

IBM z13s and z13 Innovations

*A technical walk-through of the IBM
z13s and z13 HW and SW stack*

Eldee Stephens
z Systems Nest Bringup Lead
estephen@us.ibm.com

InterConnect 2017



Please Note:

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including 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 results similar to those stated here.

What's under the covers?


A physical overview of the IBM z13 and z13s mainframes

IBM z13 Physical Overview




- **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
 - Dependent 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 per drawer
 - Up to 10 TB for System and up to 10 TB per LPAR (OS dependent)
 - 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)**

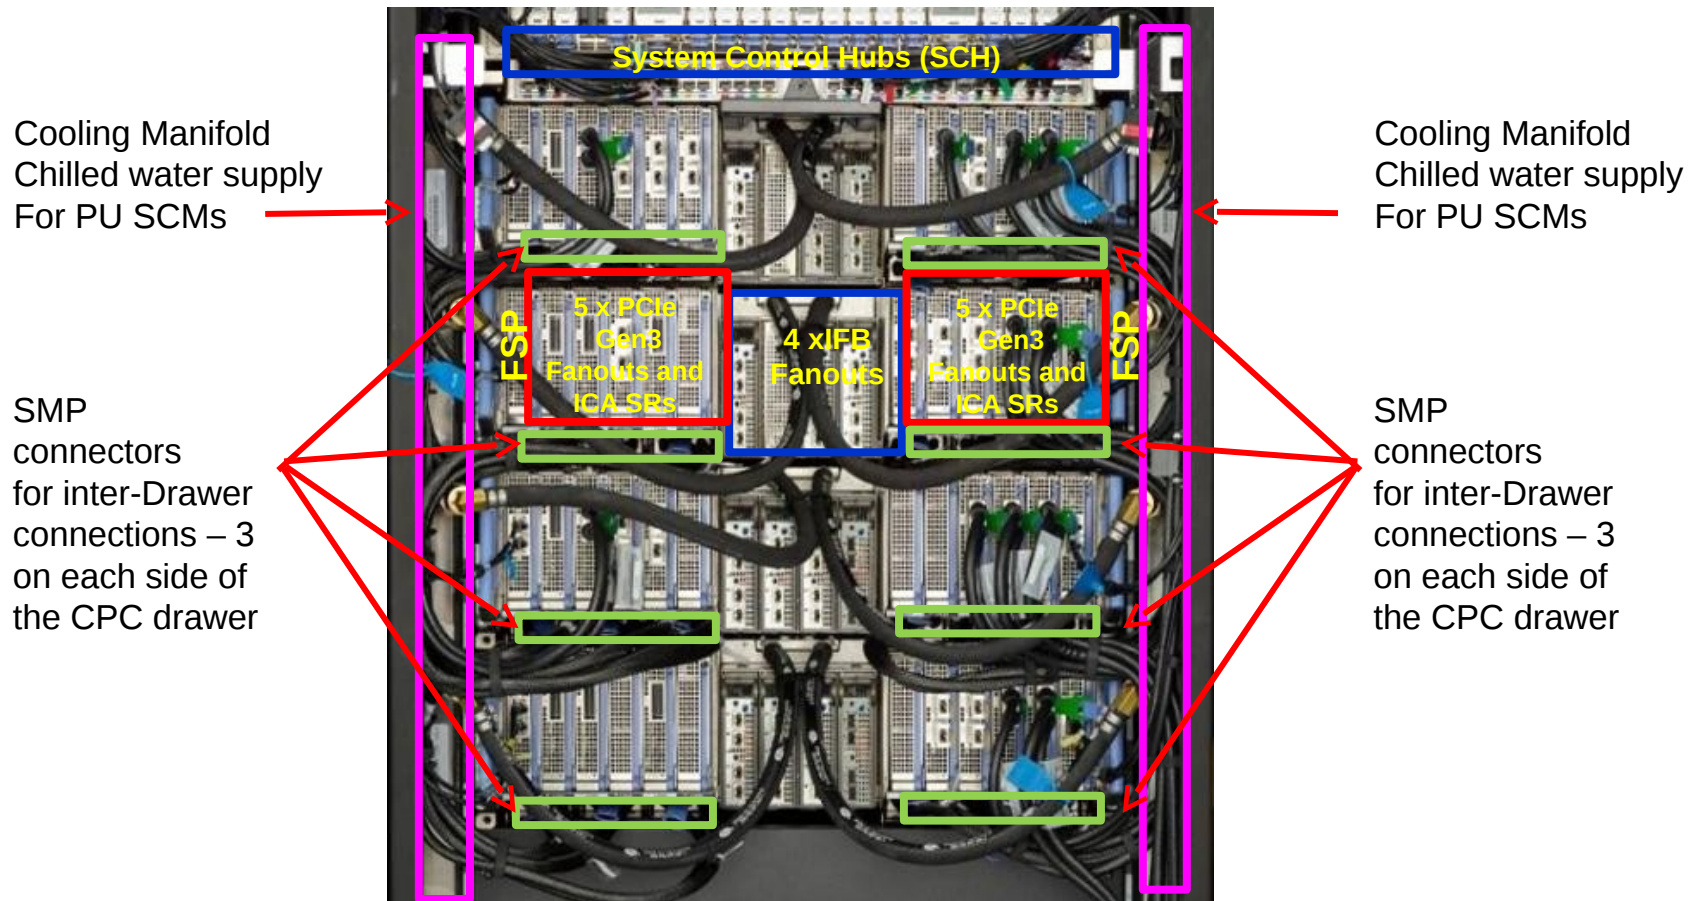
IBM z13 Walk-through: Front View (water-cooled)



IBM z13 Walk-through: Rear View (water-cooled)



IBM z13 Walk-through: CPC Drawer, Front View




IBM z13s Physical Overview



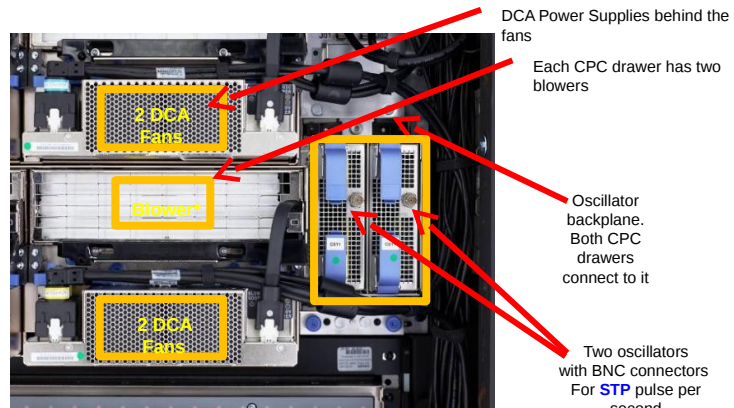
- **Machine Type**
 - 2965
- **2 Models**
 - N10 and N20
 - N20 available as one- or two- processor drawer model
 - The 2nd drawer in the N20 is driven by I/O and/or memory requirements
 - Single frame, air cooled
 - Non-raised floor option available
 - Overhead Cabling and DC Power Options
- **Processor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

Model	Customer PUs	Max Mem
N10	10	984 GB
N20 1 CPC Drawer	20	2008 GB
N20 2 CPC Drawers	20	4056 GB

IBM z13s Walk-through: CPC Drawers



Front View




*Left blower of this node is not shown in the picture


Rear View

IBM z13s Walk-through: CPC Drawers (cont.)


- Model N20 CPC drawer has two nodes
- Model N10 CPC drawer has a single node
- Each node contains:
 - One System Control (SC) chip (480 MB L4 cache)
 - Two Processing Units (PU) chips running at 4.3GHz
 - Eight-core per PU chip design
 - Up to seven active cores per PU chip (N20)
 - Six or seven per PU chip on N10
 - One memory controller per PU chip (two per node)
 - Five DDR3 DIMM slots per memory controller: 10 total per node (up to 1024GB per node)
 - Two Flexible Service Processors
 - Four PCIe fanout slots
 - Two slots for IFB fanouts or PSIFB coupling link fanouts



Two CPC Drawer System connectivity (N20-2)



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




The heart of the beast

*The most advanced microprocessor in
history powering the z13 family....*

Some history: every generation is more capable...

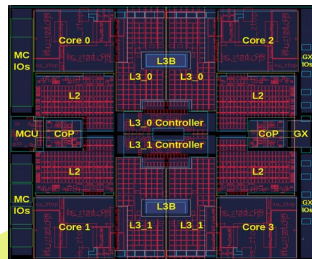
z10
2/2008



Workload Consolidation and Integration Engine for CPU Intensive Workloads

- Decimal FP
- Infiniband
- 64-CP Image
- Large Pages
- Shared Memory


z196
9/2010



Top Tier Single Thread Performance, System Capacity

- Accelerator Integration
- Out of Order Execution
- Water Cooling
- PCIe I/O Fabric
- RAIM
- Enhanced Energy Management


zEC12
8/2012



Leadership Single Thread, Enhanced Throughput

- Improved out-of-order
- Transactional Memory
- Dynamic Optimization
- 2 GB page support
- Step Function in System Capacity

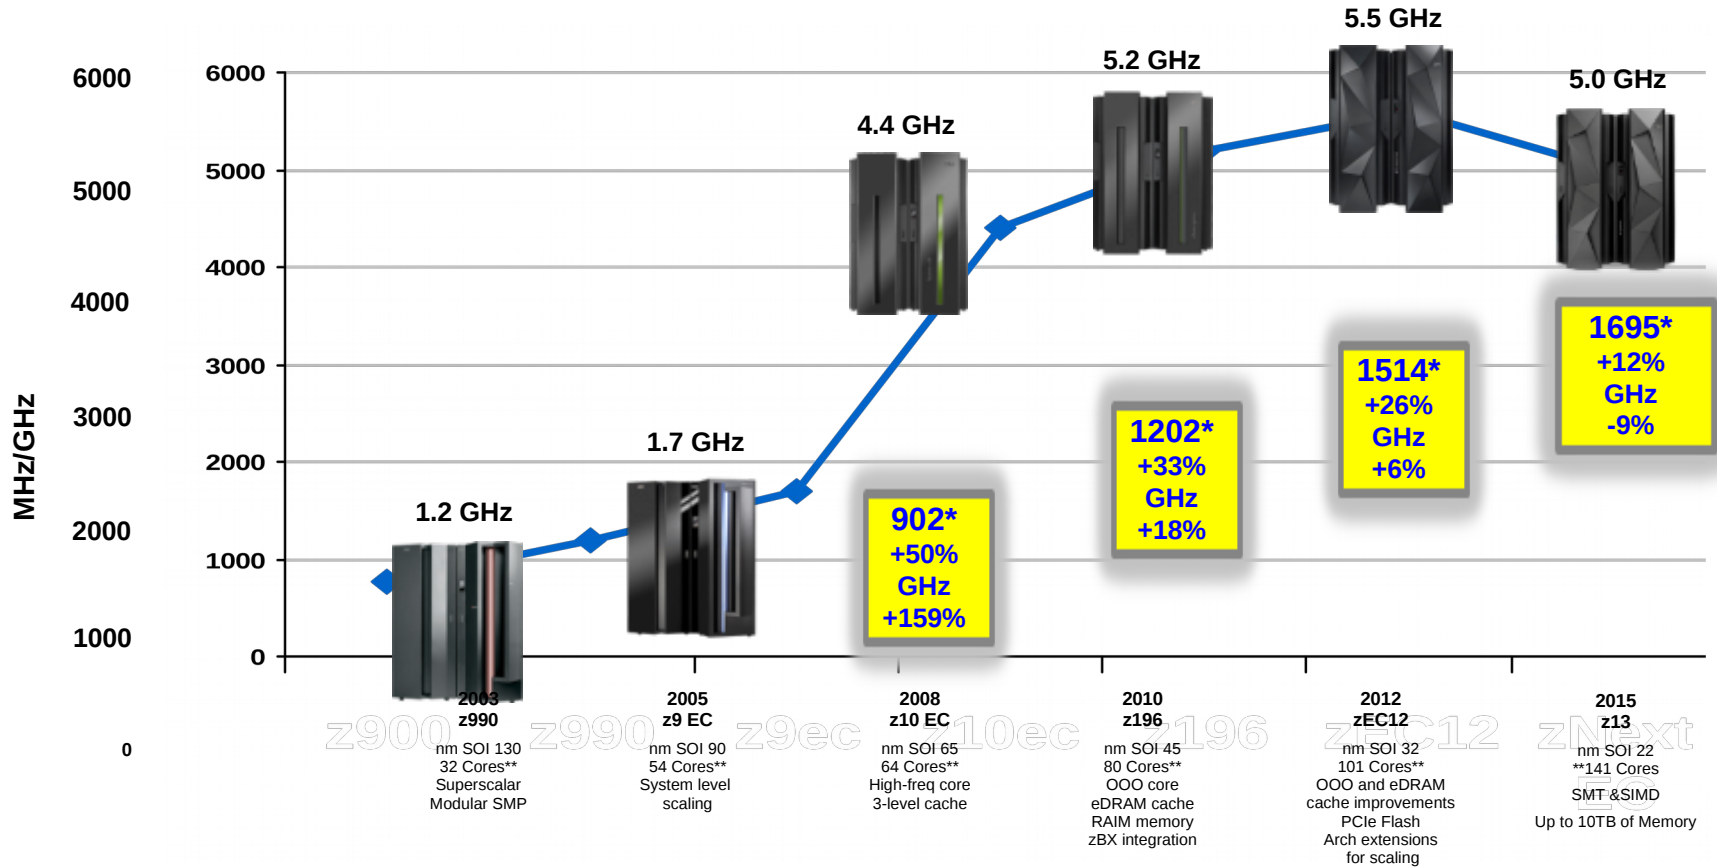
z13
1/2015



Leadership System Capacity and Performance

- Modularity & Scalability
- Dynamic SMT
- Supports two instruction threads
- SIMD
- PCIe attached accelerators
- Business Analytics Optimized

... every generation increases performance....



... every generation increases capacity....

System I/O Bandwidth: 832 GB/Sec*

384 GB/Sec*

288 GB/sec*

172.8 GB/sec*

3
TB

1.5
TB

512
GB

600

902

1202

1514

PCI for
1-way
1695

54-way

64-way

80-way

101-way

141-way Customer Processors



Memory
10 TB

*No server can fully exploit its maximum I/O bandwidth
PCI – Processor Capacity Index (IBM MIPS)*


... every generation increases overall scale

Each new range continues to deliver:

- New function
- Unprecedented capacity to meet consolidation needs
- Improved efficiency to further reduce energy consumption
- Continues to delivering flexible and simplified on demand capacity
- A mainframe that goes beyond the traditional paradigm



z13s/z13 Central Processor: Overview



▪ 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 six-seven active cores available on z13s
- Up to eight active cores (PUs) per chip on z13
 - 4.3 GHz on z13s systems
 - 5.0 GHz on z13 systems
 - 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)**
- **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
- **On chip 64 MB eDRAM L3 Cache**
 - Shared by all cores
- **I/O buses**
 - One InfiniBand I/O bus
 - Two PCIe I/O buses
- **Memory Controller (MCU)**
 - Interface to controller on memory DIMMs
 - Supports RAIM design

z13s/z13 Central Processor: Architecture Additions

- **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 zBC12
 - 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 zBC12
 - 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) ISA 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.

z13s/z13 Central Processor: SMT



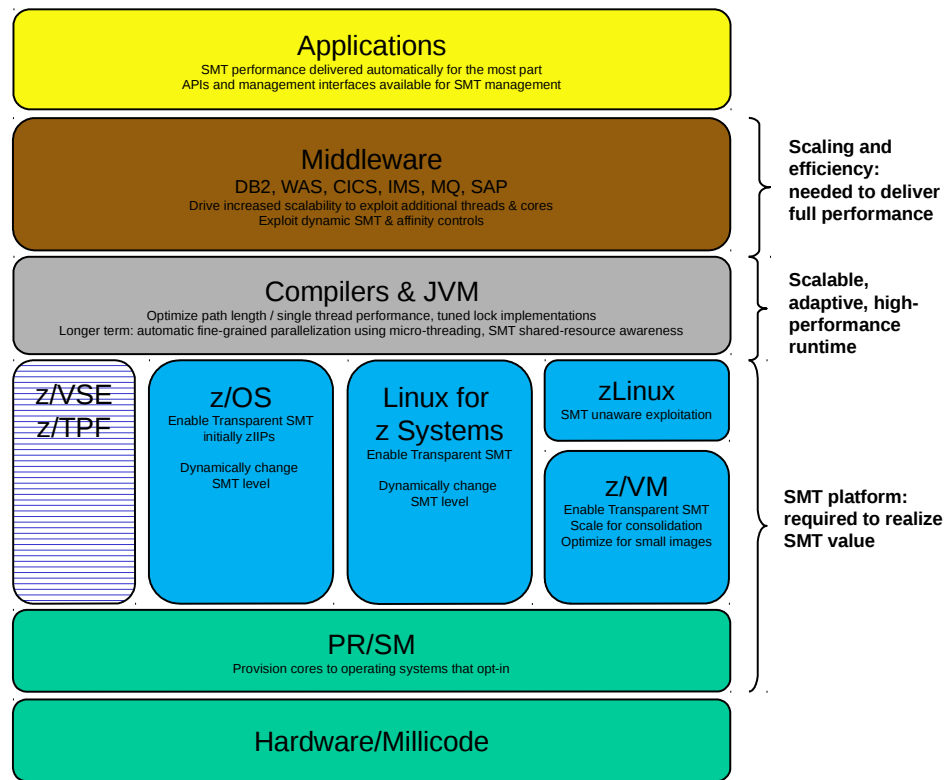
*Which approach is designed for the highest volume** of traffic?
Which road is faster?*

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

- 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:
 - ⊖ z13s zIIP is **68%** greater than a zBC12 zIIP
 - ⊖ z13s IFL is **61%** greater than a zBC12 IFL
 - ⊖ z13s zIIP is **129%** greater than a z114 zIIP
 - ⊖ z13s IFL is **119%** greater than a z114 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 individual thread) may be faster in single thread mode.
 2. Because SMT is not available for CPs (currently), LSPR ratings do not include it

z13s/z13 Central Processor: SMT (cont.)

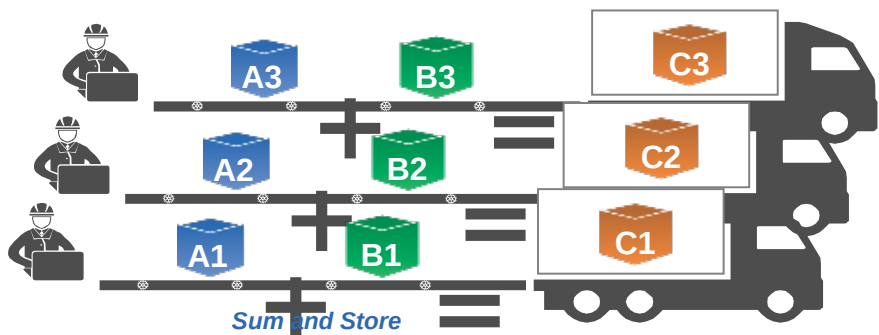
- **z13 is the first z System Processor to support SMT**
 - Enable continued scaling of per-processor capacity
 - z13 and z13s support 2 threads per core on IFLs and zIIPs
- **Increases per-core and system throughput versus single thread design**
 - More work done per unit hardware
 - Aligns with industry direction of multi-thread
 - Improves **per-core** performance comparisons vs. X86, POWER
 - Improves efficiency of IFL for Linux consolidation
- **Designed to preserve unique z System values and attributes**
 - Full support for 2-level processor virtualization
 - Full z/Architecture capability for each thread
- **Design will allow independent enablement of SMT by LPAR**
 - Operating systems must be explicitly enabled for SMT
 - Operating system may opt to run in single-thread mode
- **Processors can run in single-thread operation for workloads needing maximum thread speed**
- **Functionally transparent to middleware and applications**
 - No changes required to run in SMT partition
 - **Operating System/Hypervisor Support**
 - z/OS (for zIIPs)
 - zVM (for IFLs)
 - Linux: IBM continues working with its Linux Distribution partners to support new functions/features



z13s/z13 Central Processor: SIMD

Scalar


SINGLE INSTRUCTION, SINGLE DATA



Instruction is performed for
every data element

SIMD

SINGLE INSTRUCTION, MULTIPLE DATA



Perform instructions on
every element at once

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

z13s/z13 Central Processor: SIMD (cont.)

- **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, z/Linux, and z/VM support available or under development
 - Compiler exploitation
 - IBM Java
 - XL C/C++ on zOS or Linux on z
 - Enterprise COBOL
 - Enterprise PL/I



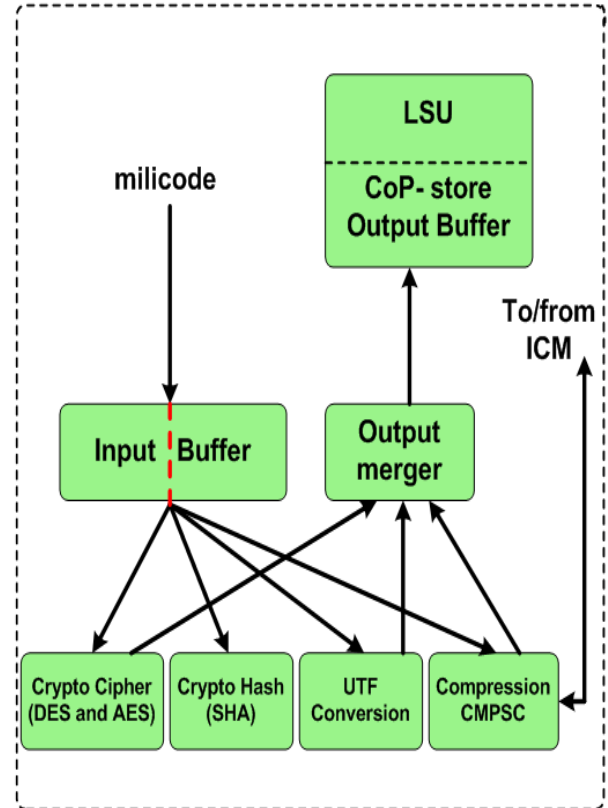
Workloads		
Java.Next	C/C++Compiler built-ins for SIMD operations (z/OS and Linux on z Systems)	MASS & ATLAS Math Libraries (z/OS and Linux on z Systems)
SIMD Registers and Instruction Set		

z13s/z13 Central Processor: Co-processors

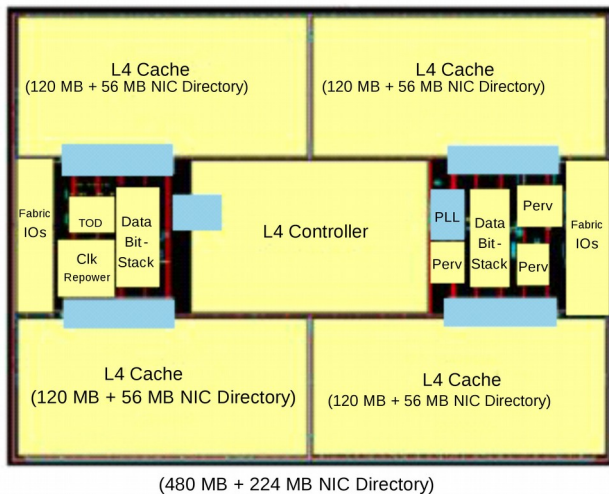
- **Coprocessor dedicated to each core (was shared by two cores on z114)**
 - Independent compression and cryptographic engines
 - Redesigned from “ground-up” for crypto, compression, hashing, UTF-conversion
 - Available to any processor type (CP, zIIP, IFL)
 - Owing processor is busy when its coprocessor is busy
 - Instructions available to any processor type; data stored direct to L1D
- **Data compression/expansion engine**
 - Static dictionary compression and expansion
- **CP Assist for Cryptographic Function**
 - Significant performance improvements for large blocks of data
 - **AES: 2x throughput compared to zBC12**
 - **TDES: 2x throughput compared to zBC12**
 - **SHA: 3.5x throughput compared to zBC12**

Exploiters of the CPACF benefit include:


- 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



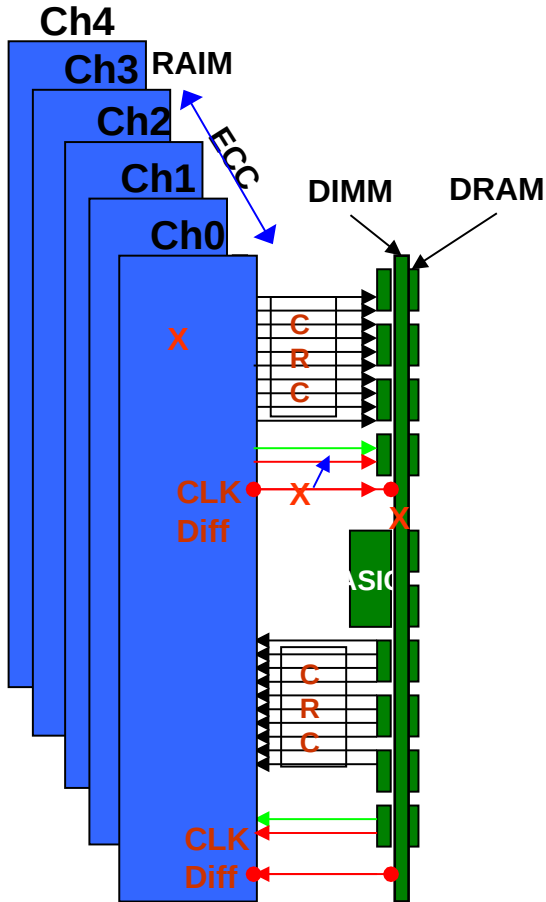
z13s/z13 System Controller



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



z13s/z13 RAIM 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


Big Performance Gains with Big Memory!

- **10TB available on z13! 4TB on z13s!**
- **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
 - Unlocks in-transaction analytics capability
- **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

The I/O monster unveiled

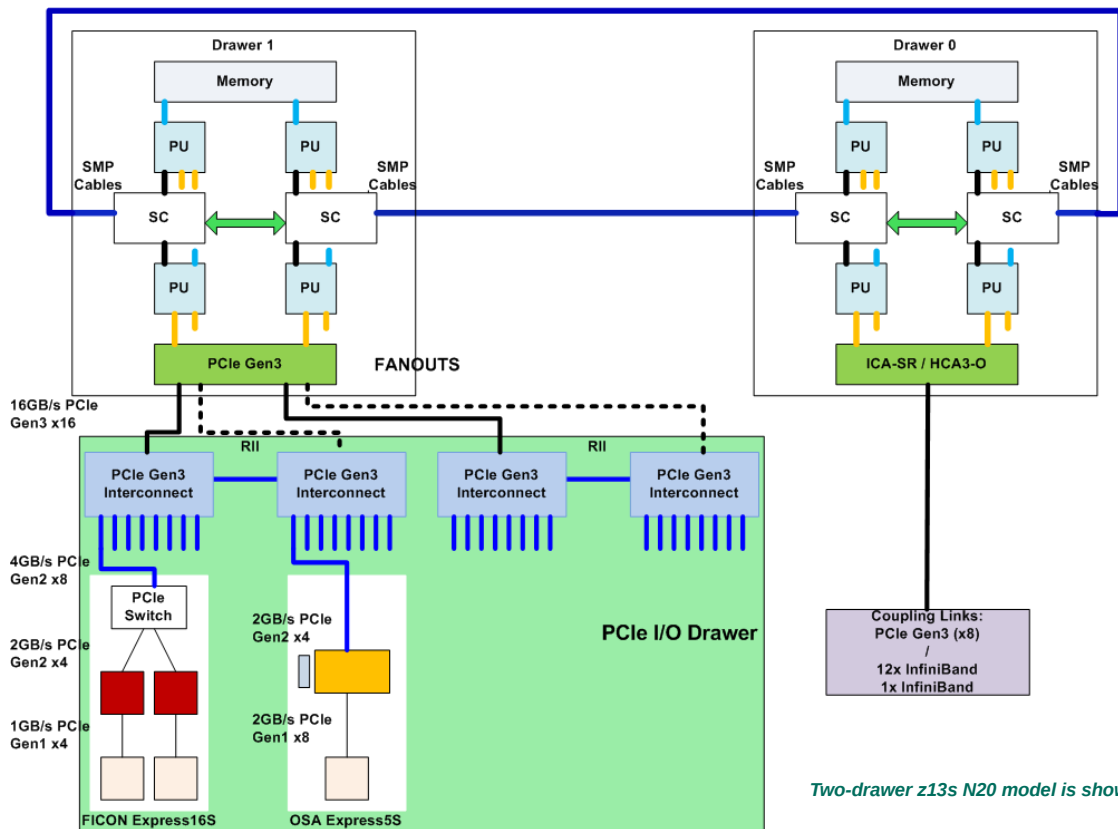
*Connectivity to millions of devices,
massive throughput, and world-class
reliability*

Massive I/O throughput via massive processing...



Dedicated I/O coprocessors bring RAS, cost savings, massive throughput and connectivity, and added compute power to workloads, allowing CPUs to focus where they should

I/O Connectivity Example: Logical View of z13s



Two-drawer z13s N20 model is shown


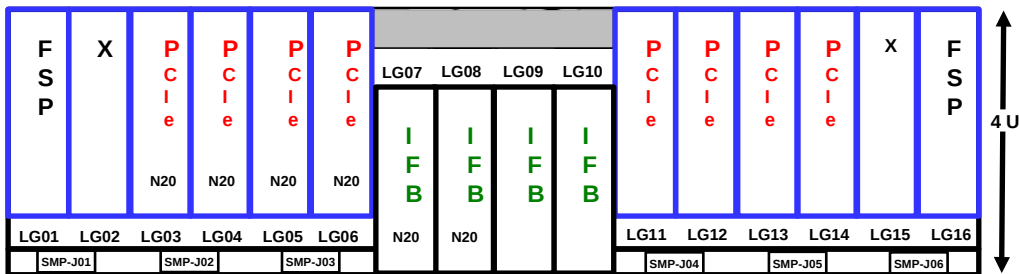
▪ The InfiniBand I/O infrastructure first introduced on the z10 is still supported...

- InfiniBand fanouts supporting the 6 GBps InfiniBand I/O interconnect
- InfiniBand I/O card domain multiplexers with Redundant I/O interconnect in:
 - The 5U, 8-slot, 2-domain I/O drawer (carry forward only)
- Selected non-PCIe I/O features
 - FICON Express8 LX (FC 3325) and FICON Express8 SX (FC 3326)

▪ ... but PCI Express Generation 3 (PCIe Gen3) I/O infrastructure is now standard

- PCI Express Generation 2 (PCIe Gen2) I/O infrastructure introduced with z196/z114
- PCIe Gen3 fanouts and PCIe Interconnect Gen3 supporting the 16 GBps PCIe I/O interconnect is the standard starting with the z13
- PCIe Interconnect Gen3 (i.e. PCIe Switch) with Redundant I/O interconnect for I/O domains in a 7U, 32-slot, 4-domain PCIe I/O drawer
- **Up to two PCIe I/O drawers available on z13s and five on z13**

I/O Connectivity Example: Physical View of z13s




- PCIe Fanout Slots: LG03 – LG06 (N20) and LG11 – LG14 (N10 and N20), can support:
 - Up to 8 one-port PCIe 16 GBps I/O fanouts to support up to 8 domains in 32-slot PCIe I/O drawers
 Note: A zBC12 Model H06 with four two-port 8 GBps PCIe fanouts supports up to 8 domains in 32-slot PCIe I/O drawers; but **a z13s CPC drawer supports double the bandwidth to each domain**
 - Up to 8 ICA (PCIe-SR) two-port coupling fanouts to support up to 16 8 GBps coupling links
- IFB Fanout Slots, LG07 – LG08 (N20) and LG09 – LG10 (N10 and N20), can support:
 - Up to four HCA3-O 12x InfiniBand coupling fanouts, 8 12x 6 GBps links – Two per fanout
 - Up to four HCA3-O LR 1x InfiniBand coupling fanouts 16 1x 5 Gbps links – Four per fanout
 Note: A zBC12 Model H06 with 4 two-port HCA3-O 12x InfiniBand coupling fanouts can support 8 12x links
 A zBC12 Model H06 with 4 four-port HCA3-O LR 1x InfiniBand coupling fanouts can support 16 1x links
 - Up to two two-port HCA2-C 6GBps I/O fanouts (2 8-slot I/O drawers) with two slots left
- Slots LG01 and LG16 always have Flexible Support Processors (FSPs)
- SMP-J01 to J06 connectors are for A-Bus cables to nodes in the other CPC drawer

Eight PCIe Fanout slots
FSP = Flexible Support Processor


Four IFB Fanout slots

I/O Connectivity Example: PCIe Cages

Front




Rear




- Supports only PCIe I/O cards
 - z13s: Up to two drawers
 - z13: Up to five drawers
- Supports 32 PCIe I/O cards, 16 front and 16 rear, vertical orientation, in four 8-card domains (shown as 0 to 3)
- Requires four 16 GBps PCIe switch cards, each connected to a 16 GBps PCIe I/O interconnect to activate all four domains.
- To support Redundant I/O Interconnect (RII) between front to back domain pairs 0-1 and 2-3 the two interconnects to each pair will be from 2 different PCIe fanouts. (All four domains in one of these drawers can be activated with two fanouts.)
- Concurrent field install and repair.
- Requires 7 EIA Units of space (12.25 inches ≈ 311 mm)

FICON Express16S: Overview

- For FICON, zHPF, and FCP environments
 - CHPID types: FC and FCP
 - Two PCHIDs/CHPIDs
- Auto-negotiates to 4, 8, or 16 Gbps
 - 2Gbps connectivity NOT supported
 - FICON Express8S will be available to order for 2Gbps connectivity
- Increased I/O Devices (subchannels) per channel for all FICON features:
 - TYPE=FC: Increased from 24k to 32k to support more base and alias devices
- Increased bandwidth compared to FICON Express8S
- 10KM LX - 9 micron single mode fiber
 - Unrepeated distance - 10 kilometers (6.2 miles)
 - Receiving device must also be LX
- SX - 50 or 62.5 micron multimode fiber
 - Distance variable with link data rate and fiber type
 - Receiving device must also be SX
- Two channels of LX or SX (no mix)
- Small form factor pluggable (SFP) optics
 - Concurrent repair/replace action for each SFP




FC 0418 – 10KM LX, FC 0419 - SX




FICON Express16S: Changed from previous cards

- **FICON Express16S - 16 Gbps Link Speeds**
 - 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 by up to 32%
- **Preserve Virtual WWPNS for NPIV configured FCP channels**
 - Designed to simplify migration to a new-build z13 or z13s
- **Support for a maximum of 32K devices per FICON channel**
 - Up to 85 Logical Partitions (40 on z13s): More flexibility for server consolidation
 - More than 10.2 million I/O device attachments supported
- **zHPF Extended I/O execution at Distance**
 - 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**
 - 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**
 - Designed to address high bit-error rate on high frequency ($\geq 8\text{Gb/s}$) links
 - Estimated equivalence to doubling optical signal power
- **Fibre Channel Read Diagnostic Parameters Extended Link Services (ELS)**
 - Supports links failure diagnostics and predictive analysis for optics, cables, and ports
- **SAN Fabric I/O Priority**
 - 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: Speeds and Feeds

zHPF and FICON Performance*



FCP Performance*



*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.

RoCE Express: 10 Gigabit Ethernet

Designed to support high performance system interconnect

- Shared Memory Communication (SMC) over Remote Direct Memory Access (RDMA) (SMC-R) Architecture exploits RDMA over Converged Ethernet (CE) - RoCE
- Shares memory between peers
- Read/write access to the same memory buffers without application changes
- Designed to increase transaction rates greatly with low latency and reduced CPU cost

Configuration

- z13/z13s - Both 10 GbE SFP+ ports enabled
- z13/z13s - Support for up to 31 Logical Partitions
- A switched connection requires an enterprise-class 10 GbE switch with SR Optics, Global Pause enabled & Priority Flow Control (PFC) disabled
- Point-to-point connection is supported
- Either connection supported to z13, z13s, zEC12 and zBC12
- Not defined as a CHPID and does not consume a CHPID number
- Up to 16 features supported on a z13/z13s/zEC12/zBC12
- Link distance up to 300 meters over OM3 50 micron multimode fiber


Exploitation and Compatibility

- z/OS V2.1
- IBM SDK for z/OS Java Technology Edition, Version 7.1
- z/VM V6.3 support for z/OS V2.1 guest exploitation
- Linux on z Systems – IBM is working with Linux distribution partners to include support in future releases*

RoCE Express links are now fully shareable between multiple z/OS images!

10 GbE SFP+

10 GbE SFP+



FC 0411 10GbE RoCE Express



OM3 fiber recommended

Crypto Express5S: Encryption Acceleration

- **Native PCIe card (FC 0890)**
 - Resides in the PCIe I/O drawer
 - Requires CPACF Enablement (FC 3863)
- **New Crypto Module**
 - Designed for **2x performance increase** over Crypto Express4S
 - Added L2 Cache, New Crypto ASIC and processor upgrade
 - Designed to support up to 85 domains for logical partitions or z/VM guests
- **Designed to Meet Physical Security Standards**
 - FIPS 140-2 level 4
 - ANSI 9.97
 - Payment Card Industry (PCI) HSM
 - Deutsche Kreditwirtschaft (DK)
- **New Functions, Standard and Compliance**
 - Drivers: NIST via FIPS standards and implementation guidance requirements; emerging banking standards: and strengthening of cryptographic standards for attack resistance
 - VISA Format Preserving Encryption (VFPE) for credit card numbers
 - Enhanced public key Elliptic Curve Cryptography (ECC) for users such as Chrome, Firefox, and Apple's iMessage
- **New Trusted Key Entry Workstation**
 - Workstation and LIC – FC 0847 with new crypto module and TKE LIC 8.0 or higher is required for new functions (see later chart for details)
 - Required: EP11 (PKCS #11) Mode, Recommended: Common Cryptographic Architecture (CCA) Mode
 - Additional Smart Cards (FC 0892) – Support for stronger encryption than previous cards



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 4 certification to meet regulations and compliance for PCI standards**

z13 Flash Express: Storage-class Memory

Four 400 GByte ($G=10^9$) SSDs support 1.4 TBytes ($T=2^{40}$) of Storage Class Memory (AES encrypted)



Cable connections to form a RAID 10 Array across a pair of Flash Express Cards.

z13 Flash Express: Details

▪ Provides Storage Class Memory

- Implemented via NAND Flash SSDs mounted in PCIe Flash Express features
- Protected by strong AES Encryption done on the features
- 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
- Enables extremely responsive paging of 4k pages to improve z/OS availability
- Enables pageable large (1 MB) pages
- ... less paging for Java and DB2!

▪ 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)
- Significantly faster, less disruptive diagnostics with shortened first failure data capture time
- Linux on z Systems
 - SLES 11 SP3 and RHEL 6.4

~25%

*Reduction in SVC
dump elapsed time*

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*

10x

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

zEnterprise Data Compression (zEDC)

	<ul style="list-style-type: none">▪ Alebra – Parallel Data Mover (PDM)<ul style="list-style-type: none">– Uses zEDC compression in lieu of Software-based compression to provide excellent qualities of service.
	<ul style="list-style-type: none">▪ ASE – OMCS<ul style="list-style-type: none">– Takes SLIKZIP and SLIKSFTP performance to a whole new level
	<ul style="list-style-type: none">▪ PKWARE – PKZIP and SecureZIP v15<ul style="list-style-type: none">– Accelerated deflate compression and automatic detection of zEDC
	<ul style="list-style-type: none">▪ Software AG – Entire Net-Work<ul style="list-style-type: none">– High performance transaction processing

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



Example Use Cases

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

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

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

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

zEDC was expressly created using industry standard APIs to encourage ISVs to leverage its high-speed compression value in applications ISVs create. With access to zEDC, ISV applications are more valuable to end users.

Integrated Coupling Adapter SC (ICA-SR)

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


- Recommended for Short Distance Coupling z13/z13s to z13/z13s, 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



Greater Connectivity

- z13s provides more ICA SR coupling fanouts per CPC drawer when compared to 12x PSIFB Coupling on either z114 or zBC12
- A single z13/z13s CPC drawer supports up to 20 ICA SR links vs 16 12x on z114/zBC12

Alleviate PSIFB Constrained Configurations

- Utilizing ICA SR frees HCA fanout slots for essential PSIFB Coupling links during migration
- For z13/z13s to z13/z13s connectivity, using ICA SR in place of PSIFB over Infiniband may enable clients to remain in the same CPC footprint as their z114 or zBC12 enterprises

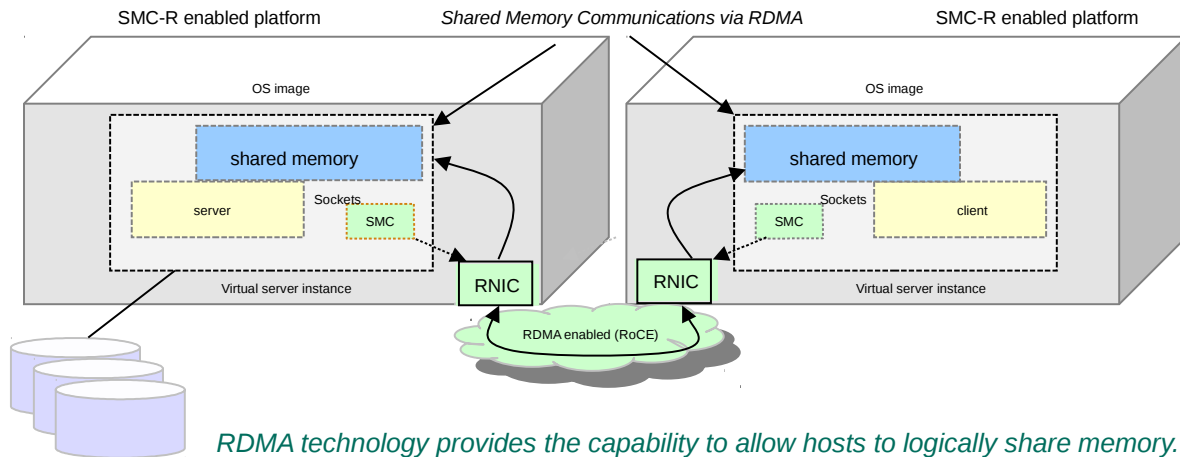
Type	Speed	Distance	Fanout
ICA SR	8 GBps	150 meters	ICA SR
12x InfiniBand	6 GBps	150 meters	HCA3-O
1x InfiniBand	5 or 2.5 Gbps	10 km	HCA3-O LR

Software Technologies Update

*Hypervisors, operating systems,
operations management and
automation, and more!*

Shared Memory Communication – RDMA (SMC-R)

- **Optimized Network Performance (leveraging RDMA technology)***
- **Transparent to (TCP socket based) application software**
- **Leverages existing 10Gbps Ethernet infrastructure (RoCE)**
- **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**




RDMA technology provides the capability to allow hosts to logically share memory. The SMC-R protocol defines a means to exploit the shared memory for communications - transparent to the applications!

Up to **50% CPU savings** for FTP file transfers across z/OS systems versus standard TCP/IP

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

z/OS V2.2 Communications Server now automatically selects between TCP/IP and RoCE



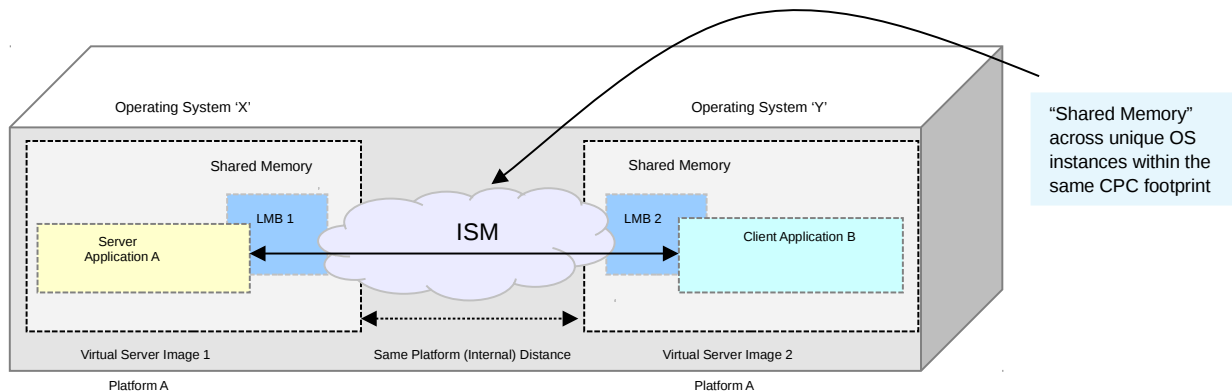
* Latency and CPU savings are based on workload type (latency focus for interactive workloads while CPU savings is on bulk traffic).

Shared Memory Communication – Direct (SMC-D)

Up to **61%** CPU savings for FTP file transfers across z/OS systems versus HiperSockets*

Up to **9x** improvement in throughput with more than a **88%** decrease in CPU consumption and a **90%** decrease in response time for streaming workloads versus using HiperSockets*

Up to **91%** improvement in throughput and up to **48%** improvement in response time for interactive workloads versus using HiperSockets*



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

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

- The Shared Memory Communications – Direct Memory Access protocol can significantly optimize intra-CEC Operating Systems communications – transparent to socket applications
- Tightly couples socket API communications/memory within the CPC
- Eliminates TCP processing in the data path
- Does not require additional hardware

* All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

Secure Service Containers

- **Secure Service Containers** provide the base infrastructure needed to host and build Software Appliances including Operating System, middleware, SDK and firmware support
- A new SSC partition mode (LPAR) enables hosting a software virtual appliance
- The first software virtual appliance will be z/VSE Network Appliance - used by z/VSE for faster TCP/IP communication to other systems. Customers see up to 3X performance with this offering....

Shared features:

- Encapsulated Operating Systems
- Services provided via Remote APIs (RESTful) and web interfaces
- Embedded monitoring and self-healing
- End-2-End tamper protection from installation, secure boot, and memory protection during runtime
- Protected Intellectual Property of appliance components: appliances cannot be altered
- Tested/Qualified by provider for a specific use case

Potential future appliances:

- Integrated Analytics – Investigate and diagnose problems faster, predict and prevent problems and optimize the systems within the z IT environment via IBM zAware
- Security – appliances that allows client to leverage z Systems cryptographic hardware across platforms; provide ultra-secure environment for Blockchain environments
- Full SDK now available from IBM to z Systems ISVs

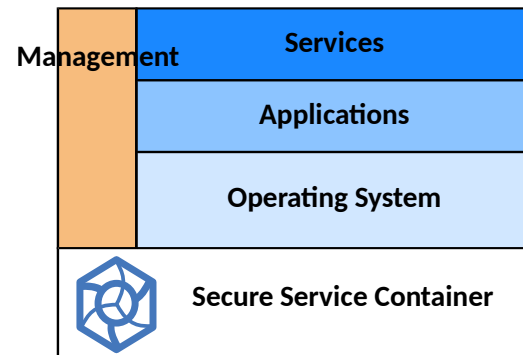
Protect Blockchain software, chain code and data!

Root users and system administrators cannot access Blockchain contents




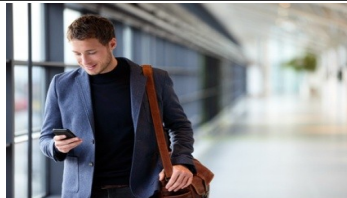
Malware cannot self-install in this container

Encryption keys are always protected

No other platform provides this capability!



z Software Stack: Tightly coupled to HW

<p>z/OS Version 2.2</p>		<ul style="list-style-type: none">• Extreme scalability with support for up to 40 LPARS (85 LPARS for z13)• Strengthened security - faster Cryptography and digitally signed audit records• Up to 4 TB memory per z/OS image• Simplified management with entitled web-based console• New support for Multi-factor Authentication Services with RACF
<p>z/VSE Version 6.1</p>		<ul style="list-style-type: none">• Faster I/O with FICON Express 16S with link rate of 16 Gbps• Performance and functional enhancements for online processing• Improved network security with firewall functionality, 3X TCP/IP performance• Better HW Encryption with Crypto Express5S• Wide portfolio using Linux on z; delivered as SSC LPAR
<p>Linux</p>		<ul style="list-style-type: none">• Multithreading with SMT may allow for per core software savings• Ability to host and manage more workloads efficiently / cost-effectively• Automatic identification of unusual messages• Integrated continuous availability & disaster recovery solution: New GDPS support via virtual appliance!• Introducing Canonical Ubuntu in addition to SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL) for more choice
<p>z/TPF</p>		<ul style="list-style-type: none">• Management of extreme transaction volumes up to hundreds of thousands of transactions per second• Fast / consistent response across predictable and unpredictable peaks• Low cost per transaction for large applications and memory tables Centralized database handling routines to effectively manage databases• Application interface enables high speed access to persistent data

Multi-Factor Authentication on z/OS

- IBM Multi-Factor Authentication on z/OS provides a way to raise the assurance level of z/OS and applications / hosting environments by extending RACF to authenticate users with multiple authentication factors.
- Support for multiple third-party authentication systems at the same time
 - RSA® Ready supporting RSA SecurID® Tokens (hardware & software based)
 - IBM TouchToken – Timed One time use Password (TOTP) generator token
 - Support PIV/CAC cards - Commonly used to authenticate in the Public Sector enterprises
- Tightly coupled with SAF & RACF
 - RACF provides the configuration point to describe multi-factor authentication requirements down to a per User ID basis
 - Deep RACF integration for configuration and provisioning data stored in RACF database allowing seamless back-up and recovery, and remote management

Typical Client Use Cases:

- **Enable higher-security user authentication** on IBM z/OS systems using RACF for security
- Provides multi-factor authentication in support of PCI-DSS requirements for personnel with non-console administrative access to the systems handling card data.



Linux + z Systems = "A match made in Heaven"

Distributions



Hypervisors



Languages



Runtimes



Management



Cloud Manager



Database



DB2

Analytics



z Systems and Virtualization

IBM z/VM

z/Proprietary Server Virtualization that is completely integrated into the full stack. Complete hardware awareness. Supported on all IBM z Systems and LinuxONE servers. z/VM will continue to be enhanced to support Linux Workloads.

KVM for IBM z Systems

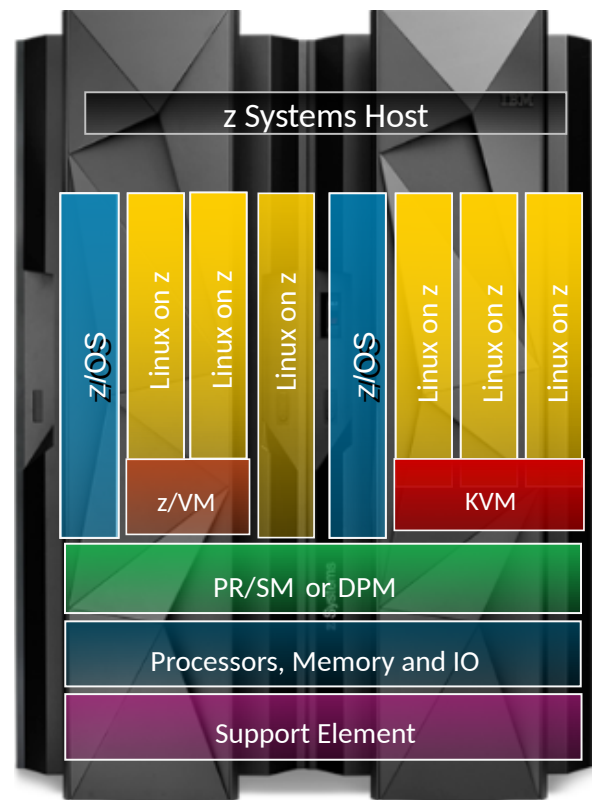
KVM for IBM z provides an open source choice for IBM z Systems and LinuxONE virtualization for Linux workloads. Best for clients that are not familiar with z/VM and are Linux centric admins

IBM Processor Resource/System Manager (PR/SM)

Divide one physical server into up to 85 logical partitions (LPAR) running a mix of multiple z/OS®, z/VM, Linux, KVM for IBM z, Transaction Processing Facility (TPF) and z/VSE® instances isolated and secured in parallel. Share resources across LPARs or dedicated to a particular LPAR. Running a mix of multiple z/OS, z/VM, Linux, TPF, KVM for IBM z and z/VSE instances isolated and secured in parallel.

IBM Dynamic Partition Manager (DPM)

DPM is a new administrative mode of PR/SM that simplifies configuration of partitions, associated resources and I/O. It allows partitions to be quickly configured, along with the management of system resources including integrated dynamic I/O management, as easily as other virtualized environments. It was developed for new-to-z users working on servers with KVM for z, z/VM and/or Linux for z Systems or LinuxONE as a partition-hosted operating system. IBM z/VM V6.4 is a supported environment using IBM Dynamic Partition Manager for Linux-only systems with SCSI storage.



z/VM 6.4: The Virtualization Gold Standard



Support for 2TB of memory and SMT means higher levels of workload consolidation, superior levels of elasticity for workload spikes, allowing considerable growth in memory-intensive applications; SMT-2 enablement means 30% or higher capacity improvement on IFLs

Operation improvements with ease-of-use enhancements from users include querying service of the running hypervisor; providing environment variables to allow client programming automation based on systems characteristics and client settings; and improved query capabilities for system shutdown

SCSI (Small Computer System Interface) improvements for guest attachment of disks and other peripherals, and host or guest attachment of disk drives:

Improve RAS capabilities within the z/VM SCSI subsystem for greater resiliency for SCSI devices behind an SVC (SAN Volume Controller)

Increase efficiency and reduce complexity by allowing Flash Systems to be directly attached for z/VM system use without the need for an SVC

Allow concurrent code loads on the SVC – and devices incorporating SVC technology – without quiescing EDEVICE I/O

Increased scalability by exploiting Guest Enhanced DAT to allow guests to take advantage of large (1MB) pages, decreasing the memory and overhead required to perform address translation

Increased efficiency with HyperPAV paging that takes advantage of IBM DS8000 features to increase the bandwidth for paging and allow for more efficient memory management of over-committed workloads

Customer Choice of Linux Distribution with planned support for **Canonical Ubuntu** distribution in addition to **Red Hat** and **SUSE**

IBM Wave for z/VM now an optional, priced feature of z/VM itself!

VM's world class industry proven virtualization technology offers the ability to host extremely large number of virtual servers on a single server

Host non-Linux environments with z/VM on IBM z Systems – z/OS, z/VSE and z/TPF

Virtual machines share system resources with very high levels of resource utilization.

Optimized for z Systems architecture multi-tenancy, capacity on demand and support for multiple types of workloads

KVM for z Systems and LinuxONE



Support new analytics workloads with **Single Instruction Multiple Data (SIMD)** for competitive advantage

Deliver higher compute capacity with support for **Simultaneous Multithreading (SMT)** to meet new business requirements

RAS support enhanced for problem determination and high availability setup to reduce down time and quickly react to business needs

Secure and protect business data with Crypto exploitation that leverages **hardware acceleration for cryptographic functions**

Provide clients with choices for flexibility based upon their protocol environment

- Connect a variety of peripherals, especially storage devices drives, with Internet Small Computer System Interface (iSCSI)
- Access files on remote hosts exactly the same way a user would access any local files with Network File System (NFS) which works across a variety of server and host architectures

Unattended installation of the KVM hypervisor simplifies administration


Customer choice of Linux Distribution with planned support for **Canonical Ubuntu** distribution in addition to **SUSE for KVM** on IBM z Systems and IBM LinuxONE

Optimized for z Systems architecture multi-tenancy and support for multiple types of workloads

Higher-level virtualization functions so critical workloads receive resources and priority based on established goals

Coexists with z/VM , Linux on IBM z, z/OS, z/VSE, z/TPF





**IBM z13 & z13s: a secure
platform for hybrid cloud
innovation**

**With the highest speeds, the most diverse workloads,
and simply the best security on the planet**

Notices and Disclaimers

Copyright © 2016 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY. IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law

Notices and Disclaimers Con't.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. IBM EXPRESSLY DISCLAIMS ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, Aspera®, Bluemix, Blueworks Live, CICS, Clearcase, Cognos®, DOORS®, Emptoris®, Enterprise Document Management System™, FASP®, FileNet®, Global Business Services®, Global Technology Services®, IBM ExperienceOne™, IBM SmartCloud®, IBM Social Business®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®, Netcool®, OMEGAMON, OpenPower, PureAnalytics™, PureApplication®, pureCluster™, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, Smarter Commerce®, SoDA, SPSS, Sterling Commerce®, StoredIQ, Tealeaf®, Tivoli®, Trusteer®, Unica®, urban{code}®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.