

Frühjahrstagung 2015
z/VSE, z/VM und Linux on z Systems
27. – 29. April 2015 in Berlin

G02

Neue Server Generation: IBM z13

Dr. Manfred Gnirss
Senior IT Specialist
IBM Client Center, Böblingen
gnirss@de.ibm.com



**Hotel The Westin Grand
Berlin**

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 z13 platform positioning

Platform Core Capabilities:

Transaction Processing

Data Serving

Mixed Workloads

Operational Efficiency

Trusted and Secure Computing


Reliable, Available, Resilient

Virtually Limitless Scale



- ***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***

z13 Continues the CMOS Mainframe Heritage Begun in 1994

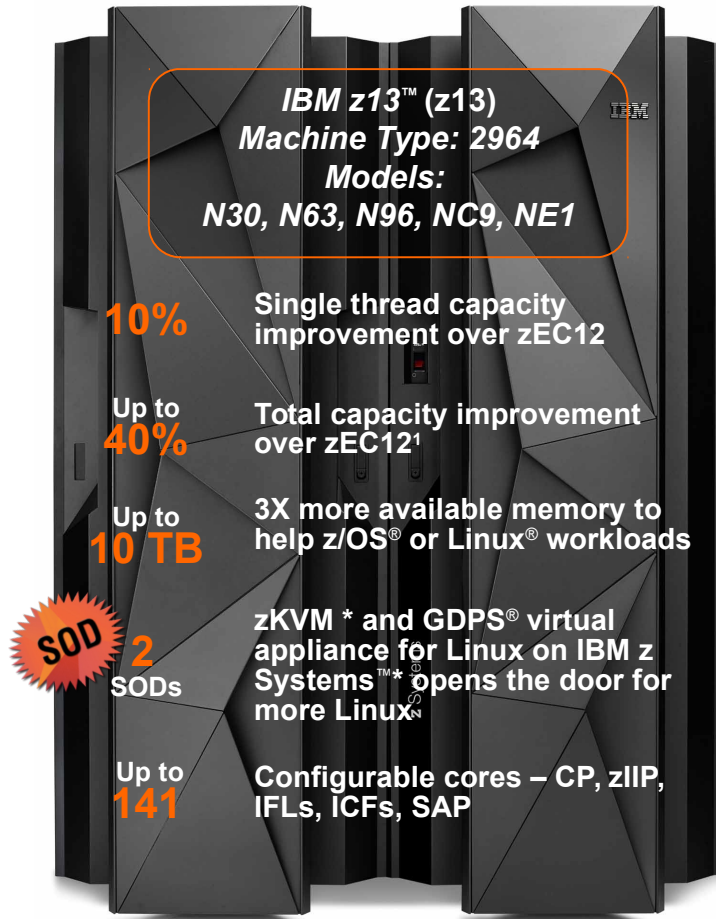


* MIPS Tables are NOT adequate for making comparisons of z Systems processors. Additional capacity planning required

** Number of PU cores for customer use

Introducing the IBM z13

The mainframe optimized for the digital era




- Performance, scale, intelligent I/O and security enhancements to support transaction growth in the mobile world
- More memory, new cache design, improved I/O bandwidth and compression help to serve up more data for analytics
- Enterprise grade Linux solution, open standards, enhanced sharing and focus on business continuity to support cloud

Upgradeable from IBM zEnterprise® 196 (z196) and IBM zEnterprise EC12 (zEC12)


¹ Based on preliminary internal measurements and projections. Official performance data will be available upon announce and can be obtained online at LSPR (Large Systems Performance Reference) website at: <https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument> . Actual performance results may vary by customer based on individual workload, configuration and software levels

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

z13 Radiator-based Air cooled – Front View (Model NC9 or NE1)



z13 Support Elements - Under the covers (Front View)




Two 1U units

Smart Card Readers
 Internal USB attached
 units to support Flash
 Express and Feature on
 Demand.



Two Displays with
 Keyboards

IBM z13: Advanced system design optimized for digital business



Accelerate Key Workloads with Special-Purpose Hardware

- **On-processor**
 - Crypto (CPACF), Compression, SIMD
 - 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**
 - RDMA over Converged Ethernet SMC-R (RoCE)
- **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 DB2 Analytics
Accelerator built on
Netezza Technology*



Operating Systems focused on exploiting hardware innovation

z/OS Version 2.1



- Improved price performance for zIIP workloads with SMT
- Support new analytics workloads with SIMD
- New Crypto capabilities for faster encryption
- Large memory to improve performance and enable new applications



z/VM Version 6.3



- Improved price performance with support for multi-threading technology
- Support for twice as many processors (z13 only)
- Improved systems management and economics
- Embracing Open Standards and Open Source Interoperability
- Supports more virtual servers than any other platform in a single footprint



z/VSE Version 5.1



- Reduced risk of access from unauthorized users
- Reduced memory constraints
- Wide portfolio using Linux on z Systems
- Continued system usability enhancements with CICS Explorer
- More efficient communications



Linux on z System



- Multithreading allows for per core software savings
- Ability to host and manage more workloads efficient and cost-effective
- Automatic identification of unusual messages
- Integrated continuous availability & disaster recovery solution



z13 Processor and Memory subsystem

z13 z/Architecture Extensions


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

z13 SCM Vs zEC12 MCM Comparison


z13 Single Chip Modules (SCMs)

- **Processor Unit (PU) SCM**
 - 68.5mm x 68.5mm – fully assembled
 - PU Chip area 678 mm²
 - Eight core chip with 6, 7 or 8 active cores
- **System Controller (SC) SCM**
 - 68.5mm x 68.5mm – fully assembled
 - SC Chip area 678 mm²
 - 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

Single PU Chip
without
Module Thermal Cap




Single SC Chip
without
Module Thermal Cap

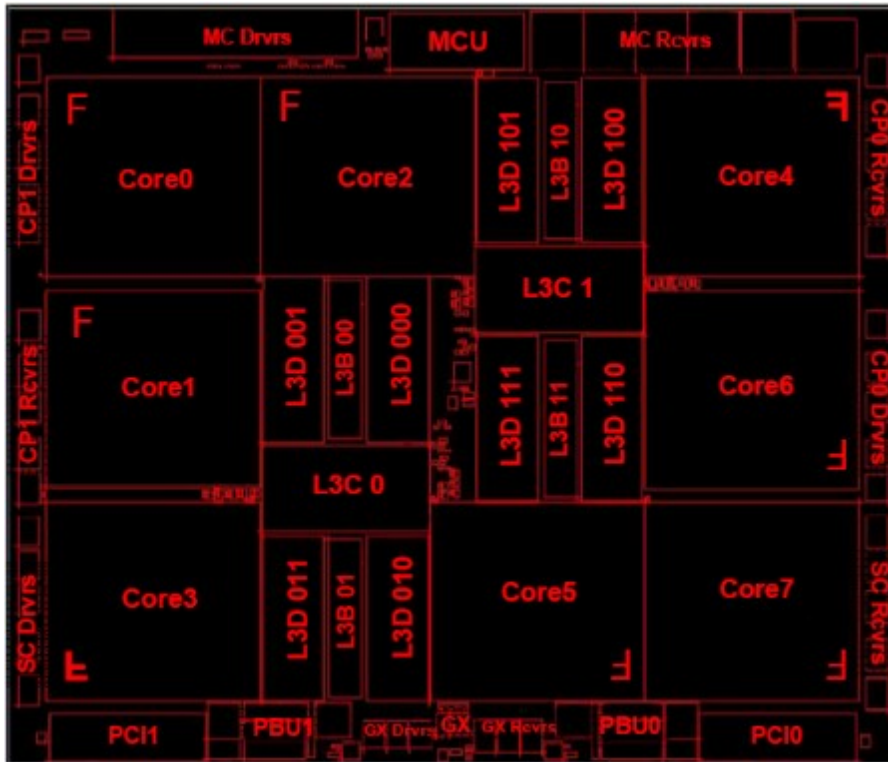


zEC12 Multi Chip Module (MCM)

- **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 System Controller (SC) chips per MCM**
 - 192 MB 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 8-Core Processor Unit (PU) Chip Detail

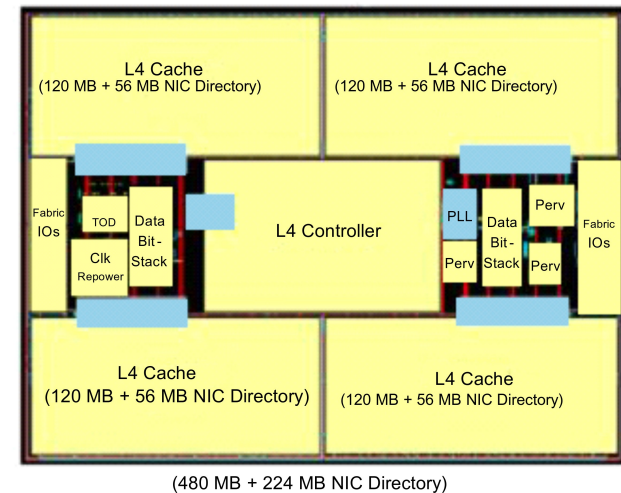



- **14S0 22nm SOI Technology**
 - 17 layers of metal
 - 3.99 Billion Transistors
 - 13.7 miles of copper wire
- **Chip Area**
 - 678.8 mm²
 - 28.4 x 23.9 mm
 - 17,773 power pins
 - 1,603 signal I/Os


- **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 InfiniBand I/O bus
 - Two PCIe I/O buses
- **Memory Controller (MCU)**
 - Interface to controller on memory DIMMs
 - Supports RAIM design

z13 System 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)
 - 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**



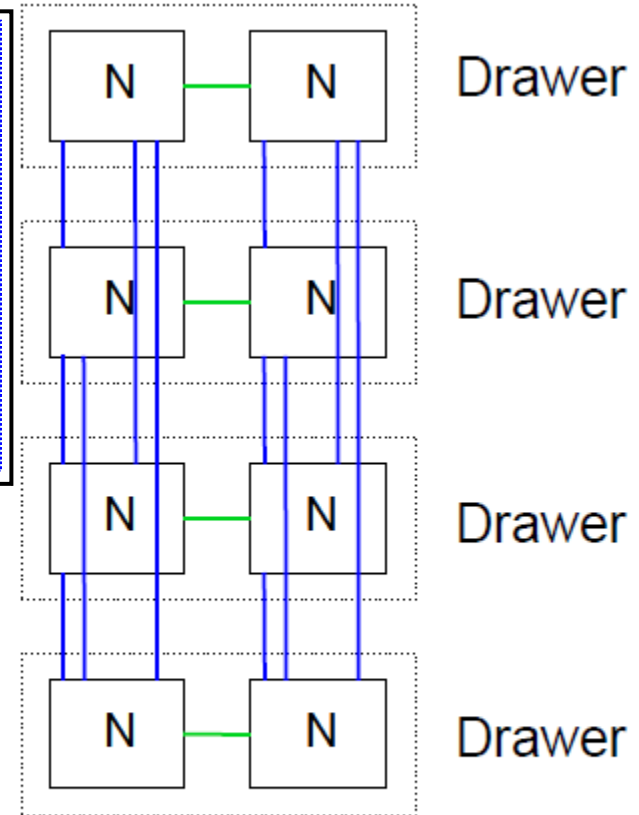
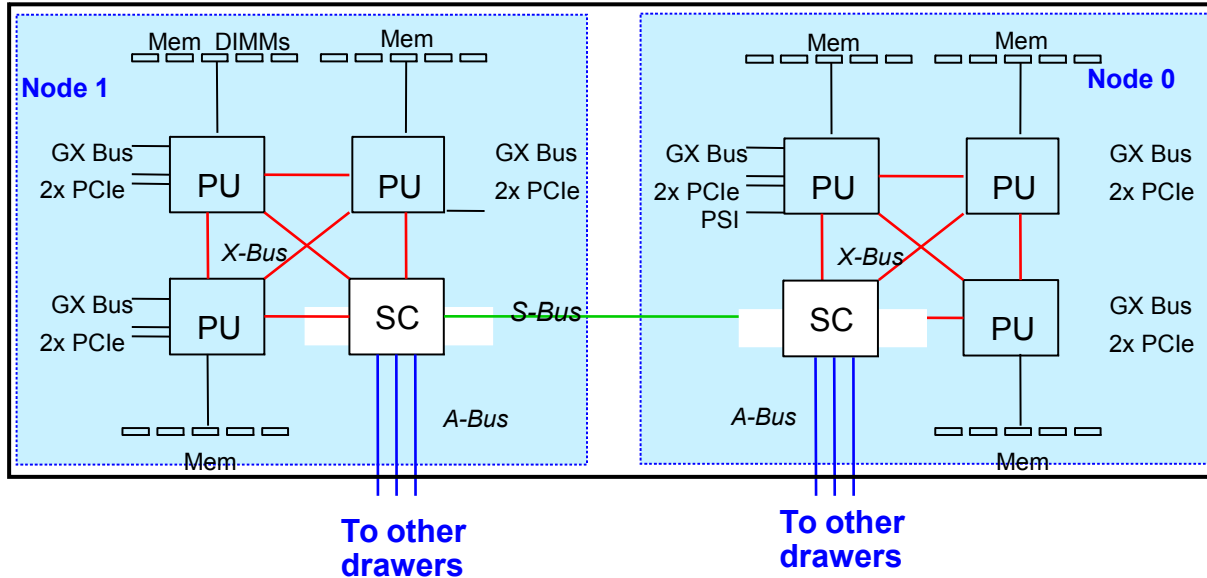
z13 Processor Drawer (Top View)



- **Two physical nodes, left and right**
- **Each logical node:**
 - One SC chip (480 MB L4 cache)
 - Three PU chips
 - Three Memory Controllers: One per CP Chip
 - Five DDR3 DIMM slots per Memory Controller: 15 total per logical node
- **Each drawer:**
 - Six PU Chips: 39 active PUs (42 in z13 Model NE1)
 - Two SC Chips (960 MB L4 cache)
 - Populated DIMM slots: 20 or 25 DIMMs to support up to 2,560 GB of addressable memory (3,200 GB RAIM)
 - Water cooling for PU and SC chips
 - 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 Drawer Structure and Interconnect

Fully Populated Drawer

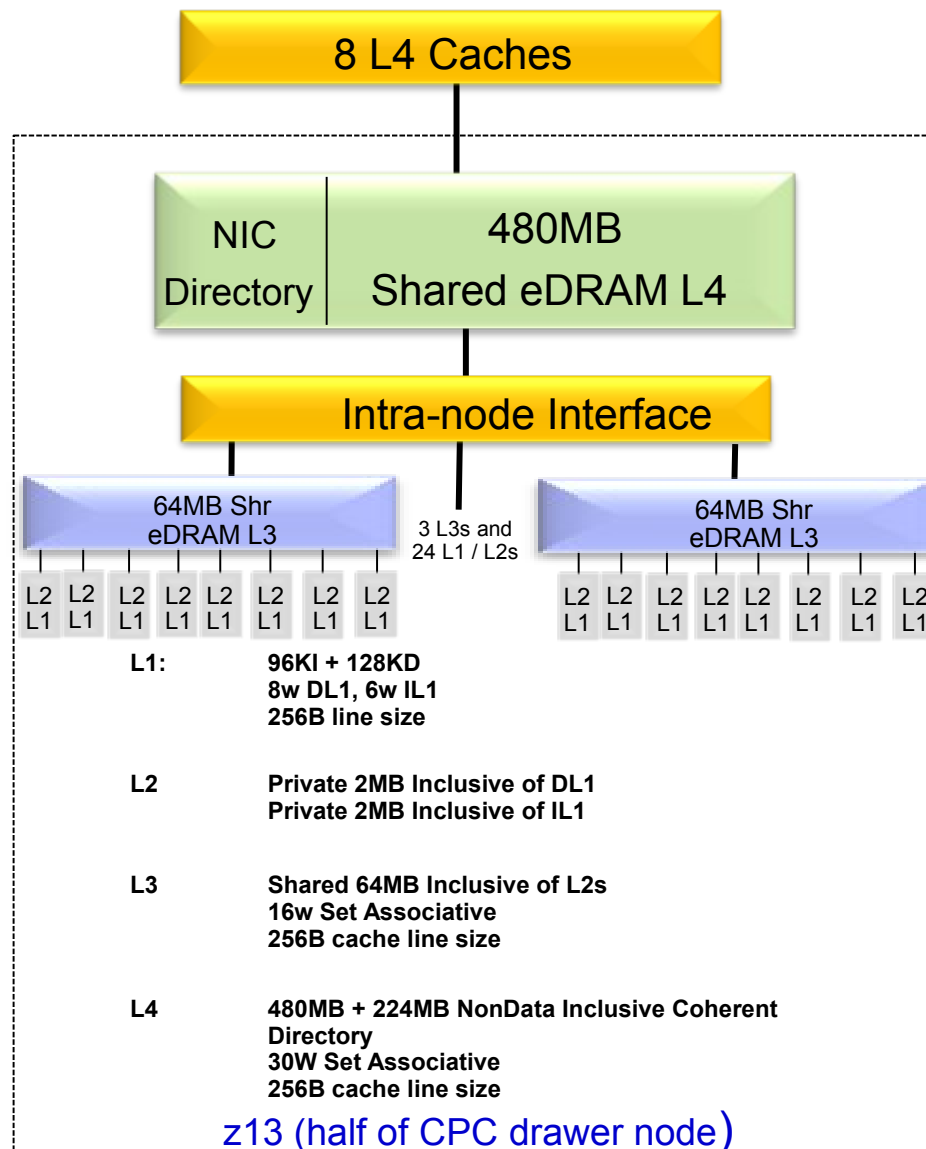
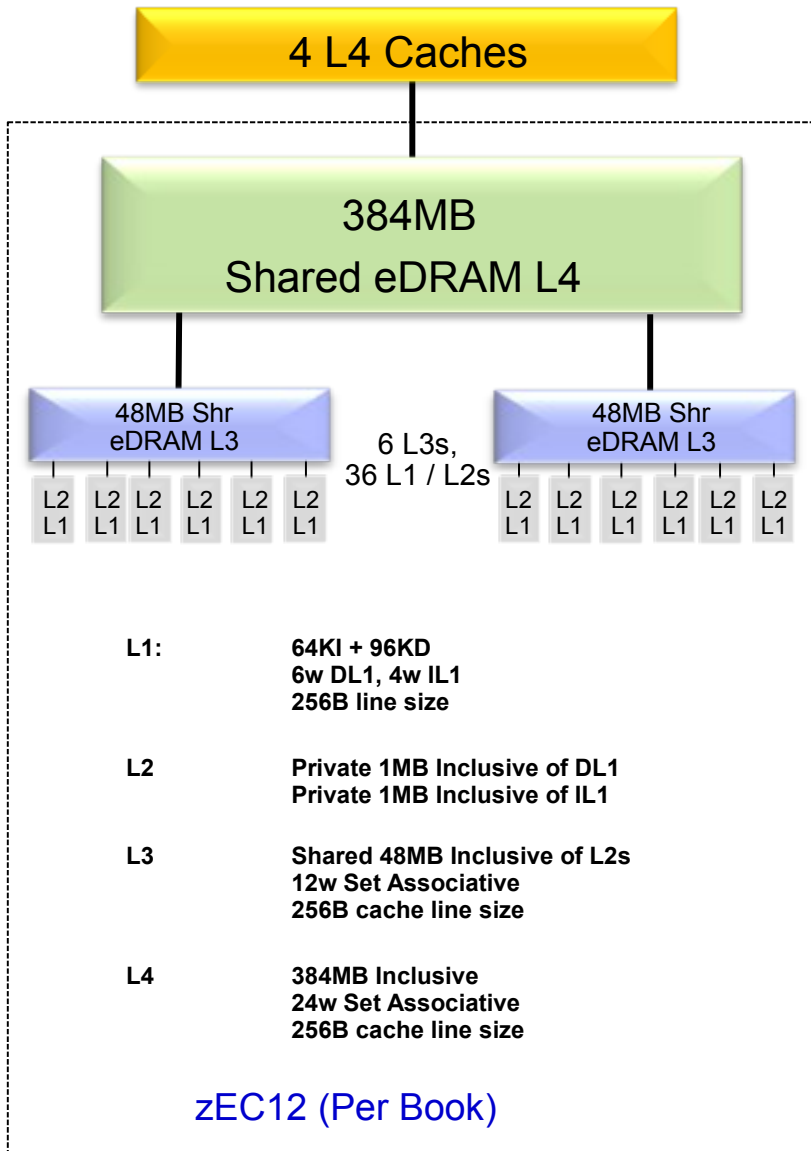


4 Drawer System Interconnect

Physical node: (Two per drawer)

- **Chips**
 - Three PU chips
 - One SC chip (480 MB L4 cache)
- **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 (up to 2.5TB)
- **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**

z System Cache Topology – zEC12 vs. z13 Comparison



Large Memory Value – Potential Performance Gains

- **Potential Latency Reduction for OLTP workloads**
 - Response time reductions
 - Increased transaction rates
- **Enables In-Memory Databases**
 - Dramatic reduction in response time by avoiding I/O wait
- **Batch Window Reduction**
 - More concurrent Workloads
 - Shorter elapsed times for Jobs
- **Reducing time to insight for analytic workloads**
 - Process data more efficiently - keep pace with influx of data
 - Reduces time to get from *raw data* to *business insight*
- **CPU performance improvements**
 - Improves response time and shrinks batch windows
 - Reduce the need for application/system redesign to meet service goals
 - Reduction in CPU time per transaction
 - Run more work at the same HW and SW MSU rating **-or-**
 - Run the same workload with lower HW and SW MSU rating

Simultaneous Multithreading (SMT)


Simultaneous Multithreading - Background

- **SMT enables to run multiple threads on a single core**
 - Other processor families (i.e. x86, etc.) already have similar support
 - Each thread runs slower than a non-SMT core, but the combined ‘threads’ throughput is higher. The overall throughput benefit depends on the workload
- **Hardware support**
 - Single thread (ST) operation
 - SMT operation with seamless transition between ST and SMT
 - Precise metering of SMT utilization => Monitors Dashboard
- **Software must actually enable the use of SMT operation**
 - You must have software at levels that can exploit SMT.
 - Use of SMT is on a per-LPAR basis
 - The support is present in the z13
 - The OS(es) must actually issue instructions to switch into SMT mode
 - Enabling the use of SMT is unidirectional
 - Once the OS switches, the only way back to ST mode is via a disruptive action (re-activate the partition or to re-IPL it).
 - With the SMT enabled mode it is possible to dynamically switch between MT-1 (multi thread) and MT-2 mode for the processor types that support MT-2

Simultaneous Multithreading Value Example




Two tasks, one core



Two tasks, two threads

Additional capacity

Elapsed Time



Task A

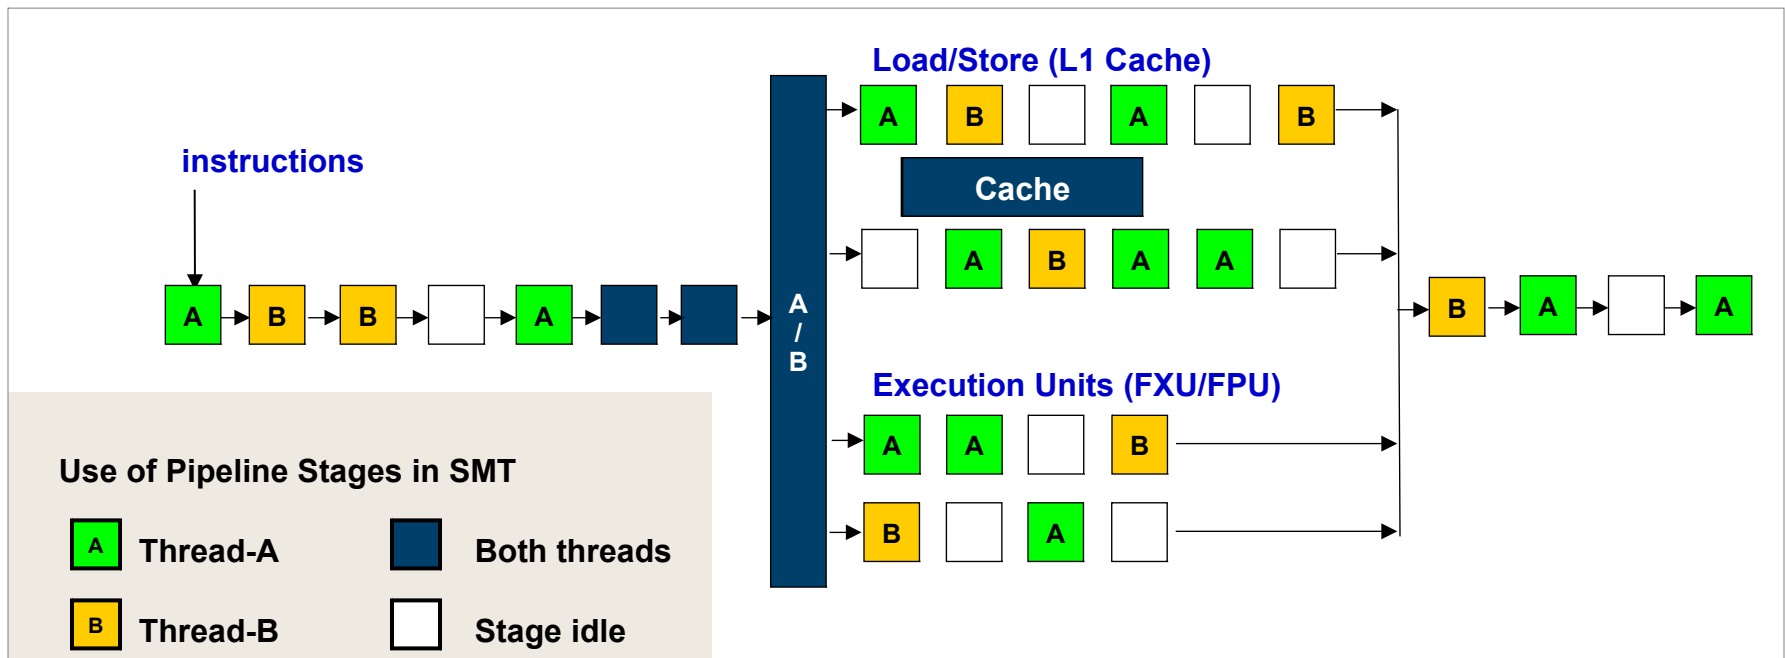


Task B

(assumes one thread delivers 70% of a core)

Simultaneous Multithreading – The Technology


- **Simultaneous Multithreading (SMT) technology**
 - Multiple programs (software threads) run on the same processor core
 - More efficient use of the core hardware
- **Active threads share core resources**
 - In space: data and instruction caches, TLBs, branch history tables, etc.
 - In time: pipeline slots, execution units, address translator, etc.
- **Increases overall throughput per core when SMT is active**
 - Amount that increase, varies widely with workload – typically 1.X-1.Y >1
 - Each thread runs more slowly than on a single-thread core



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/IFL with exploitation of the SMT option is:
 - zIIP is **38%** greater than a zEC12 zIIP
 - IFL is **32%** greater than a zEC12 IFL
 - zIIP is **72%** greater than a z196 zIIP
 - IFL is **65%** greater than a z196 IFL
- SMT exploitation: z/VM V6.3 + PTFs for IFLs and z/OS V2.1 + PTFs in an LPAR for zIIPs
- The use of SMT mode can be enabled on an LPAR by LPAR basis via operating system parameters.
 - When enabled, z/OS can transition dynamically between MT-1 (multi thread) and MT-2 modes with operator commands.
- Notes:
 1. SMT is designed to deliver better overall capacity (throughput) for many workloads. Thread performance (instruction execution rate for an individual thread) may be faster running in single thread mode.
 2. Because SMT is not available for CPs, LSPR ratings do not include it

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



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

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

Single Instruction Multiple Data (SIMD)

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




Sum and Store

Instruction is performed for every data element

SIMD

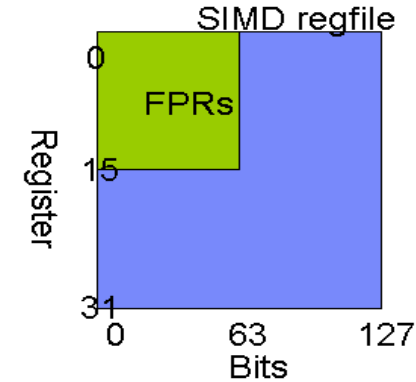

SINGLE INSTRUCTION, MULTIPLE DATA



Sum and Store

Perform instructions on every element at once

z System SIMD Hardware Accelerator

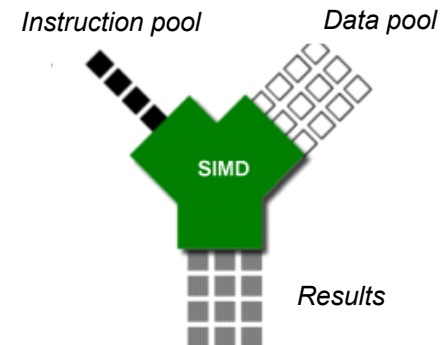



Operates on three distinct data types:

Integer	String	Floating-point
<p>16 x Byte, 8 x HW, 4xW, 2xDW, 1xQW</p> <ul style="list-style-type: none"> ▪ Byte to QuadWord add, sub, compare ▪ Byte to DoubleWord min, max, ave. ▪ Byte to Word multiply, multiply/add 4 - 32 x 32 multiply/adds ▪ Logical ops, shifts, ▪ CRC (GF multiply up to 64b), Checksum (32b), ▪ Loads efficient with 8B alignment though minor penalties for byte alignment ▪ Gather by Step 	<ul style="list-style-type: none"> ▪ Find 8b, 16b, 32b, equal or not equal with zero character end ▪ Range compare ▪ Find any equal ▪ Load to block boundary, load/store with length 	<p>BFP DP only 32 x 2 x 64b</p> <ul style="list-style-type: none"> ▪ 2 BFUs with an increase in architected registers ▪ Exceptions suppressed

Single Instruction Multiple Data (SIMD) Vector Processing

- **Single Instruction Multiple Data (SIMD)**
 - A type of data parallel computing that can accelerate code with integer, string, character, and floating point data types
- **Provide optimized SIMD math & linear algebra libraries that will minimize the effort on the part of middleware/application developers**
- **Provide compiler built-in functions for SIMD that software applications can leverage as needed (e.g. for use of string instructions)**
- **OS/Hypervisor Support:**
 - z/OS: 2.1 SPE available at GA
 - Linux: IBM is working with its Linux Distribution partners to support new functions/features
 - No z/VM Support for SIMD
 - Compiler exploitation
 - IBM Java => 1Q2015
 - XL C/C++ on zOS => 1Q2015
 - XL C/C++ on Linux on z => 2Q2015
 - Enterprise COBOL => 1Q2015
 - Enterprise PL/I => 1Q2015




Workloads		
Java.Next	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


z13 I/O update

z Systems I/O Subsystem Internal Bus Interconnect Speeds (GBps)



STI: Self-Timed Interconnect


z13 Processor Drawer Connectivity for I/O and Coupling



Ten PCIe Fanout slots
FSP = Flexible Support Processor

▪ Ten PCIe fanout slots per drawer (40 maximum per System))

- ICA (ICA SR) two-port 8 GBps PCIe Gen3 fanout 150 meter fiber optic coupling link





- PCIe Gen3 one-port 16 GBps PCIe fanout connects to a switch card for an 8-slot PCIe I/O domain (plugs in pairs)



▪ Four IFB HCA fanout slots per drawer (16 maximum on a four drawer system)

- HCA2-C 2-port 6 GBps I/O drawer fanout (plugs in pairs, maximum two per System)
- HCA3-O 2-port 12x IFB Coupling Link fanout
- HCA3-O LR 4-port 1x IFB Coupling Link fanout



Carry forward
(One pair only)

Carry forward or
New Build


Connectivity for Storage

New FICON Function for z13


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

FICON Express16S – SX and 10KM

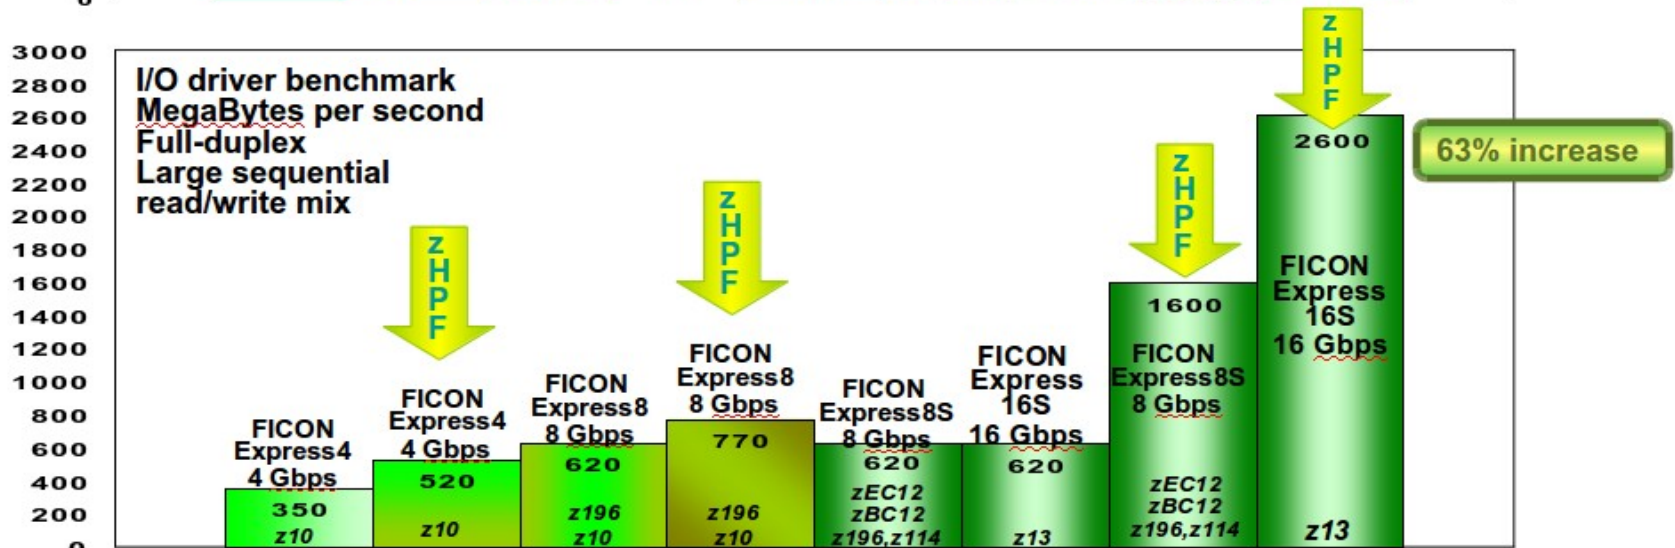
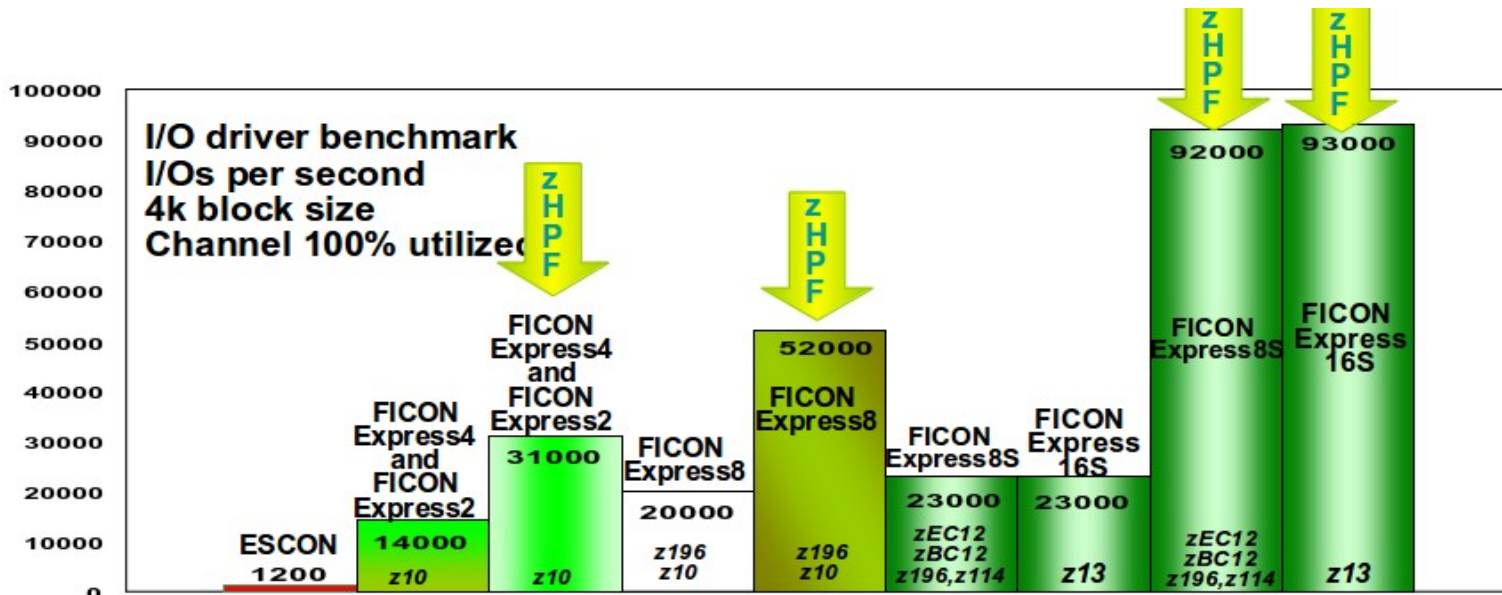
- **For FICON, zHPF, and FCP environments**
 - CHPID types: FC and FCP
 - 2 PCHIDs/CHPIDs
- **Auto-negotiates to 4, 8, or 16 Gbps**
 - 2Gbps connectivity **NOT** supported
 - FICON Express8S will be available to order for 2Gbps connectivity
- **Increased I/O Devices (subchannels) per channel for all 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
- **2 channels of LX or SX (no mix)**
- **Small form factor pluggable (SFP) optics**
 - Concurrent repair/replace action for each SFP




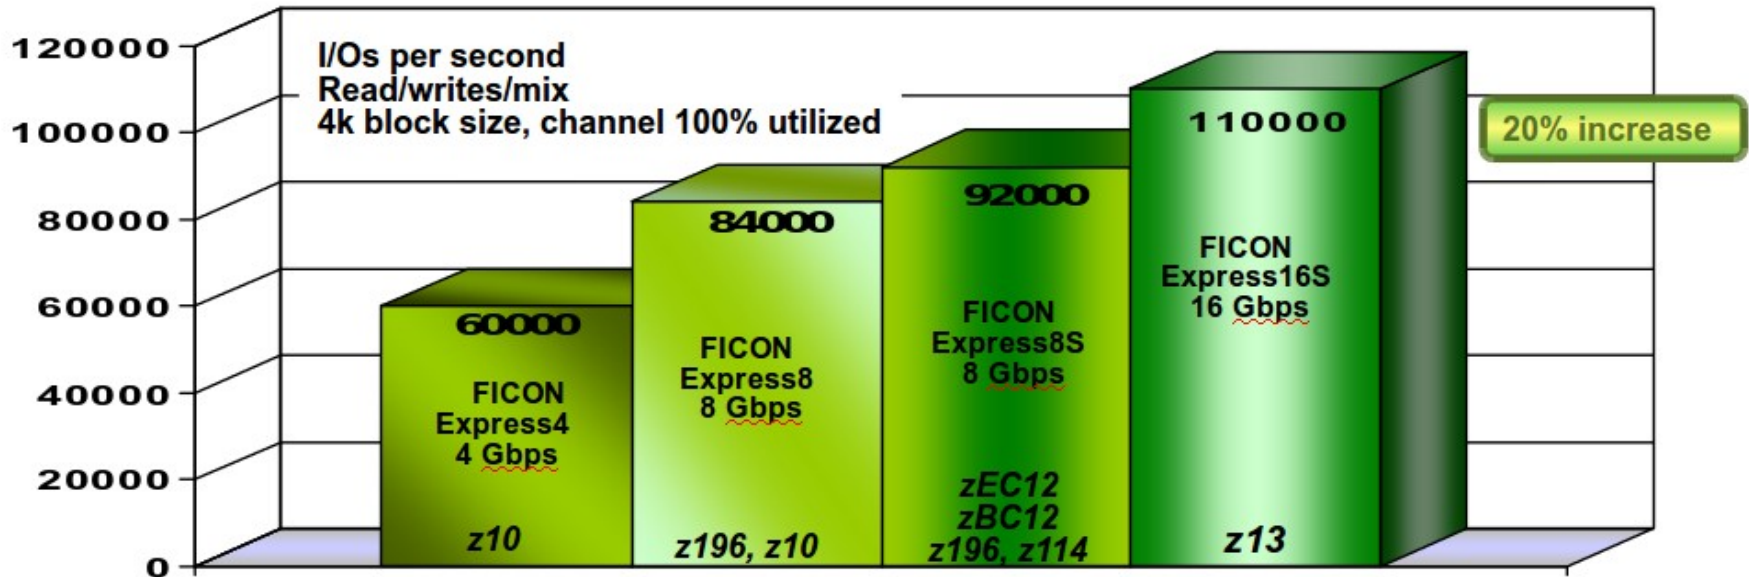
FC 0418 – 10KM LX, FC 0419 – SX



zHPF and FICON Performance* z13



FCP Performance* for z13



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

Connectivity for Networking

OSA-Express5S 1000BASE-T Ethernet Feature - PCIe I/O Drawer

- **PCI-e form factor feature supported by PCIe I/O drawer**
 - One two-port CHPID per feature
 - Half the density of the OSA-Express3 version
- **Small form factor pluggable (SFP+) transceivers**
 - Concurrent repair/replace action for each SFP
- **Exclusively Supports: Auto-negotiation to 100 or 1000 Mbps and full duplex only on Category 5 or better copper**
- RJ-45 connector
- Operates at “line speed”
- CHPID TYPE Support:



FC 0417



Connector = RJ-45

Mode	TYPE	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, OS system console operations
QDIO	OSD	TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
Non-QDIO	OSE	TCP/IP and/or SNA/APPN/HPR traffic
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN)
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CCL)

Note: OSA-Express5S feature are designed to have the same performance and to require the same software support as equivalent OSA-Express4S features.



OSA-Express5S Fiber Optic Features – PCIe Drawer

▪ 10 Gigabit Ethernet (10 GbE)

- CHPID types: OSD, OSX
- Single mode (LR) or multimode (SR) fiber
- One port of LR or one port of SR
 - 1 PCHID/CHPID
- Small form factor pluggable (SFP+) optics
 - Concurrent repair/replace action for each SFP
- LC duplex

▪ Gigabit Ethernet (1 GbE)

- CHPID types: OSD (**OSN not supported**)
- Single mode (LX) or multimode (SX) fiber
- Two ports of LX or two ports of SX
 - 1 PCHID/CHPID
- Small form factor pluggable (SFP+) optics
 - Concurrent repair/replace action for each SFP
- LC Duplex




Note: OSA-Express5S features are designed to have the same performance and to require the same software support as equivalent OSA-Express4S features.

Shared Memory Communication – Remote Direct Memory Access (RDMA) utilizing the 10GbE RoCE Express (SMC-R)

SMC-R Key Attributes

- **Optimized Network Performance (leveraging RDMA technology)¹**
 - **Transparent to (TCP socket based) application software**
 - **Preserves existing network security model**
 - **Resiliency (dynamic failover to redundant hardware)**
 - **Transparent to Load Balancers**
 - **Preserves existing IP topology and network administrative and operational model**
- ✓ ¹*Latency and CPU savings are based on workload type (latency focus for interactive workloads while CPU savings is on bulk traffic).*



Network latency reduced up to 80% for z/OS TCP/IP multi-tier OLTP workloads such as web based claims and payment systems *

* Based on internal IBM benchmarks of modeled z/OS TCP sockets-based workloads with request/response traffic patterns using SMC-R vs TCP/IP. The actual throughput that any user will experience will vary.

z13 Shared RoCE (SR-IOV) Overview

- **10GbE RoCE Express feature becomes sharable among multiple z/OS instances (LPARs or z/VM guest virtual machines)**
- **Multiple RoCE PFIDs (PCIe Function IDs) with unique Virtual Function IDs are configured for each physical adapter (PCHID) in HCD (IOCDs)**
- **Up to 31 PFIDs supported per physical adapter**
- **Each z/OS instance (LP or z/VM guest) sharing the adapter consumes a unique (at least one) PFID (each PFID has a corresponding Virtual Function ID / number)
Note. An OS instance and VFs are generally (but not always) one to one**
- **Up to 16 physical adapters per CPC (no change)**
- **Adapter virtualization is transparent to upper layers (stack (transport layer) and application software)**
- **Dedicated RoCE (zEC12 and zBC12) interoperates with (transparent to) SR-IOV RoCE (z13)... shared RoCE environment is transparent to peer hosts**
- **No changes in sub-net'ing compared to zEC12 (all endpoints must be in same layer2-subnet) – for more information see:**

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

Connectivity for Coupling Links and STP

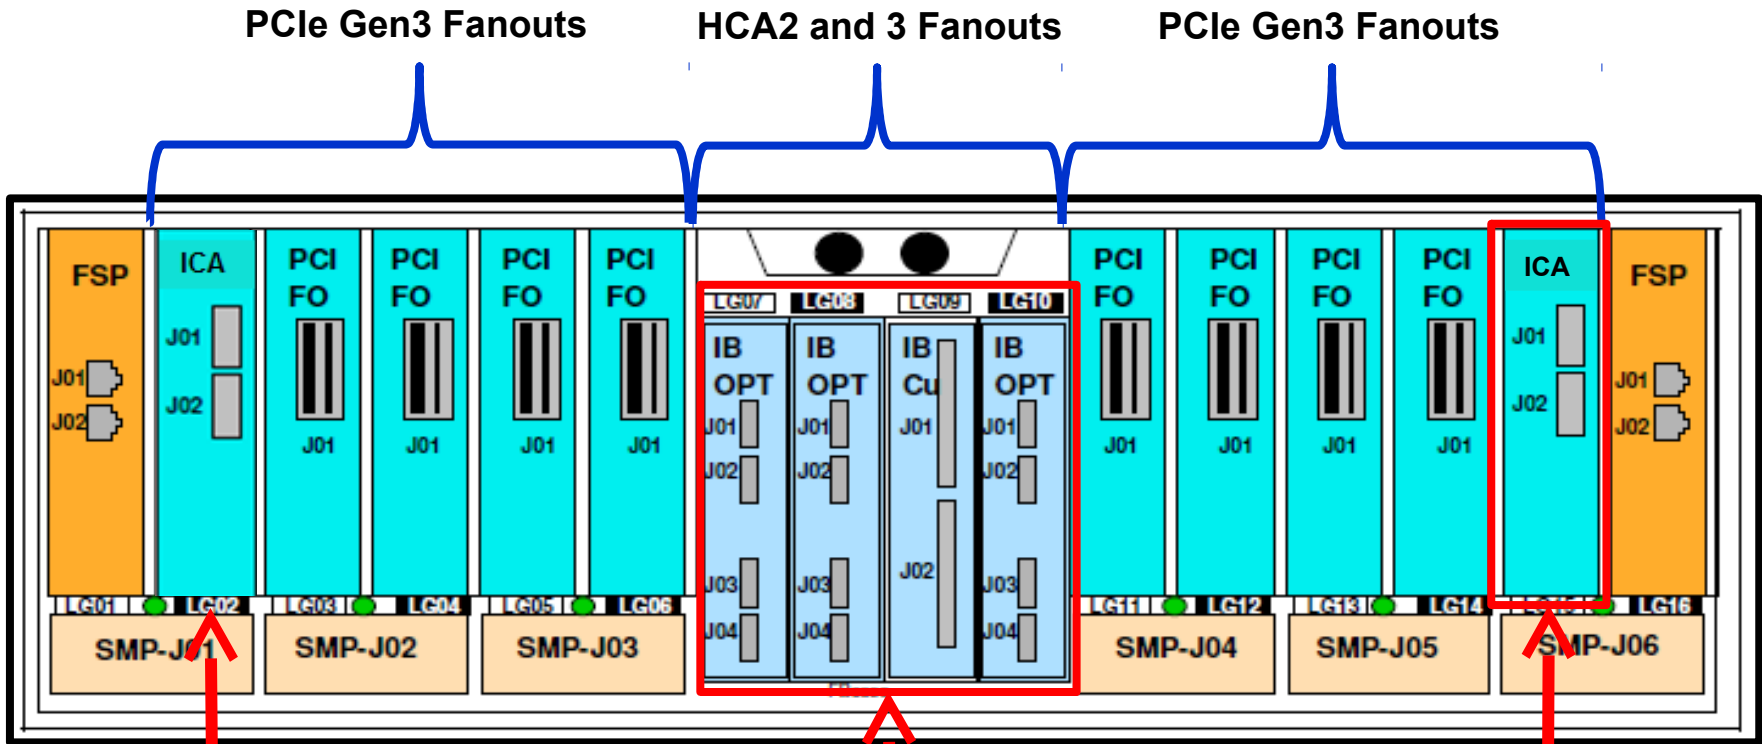
z13 supports two different internal coupling infrastructures

- **PCI Express Generation 3 (PCIe Gen3) I/O infrastructure introduced with z13**
 - New Build for PCIe Gen3 Integrated Coupling Adapter (ICA) SR Coupling Fanout

- **Host Channel Adapter I/O Fanouts (Infiniband coupling links)**
 - New build and Carry Forward for HCA3-O only, 12x InfiniBand (#0171 - HCA3-O fanout) or 1x InfiniBand (#0170 - HCA3-O LR fanout) coupling links.
 - HCA2-O fanout for 12x IFB coupling links (#0163) and HCA2-O LR fanout for 1x IFB coupling links (#0168) New Build and Carry Forward are NOT SUPPORTED

ISC3 Links - FC 0217, 0218 and 0219, New Build and Carry Forward are NOT supported

CPC Drawer Front View – Coupling Links



HCA2-C (I/O Drawer)
or
HCA3 (1X or 12X PSIFB Links)

ICA SR Coupling Link
(Integrated Coupling Adapter)

Integrated Coupling Adapter (ICA SR)

Integrated Coupling Adapter SR (ICA SR) Fanout in the CPC drawer

- Recommended for Short Distance Coupling z13 to z13, not available on older servers
- No performance degradation compared to Coupling over Infiniband 12X IFB3 protocol

Hardware Details

- Short reach adapter, distance up to 150 m
- Up to 32 ports maximum
- IOCP Channel Type = CS5
- Feature code 0172, 2 ports per adapter
 - Up to 4 CHPIDs per port, 8 per feature, 7 buffers (i.e. 7 subchannels) per CHPID
- ICA requires new cabling for single MTP connector
 - Differs from 12X Infiniband split Transmit/Receive connector

Requirements

- CF: z13; z/OS: z13
- z/OS V2.1, V1.13, or V1.12 with PTFs for APARs OA44440 and OA44287




Coupling links on z13

Type	Speed	Distance	Fanout
ICA SR	8 GBps	150 meters	ICA SR
12x InfiniBand	6 GBps	150 meters	HCA3-O
1x InfiniBand	5 or 2.5 Gbps	10 km	HCA3-O LR


Up to 4 CHPIDs – per port



Up to 16 CHPIDs – across 2 ports



Up to 16 CHPIDs – across 2 ports*



Up to 16 CHPIDs – across 4 ports*

- Ports exit from the front of a CPC drawer with HCA3s or ICA SRs.
- ICA SR
 - 8 GBps
- 12x InfiniBand
 - 6 GBps
- 1x InfiniBand
 - 5 Gbps (Server to Server and with DWDM)
 - 2.5 Gbps (with DWDM)

* Performance considerations may reduce the number of CHPIDs per port

HCA3 for Coupling Links

HCA3-O for 12x IFB & 12x IFB3



Up to 16 CHPIDs – across 2 ports*

HCA3-O LR for 1x IFB



Up to 16 CHPIDs – across 4 ports*

- **12x InfiniBand and 1x InfiniBand fanout features**
 - **HCA3-O fanout for 12x InfiniBand coupling links**
 - CHPID type – CIB
 - Improved service times with 12x IFB3 protocol
 - Two ports per feature
 - Fiber optic cabling – 150 meters
 - Supports connectivity to HCA2-O
 - Link data rate of 6 GBps
 - **HCA3-O LR fanout for 1x InfiniBand coupling links**
 - CHPID type – CIB
 - Four ports per feature
 - Fiber optic cabling
 - 10 km without repeaters
 - Extended distances of up to 100 km are also possible using Dense Wavelength Division Multiplexors (DWDMs)
 - For distances > 10 km WITHOUT repeaters, based on the z System, following RPQ is required:
 - z196/z114 RPQ 8P2340
 - zEC12/zBC12 RPQ 8P2781
 - z13 RPQ 8P2981
 - Supports connectivity to HCA2-O LR
 - Link data rate server-to-server 5 Gbps
 - Link data rate with DWDM; 2.5 or 5 Gbps


* Performance considerations may reduce the number of CHPIDs per port

Note: The InfiniBand link data rates of 6 GBps, 3 GBps, 2.5 Gbps, or 5 Gbps do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

z13 Parallel Sysplex Connectivity

z196 and z114

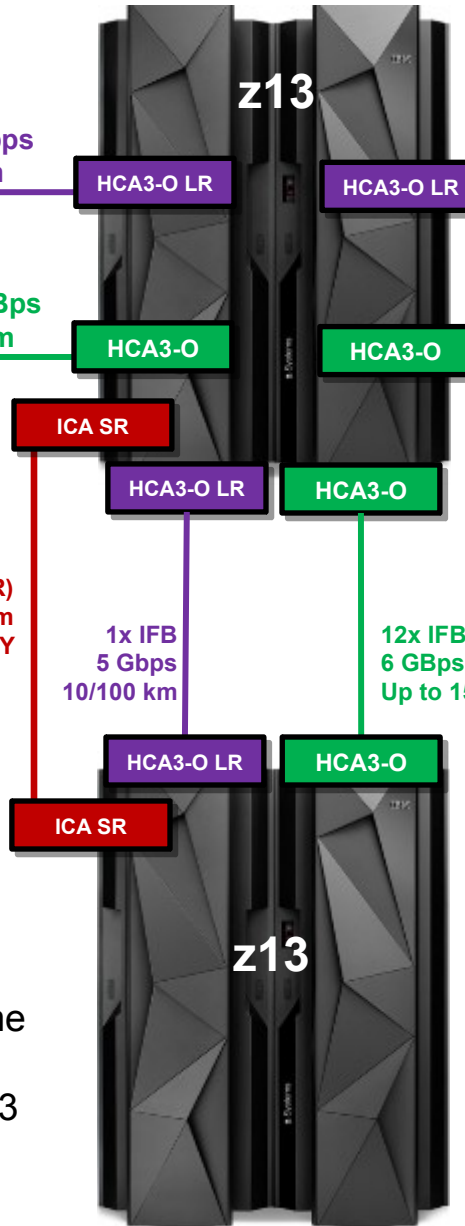
12x IFB, 12x IFB3, 1x IFB



1x IFB, 5 Gbps
10/100 km

12x IFB, 6 GBps
Up to 150 m

z13



ICA SR

1x IFB
5 Gbps
10/100 km

12x IFB
6 GBps
Up to 150 m

zEC12 and zBC12

12x IFB, 12x IFB3, 1x IFB



1x IFB, 5 Gbps
10/100 km

12x IFB, 6 GBps
Up to 150 m

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

z10, z9 EC, z9 BC,
z890, z990
Not supported in same
Parallel Sysplex
or STP CTN with z13

IC (Internal Coupling Link):
Only supports IC-to-IC connectivity

HCA2-O and HCA2-O LR are NOT supported on z13 or future High End z enterprises as Per SOD

ISC-3 is not supported on z13 even if I/O Drawer is Carried Forward for FICON Express8

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



Flash Express

Why Flash Express on z13?

▪ Provides Storage Class Memory

- Implemented via NAND Flash SSDs (Solid State Drives) mounted in PCIe Flash Express features
- Protected by strong AES Encryption done on the features
- Not defined as I/O devices or with PCIe FUNCTIONS
- Assigned to partitions similarly to Main Memory; but, not in the partition Image Profile. Reconfigurable.
- Accessed using the new z System architected EADM (Extended Asynchronous Data Mover) Facility
- Designed to enable extremely responsive paging of 4k pages to improve z/OS availability
- Enables pageable large (1 MB) pages

▪ Flash Express Exploitation

- z/OS V2.1, V1.13 + PTFs and RSM Enablement Offering
 - With z/OS Java SDK 7 SR3: CICS TS V5.1, WAS Liberty Profile V8.5, DB2 V11, IMS 12 and higher, SOD: Traditional WAS 8.0.0x*
 - CFCC Level 19 with WebSphere MQ for z/OS Version 7 MQ Shared Queue overflow support (March 31, 2014)
- Linux on z Systems
 - SLES 11 SP3 and RHEL 6.4

10x

Faster response time and 37% increase in throughput compared to disk for morning transition

28%

Improvement in DB2 throughput leveraging Flash Express with Pageable Large Pages (PLP)

19%

Reduction in total dump time for a 36 GB standalone dump

~25%

Reduction in SVC dump elapsed time

***Note: 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.**

IBM zEnterprise Data Compression (zEDC)

Compression Coprocessor (CMPSC) vs. zEDC

Using the right hardware compression acceleration for each of your workloads

Compression Coprocessor

z Enterprise Data Compression

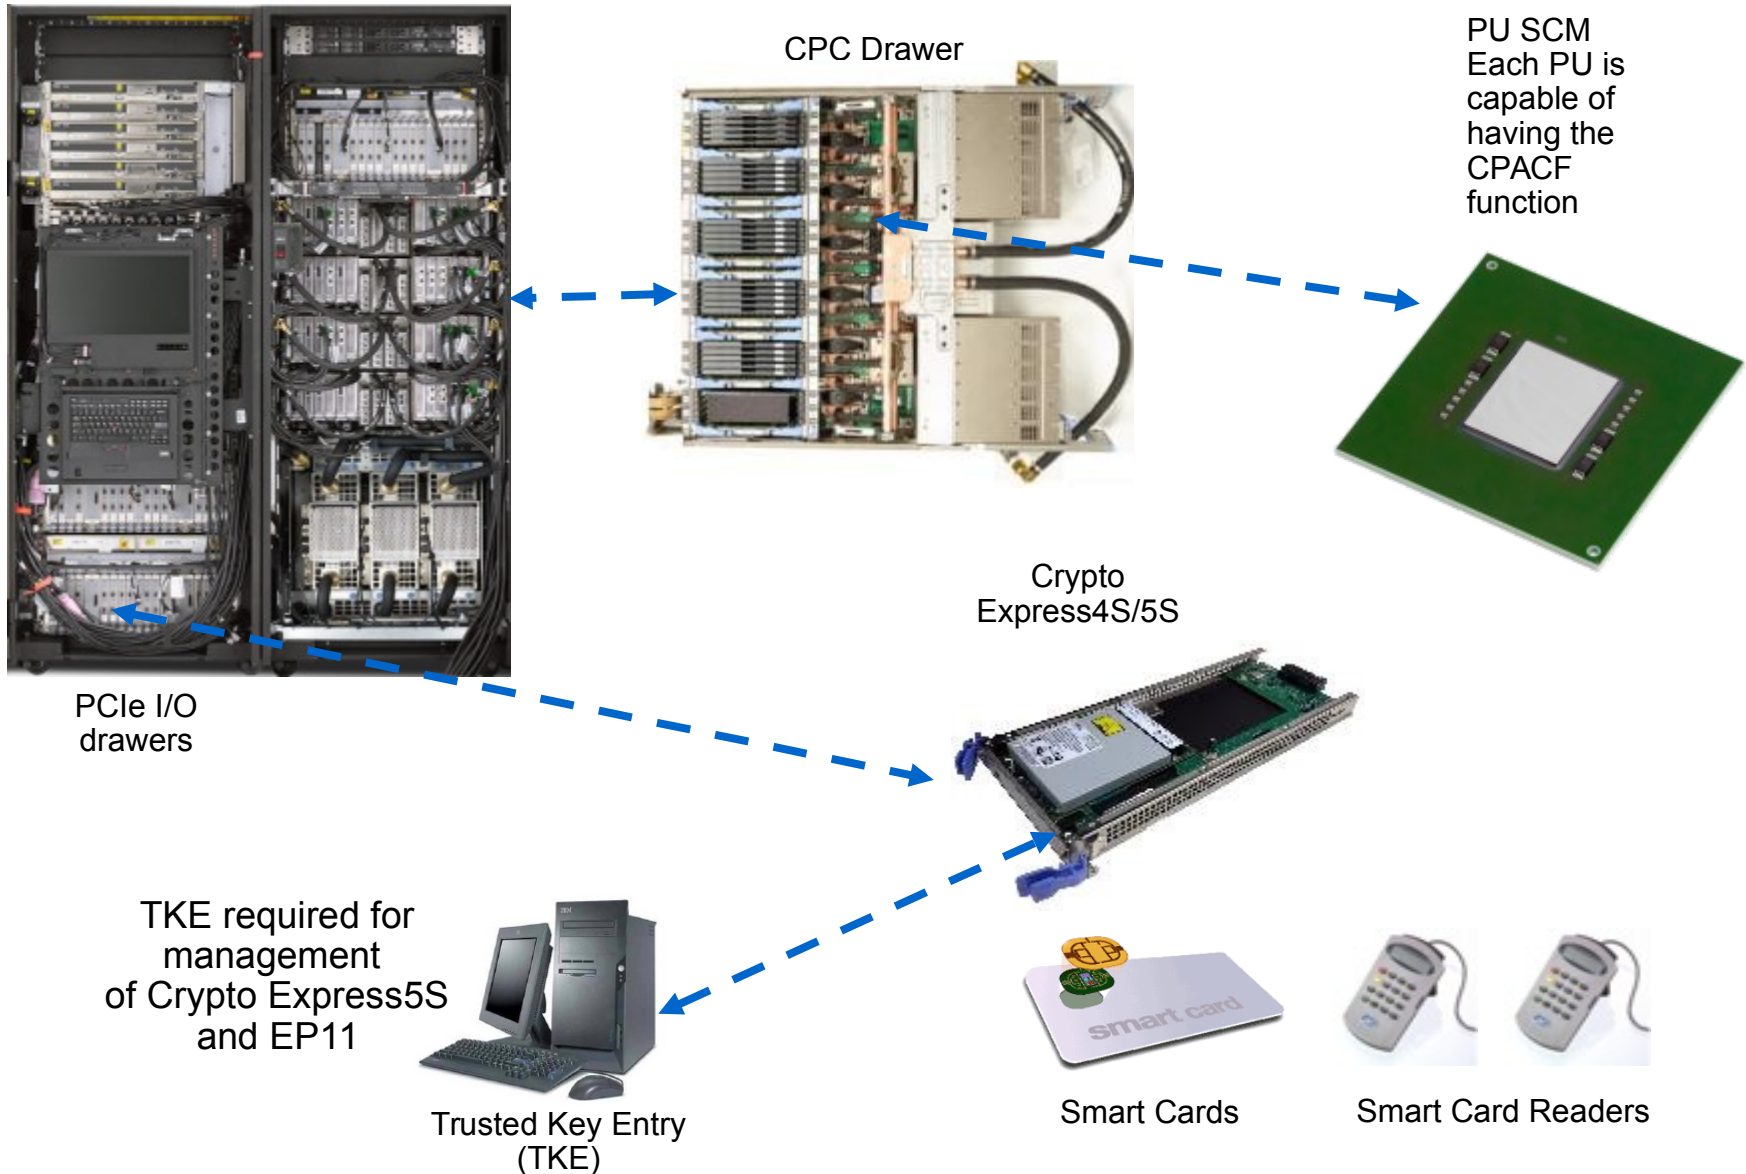
<p>On Chip</p> <p>In every IBM eServer™ zSeries® today (and tomorrow)</p> <p>Mature: Decades of use by Access Methods and DB2®</p> <p>Work is performed jointly by CPU and Coprocessor</p> <p>Proprietary Compression Format</p>	<p>PCIe Adapter</p> <p>Introduced with IBM zEnterprise® EC12 GA2 and IBM zEnterprise BC12 and z13</p> <p>Mature: Industry Standard with decades of software support</p> <p>Work is performed by the PCIe Adapter</p> <p>Standards Compliant (RFC1951)</p>
--	---

Use Cases

<p><u>Small object compression</u></p> <ul style="list-style-type: none"> ▪ Rows in a database 	<p><u>Large Sequential Data</u></p> <ul style="list-style-type: none"> ▪ QSAM/BSAM Online Sequential Data ▪ Objects stored in a data base 	<p><u>Industry Standard Data</u></p> <ul style="list-style-type: none"> ▪ Cross Platform Data Exchange
<p><u>Users</u></p> <ul style="list-style-type: none"> ▪ VSAM for better disk utilization ▪ DB2 for lower memory usage ▪ The majority of customers are currently compressing their DB2 rows 	<p><u>Users</u></p> <ul style="list-style-type: none"> ▪ QSAM/BSAM for better disk utilization and batch elapsed time improvements ▪ SMF for increased availability and online storage reduction 	<p><u>Users</u></p> <ul style="list-style-type: none"> ▪ Java for high throughput standard compression via java.util.zip ▪ Encryption Facility for z/OS for better industry data exchange ▪ IBM Sterling Connect: Direct® for z/OS for better throughput and link utilization ▪ ISV support for increased client value


z Systems Crypto

Overview – HW Crypto support in z Systems



z13 CPACF

- **CP Assist for Cryptographic Function Co-processor redesigned from "ground up"**
- **Enhanced performance over zEC12**
 - Does not include overhead for COP start/end and cache effects
 - Enhanced performance for large blocks of data
 - AES: 2.5x throughput vs. zEC12
 - TDES: 2.5x throughput vs. zEC12
 - SHA: 3.5x throughput vs. zEC12
- **Exploiters of the CPACF benefit from exploited by the throughput improvements of z13's CPACF such as:**
 - DB2/IMS encryption tool
 - DB2® built in encryption
 - z/OS Communication Server: IPsec/IKE/AT-TLS
 - z/OS System SSL
 - z/OS Network Authentication Service (Kerberos)
 - DFDSS Volume encryption
 - z/OS Java SDK
 - z/OS Encryption Facility
 - Linux on z Systems; kernel, openssl, openCryptoki, GSKIT



Crypto Express5S

- One PCIe adapter per feature
 - Initial order – two features
- Designed to be FIPS 140-2 Level 4
- Installed in the PCIe I/O drawer
- Up to 16 features per server
- Prerequisite: CPACF (#3863)
- Designed for 2X performance increase over Crypto Express4S

Three configuration options for the PCIe adapter

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



Accelerator

CCA Coprocessor

EP11 Coprocessor


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

Business Value

- High speed advanced cryptography; intelligent encryption of sensitive data that executes off processor saving costs
- PIN transactions, EMV transactions for integrated circuit based credit cards(chip and pin), and general-purpose cryptographic applications using symmetric key, hashing, and public key algorithms, VISA format preserving encryption(VFPE), and simplification of cryptographic key management.
- Designed to be FIPS 140-2 Level certification to meet regulations and compliance for PCI standards

IBM zAware

IBM zAware V2.0 - Analyze z/OS and Linux on z Systems



IBM z Systems Advanced Workload Analysis Reporter:

- Identify unusual system behavior of Linux on system z images
- Monitors **syslog*** from guest or native image in real time
- Improved analytics for z/OS message logs
- Upgraded internal database for improved RAS
- Completely rewritten UI, including heat map views

IBM zAware Background – Heat Map – All Sstems in a group

▪ UI with Drill down system list (ModelGroup)

Firefox | Analysis - IBM zAware

IBM zAware | admin | IBM

Analysis

Date (UTC): October 11, 2013

Analysis Source: Charge Source | Previous Group Selection

All systems in MEL*5SH3

Actions | Zoom: 24 Hrs | View: Heat Map Table | Filter

No filter applied

System group	System	24 Hour Peak	Peak Anomaly Score Per Hour																							
			0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
MEL*5SH3	MEL15SHB.OSI1	101	96.6	99.8	81.2	69.2	95.3	94.5	95.3	96.8	94.5	90.1	101	101	100	100	95.9	94.7	93.4	93.9	91.5	93.1	97.1	89.7	99.7	94.5
MEL*5SH3	MEL15SHB.OSI2	100	100	95.1	75	57.2	90.7	84.3	90.3	83.7	90.4	81.6	98.6	100	100	99	84.3	94.1	93.2	99	84.5	73.3	87.1	89.7	97.8	57
MEL*5SH3	MEL15SHB.OSI1	98.9	73.5	98.6	74.7	82	65.4	76.5	88	94.4	97.1	93.9	97.7	97.2	95.6	95.6	96.3	97.8	97.7	95.1	87.1	85.7	95.8	93.7	83.5	98.9
MEL*5SH3	MEL15SHB.OSI2	98.1	74.5	98.1	76.9	80.2	80.7	78.1	80.8	92.7	91.6	91.5	95.1	95.3	87.3	93.6	96.7	82.8	95.7	91.4	93	75.4	86.7	76.8	75.5	96.9

▼ Details

MEL15SHB.OSI1

Timeline (UTC)

z13 Models & Upgrades, Capacity and Performance Planning

z13 Full and Sub-Capacity CP Offerings

CP Capacity

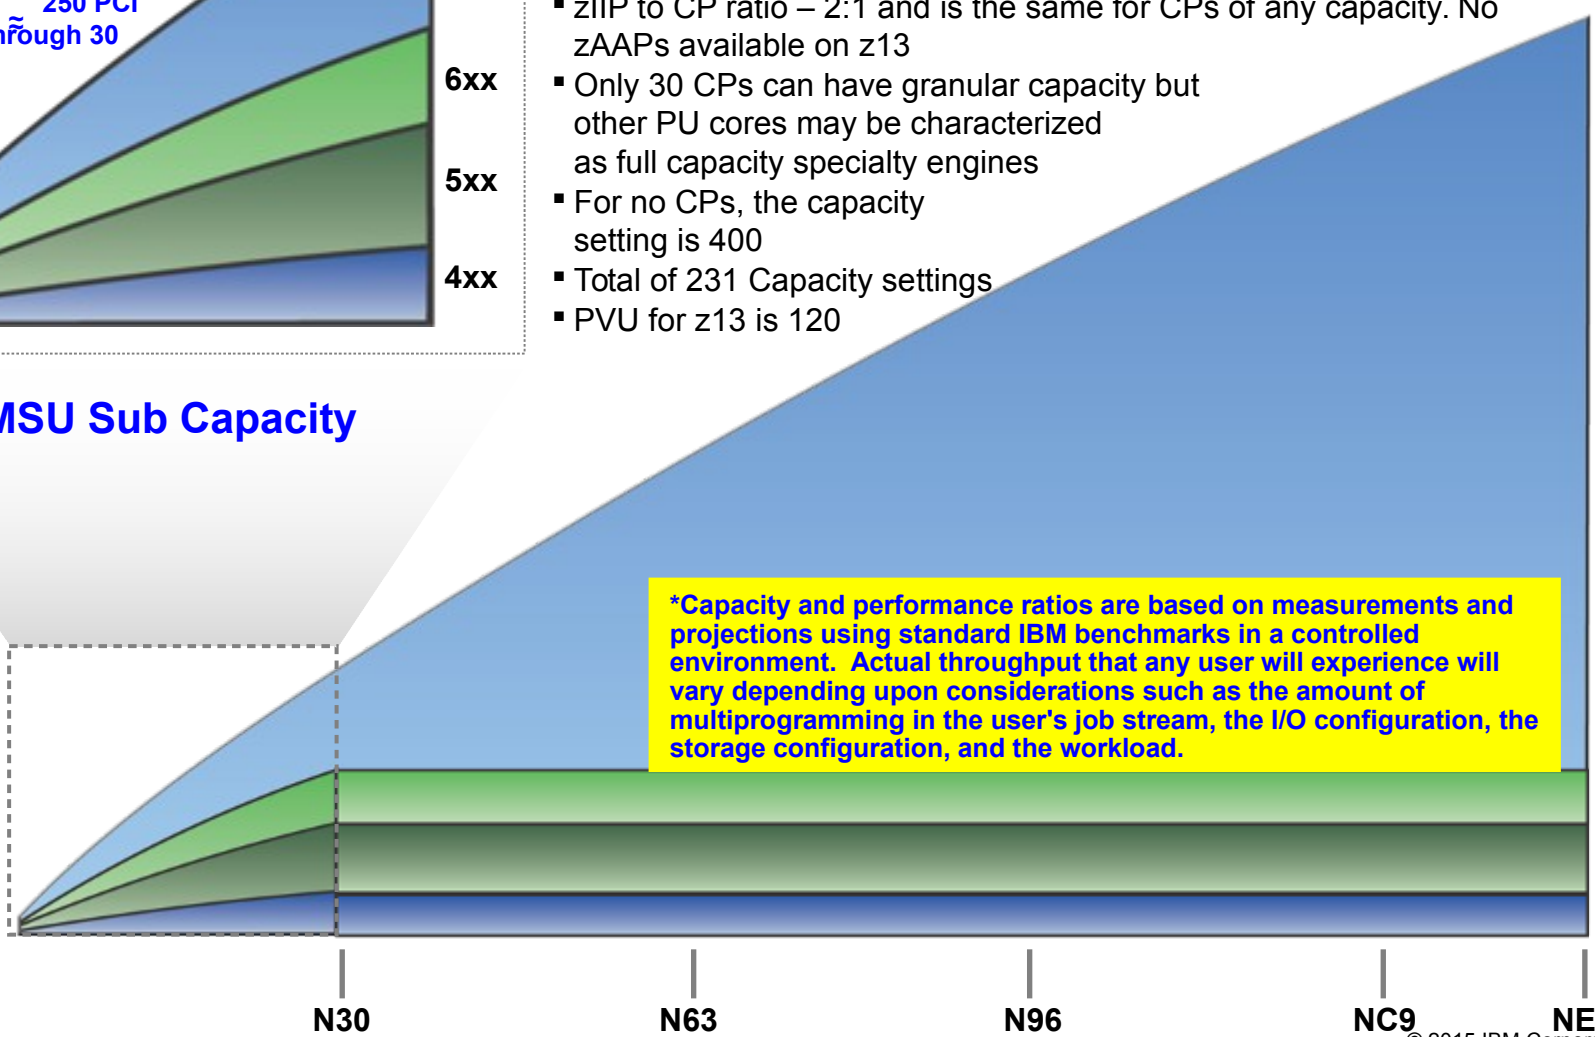
Relative to Full Capacity Uni

7xx	=	100%	≈	1695 PCI
6xx		63%	≈	1068 PCI
5xx		44%	≈	746 PCI
4xx		15%	≈	250 PCI
xx		1/4		01 Through 30

7xx
6xx
5xx
4xx

- Subcapacity CPs, up to 30, may be ordered on ANY z13 model. If 31 or more CPs are ordered all must be full 7xx capacity
- All CPs on a z13 CPC must be the same capacity
- All specialty engines are full capacity.
- zIIP to CP ratio – 2:1 and is the same for CPs of any capacity. No zAAPs available on z13
- Only 30 CPs can have granular capacity but other PU cores may be characterized as full capacity specialty engines
- For no CPs, the capacity setting is 400
- Total of 231 Capacity settings
- PVU for z13 is 120

MSU Sub Capacity



N30

N63

N96


NC9

NE1

z13 Vs zEC12 Vs z196 Vs z10 EC Vs z9 EC capacity comparison

- z9 EC up to 54-way: 193 – 18,505 PCIs.
- z10 EC up to 64-way: 214 – 31,826 PCIs
- z196 up to 80-way: 240 – 52,286 PCIs
- zEC12 up to 101-way: 240 – 78,426 PCIs
- z13 up to 141-way: 250 – up to 111,556 PCIs

Note: For z13 and zEC12 - z/OS 2.1 and for rest z/OS V1.13 LSPR numbers quoted.



Statements of Direction

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 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

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

- **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).

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

- **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.

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

- **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 server 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.

Herzlichen Dank

Anhang

z13 Overview

z13



Model	Customer PUs	Max Memory
NE1	141	10 TB
NC9	129	10 TB
N96	96	7.5 TB
N63	63	5 TB
N30	30	2.5 TB

- **Machine Type**
 - 2964
- **5 Models**
 - N30, N63, N96, NC9 and NE1
- **Processor Units (PUs)**
 - 39 (42 for NE1) PU cores per CPC drawer
 - Up to 24 SAPs per system, standard
 - 2 spares designated per system
 - Dependant on the H/W model - up to 30, 63, 96, 129, 141 PU cores available for characterization
 - Central Processors (CPs), Internal Coupling Facility (ICFs), Integrated Facility for Linux (IFLs), IBM z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs) and Integrated Firmware Processor (IFP)
 - 85 LPARs, increased from 60
 - Sub-capacity available for up to 30 CPs
 - 3 sub-capacity points
- **Memory**
 - RAIM Memory design
 - System Minimum of 64 GB
 - Up to 2.5 TB GB per drawer
 - Up to 10 TB for System and up to 10 TB per LPAR (OS dependant)
 - LPAR support of the full memory enabled
 - 96 GB Fixed HSA, standard
 - 32/64/96/128/256/512 GB increments
 - Flash Express
- **I/O**
 - 6 GBps I/O Interconnects – carry forward only
 - Up to 40 PCIe Gen3 Fanouts @ 16 GBps each and Integrated Coupling Adapters @ 2 x 8 GBps per System
 - 6 Logical Channel Subsystems (LCSSs)
 - 4 Sub-channel sets per LCSS
- **Server Time Protocol (STP)**

IBM z13 Redbooks



- **IBM z13 Technical Introduction, SG24-8250:** This publication provides concepts, positioning, and a business value view of IBM z13 capabilities, hardware functions/features, and associated software support. It is intended for IT Managers, consultants, IT Architects and Specialists, and anyone who wants to understand the basic elements of the IBM z13.
- **IBM z13 Technical Guide, SG24-8051:** This publication provides specific information about the IBM z Systems z13 (z13) and its functions, features, and associated software support. Greater detail is offered in areas relevant to technical planning. It is intended for systems engineers, system programmers (IT Specialists), planners, and anyone wanting to understand the z13 functions and plan for their usage.
- **IBM z Systems Connectivity Handbook, SG24-5444:** This publication highlights the hardware and software components, typical uses, coexistence, and relative merits of the z System I/O features. It is intended for data center planners, IT Specialists, system engineers, technical sales staff, and network planners who are involved in planning connectivity solutions for z System servers.
- **IBM z13 Configuration Setup, SG24-8260:** This publication helps you install, configure, and maintain the IBM z13. This book is intended for systems engineers, hardware planners, and anyone who needs to understand IBM z Systems® configuration and implementation. Readers should be generally familiar with current IBM z Systems technology and terminology. For details about the z13, see IBM z13 Technical Introduction, SG24-8250, and IBM z13 Technical Guide, SG24-8251.
- **The z13 IBM Redbooks launch page will be:**
<http://www.redbooks.ibm.com/redbooks.nsf/pages/z13?Open>

z13 Functional Comparison to zEC12

Performance and Scale	<ul style="list-style-type: none"> ▪ Uniprocessor Performance ▪ System Capacity ▪ SMT ▪ SIMD ▪ Cache ▪ Models ▪ Processing cores ▪ Granular Capacity ▪ Memory ▪ Fixed HSA ▪ Compression ▪ Internal I/O Bandwidth 	<ul style="list-style-type: none"> ▪ Up to 10% performance improvement over zEC12¹ ▪ Up to 40% system total z/OS capacity performance improvement over zEC12¹ ▪ SMT delivers up to 32% price performance improvement for Linux on z Systems and up to 38% price performance improvement for zIIP workloads versus single threaded only on zEC12 ▪ Vector processing (SIMD) model provides construction of richer, complex analytics models, increased programmer productivity and faster mathematical modeling versus no SIMD on zEC12 ▪ z13 has 2x the cache versus zEC12 ▪ Five models with up to 4 CPC drawers (zEC12 has five models with up to 4 books) ▪ Up to 141 cores to configure, up to 101 on zEC12 ▪ Up to 231 capacity settings versus 161 on the zEC12 ▪ Up to 10 TB RAIM memory versus 3 TB RAIM memory on zEC12 ▪ Up to 3x HSA (96 GB) versus zEC12 (32 GB) fixed HSA ▪ Continued support of zEDC Express and hardware compression on coprocessor ▪ 832 GB/sec I/O bandwidth versus 384 on zEC12 (Note: servers exploit a subset of its designed capacity)
Virtualization	<ul style="list-style-type: none"> ▪ LPAR virtualization ▪ LPAR memory support ▪ RoCE adapter virtualization 	<ul style="list-style-type: none"> ▪ 85 partitions versus 60 on zEC12 ▪ Full memory support per LPAR (10 TB) versus 1 TB on zEC12 (OS exploitation varies) ▪ 10 GbE RoCE Express shared across LPARs versus dedicated 10GbE RoCE Express on zEC12
Infrastructure Efficiency	<ul style="list-style-type: none"> ▪ Networking ▪ FICON ▪ zHPF ▪ ROCE ▪ Forward Error Correction ▪ Fabric I/O priority ▪ FICON dynamic routing ▪ LCSS/Subchannel sets ▪ WWPN 	<ul style="list-style-type: none"> ▪ New OSA-Express5S supported by both ▪ FICON Express16S plus increased FICON subchannels to 32K versus FICON Express 8S and 24K on zEC12 ▪ zHPF extended distance II offers faster remote site recovery with improved I/O service time improvement when writing data remotely (GDPS HyperSwap) versus zHPF only on z13 ▪ 10GbE RoCE Express supported on both – enhancement to share feature only on z13 ▪ Industry standard FEC for optical connections for substantially reduced I/O link errors not on zEC12 ▪ z/OS WLM extended to SAN fabric leveraging capabilities of the SAN vendors not available on zEC12 ▪ Dynamic Routing allows for sharing of switches between FICON and FCP without creating separate virtual switches not available on zEC12 ▪ Up to six LCSS versus four on zEC12 and 4 Subchannel sets versus 3 on zEC12 ▪ I/O serial number migration allows keeping same serial number on replacement server not on zEC12
Resiliency and Availability	<ul style="list-style-type: none"> ▪ Coupling – HCA-3 ▪ Coupling – ICA SR ▪ STP ▪ Sparing ▪ IBM zAware ▪ Environmentals 	<ul style="list-style-type: none"> ▪ Coupling with HCA-3 InfiniBand Coupling Links – long and short distance – same as zEC12 ▪ New short distance coupling with PCIe-based links versus not available on zEC12 ▪ Simplified sysplex management with STP enhancements not available on zEC12 ▪ Enhanced integrated sparing on z13 reducing the number of on site service and maintenance events ▪ IBM zAware offers high speed analytics to consume large quantities of message logs for smarter monitoring – available for both z/OS and Linux – IBM zAware for z/OS only supported on zEC12 ▪ Optional non raised floor, overhead cabling, water cooling and DC power – same on zEC12
Security	<ul style="list-style-type: none"> ▪ Cryptographic Coprocessor ▪ Crypto Express ▪ Crypto Key Management 	<ul style="list-style-type: none"> ▪ CPACF improves performance by 2.5x for AES/TDES and 3.5x for SHA versus zEC12 ▪ Crypto Express 5S with performance increase plus new algorithms for elliptic curve, SHA, VISA FPE versus zEC12 Crypto Express4S ▪ Cryptographic key support for management, simplification and compliance

¹ For average LSPR workloads running z/OS 1.13. Official performance data are available and can be obtained online at LSPR (Large Systems Performance Reference) website at: <https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lspindex?OpenDocument>. Actual performance results may vary by customer based on individual workload, configuration and software levels.