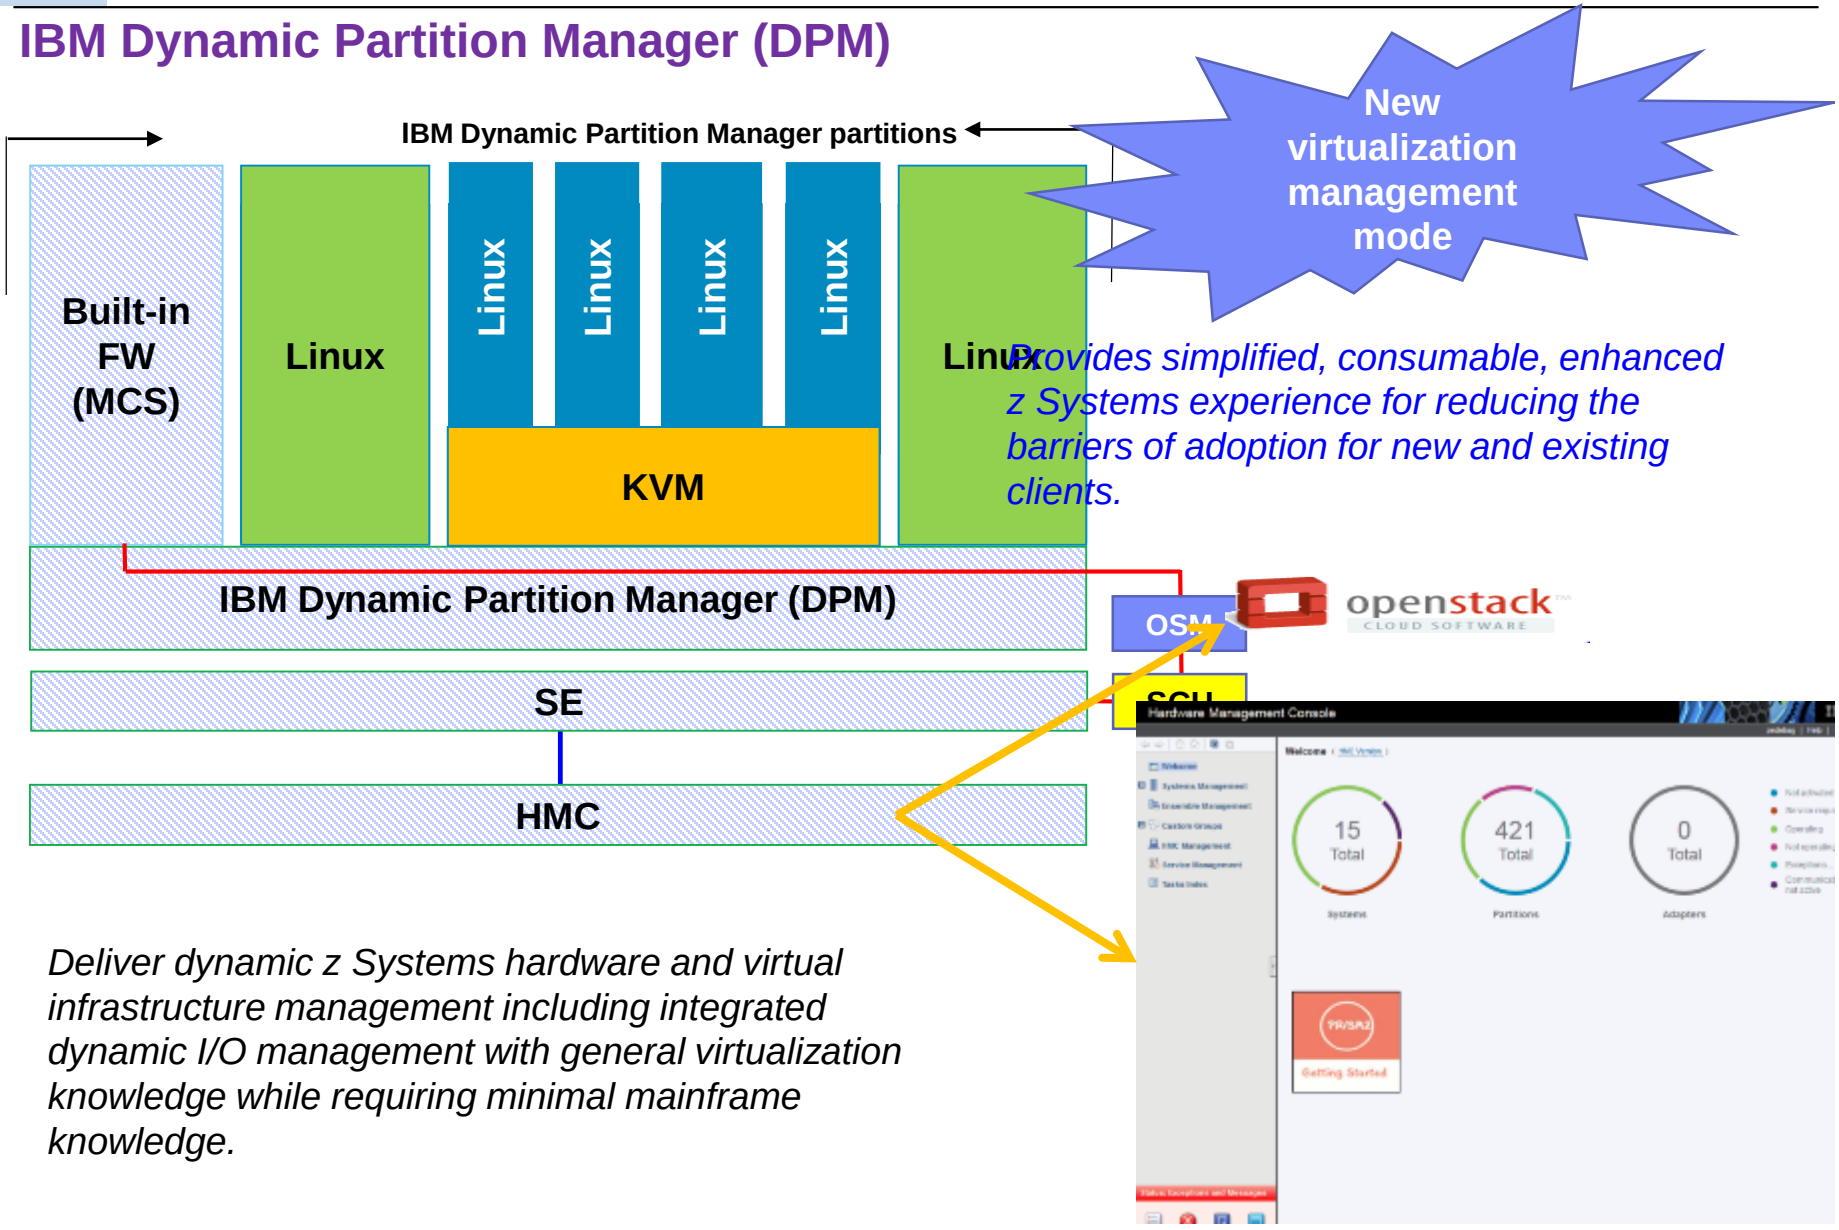

New virtualization management mode

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

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

What is GDPS Virtual Appliance*

- **Fully integrated Continuous Availability & Disaster Recovery solution for Linux on z Systems customers with no or little z/OS skills**
  - It is an image comprising an operating system, the application components, an appliance management layer which makes the image self-containing, and APIs / UIs for customization, administration, and operation tailored to the appliance function.
  - It improves both consumability and time-to-value for customers.

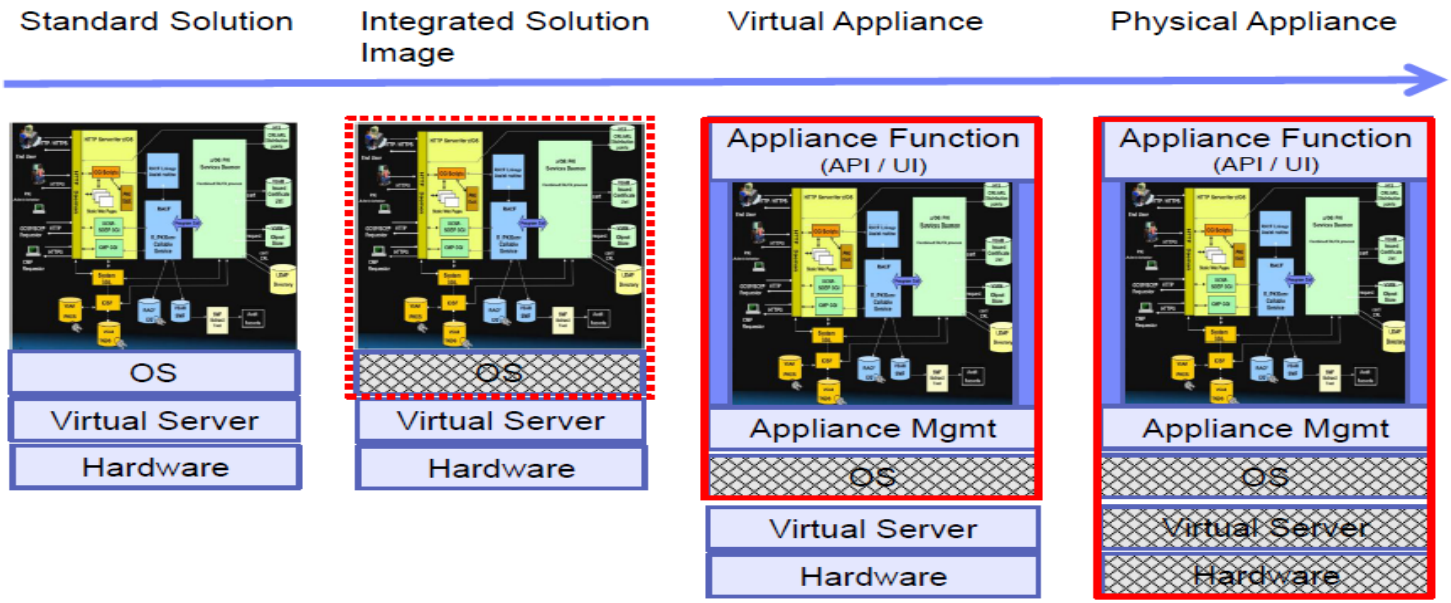

***See the *IBM z13s Technical Guide*, SG24-8294 for more details**

Operating System Support for z13s

- Currency is key to operating system support and exploitation of future servers
- The following releases of operating systems will be supported on z13s
  (Please refer to PSP buckets for any required maintenance)

Operating System	Supported levels
z/OS	▪ z/OS V2.2 with PTFs (Exploitation) ▪ z/OS V2.1 with PTFs (Exploitation) – End of service planned* 9/30/2018 ▪ z/OS V1.13 with PTFs (Limited exploitation) – End of service planned* 9/30/2016 ▪ z/OS V1.12 with PTFs – Service support ended 9/30/2014 ▪ Note: TSS Service Extension for z/OS V1.12 Defect Support: Offered 10/1/14 – 9/30/17 will be required for z13/z13s and compatibility at GA2
z/VSE	▪ z/VSE V6.1 ▪ z/VSE V5.2 with PTFs (Compatibility plus Crypto Express5S – up to 85 LPARs) ▪ z/VSE V5.1 with PTFs (Compatibility) – End of service 6/30/2016
z/TPF	▪ z/TPF V1.1 with PTFs (Compatibility)
z/VM	▪ z/VM V6.3 with PTFs (Exploitation) ▪ z/VM V6.2 with PTFs (Compatibility plus Crypto Express5S support) ▪ Note: z/VM 5.4 – NOT (Compatible) – even though still in service until 12/31/2016
Linux for z Systems	▪ SUSE SLES 12 and 11 (Later releases: GA support TBD by SUSE) ▪ Red Hat RHEL 7 and 6 (Later releases: GA support TBD by Red Hat)
KVM on z Systems	▪ KVM on z Systems V1.1.1 Linux guest support announced: SLES 12 SP1

- Note: Red Hat RHEL V5.11: Support was NOT announced but has been tested.
  For the latest status of Linux on z Systems check this site: Linux on z Systems Tested Platforms

*All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Herzlichen Dank

Dr. Manfred Gnirss

IBM Client Center –
Systems and Software –
z ATS

IBM Germany Lab

Schoenaicher Str. 220
D-71032 Boeblingen
Phone +49 (0) 7031 16-4093
gnirss@de.ibm.com