The IBM z13 Part I: Processor Design, Server Structure, z/Architecture Enhancements, and Operating System Support (17434)

Harv Emery
Executive IT Specialist, z Systems Hardware
IBM Washington System Center

Please evaluate sessions:
www.share.org/Orlando-Eval

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®, GDPS®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER7®, Power Architecture®, PowerVM®, S/390®, ServerProven®, Sysplex Timer®, System p®, System x®, System z®, System z9®, System z10®, Tivoli®, WebSphere®, X-Architecture®, z Systems, z9®, z10, z13, 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.
Oracle®, Java and all Java-based trademarks are trademarks of Oracle Corporation and/or its affiliates 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.

IBM z13 Launch

IBM z13 platform positioning

Platform Core Capabilities:

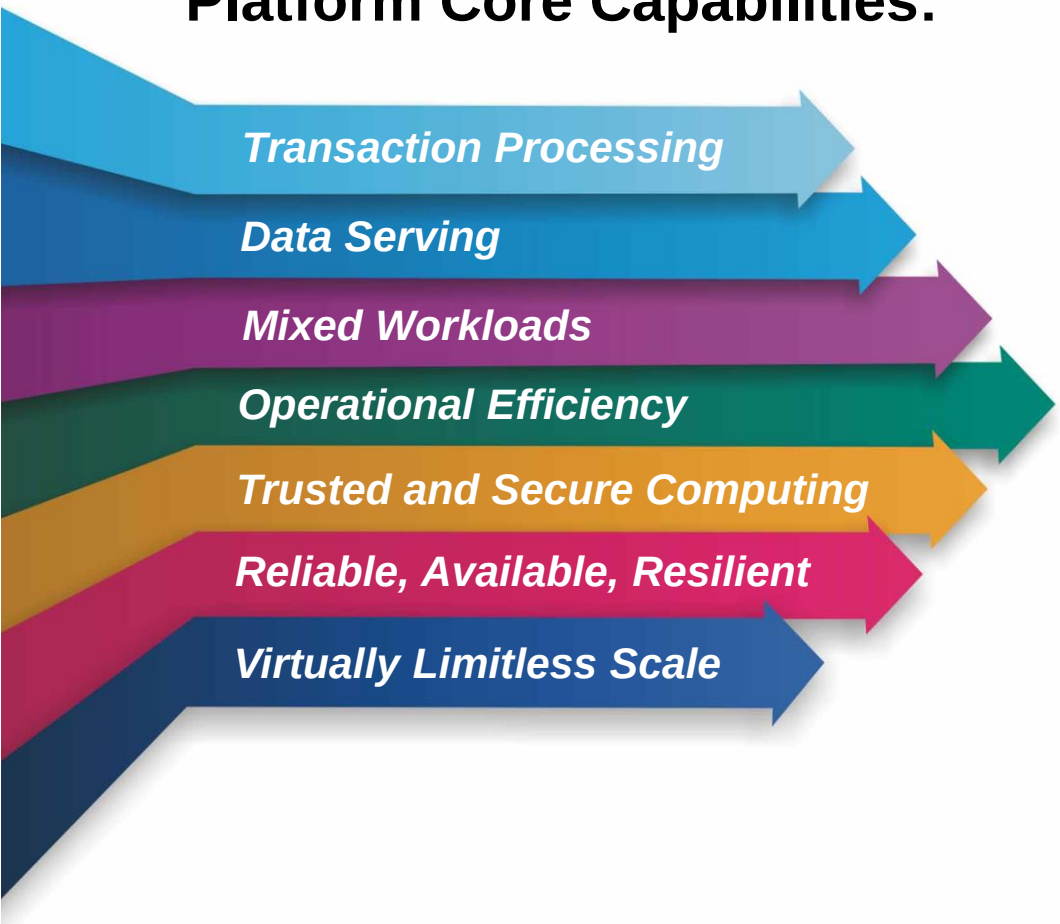

- *The world's premier transaction and data engine now enabled for the **mobile** generation*
- *The integrated transaction and **analytics** system for right-time insights at the point of impact*
- *The world's most efficient and trusted **cloud** system that transforms the economics of IT*

IBM z Systems Server Generations

z196 "N-2"

- Announced: 7/22/2010
- End of mkt: **6/30/2014 (LIC 2015)**
- Chip: 4 core, *5.2 GHz*
- Key z/Arch: *Out of order execution*
- Up to 80 client cores
- CP, IFL, ICF, zAAP, zIIP
- Single thread
- zIIP-zAAP to CP ratio 1x1
- Uni MIPS: 1,202
- Max MIPS: 52,286
- Max mem *3 TB (RAIM)*
- Max mem/LPAR: 1 TB
- LCSS: 4, LPARs: 60
- Subchannel Sets: 3/LCSS
- Max I/O Slots: 160*
- IO buses: 6 GBps (*8 GBps GA2*)
- Max FICON Channels: 320
- FICON Express8 (*8S GA2*)
- Max OSA Ports: 96
- OSA-Express4S (GA2)
- Crypto Express3
- Coupling: *128 CHPIDs*, ISC3, PSIFB DDR: 12x, 1x
- CTN: STP Mixed
- *zBX Model 2 (Attached)*

zE12 "N-1"

- Announced: 8/28/2012
- End of mkt: **Not yet announced**
- Chip: *6 core, 5.5 GHz*
- Key z/Arch: *Transactional Exec*
- Up to *101* client cores
- CP, IFL, ICF, zAAP, zIIP
- Single thread
- zIIP-zAAP to CP *ratio 2x1 (GA2)*
- Uni MIPS: 1,514
- Max MIPS: 78,426
- Max mem 3 TB (RAIM)
- Max mem/LPAR: 1 TB
- LCSS: 4, LPARs: 60
- Subchannel Sets: 3/LCSS
- Max I/O Slots: 160*
- I/O buses: 8 GBps
- Max FICON Channels: 320
- FICON Express8S
- Max OSA Ports: 96
- OSA-Express4S *(5S GA2)*
- *Crypto Express4S*
- Coupling: 128 CHPIDs, PSIFB DDR: 12x, 1x
- CTN: STP mixed
- *Flash Express*
- *Native PCIe: zEDC, RoCE (GA2)*
- *zBX Model 3 (Attached)*

z13 "N"

- Announced 1/14/2015
- End of mkt: Current generation
- *Chip: 8 core,* 5.0 GHz
- Key z/Arch*: SMT, Vector (SIMD)*
- Up to *141* client cores
- CP, IFL, ICF, zIIP *(No zAAPs)*
- *SMT Support: zIIP, IFL*
- zIIP to CP ratio 2x1
- Uni MIPS: *1,695*
- Max MIPS: *111,566*
- Max mem *10 TB* (RAIM)
- Max LPAR: *10 TB (z/OS: 4TB)*
- *LCSS: 6, LPARs: 85*
- Subchannel Sets: *4/LCSS*
- Max I/O Slots: 160*
- I/O buses: *16 GBps*
- Max FICON Channels: 320
- *FICON Express16S**
- Max OSA Ports: 96
- *OSA-Express5S**
- *Crypto Express5S*
- Coupling: *256 CHPIDs*, *ICA 8 GBps*, PSIFB DDR: 12x, 1x
- *CTN: STP Only*
- *Flash Express (refresh)*
- Native PCIe: zEDC, *RoCE**
- *zBX Model 4 (Independent)*

** Major functional enhancements*

IBM z13 Key Planning and Support Dates

- **January 14, 2015 – Announcement Day**
  - IBM United States Hardware Announcement Letter 115-001
- **March 9, 2015 – General Availability** ✓
- **June 26, 2015** ✓ **(Bundle S14, Driver 22. Released June 10)**
  - Field install of MES hardware features for z13 Models N30, N63, N96, NC9, and NE1
  - z/VM V6.3 support for Multi-VSwitch Link Aggregation
  - Support for 256 Coupling CHPIDs
  - HMC STP Panel Enhancements: Initialize Time, Set Date and Time, Time Zone, View-Only Mode
  - Fibre Channel Protocol (FCP) channel configuration discovery and debug
  - Improved High Performance FICON for z Systems (zHPF) I/O Execution at Distance
  - IBM zAware support for Linux on z Systems
- **September 25, 2015**
  - FICON Dynamic Routing
  - Forward Error Correction (FEC) for FICON Express16S
  - Storage Area Network (SAN) Fabric I/O Priority

IBM z13 Model Structure and Performance

z13 System Offering Overview

zEC12 zBX Model 3

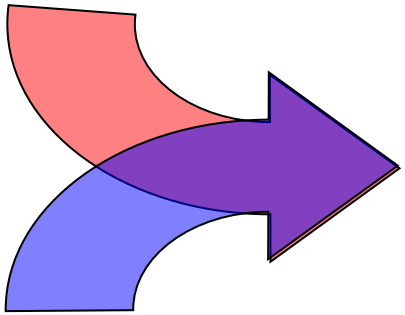

z196 zBX Model 2

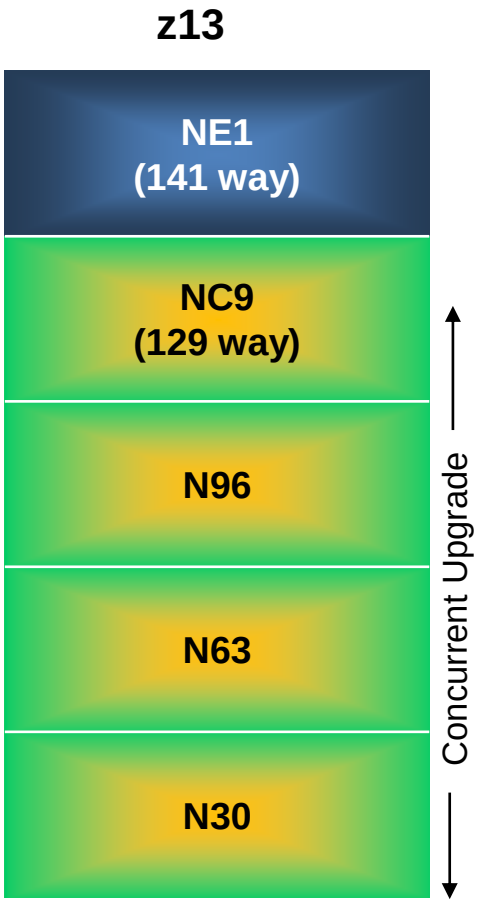

Machine Type for z13

- 2964

Processors

- 39 PUs per drawer (42 in NE1)
- Sub-capacity available up to 30 CPs
- 2 standard spare PUs per system

Memory

- System minimum = 64 GB with separate 96 GB HSA
- Maximum: ~10 TB / ~2.5TB per drawer
- RAIM memory design
- Purchase Increments – 32 to 512 GB

I/O

- Up to 14 fanouts per drawer
  - **Up to 10 PCIe Gen 3 fanouts: 1-port 16 GBps I/O or 2-port 8 GBps PCIe coupling**
  - Up to 4 IFB HCA fanouts: 2-port 6 GBps I/O, 2-port 12x PSIFB, or 4-port 1x PSIFB

On upgrade from zEC12 or z196

- Detach zBX Model 3 or 2 and upgrade it to zBX Model 4 (Option: Move zBX Model 3)
- Feature convert installed zAAPs to zIIPs (default) or another processor type
- For installed On Demand Records, change temporary zAAPs to zIIPs. Stage the record

z13 Continues the CMOS Mainframe Heritage Begun in 1994

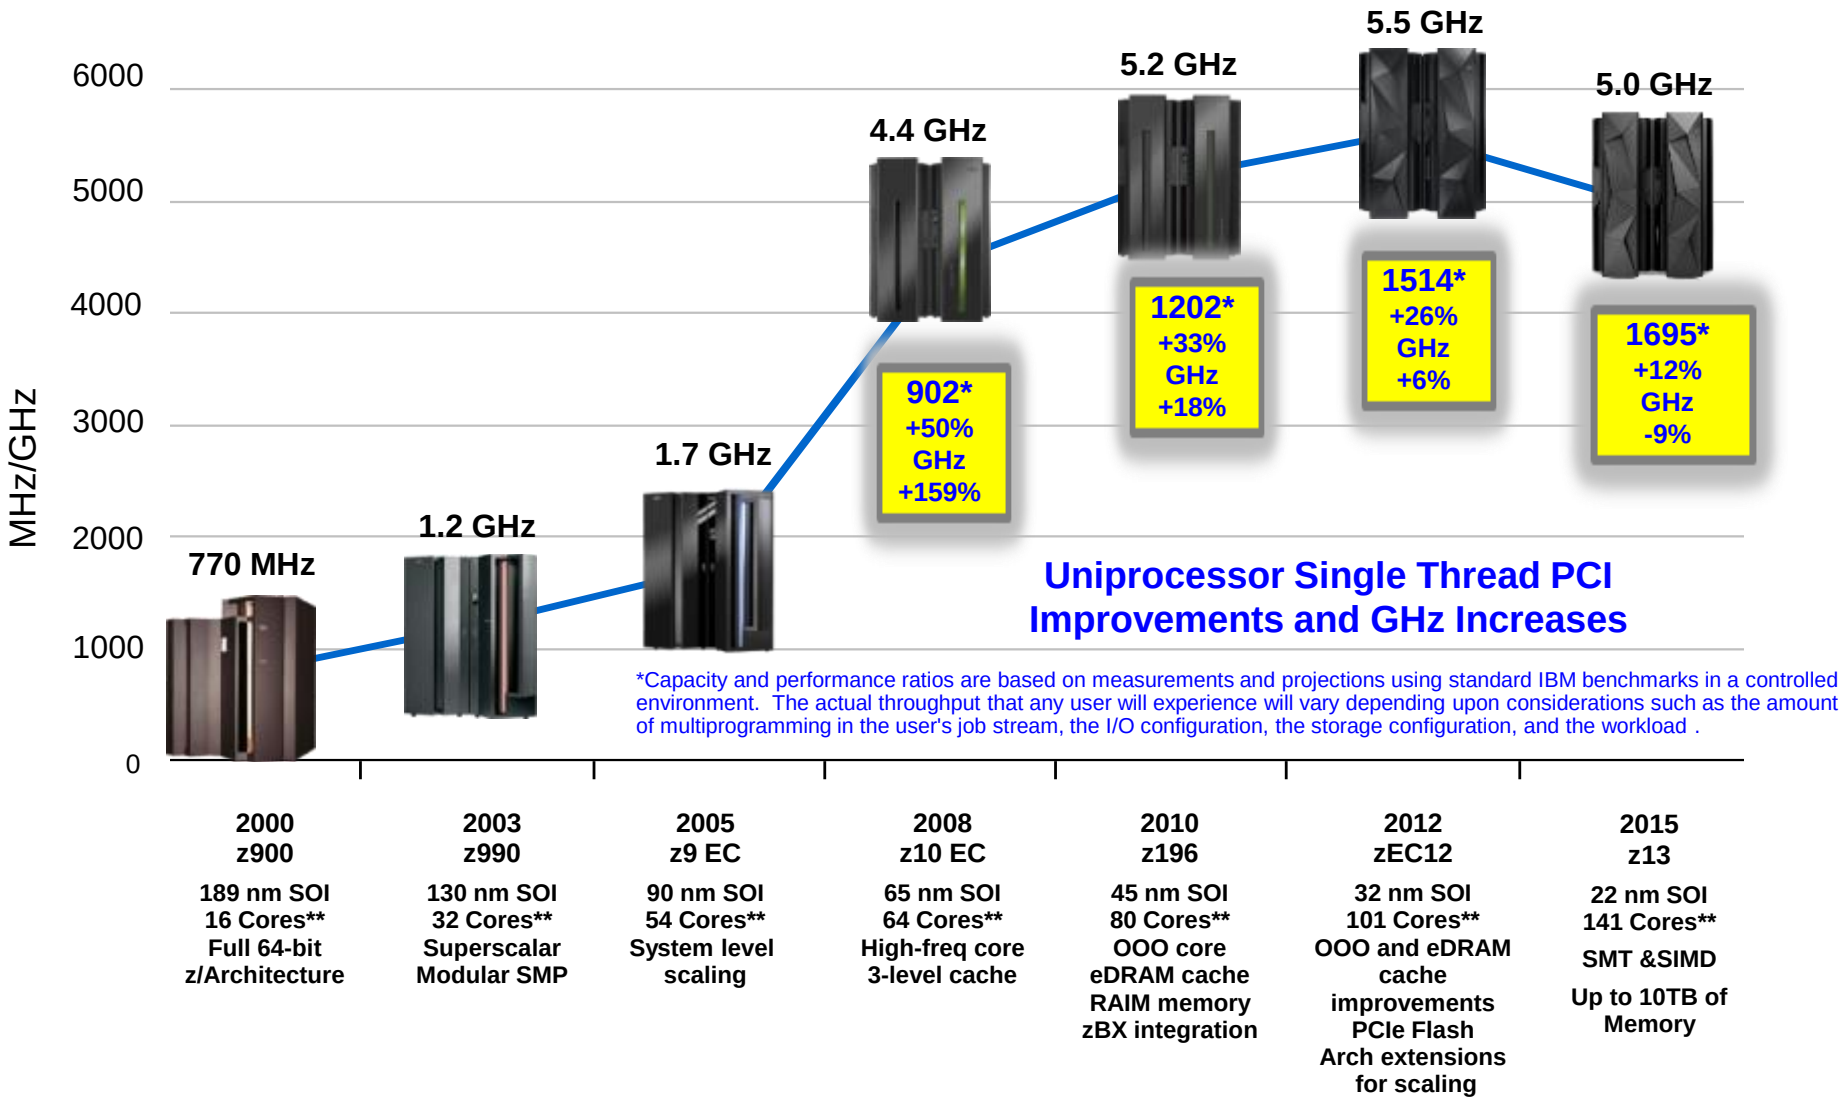

*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload .

2000 z900	2003 z990	2005 z9 EC	2008 z10 EC	2010 z196	2012 zEC12	2015 z13
189 nm SOI 16 Cores** Full 64-bit z/Architecture	130 nm SOI 32 Cores** Superscalar Modular SMP	90 nm SOI 54 Cores** System level scaling	65 nm SOI 64 Cores** High-freq core 3-level cache	45 nm SOI 80 Cores** OOO core eDRAM cache RAIM memory zBX integration	32 nm SOI 101 Cores** OOO and eDRAM cache improvements PCIe Flash Arch extensions for scaling	22 nm SOI 141 Cores** SMT &SIMD Up to 10TB of Memory

* PCI (MIPS) Tables are NOT adequate for making comparisons of z Systems processors. Use IBM Capacity Planning Tools!
** Number of PU cores for customer use

z13 Full and Sub-Capacity CP Offerings

CP Capacity* – (Measured with z/OS V2.1)
Relative to Full Capacity Uniprocessor
701 = 100% = 1695 PCI (IBM MIPS)
601 ≈ 63% = 1068 PCI
501 ≈ 44% = 746 PCI
401 ≈ 15% = 250 PCI

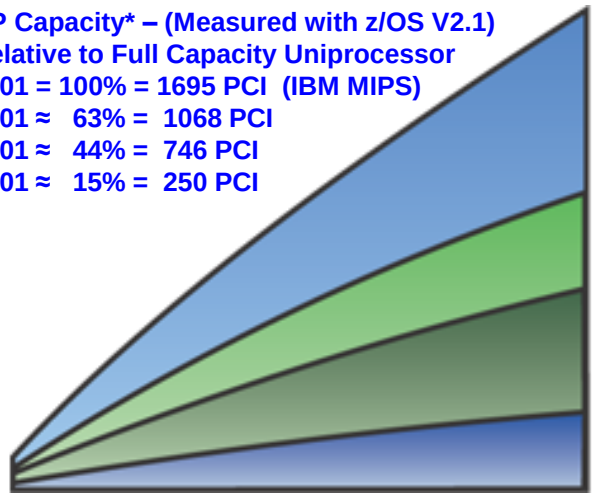

- **Up to 30 subcapacity CPs may be ordered on ANY z13 model. If 31 or more CPs are ordered all must be full 7xx capacity**
- **232 CP capacity settings including 400 for no CPs**
- **Specialty engines run at full capacity, SMT optional: IFL and zIIP**
- **SMT *average* capacity* benefit: 25% for zIIP, 20% for IFL**
- **z13 NE1 maximum CP capacity is 1.4 times the zEC12 HA1**
- **IFL PVU rating remains at 120, the same as zEC12, z196 and z10 EC, no increase for SMT**
- **Entitlement to purchase two zIIPs for each CP purchased.**

MSU Sub Capacity

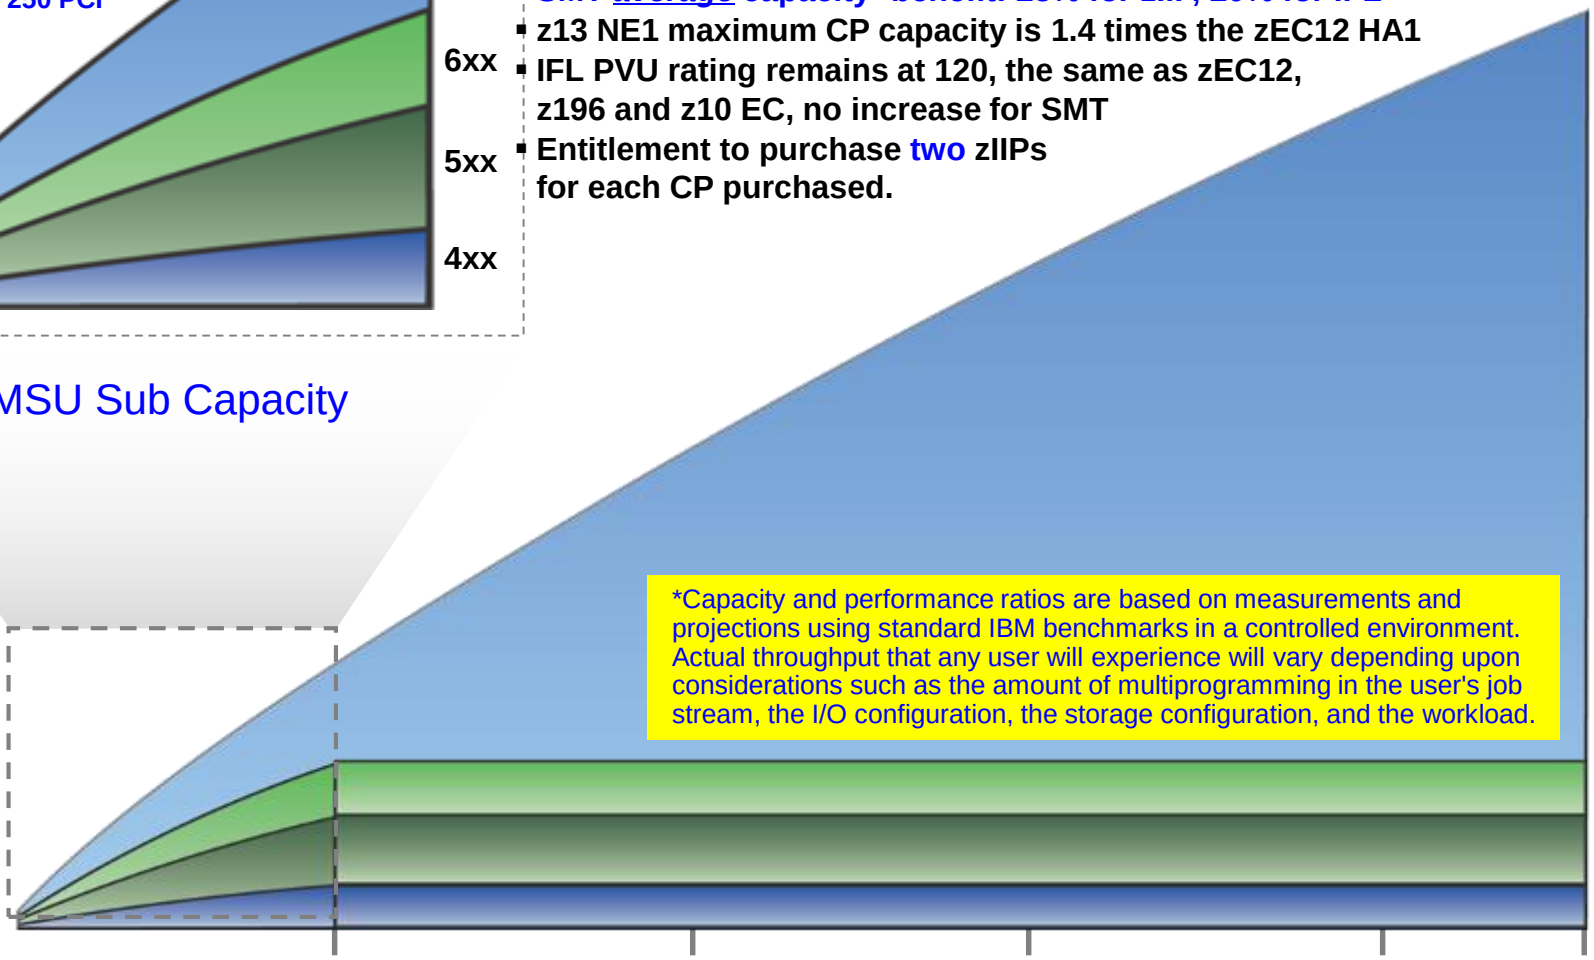

*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. Actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload.

z13 z/Architecture Extensions

- **Two-way 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. Changes made to support SMT are designed to improve both core capacity in SMT mode and single thread performance.
  - PR/SM dispatches online logical processors to 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 CPU Measurement Facility monitoring support will measure thread usage, capacity and performance.
- **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.

Simultaneous Multithreading (SMT)

- 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 is **40%** 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
- SMT can be turned on or off on an LPAR by LPAR basis by operating system parameters. z/OS can also do this dynamically 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

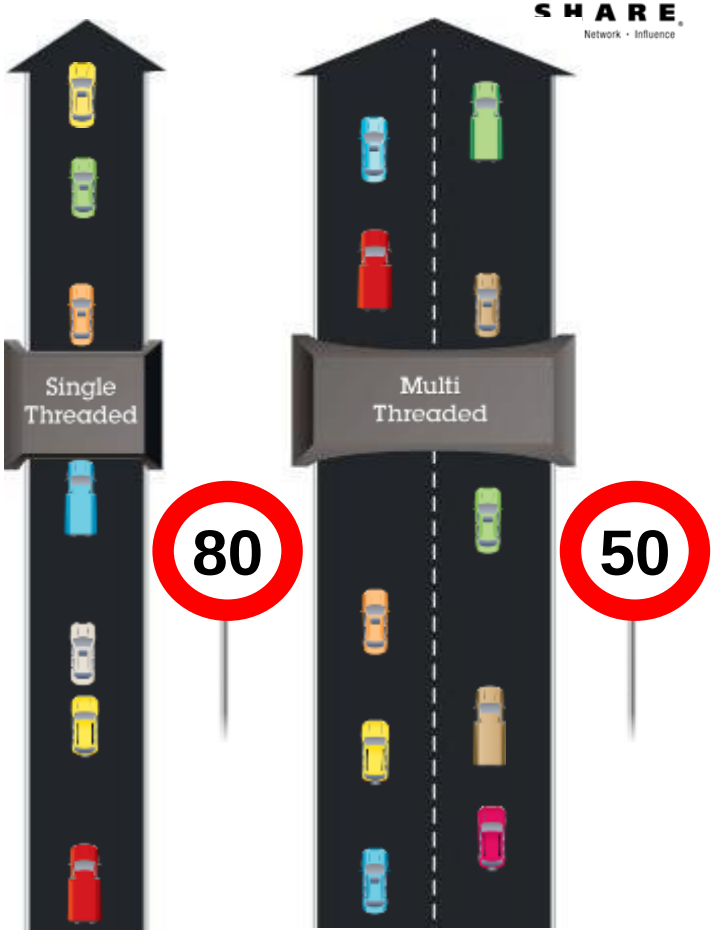

*Which approach is designed for the highest volume** of traffic? Which road is faster?*

***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 (Single Instruction Multiple Data) Instructions

Increased parallelism to enable analytics processing

- Fewer instructions helps to 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

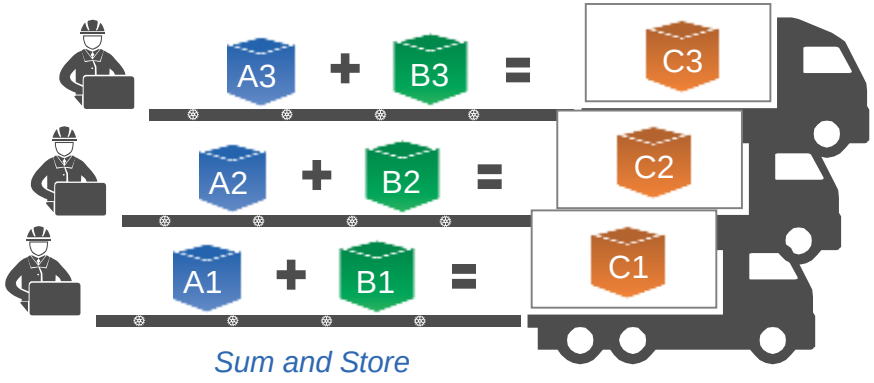

Instruction is performed for *every data element*

SIMD
SINGLE INSTRUCTION, MULTIPLE DATA

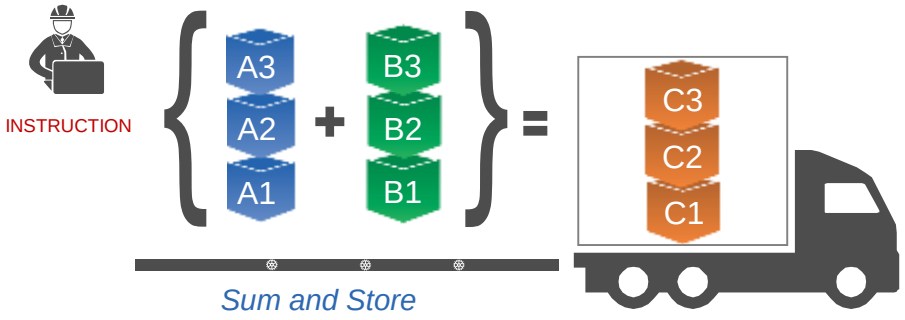

Perform instructions on *every element at once*

IBM z13 Technology

z13 8-Core Processor Chip Detail

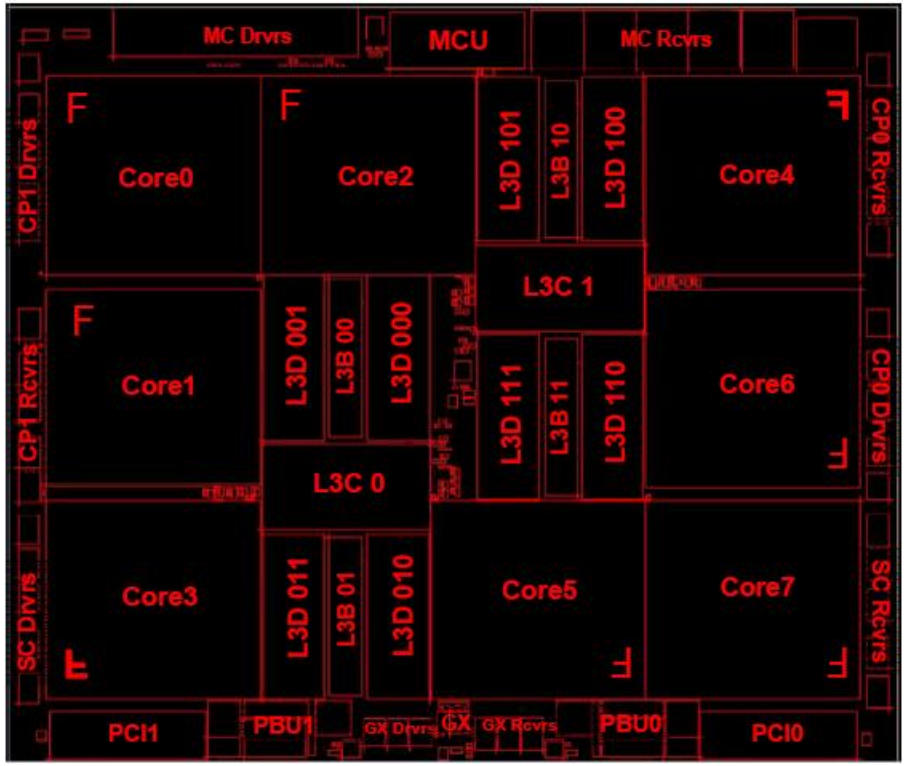

- **14S0 22nm SOI Technology**
  - 17 layers of metal
  - 3.99 Billion Transistors
  - 13.7 Miles copper wire
- **Chip Area**
  - 678.8 mm²
  - 28.4 x 23.9 mm
  - 17,773 power pins
  - 1,603 signal I/Os

- **Up to eight active cores (PUs) per chip**
  - **5.0 GHz (v5.5 GHz zEC12)**
  - L1 cache/ core
    - 96 KB I-cache
    - 128 KB D-cache
  - L2 cache/ core
    - 2M+2M 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 GX++ I/O bus
  - Two PCIe I/O buses
- **Memory Controller (MCU)**
  - Interface to controller on memory DIMMs
  - Supports RAIM design

z13 Storage Control (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)
  - 224 MB L3 NIC Directory
  - 2 SCs = 960 MB L4 per z13 drawer
- **Interconnects (L4 – L4)**
  - **3 to CPs in node**
  - **1 to SC (node – node) in drawer**
  - **3 to SC nodes in remote drawers**
- **6 Clock domains**

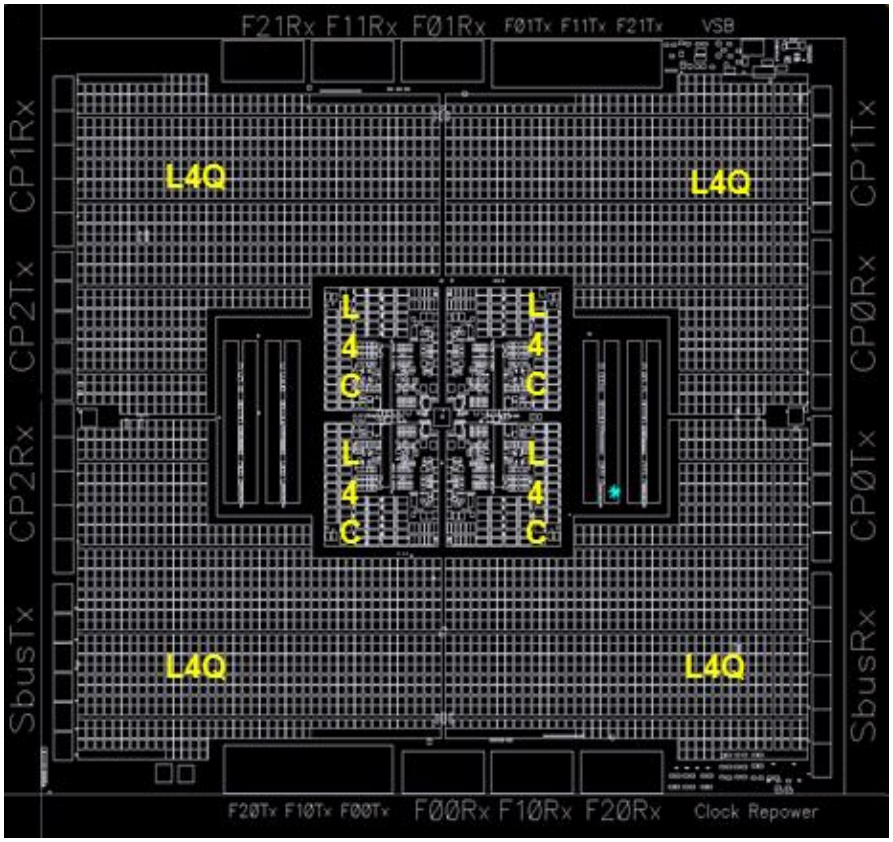

z13 SCM Vs zEC12 MCM Comparison

z13 Single Chip Modules (SCMs)

- **Processer Unit (PU) SCM**
  - 68.5mm x 68.5mm – fully assembled
  - PU Chip area 678 mm^2
  - Eight core chip with 6, 7 or 8 active cores
- **Storage Control (SC) SCM**
  - 68.5mm x 68.5mm – fully assembled
  - SC Chip area 678 mm^2
  - 480 MB on-inclusive L4 cache per SCM
  - Non-Data Integrated Coherent (NIC) Directory for L3
- **Processor Drawer – Two Nodes**
  - Six PU SCMs for 39 PUs (42 PUs in Model NE1)
  - Two SC SCMs (960 MB L4)
  - N30: One Drawer, N63: Two Drawers, N96: Three Drawers, NC9 or NE1: Four Drawers

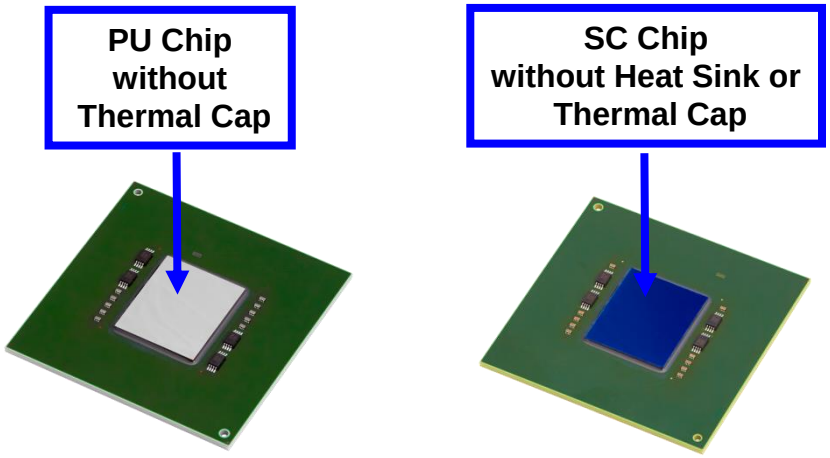

zEC12 Multi Chip Module (MCM)

- **Technology**
  - 96mm x 96mm with 102 glass ceramic layers
  - 7,356 LGA connections to 8 chip sites
- **Six 6-core Processor (PU) chips**
  - Each with 4, 5 or 6 active cores
  - 27 active processors per MCM (30 in Model HA1)
  - PU Chip size 23.7 mm x 25.2 mm
- **Two Storage Control (SC) chips per MCM**
  - 192 MB Inclusive L4 cache per SC, 384 MB per MCM
  - SC Chip size 26.72 mm x 19.67 mm
- **One MCM per book, up to 4 books per System**

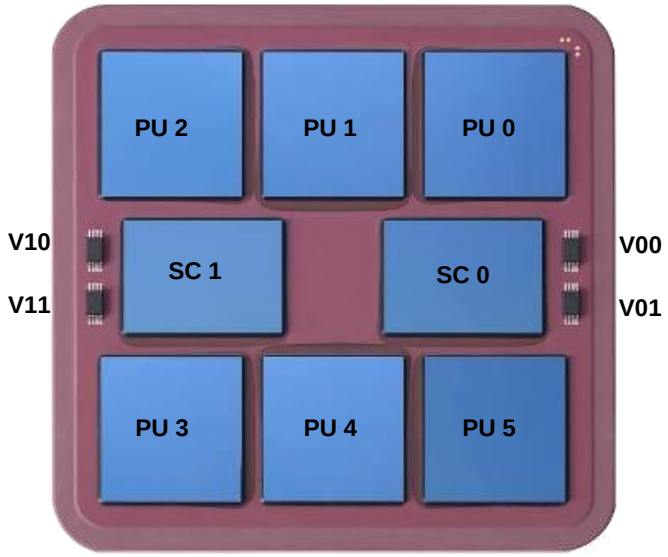

z13 Logical Drawer Structure and Interconnect

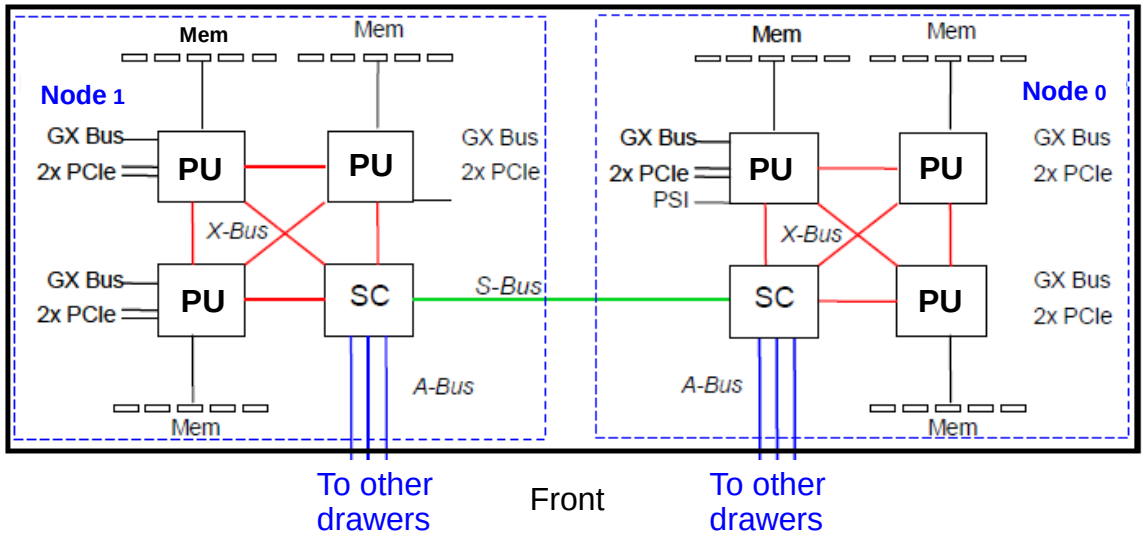

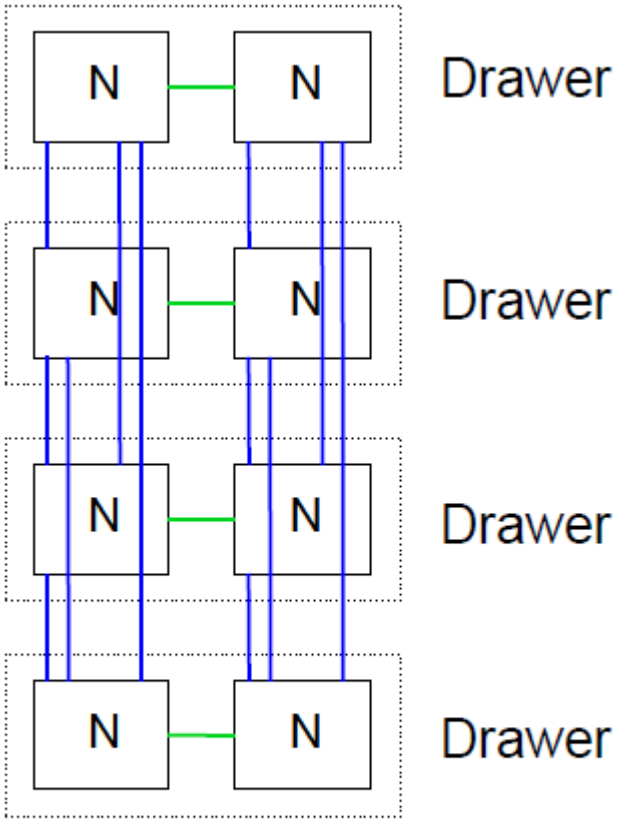

4 Drawer System Interconnect

Physical node: (Two per drawer)

- **Chips**
  - Three PU chips
  - One SC chip (480 MB L4 cache + 224 MB NIC Directory)
- **RAIM Memory**
  - Three Memory Controllers: One per CP Chip
  - Five DDR3 DIMM slots per Controller: 15 total per logical node
  - Populated DIMM slots: 20 or 25 per drawer
- **SC and CP Chip Interconnects**
  - **X-bus: SC and CPs to each other (same node)**
  - **S-bus: SC to SC chip in the same drawer**
  - **A-bus: SC to SC chips in the remote drawers**

zEC12 Book (Left) to z13 Node (Right) Cache Comparison

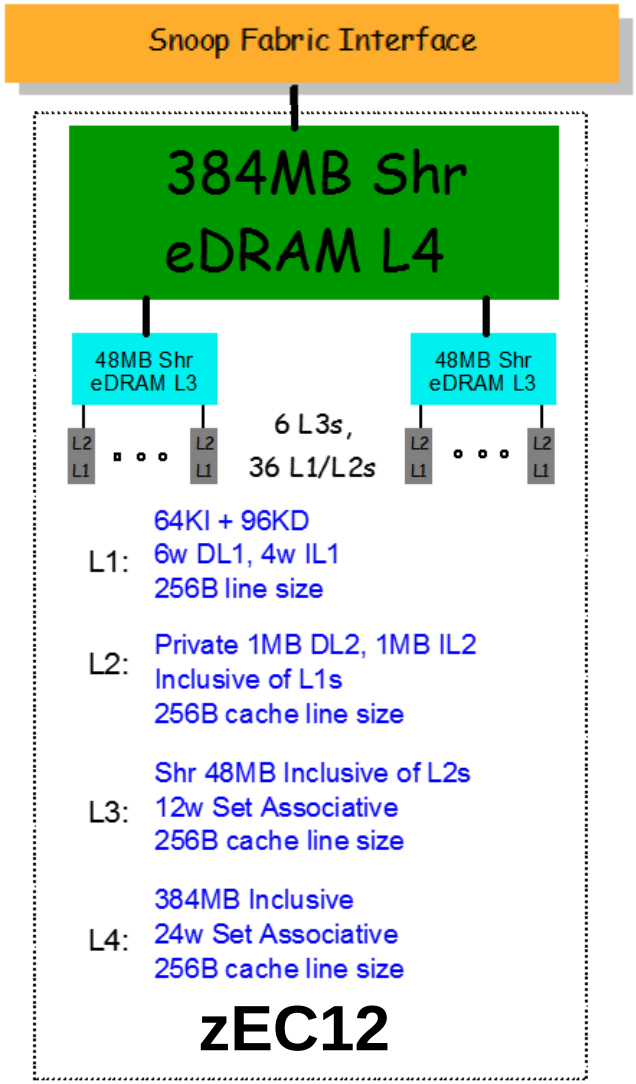

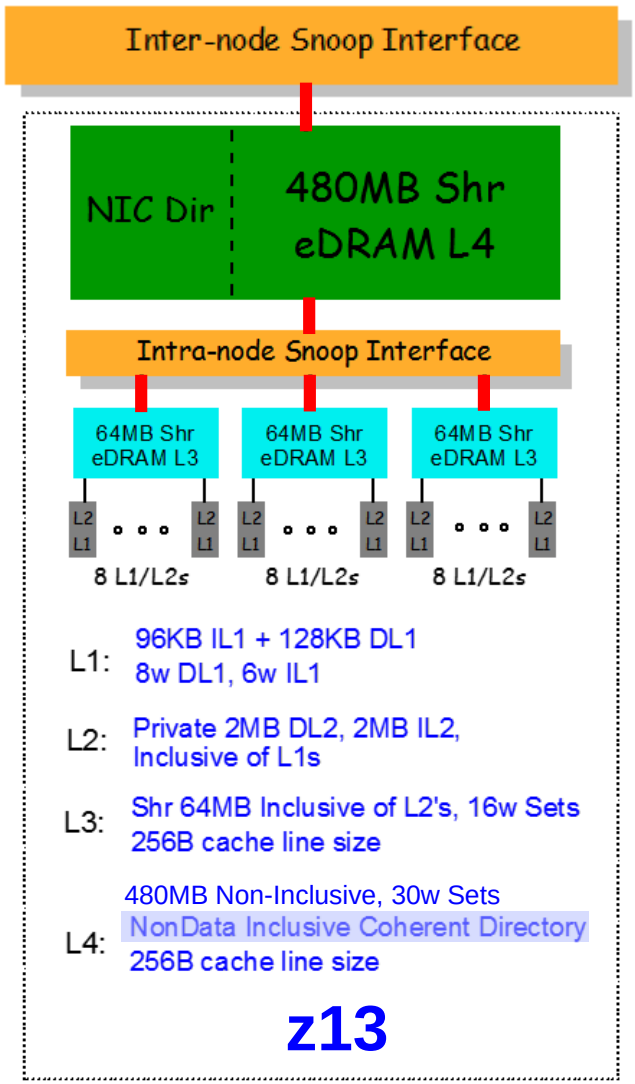

z13 CPC Drawer (Top View)

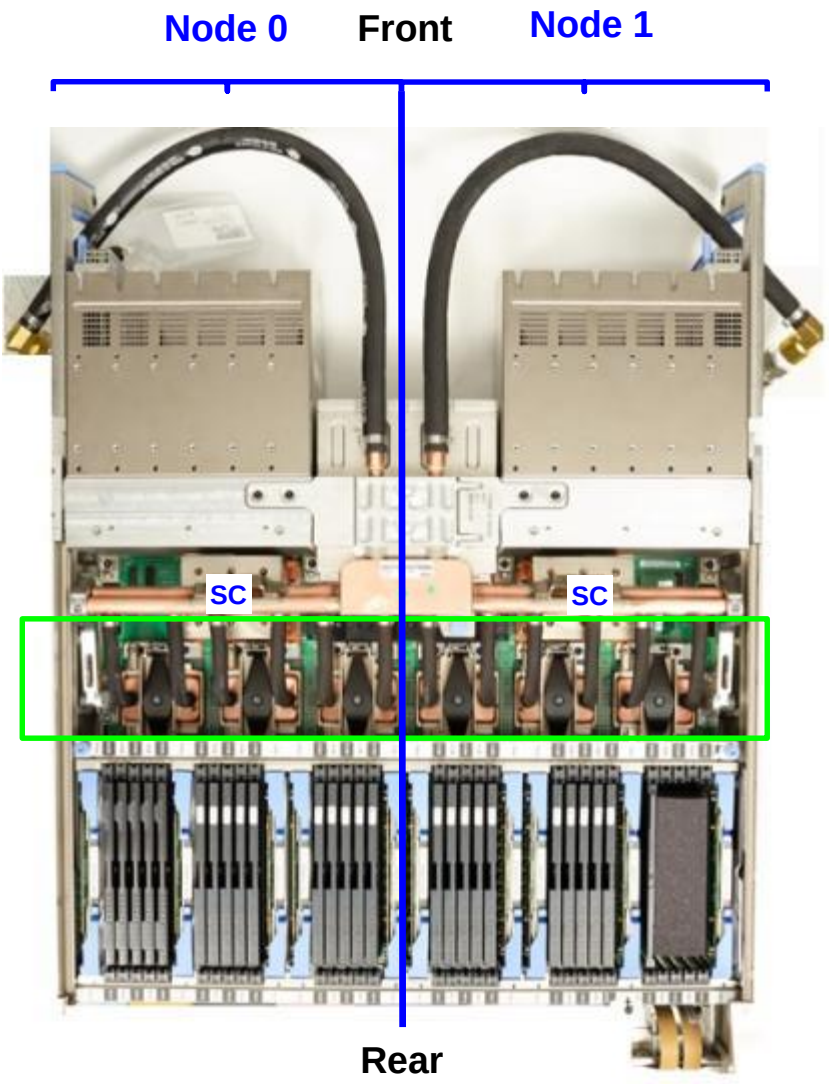

- **Two physical nodes, left and right**
- **Each logical node:**
  - Three PU chips
  - One **SC** chip (480 MB L4 cache)
  - Three Memory Controllers: One per CP Chip
  - Five DDR3 DIMM slots per Memory Controller: 15 total per node (One bank of 5 never used in Node 1)
- **Each drawer:**
  - Six PU Chips with water cooling for 39 active PUs (42 in Model NE1)
  - Two **SC** Chips (960 MB L4 cache) with heat sink for air cooling
  - Populated DIMM slots: 20 or 25 DIMMs to support up to 2,560 GB of addressable memory (3,200 GB RAIM)
  - Two Flexible Support Processors
  - Ten fanout slots for PCIe I/O drawer fanouts or PCIe coupling fanouts
  - Four fanout slots for IFB I/O drawer fanouts or PSIFB coupling link fanouts

z13 Model NE1 or NC9 Radiator (Air) Cooled – Under the covers (Front View)

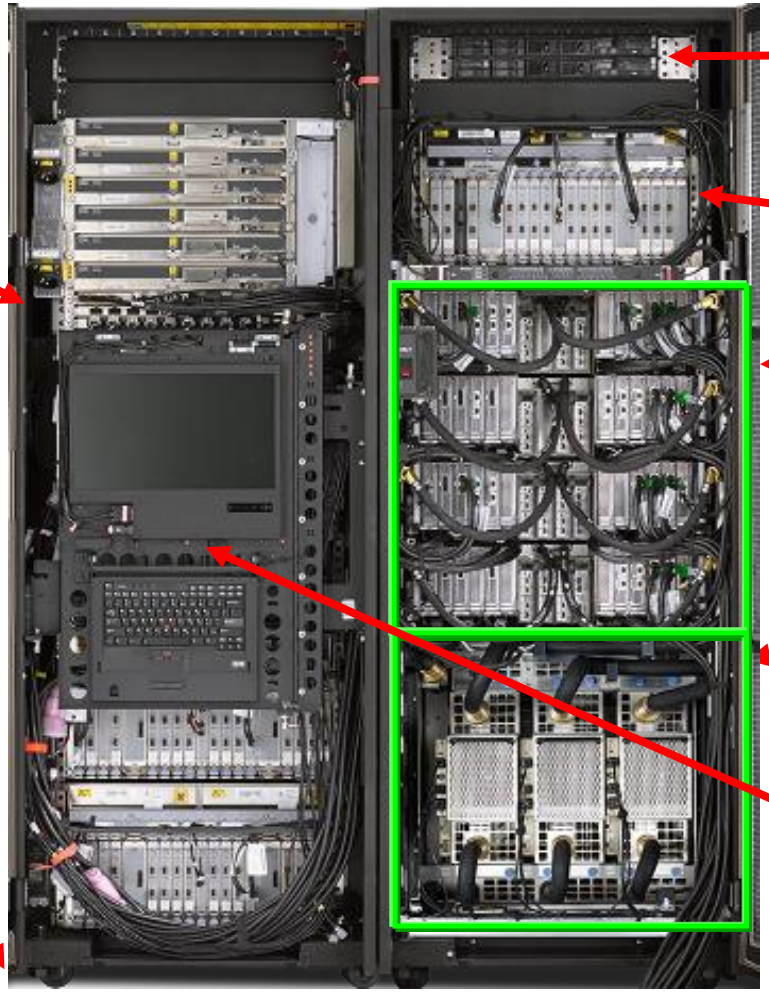

Space for Optional Integrated Battery Features (IBFs)

Power Components

Space for the first four I/O drawers. The top two can be 8-slot for carried forward FICON Express8. All can be PCIe I/O drawers

Two 1U Support Element (SE) System Units

Last PCIe I/O drawer (5th)

Processor Drawers (1st bottom to 4th top) with Flexible Support Processors (FSPs), and I/O fanouts

N+2 Pumps and Blowers for Radiator Air Cooling Unit

2 SE Displays with Keyboards

z13 CPC Drawer and I/O Drawer Locations

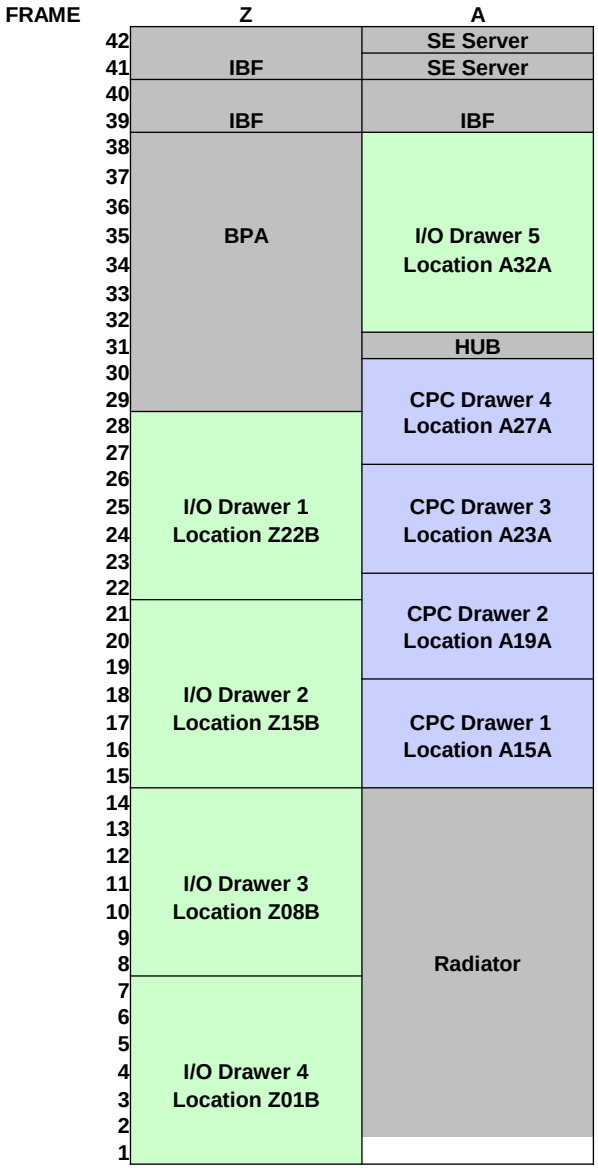

- Drawer locations are based on the front view of the machine: Frame A (right), Frame Z (left) and EIA Unit location of the lower left of drawer corner
- Locations are reported in eConfig “AO Data” reports along with PCHIDs for I/O definition
- CPC Drawers are populated from bottom to top
  - Drawer 1: A15A – N30, N63, N96, NC9 and NE1
  - Drawer 2: A19A – N63, N96, NC9 and NE1
  - Drawer 3: A23A – N96, NC9 and NE1
  - Drawer 4: A27A – NC9 and NE1
- Old technology 8-slot I/O drawers (if present) populate top down in Frame Z
  - Drawer 1: Z22B, Drawer 2: Z15B
- PCIe 32-slot I/O Drawers populate in remaining locations:
  - PCIe I/O Drawer 1: Z22B, Z15B or Z08B
  - PCIe I/O Drawer 2: Z15B, Z08B, or Z01B
  - PCIe I/O Drawer 3: Z08B, Z01B, or A32A
  - PCIe I/O Drawer 4: Z01B
  - PCIe I/O Drawer 5: A32A

z13 Processor Features – zIIP to CP 2:1 ratio

Model	Drawers /PUs	CPs	IFLs uIFLs	*zIIPs*	ICFs	Std SAPs	Optional SAPs	Std. Spares	IFP
N30	1/39	0-30	0-30 0-29	*0-20*	0-30	6	0-4	2	1
N63	2/78	0-63	0-63 0-62	*0-42*	0-63	12	0-8	2	1
N96	3/117	0-96	0-96 0-95	*0-64*	0-96	18	0-12	2	1
NC9	4/156	0-129	0-129 0-128	*0-86*	0-129	24	0-16	2	1
NE1	4/168	0-141	0-141 0-140	*0-94*	0-141	24	0-16	2	1

- z13 Models N30 to NC9 use drawers with 39 cores. The Model NE1 has 4 drawers with 42 cores.
- The maximum number of logical ICFs or logical CPs supported in a CF logical partition is 16
- **The integrated firmware processor (IFP)** is used for **PCIe I/O support functions**
- **Concurrent Drawer Add (CDA)** is available to upgrade in steps from model N30 to model NC9

1. At least one CP, IFL, or ICF must be purchased in every machine
2. Two zIIPs may be purchased for each CP purchased if PUs are available. This remains true for sub-capacity CPs and for “banked” CPs.
3. **On an upgrade from z196 or zEC12, installed zAAPs are converted to zIIPs by default. (Option: Convert to another engine type)**
4. “uIFL” stands for Unassigned IFL
5. The IFP is conceptually an additional, special purpose SAP

IBM z13 Memory

z13 5-Channel RAIM Memory Controller Overview
(RAIM = Redundant Array of Independent Memory)

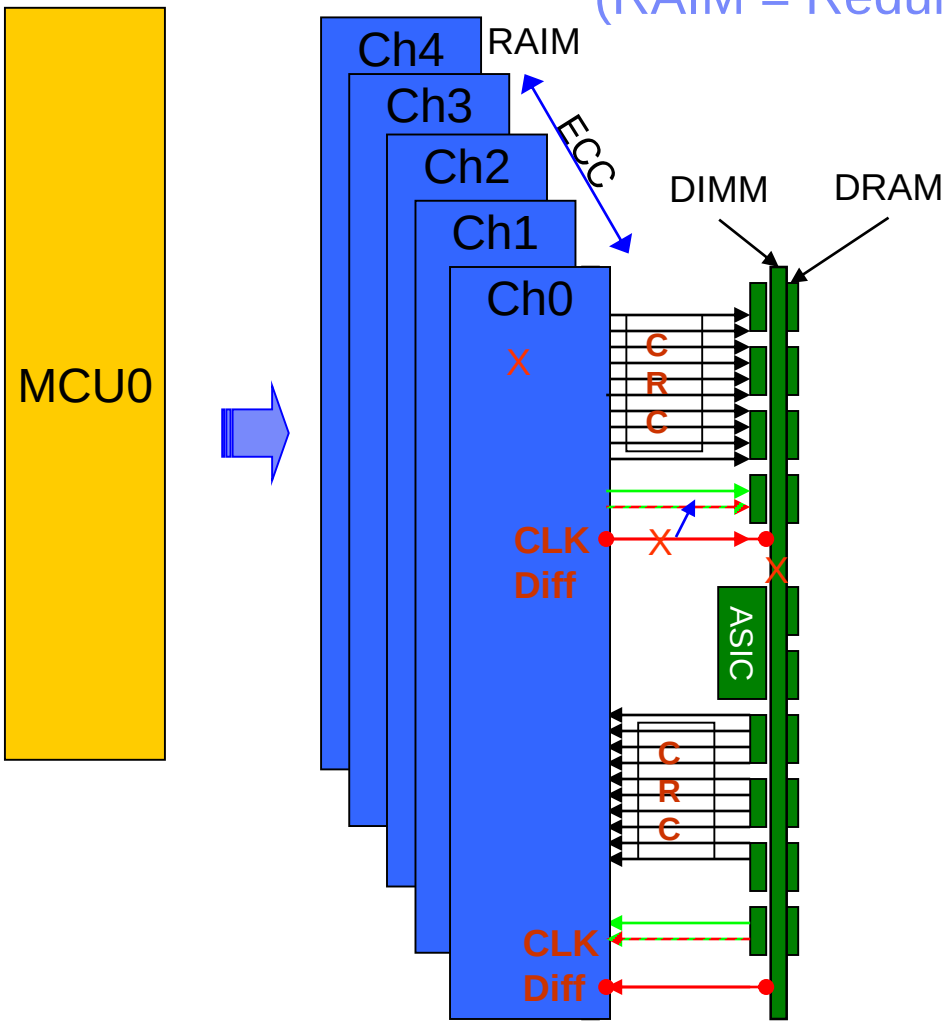

Layers of Memory Recovery

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

z13: Each memory channel supports only one DIMM

z13 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 z13 Model N30. They are always done without a POR on z13 models with multiple drawers using Enhanced CPC 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. All LPARs, if **Flexible Memory** is selected.
- **eConfig user TIP**: To determine the size of the largest LIC CC Customer Memory upgrade possible, examine the configurator default "**Memory Plan Ahead Capacity**" field. If the customer requires a LIC CC upgrade larger that that, configure Plan Ahead Memory by selecting a larger "Memory Plan Ahead Capacity" target value.

***HSA size is 96 GB on z13**

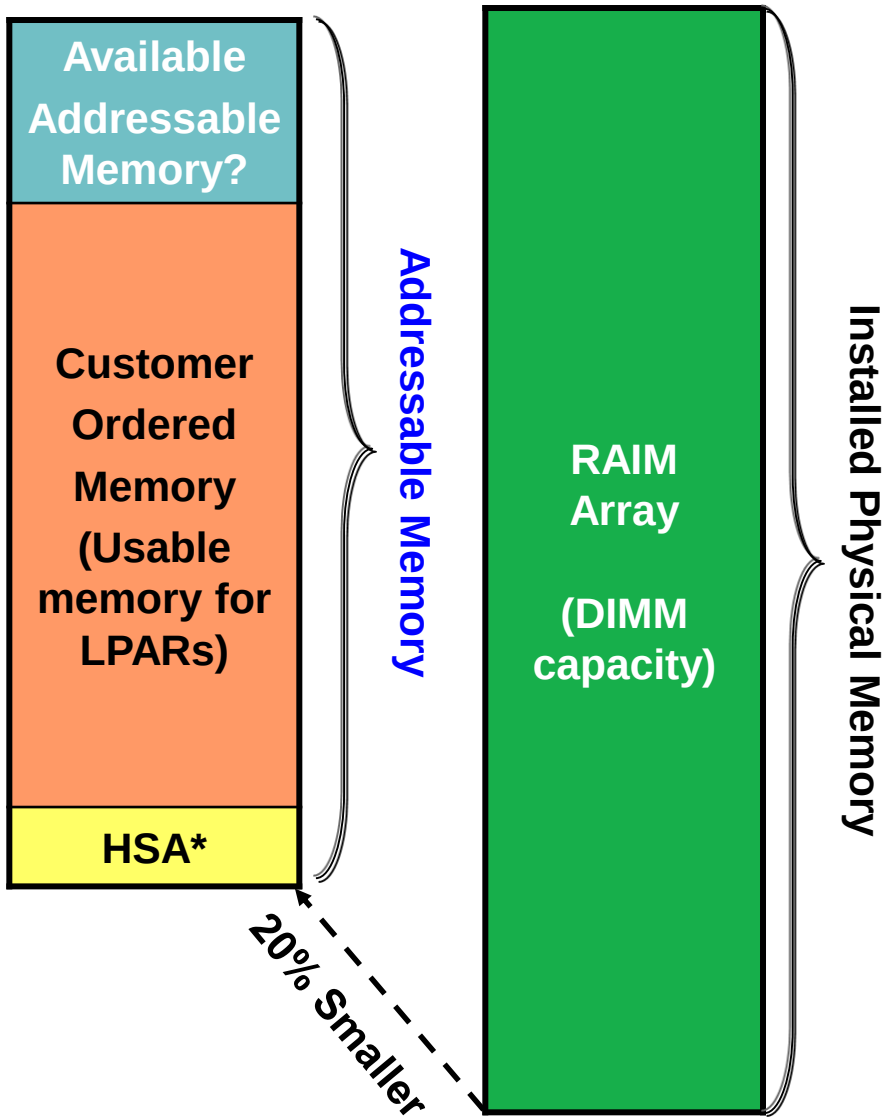

z13 Purchasable Addressable Memory Ranges

Model	Standard Memory GB	Flexible Memory GB
N30	64 - 2464	NA
N63	64 - 5024	64 - 2464
N96	64 - 7584	64 - 5024
NC9	64 - 10144	64 - 7584
NE1	64 - 10144	64 - 7584

- ***Purchased Memory -*** **Memory available for assignment to LPARs**
- ***Hardware System Area*** **– Standard 96 GB of addressable memory for system use outside customer memory**
- ***Standard Memory*** **- Provides minimum physical memory required to hold customer purchase memory plus 96 GB HSA**
- ***Flexible Memory*** **- Provides additional physical memory needed to support activation base customer memory and HSA on a multiple drawer z13 with one drawer out of service.**
- ***Plan Ahead Memory*** **– Provides additional physical memory needed for a concurrent upgrade (LIC CC change only) to a preplanned target customer memory**

z13 Standard and Flexible Addressable Memory Offerings

Memory Increment (GB)	Offered Memory Sizes (GB)	Memory Maximum Notes (GB)
32	64, 96, 128, 160,192	
64	256, 320, 384, 448	
96	544, 640, 736, 832, 928	
128	1056, 1184, 1312, 1440	
256	1696, 1952, 2208, **2464**, 2720, 2976, 3232, 3488, 3744, 4000, 4256, 4512, 4768, **5024**, 5280, 5536, 5792, 6048	**2464** – N30 Standard, N63 Flexible **5024** – N63 Standard, N96 Flexible
512	6560, 7072, **7584**, 8096, 8608, 9120, 9632, **10144**	**7584** – N96 Standard, NC9 and NE1 Flexible **10144** – NC9 and NE1 Standard

On Demand

z13 Basics of Capacity on Demand

Upgrade to z13: **Installed** On Demand records: zAAPs are converted to zIIPs and the record is migrated staged.
Upgrade to z13: **Staged** On Demand records: Records with zAAPs are rejected. Others are migrated staged.

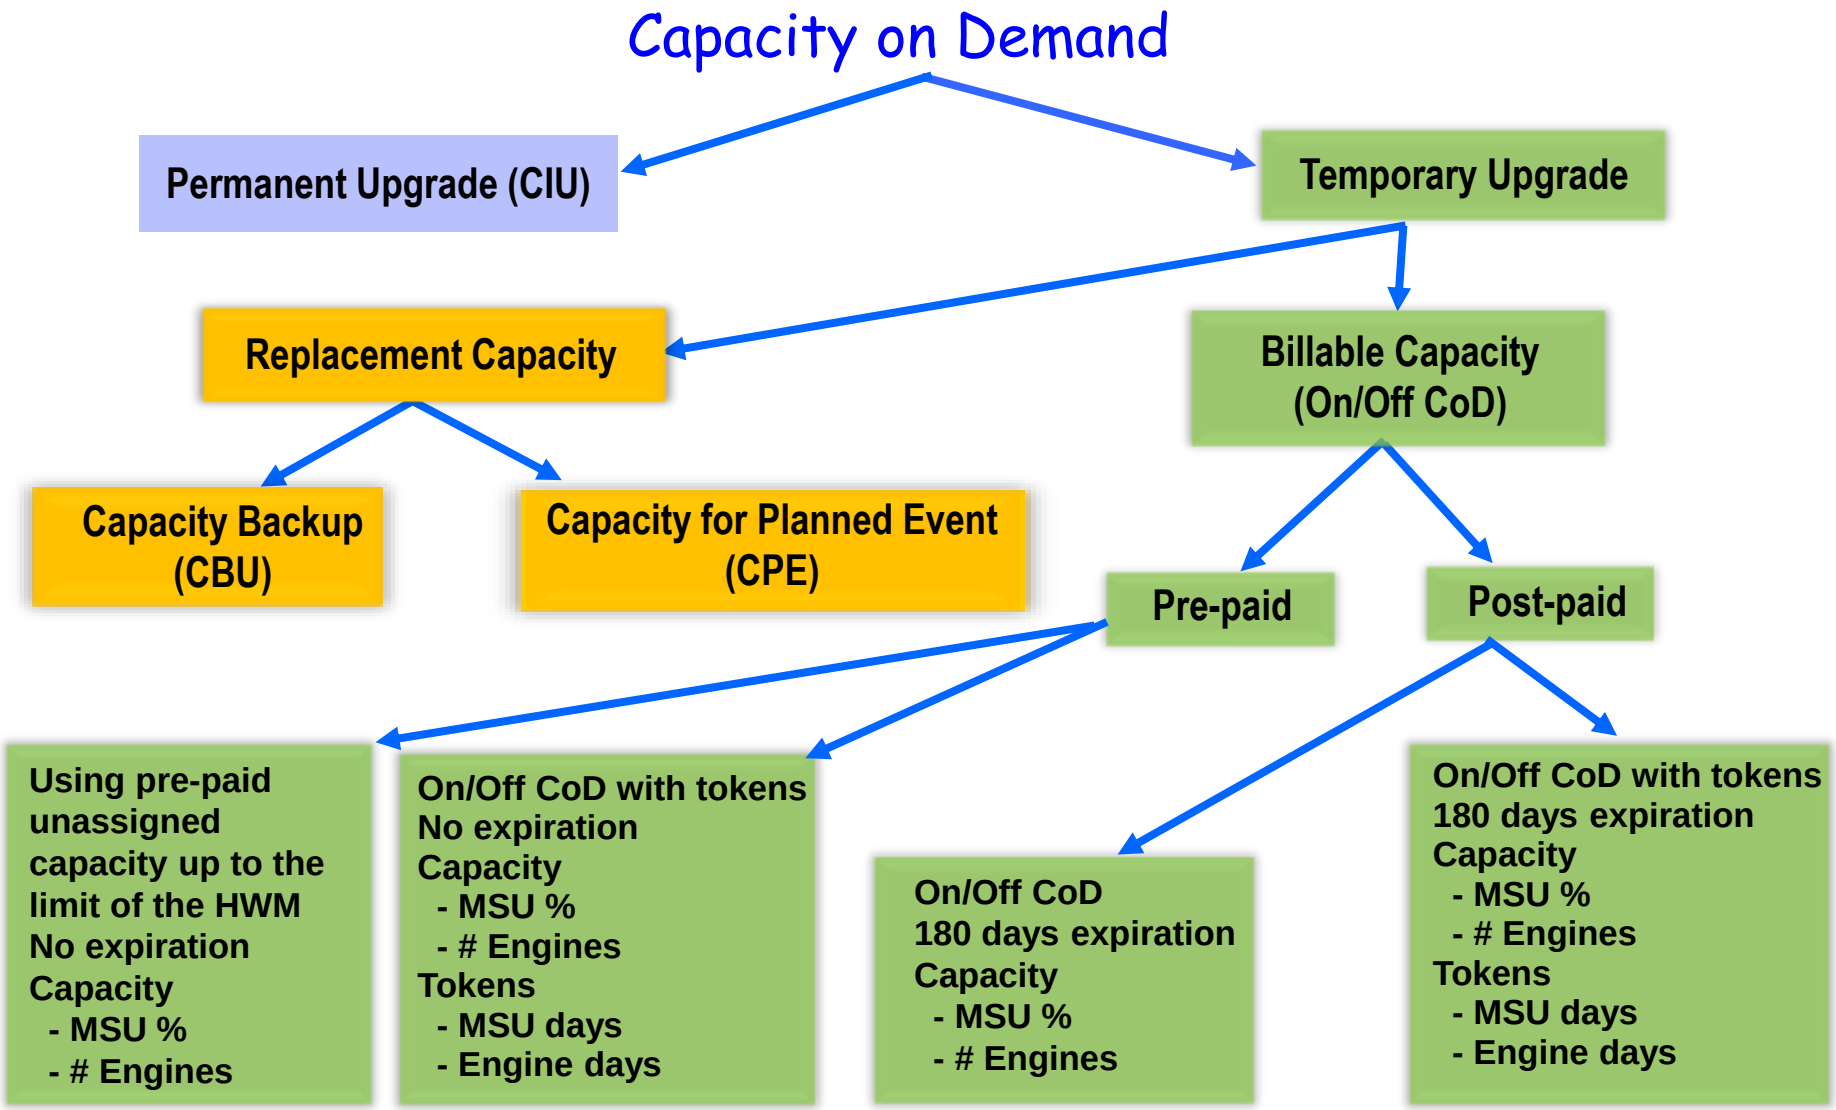

IBM z13
I/O Subsystem
Introduction

IBM z13: Optimized as a system of insight for digital business

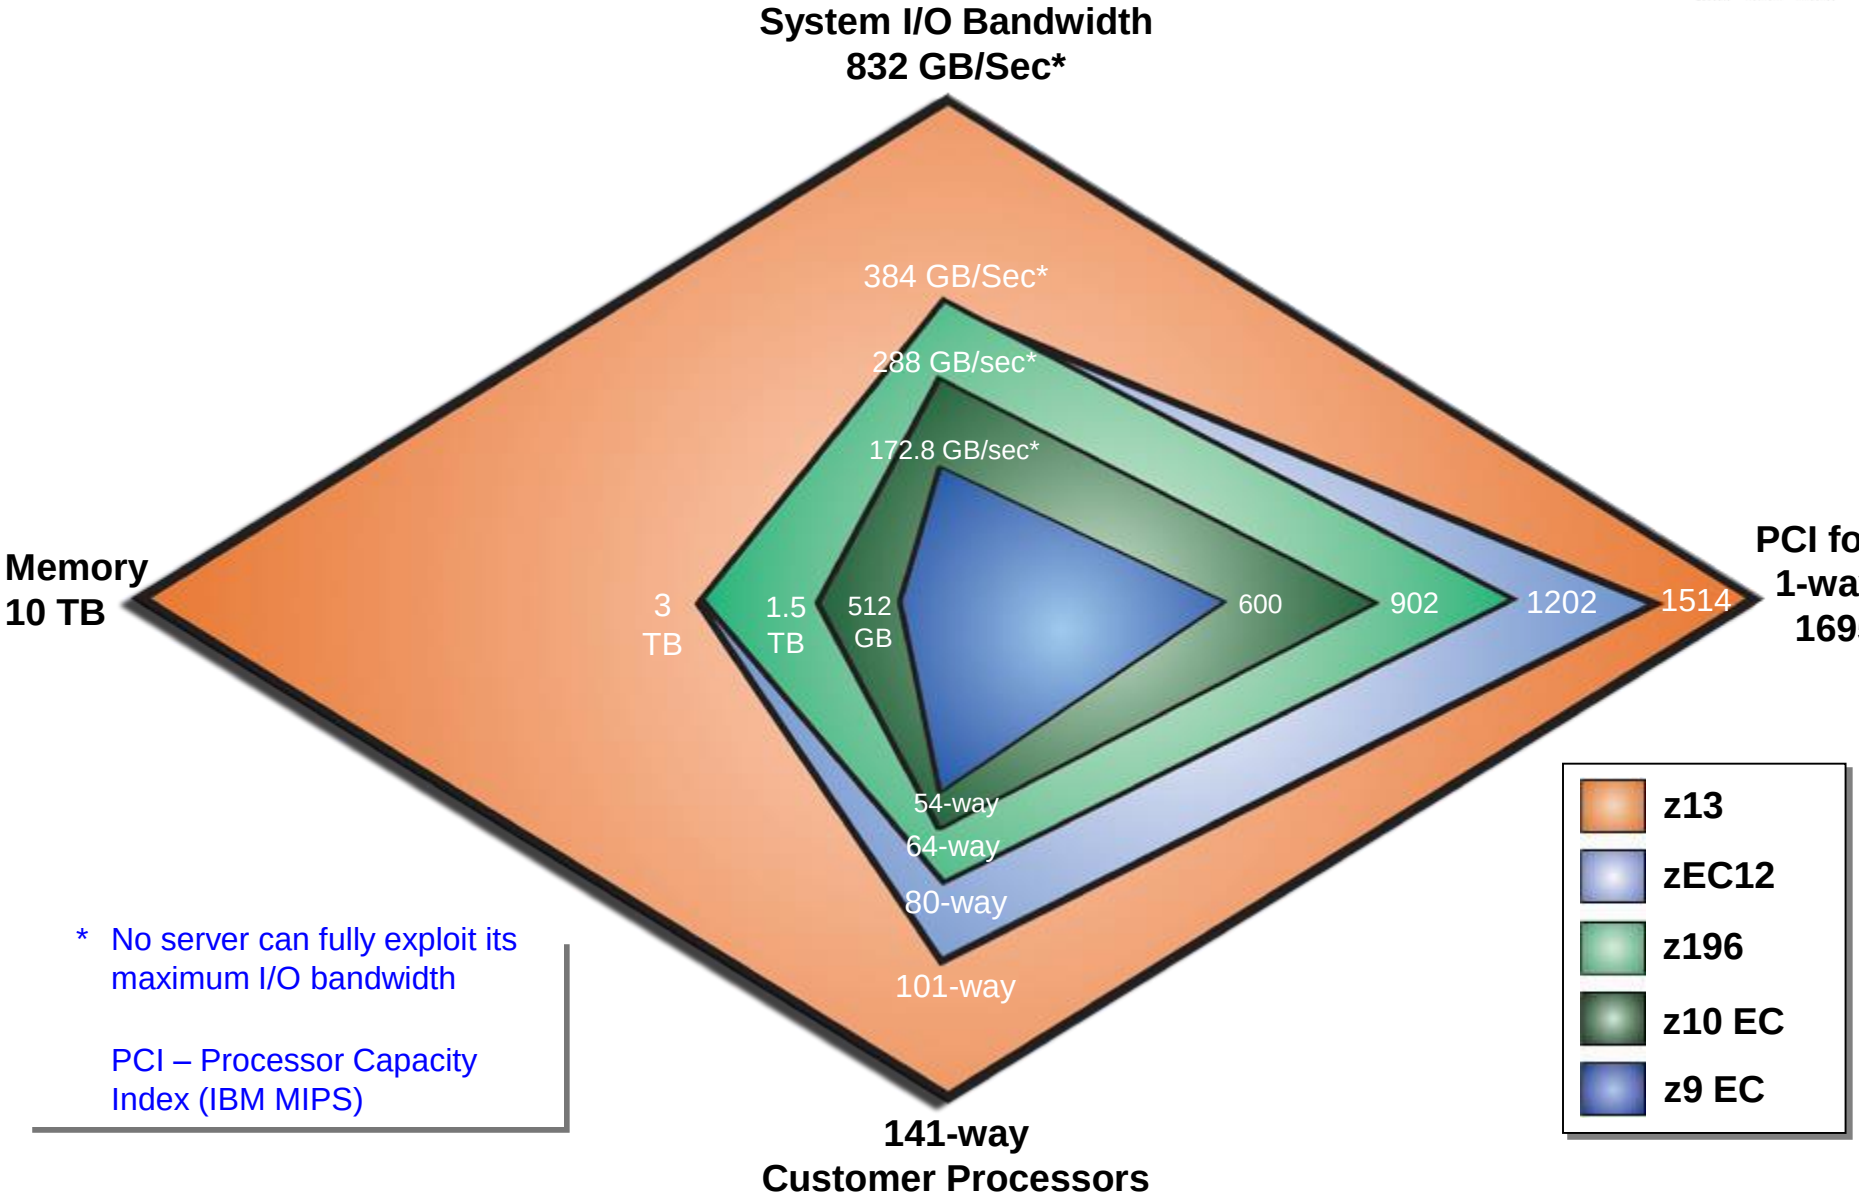

* No server can fully exploit its maximum I/O bandwidth

PCI – Processor Capacity Index (IBM MIPS)

Supported I/O Features

z13 "New Build" I/O and MES Features Supported

Note - "Plan Ahead" for I/O drawers is not offered on z13

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

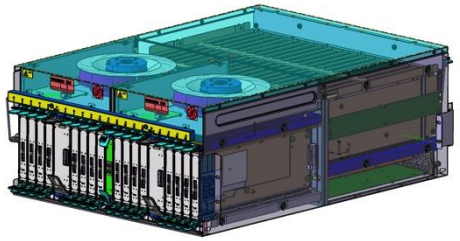

32 I/O slots

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

z13 "Carry Forward" I/O Features Supported

Note – "Plan Ahead" for I/O drawers is not offered on z13

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 z13)*
  - **zEDC Express**
  - **Flash Express**
  - **Not Supported: Crypto Express4S**

PCIe I/O drawer 32 I/O slots

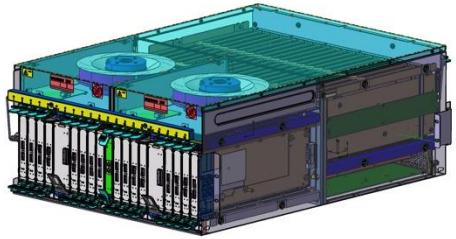

- **Features – I/O drawer (No MES adds)**
  - **FICON Express8** (SX and LX, 4 SFPs, 4 CHPIDs)
    SoD: IBM plans not to support FICON Express8 on the next high end z Systems server.
  - **Not Supported: ESCON, FICON Express4, OSA-Express3, ISC-3, and Crypto Express3**

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**

Operating System Support and Statements of Direction

Operating System Support for z13

- Currency is key to operating system support and exploitation of future servers
- The following releases of operating systems will be supported on z13
  (Please refer to PSP buckets for any required maintenance):

Operating System	Supported levels
z/OS	▪ z/OS V2.2 (Exploitation) – GA planned* September 30, 2015 ▪ z/OS V2.1 with PTFs (Exploitation) ▪ 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 z/OS V1.12 Defect Support: Offered 10/1/14 – 9/30/17)
Linux on 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.)
z/VM	▪ z/VM V6.3 with PTFs – Exploitation support ▪ 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
z/VSE	▪ 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

- 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.

A new hypervisor choice for Linux on z Systems

Statement of Direction for KVM on z Systems*

- *Expanded audience* for Linux on z Systems
  - KVM on z System will co-exist with z/VM
  - Attracting new clients with in house KVM skills
  - Simplified startup with standard KVM interfaces
- Support of modernized *open source* KVM hypervisor for Linux
  - Provisioning, mobility, memory over-commit
  - Standard management and operational controls
  - Simplicity and familiarity for Intel Linux users
- *Optimized for z System* scalability, performance, security and resiliency
  - Standard software distribution from IBM
- Flexible *integration to cloud* offerings
  - Standard use of storage and networking drivers (including SCSI disk)
  - No proprietary agent management
  - Off-the-shelf OpenStack and cloud drivers
  - Standard enterprise monitoring and automation (i.e. GDPS)

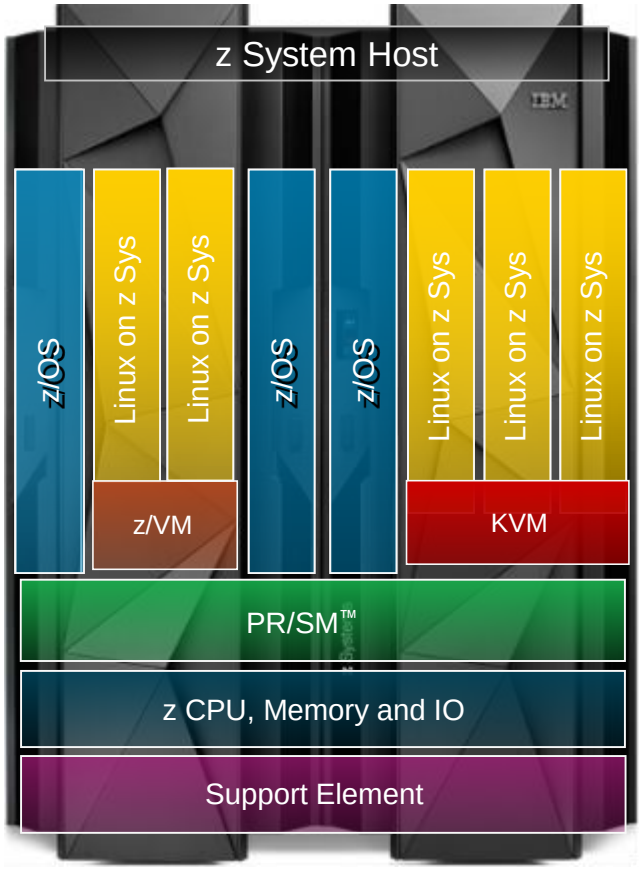

*All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Statements of Direction: z/VM

- z/VM support for Single Instruction Multiple Data (SIMD): In a future deliverable IBM intends to deliver support to enable z/VM guests to exploit the Vector Facility for z/Architecture (SIMD).
- Removal of support for Expanded Storage (XSTORE): z/VM V6.3 is the last z/VM release that will support Expanded Storage (XSTORE) for either host or guest usage. The IBM z13 server family will be the last z Systems server to support Expanded Storage (XSTORE). (Hint: Use large central memory instead)
- Product Delivery of z/VM on DVD/Electronic only: z/VM V6.3 will be the last release of z/VM that will be available on tape. Subsequent releases will be available on DVD or electronically.
- Stabilization of z/VM V6.2 support: The IBM z13 server family is planned to be the last z Systems server supported by z/VM V6.2 and the last z systems server that will be supported where z/VM V6.2 is running as a guest (second level). This is in conjunction with the statement of direction that the IBM z13 server family will be the last to support ESA/390 architecture mode, which z/VM V6.2 requires. z/VM V6.2 will continue to be supported until December 31, 2016, as announced in announcement letter # 914-012.

THANK YOU

Statements of Direction and Backup

Statements of Direction

- IBM plans to accept for review certification requests from cryptography providers by the end of 2015, and intends to support the use of cryptography algorithms and equipment from providers meeting IBM's certification requirements in conjunction with z/OS and z Systems processors in specific countries. This is expected to make it easier for customers to meet the cryptography requirements of local governments.
- KVM offering for IBM z Systems: In addition to the continued investment in z/VM, IBM intends to support a Kernel-based Virtual Machine (KVM) offering for z Systems that will host Linux on z Systems guest virtual machines. The KVM offering will be software that can be installed on z Systems processors like an operating system and can co-exist with z/VM virtualization environments, z/OS, Linux on z Systems, z/VSE and z/TPF. The KVM offering will be optimized for z Systems architecture and will provide standard Linux and KVM interfaces for operational control of the environment, as well as providing the required technical enablement for OpenStack for virtualization management, allowing enterprises to easily integrate Linux servers into their existing infrastructure and cloud offerings.
- In the first half of 2015, IBM intends to deliver a GDPS/Peer to Peer Remote Copy (GDPS/PPRC) multiplatform resiliency capability for customers who do not run the z/OS operating system in their environment. This solution is intended to provide IBM z Systems customers who run z/VM and their associated guests, for instance, Linux on z Systems, with similar high availability and disaster recovery benefits to those who run on z/OS. This solution will be applicable for any IBM z Systems announced after and including the zBC12 and zEC12

Statements of Direction

- Enhanced RACF password encryption algorithm for z/VM: In a future deliverable an enhanced RACF/VM password encryption algorithm is planned. This support will be designed to provide improved cryptographic strength using AES-based encryption in RACF/VM password algorithm processing. This planned design is intended to provide better protection for encrypted RACF password data in the event that a copy of RACF database becomes inadvertently accessible.

- IBM intends that a future release of IBM CICS Transaction Server for z/OS will support 64-bit SDK for z/OS, Java Technology Edition, Version 8 (Java 8). This support will enable the use of new facilities delivered by IBM z13 which are exploited by Java 8, including Single Instruction Multiple Data (SIMD) instructions for vector operations and simultaneous multithreading (SMT).

- z/VM support for Single Instruction Multiple Data (SIMD): In a future deliverable IBM intends to deliver support to enable z/VM guests to exploit the Vector Facility for z/Architecture (SIMD).

- Removal of support for Expanded Storage (XSTORE): z/VM V6.3 is the last z/VM release that will support Expanded Storage (XSTORE) for either host or guest usage. The IBM z13 server family will be the last z Systems server to support Expanded Storage (XSTORE).

Statements of Direction

- The IBM z13 will be the last z Systems server to support running an operating system in ESA/390 architecture mode; all future systems will only support operating systems running in z/Architecture mode. This applies to operating systems running native on PR/SM as well as operating systems running 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 contracts, and they will not be usable on systems beyond IBM z13. However, 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.

- Stabilization of z/VM V6.2 support: The IBM z13 server family is planned to be the last z Systems server supported by z/VM V6.2 and the last z systems server that will be supported where z/VM V6.2 is running as a guest (second level). This is in conjunction with the statement of direction that the IBM z13 server family will be the last to support ESA/390 architecture mode, which z/VM V6.2 requires. z/VM V6.2 will continue to be supported until December 31, 2016, as announced in announcement letter # 914-012.

- Product Delivery of z/VM on DVD/Electronic only: z/VM V6.3 will be the last release of z/VM that will be available on tape. Subsequent releases will be available on DVD or electronically.

- Removal of support for Classic Style User Interface on the Hardware Management Console and Support Element: The IBM z13 will be the last z Systems server to support Classic Style User Interface. In the future, user interface enhancements will be focused on the Tree Style User Interface.

Statements of Direction

- Removal of support for the Hardware Management Console Common Infrastructure Model (CIM) Management Interface: IBM z13 will be the last z Systems server to support the Hardware Console Common Infrastructure module (CIM) Management Interface. The Hardware Management Console Simple Network Management Protocol (SNMP), and Web Services Application Programming Interfaces (APIs) will continue to be supported.

- The IBM z13 will be the last z Systems server to support FICON Express8 channels: IBM z13 will be the last high-end server to support FICON Express8. Enterprises should begin migrating from FICON Express8 channel features (#3325, #3326) to FICON Express16S channel features (#0418, #0419). FICON Express8 will not be supported on future high-end z Systems servers as carry forward on an upgrade.

- The IBM z13 will be the last z Systems server to offer ordering of FICON Express8S channel features. Enterprises that have 2 Gb device connectivity requirements must carry forward these channels.

- Removal of an option for the way shared logical processors are managed under PR/SM LPAR: The IBM z13 will be the last high-end server to support selection of the option to "Do not end the timeslice if a partition enters a wait state" when the option to set a processor run time value has been previously selected in the CPC RESET profile. The CPC RESET profile applies to all shared logical partitions on the machine, and is not selectable by logical partition.

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.

Statements of Direction

- IBM intends to provide support for the Read Diagnostic Parameters Extended Link Service command for fiber channel SANs as defined in the T11.org FC-LS-3 draft standard. Support for the Read Diagnostic Parameters Extended Link Service command is intended to improve SAN reliability and fault isolation.

- The IBM z13 will be the last generation of z Systems hardware servers to support configuring OSN CHPID types. OSN CHPIDs are used to communicate between an operating system instance running in one logical partition and the IBM Communication Controller for Linux on z Systems (CCL) product in another logical partition on the same CPC. See announcement letter #914-227 dated 12/02/2014 for details regarding withdrawal from marketing for the CCL product.

IBM z13 and zBX Model 004

IBM z13 (2964)

- Available – March 9, 2015
- 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 CSSs, up to 85 LPARs
  - 4 Subchannel Sets per CSS
  - FICON Express16S or 8S (8 Carry forward)
  - OSA Express5S (4S carry forward)
  - HiperSockets – up to 32
  - Flash Express
  - zEnterprise Data Compression
  - RDMA over CE (RoCE) with SR-IOV Support
- Crypto Express5S
- Parallel Sysplex clustering, PCIe Coupling, Internal Coupling and InfiniBand Coupling Links
- IBM zAware: z/OS and Linux on z Systems
- Operating Systems
  - z/OS, z/VM, z/VSE, z/TPF, Linux on z Systems

IBM zBX Model 4 (2458-004)

- Available – March 9, 2015
- 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 XI50z with Firmware 7.0
- Operating Systems
  - AIX 5.3 and higher
  - Linux on System x
  - Microsoft Windows Server on System x
- Hypervisors
  - KVM Hypervisor on System x
  - PowerVM Enterprise Edition

z13 Node L4 Cache Design with Non-Data Inclusive Coherent (NIC) Directory, Intra-Node Snoop Interface and Inter-Node Snoop Interface

zEC12 Book Inclusive L4 Design
Two 192 MB L4s
384 MB Inclusive L4 per Book

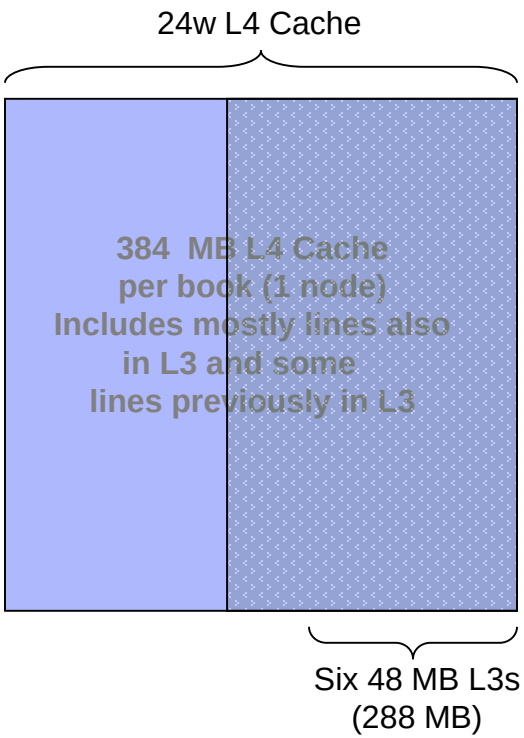

z13 Non-Inclusive L4 Design
Two 480 MB L4s with L3 NIC Directories
≈ 1344 MB of Inclusive L4 per Drawer
(960 L4 + 384 L3)

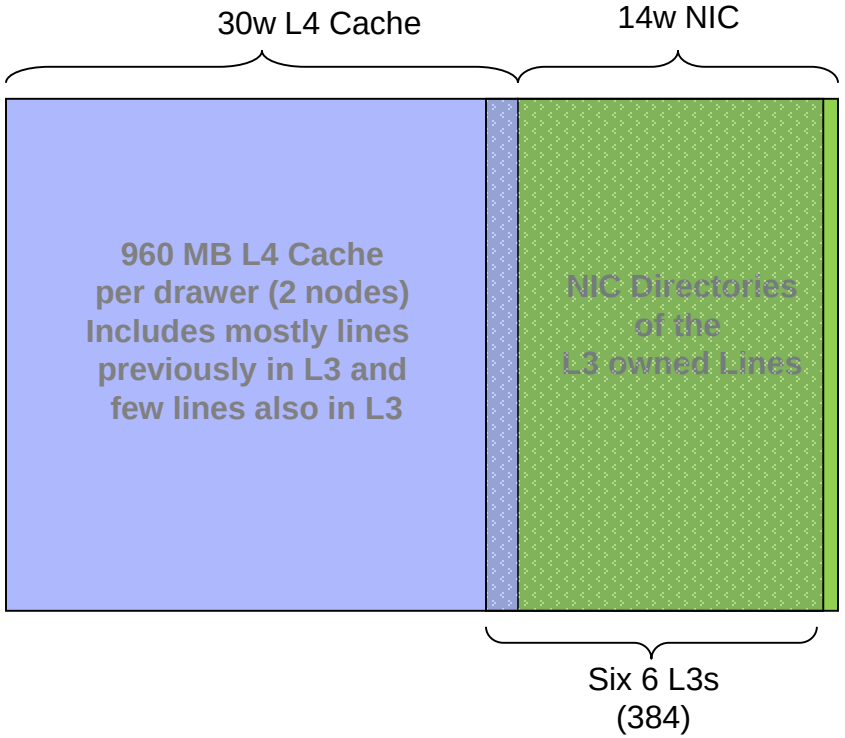

On z13, L3 lines in a node can be accessed over the **X-bus L3 – L3** using the Intra-Node Snoop Interface without duplicating them in L4. Inter-Node snoop traffic can still be handled effectively.

z13 Compression and Cryptography Accelerator

- Coprocessor dedicated to each core
  (Was shared by two cores on z196)
  - Independent compression engine
  - Independent cryptographic engine
  - Available to any processor type (CP, zIIP, IFL)
  - Owning processor is busy when its coprocessor is busy
  - Instructions available to any processor type
- Data compression/expansion engine
  - Static dictionary compression and expansion
- CP Assist for Cryptographic Function
  - Supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on z Systems
  - DES, TDES – Clear and Protected Key
    TDES: Up to double the throughput of zEC12 CPACF
  - AES128, 192, 256 – Clear and Protected Key
    AES: Up to double the throughput of zEC12 CPACF
  - SHA-1 (160 bit) – Clear Key
  - SHA-256, -384, -512 – Clear Key
    SHA: Up to 3.5 times the throughput of zEC12 CPACF
  - PRNG – Clear Key
  - DRNG – Clear Key
  - CPACF FC 3863 (No Charge – Export Control) is required to enable some functions and to support Crypto Express5S

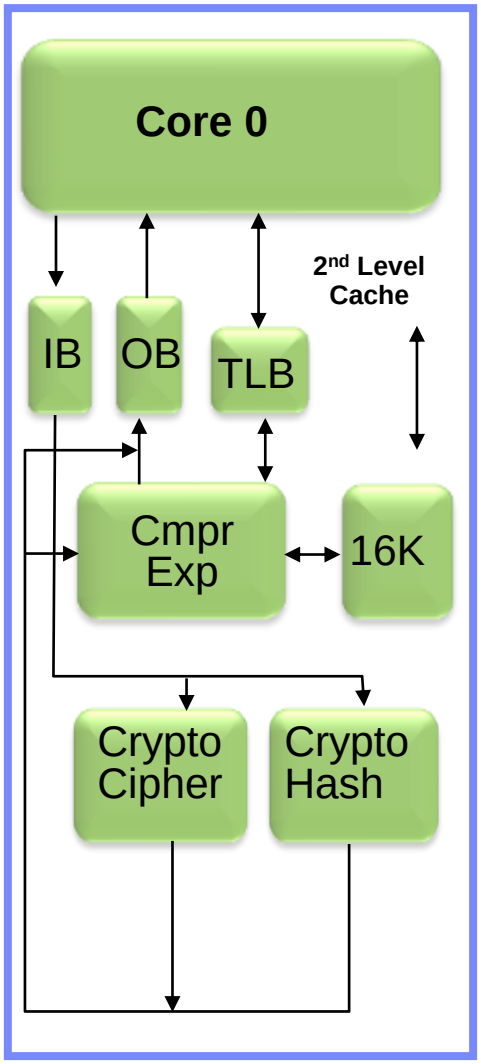

Processor Unit (Core) Locations: Customer, SAP, IFP and Spare

z13		1st Drawer				2nd Drawer				3rd Drawer				4th Drawer			
Model	Cust PUs	Cust PUs	SAPs	IFP	Spare	Cust PUs	SAPs	IFP	Spare	Cust PUs	SAPs	IFP	Spare	Cust PUs	SAPs	IFP	Spare
NE1	141	34	6	1	1	35	6	0	1	36	6	0	0	36	6	0	0
NC9	129	31	6	1	1	32	6	0	1	33	6	0	0	33	6	0	0
N96	96	31	6	1	1	32	6	0	1	33	6	0	0				
N63	63	31	6	1	1	32	6	0	1								
N30	30	30	6	1	2												

- **PUs can be configured as CPs, IFLs, Unassigned IFLs, zIIPs, ICFs or Additional SAPs**
  - zAAPs discontinued as per SOD
  - zIIP to CP purchase ratio is 2:1
  - Additional SAPs + Permanent SAPs may not exceed 32
  - Any unconfigured PU can act as an additional Spare PU
- **Upgrades available between any models**
  - Achieved via concurrent *drawer* add from model N30 to model NC9
  - Achieved via combination of *drawer* add and *drawer* replacement to model NE1

z13 Memory DIMMs and Plugging

- **z13 Memory Plugging**
  - Six memory controllers per drawer, one per PU chip, three per node
  - Each memory controller supports five DIMM slots
  - Four or five memory controllers per drawer will be populated (20 or 25 DIMMs)
  - Different memory controllers may have different size DIMMs
- **Maximum Client Memory Available**
  - Remember **RAIM** – 20% of DIMM memory is used only for error recovery
  - Minimum memory per drawer: 320 GB RAIM = 256 GB addressable
  - Maximum memory per drawer: 3200 GB RAIM = 2560 GB addressable
  - To determine maximum possible customer memory from the DIMM configuration: Calculate addressable memory, subtract 96 GB, and **round down if necessary to an offered memory size**

DIMM Size	z13 Feature (5 DIMMs) RAIM and Addressable Size
16 GB	#1610 = 80 GB RAIM, 64 GB Addressable Memory
32 GB	#1611 = 160 GB RAIM, 128 GB Addressable Memory
64 GB	#1612 = 320 GB RAIM, 256 GB Addressable Memory
128 GB	#1613 = 640 GB RAIM, 512 GB Addressable Memory

PR/SM Partition Logical Processor and Memory Assignment

- **System z9 EC to zEnterprise EC12**
  - Memory Allocation Goal: Stripe across all the available books in the machine
    Advantage: Exploit fast book interconnection; spread the memory controller work; smooth performance variability
  - Processor Allocation Goal: Assign all logical processors to one book; packed into chips of that book. Cooperate with operating system use of HiperDispatch
    Advantage: Optimal shared cache usage
- **z13**
  - Memory Allocation Goal: Assign all memory in one drawer striped across the two nodes.
    Advantage: Lower latency memory access in drawer; smooth performance variability across nodes in the drawer
  - Processor Allocation Goal: Assign all logical processors to one drawer; packed into chips of that drawer. Cooperate with operating system use of HiperDispatch
  - Reality: Easy for any given partition. Complex optimization for multiple logical partitions because some need to be split among drawers.

z13 Processor and Memory Assignment and Optimization

- **Default processor assignments by POR, MES adds, and On Demand activation**:
  - Assign IFLs and ICFs to cores on chips in "high" drawers working down
  - Assign CPs and zIIP in low drawers working up.
  - Objective: Keep "Linux Only", "IBM zAware" and "Coupling Facility" using IFLs and ICFs "away" from "ESA/390" partitions running z/OS on CPs and zIIPs and in different drawers if possible.
- **PR/SM makes optimum available memory and logical processor assignment at activation**
  - Logical Processors specified in the Image Profile, are assigned a ***core*** if Dedicated or a ***"home" drawer, node and chip*** if Shared. Later, if it becomes a HiperDispatch "Vertical High", a Shared Logical Processor is assigned a specific core.
  - Ideally assign all memory in one drawer with the processors if everything "fits"
  - With memory striped across drawers with processors if memory or processors must be split
- **PR/SM optimizes resource assignment when triggered**
  - **Triggers:** Available resources changes: partition activation or deactivation or significant processor entitlement changes, dynamic memory increases or processor increases or decreases (e.g. by CBU) or MES change.
  - Examines partitions in priority order by the size of their "processor entitlement" (dedicated processor count or shared processor pool allocation by weight) to determine priority for optimization
  - Changes logical processor "home" drawer/node/chip assignment
  - Moves processors to different chips, nodes, drawers (**LPAR Dynamic PU Reassignment)**
  - Relocates partition memory to active memory in a different drawer or drawers using the newly optimized **Dynamic Memory Relocation (DMR)**, also exploited by Enhanced Drawer Availability (**EDA**).
  - If available but inactive memory hardware is present (e.g. hardware driven by Flexible or Plan Ahead) in a drawer where more active memory would help: activate it, reassign active partition memory to it, and deactivate the source memory hardware, again using DRM.
    (PR/SM can use all memory hardware but concurrently enables no more memory than the client has paid to use.)

Last Slide