IBM z Systems Hardware - 2016 Technical Overview

Session B01 / C15

Walter Kläy, IMS SWAT Team
walter.klaey@ch.ibm.com

with support from Kevin Hite (SVL) and Silvia Mueller (BOB)

Innovation never stops.
…2016…

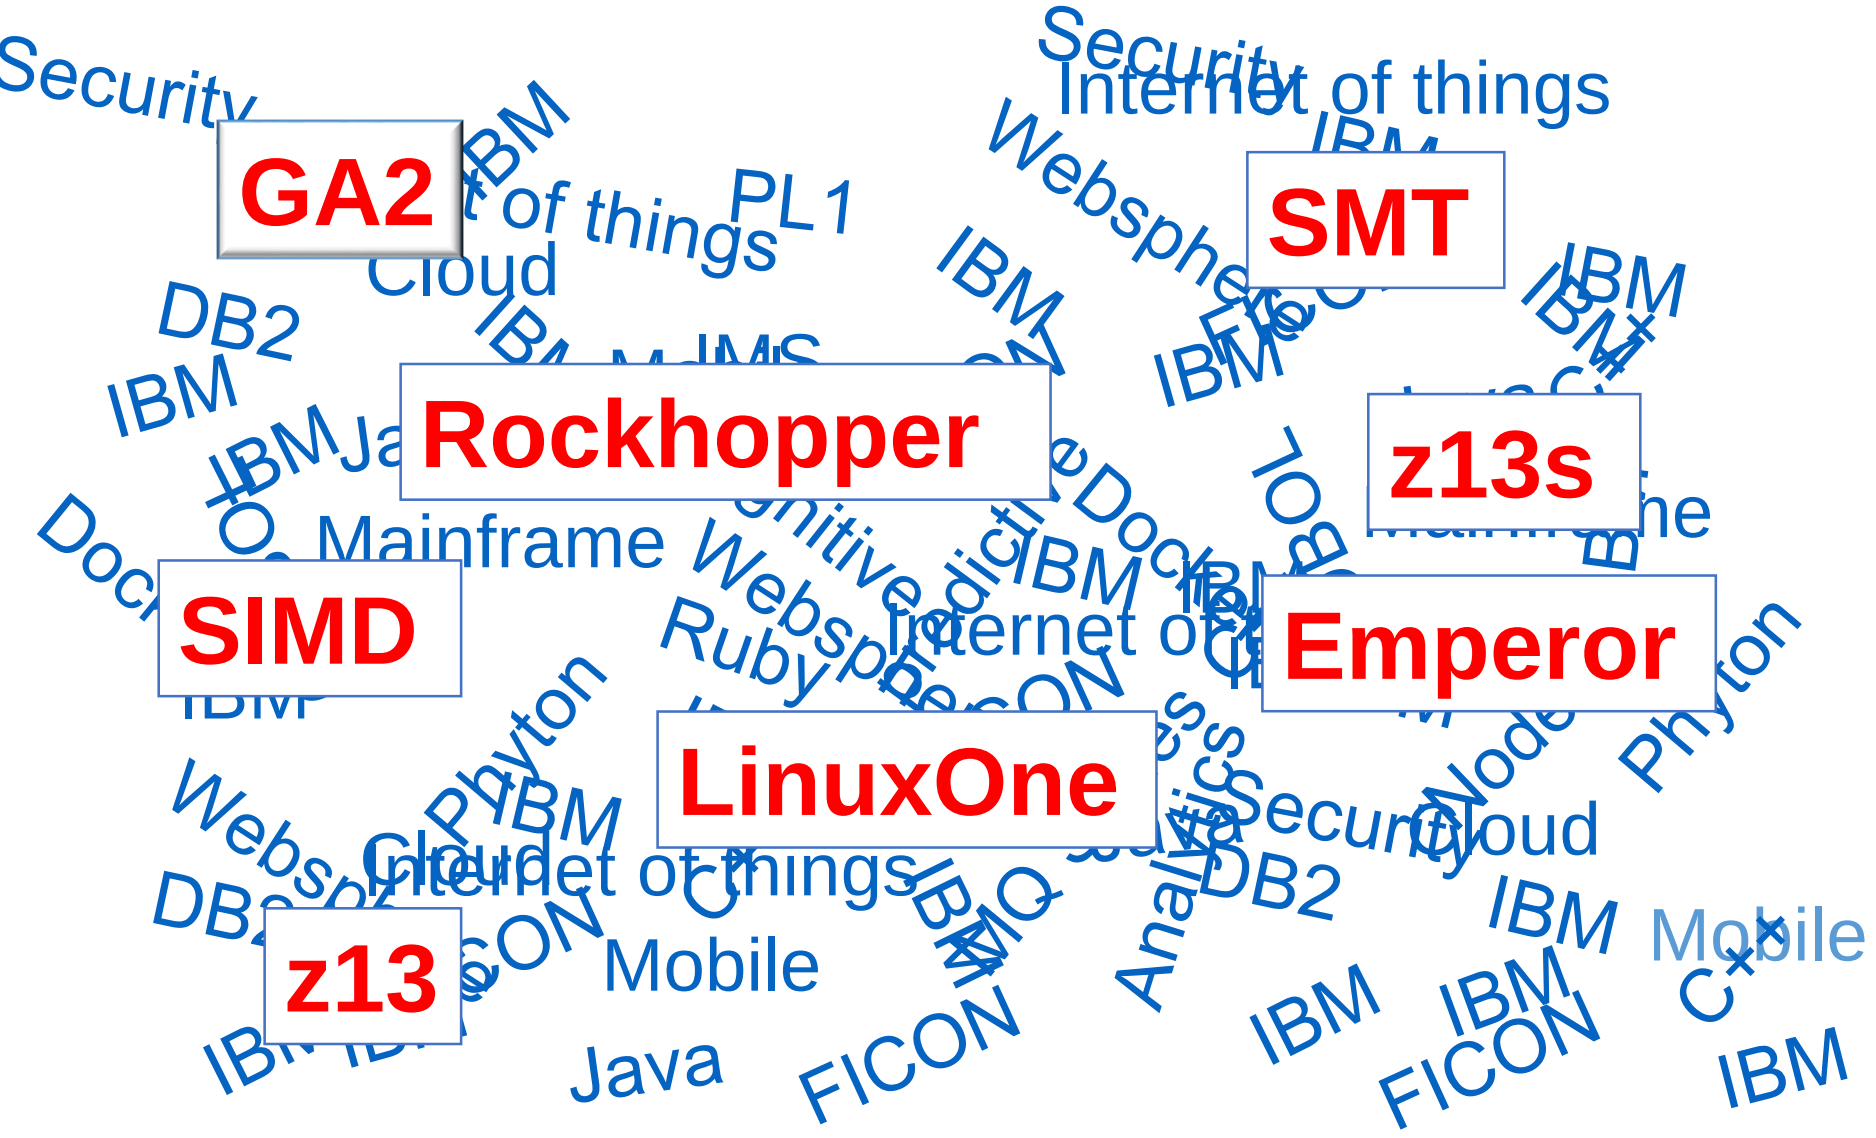

Digital Revolution
World becoming smarter

Transform interactions
Personalized everything
In the moment right now

What is happening?

16 billion connected devices
75 billion devices by 2020
7 billion smart phones

Infrastructure of the company
Infrastructure of the city
Infrastructure of the world

Respect and protect security and privacy

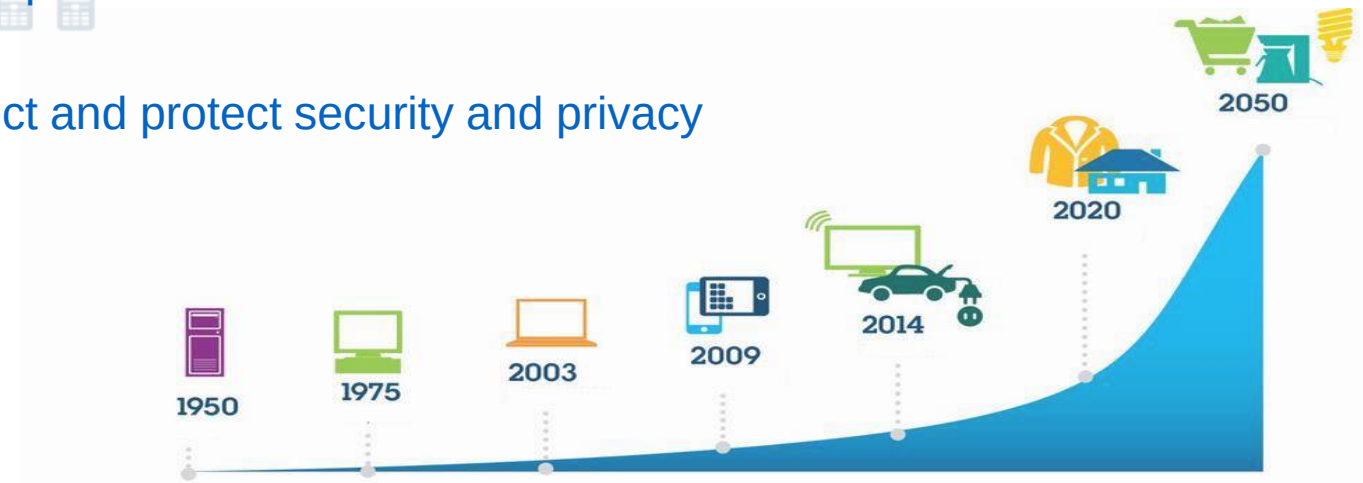

Figure 1: Each inflection point in the history of computing has triggered an explosion in the number of computing devices

The market is moving, forcing businesses to transform

Explosion in transaction growth

driven by mobility and the Internet of Things

Analytics is moving to real time

to capture new opportunities at the point of impact

Hybrid cloud is the new standard

for delivering service, agility, trust and efficiency

z13 Overview

IBM z Systems Evolution

Jan 12, 2015 – z13

Feb 16, 2016 – z13s / z13 GA2

New Brand: LinuxOne

Au – Emperor

Jan 25, 2016 – Rockhopper

z13(s) Functions and Features (DGA Driver Level 22)

System, Processor, Memory
Five hardware models
Eight core 22nm PU SCM
Up to 141 processors configurable as CPs, zIIPs, IFLs, ICFs, or optional SAPs
Increased Uni processor capacity
Up to 30 sub capacity CPs at capacity settings 4, 5, or 6
CPC Drawers and backplane Oscillator
SMT (for IFLs and zIIPs only) and SIMD
Enhanced processor/cache design with bigger cache sizes
Up to 10 TB of Redundant Array of Independent Memory (RAIM)
CPC Drawer/Memory Affinity
LPARs increased from 60 to 85

I/O Subsystem, Parallel Sysplex, STP, Security
New PCIe Gen3 I/O fanouts with 16 GBps Buses
LCSS increased from 4 to 6
4th Subchannel Set per LCSS
Maximum number of I/O Devices (subchannels) per channel increased from 24K to 32K for all z13 FICON features
FICON Enhancements
SR-IOV support for RoCE
New Integrated Coupling Adapter (ICA SR) for coupling links
Support for up to 256 coupling CHPIDs per CPC
CFCC Level 20
Crypto Express5S and Cryptographic enhancements with support for 85 Domains
STP Enhancements

RAS, Other Infrastructure Enhancements	
IBM zAware for Linux on z Systems (June 23, 2015)	System Control Hub (SCH). Replaces BPH
New N+2 ‘radiator’ design for Air Cooled System	Rack-Mounted Support Elements in the CPC
Key Locks for doors	Rack-mounted HMCs for customer supplied rack
Support for ASHRAE Class A2 datacenter	TKE 8.0 LICC

z13(s) Functions and Features (GA2 Driver Level 27)

System, Processor, Memory
Dynamic Partition Manager for Linux and KVM*
z Appliance Container Infrastructure (zACI)*
z/VSE Network Appliance using the z Appliance Container Infrastructure (zACI)**
IBM zAware in zACI Partition Mode*
LPAR Group Absolute Capping*
CPUMF sampler w/o PEMode enablement*
SNMP/BCPii performance enhancements*

RAS, Other Infrastructure Enhancements
STP Enhanced Console Assisted Recovery*
TKE 8.1 LICC and Rack-mounted TKE*

I/O Subsystem, Parallel Sysplex, STP, Security
New FICON functions* • Export/Import physical port WWPNs for FCP • Fiber channel read diagnostic parameters (RDP) Extended Link Services (ELS)
OSA Enhancements* • OSA ICC Secure Socket Layer (SSL) support • OSA ICC Concurrent MCL's for OSC CHPIDs
Shared Memory Communications–Direct (SMC-D)*
Enhanced Flash Express (R/W Cache 4 GB)*
CFCC Level 21* • Maintain Entry and Element Counts • CFCC Dump Reasons • CFCC set SLCP Event Type 11
Crypto Express5S* • Format Preserving Encryption (FPE) in FPGA Stage 2 • DK Phase 4 • CCA verb algorithm currency & interoperability • EMF simplification support • EP11 Stage 3, PKCS 11
• Regional Crypto Enablement (RCE) for Greater China Group (GCG) only**

Notes:

- *(*) New functional items available on z13 GA2 and z13s*
- *(**) Availability date update later than GA*

IBM z Systems naming for IBM z13 (z13s)

Brand Name:	**IBM**
Product Class:	**IBM mainframe**
Family Name:	**IBM z Systems™**
Family Short Name:	**z Systems**
Product Line Name:	**IBM z Systems™**
Product Line Short Name:	**z Systems**
Product Name:	**IBM z13™**
Short Name:	**Z13** **z13s**
Models:	**N30, N63, N96, NC9, NE1** **N10, N20**
Machine Type:	**2964** **2965**
Workload Optimizing Attachments:	**IBM z BladeCenter® Extension (zBX) Model 004** **IBM DB2® Analytics Accelerator for z/OS® Version 5**
Management Firmware:	**IBM z Unified Resource Manager**
Management Firmware Short Name:	**Unified Resource Manager or zManager**

IBM z Systems Generations

N-4

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

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

N-3

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

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

N-2

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

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

N-1

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

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

N

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

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

z13 Continues the CMOS Mainframe Heritage Begun in 1994

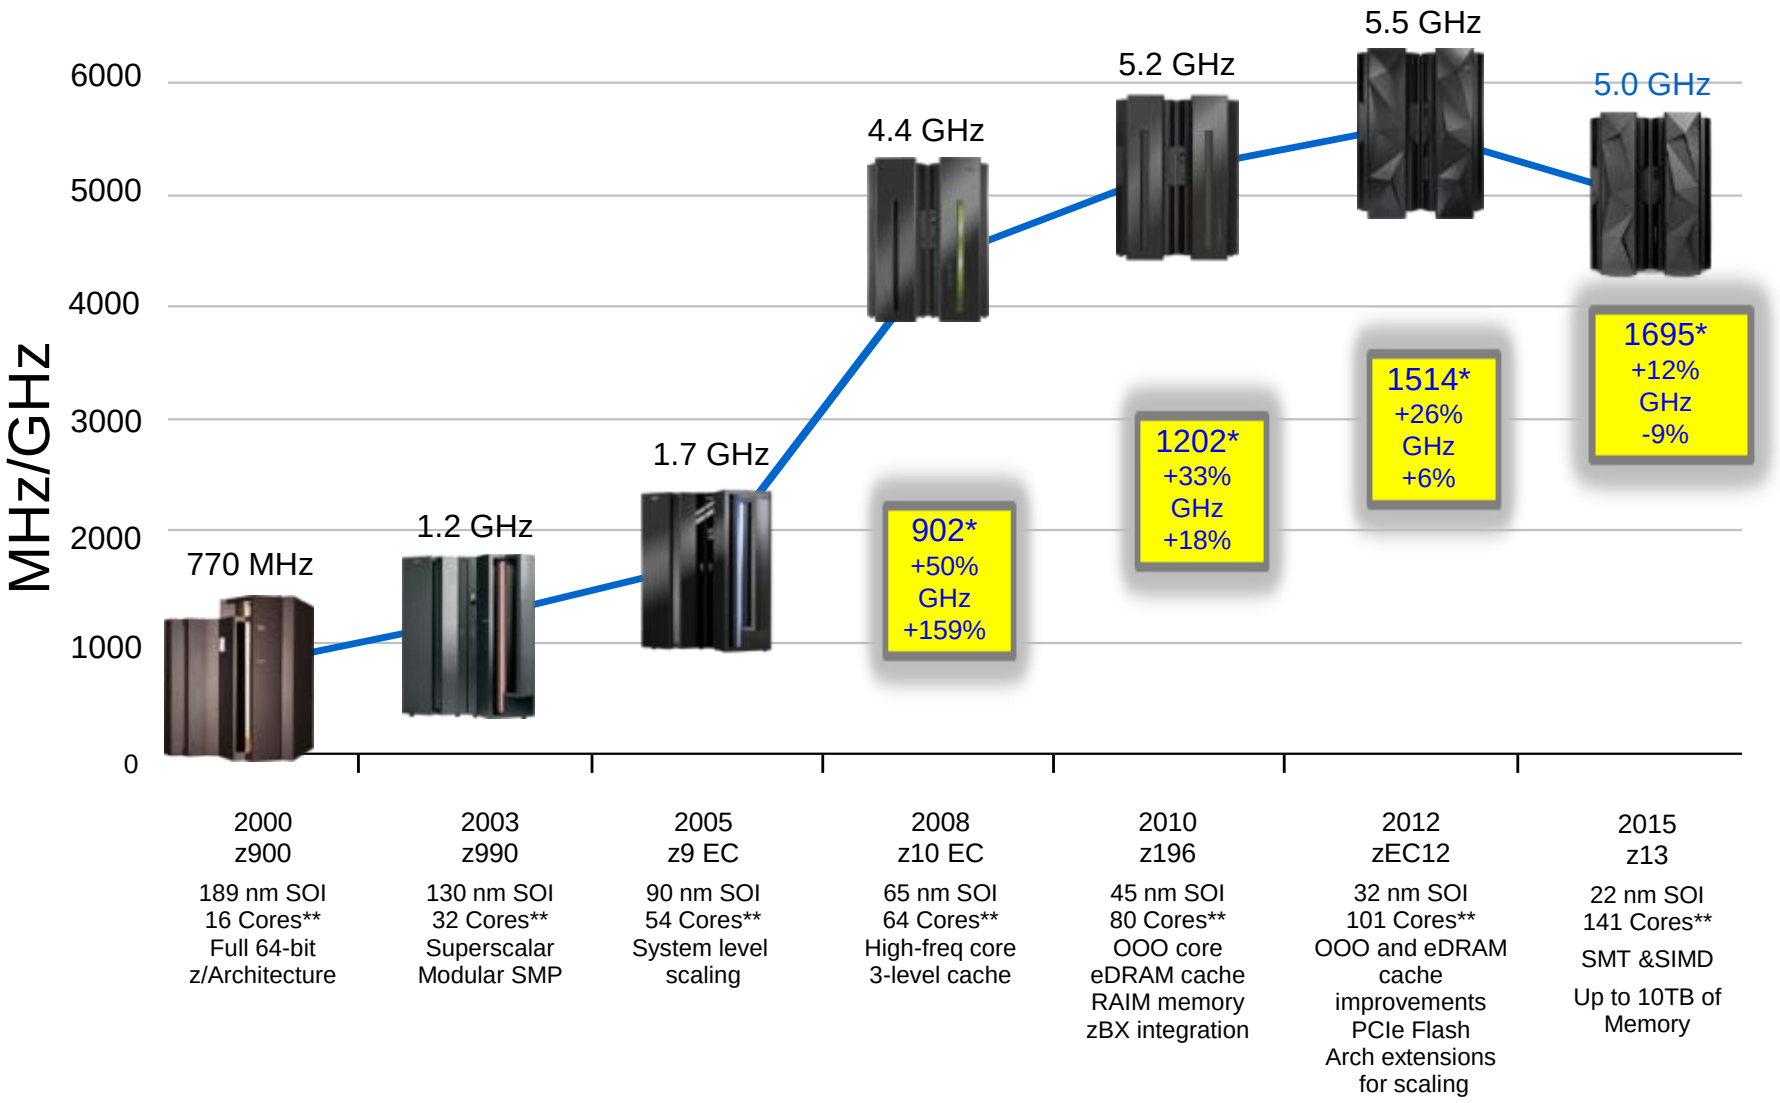

* MIPS Tables are NOT adequate for making comparisons of z Systems processors. Additional capacity planning required

** Number of PU cores for customer use

IBM z13: Advanced system design optimized for digital business

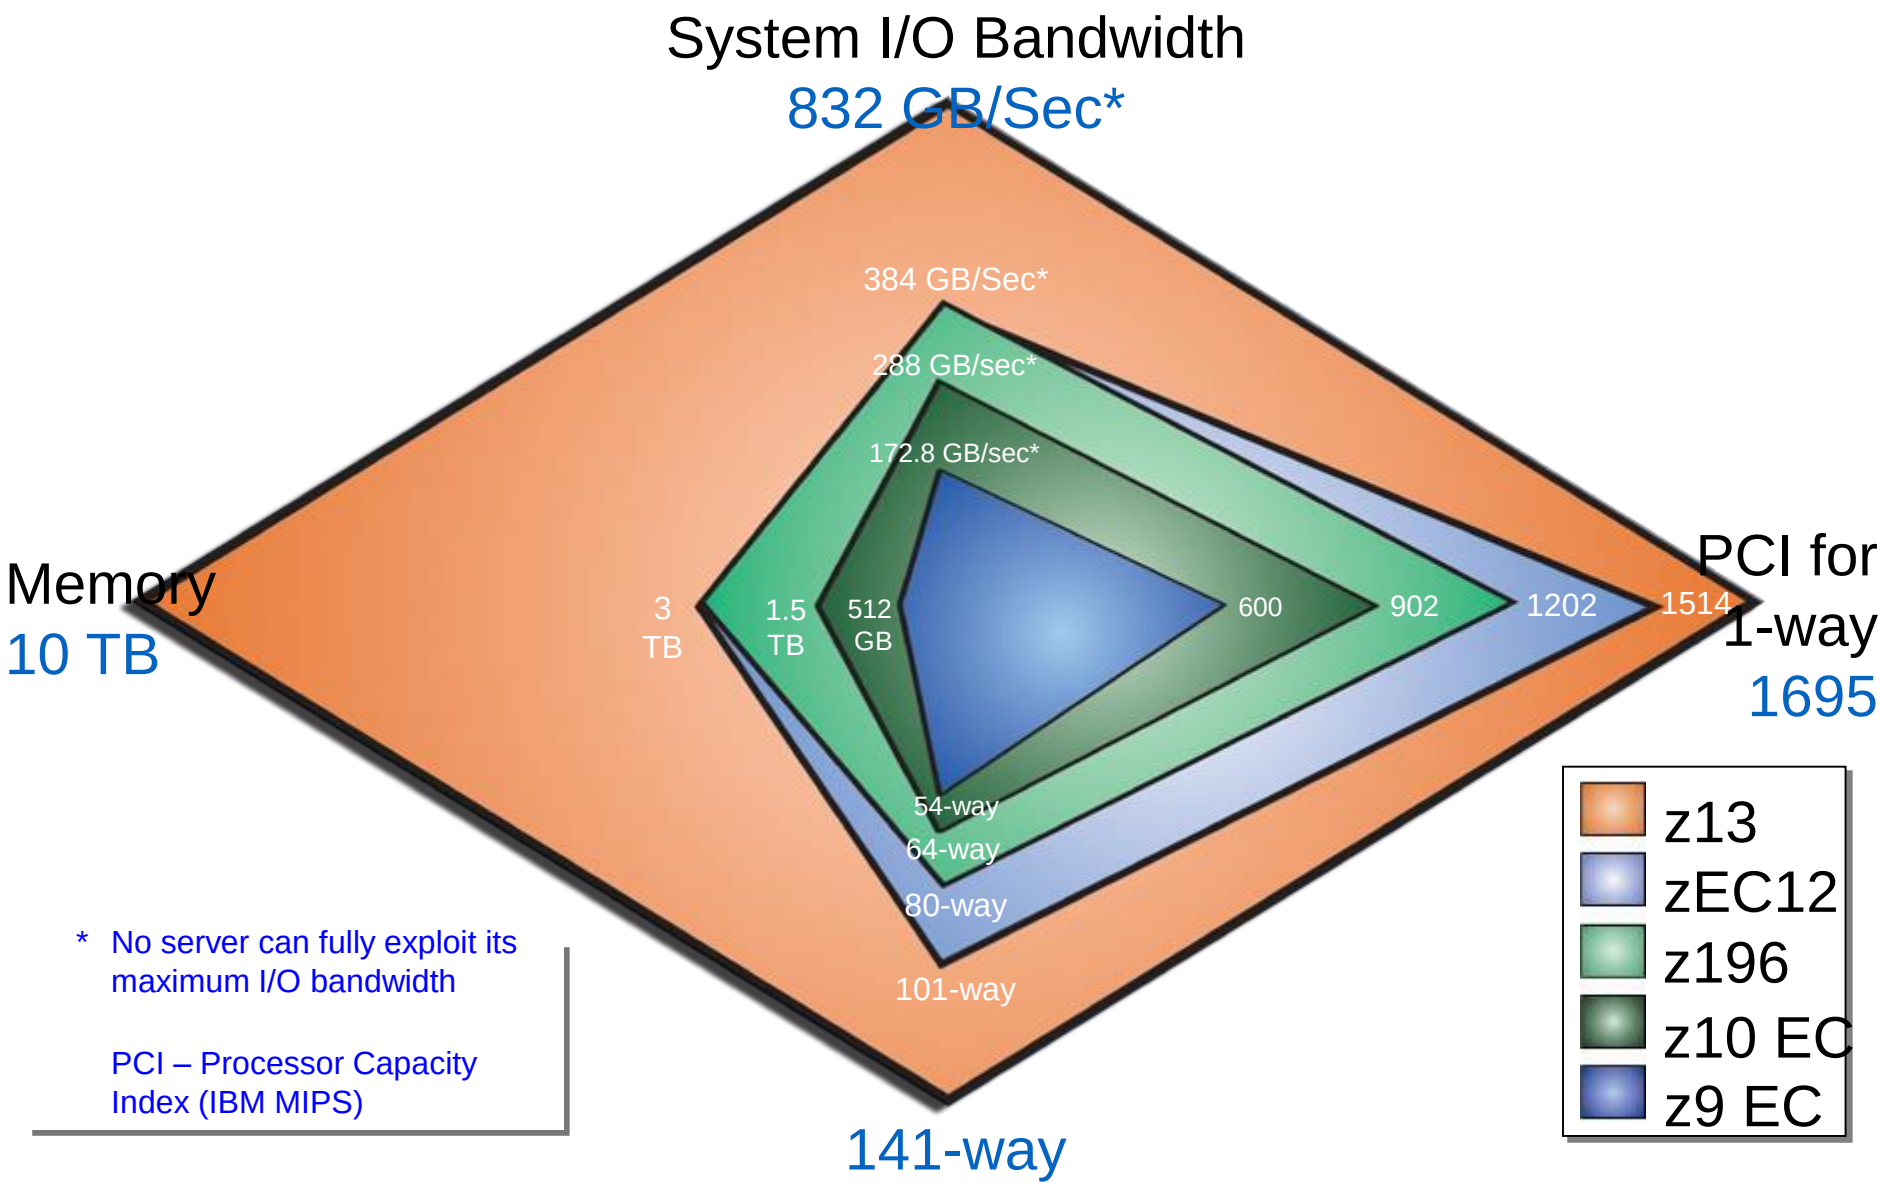

* No server can fully exploit its maximum I/O bandwidth

PCI – Processor Capacity Index (IBM MIPS)

z13 Processor Unit Allocation/Usage

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

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

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

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

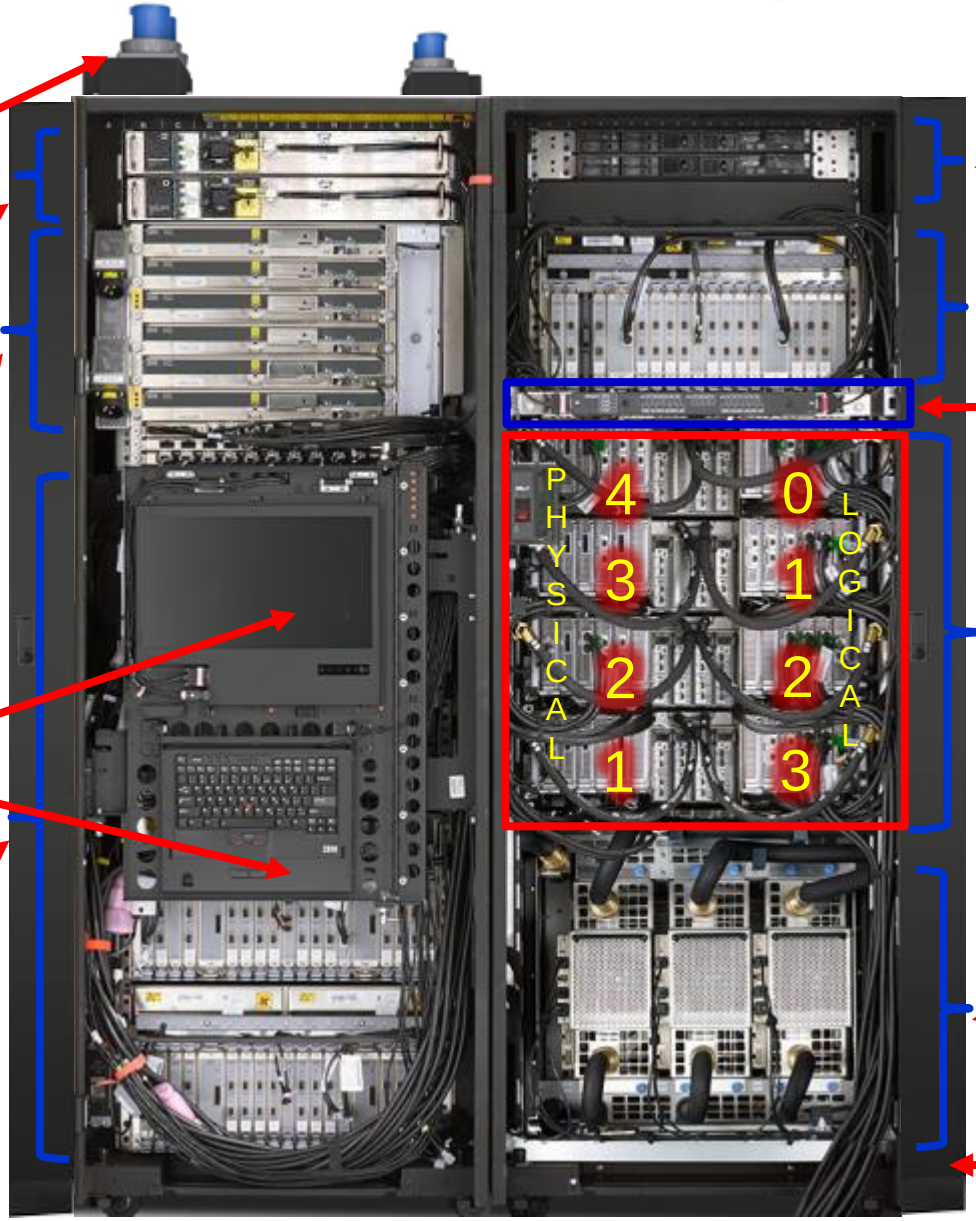

Overhead Power Cables (option)

Internal Batteries (optional)

Power Supplies

Displays and keyboards for Support Elements

PCIe I/O drawers numbers 1 to 4 (Note: for an upgraded System, drawer slots 1 and 2 are used for the I/O Drawer

2 x 1U Support Elements

PCIe I/O drawer number 5

System Control Hubs (used to be BPHs)

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

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

N+2 Radiator Pumps

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

z13 Water cooled – Rear View (Model NC9 or NE1)

z13s continues the CMOS Mainframe Heritage

**Full capacity uniprocessor*

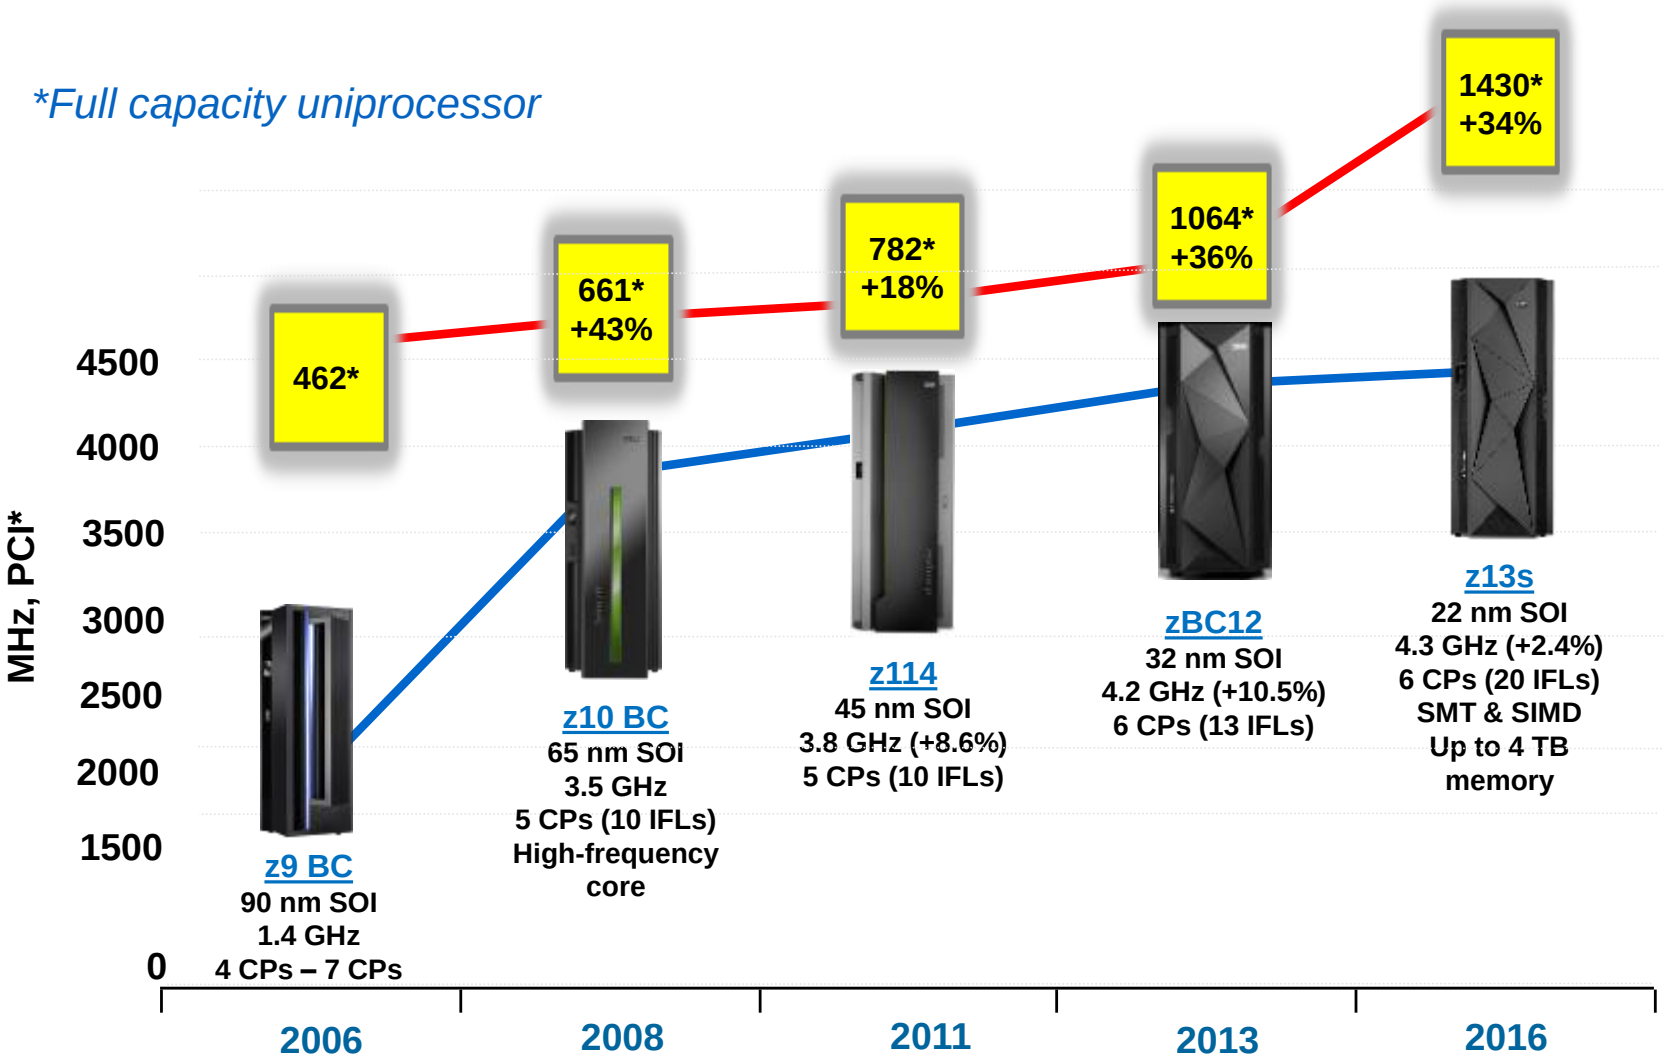

*NOTE: MIPS Tables are NOT adequate for making comparisons of z Systems processors in proposals

IBM z13s Advanced System Design Optimized for Digital Business

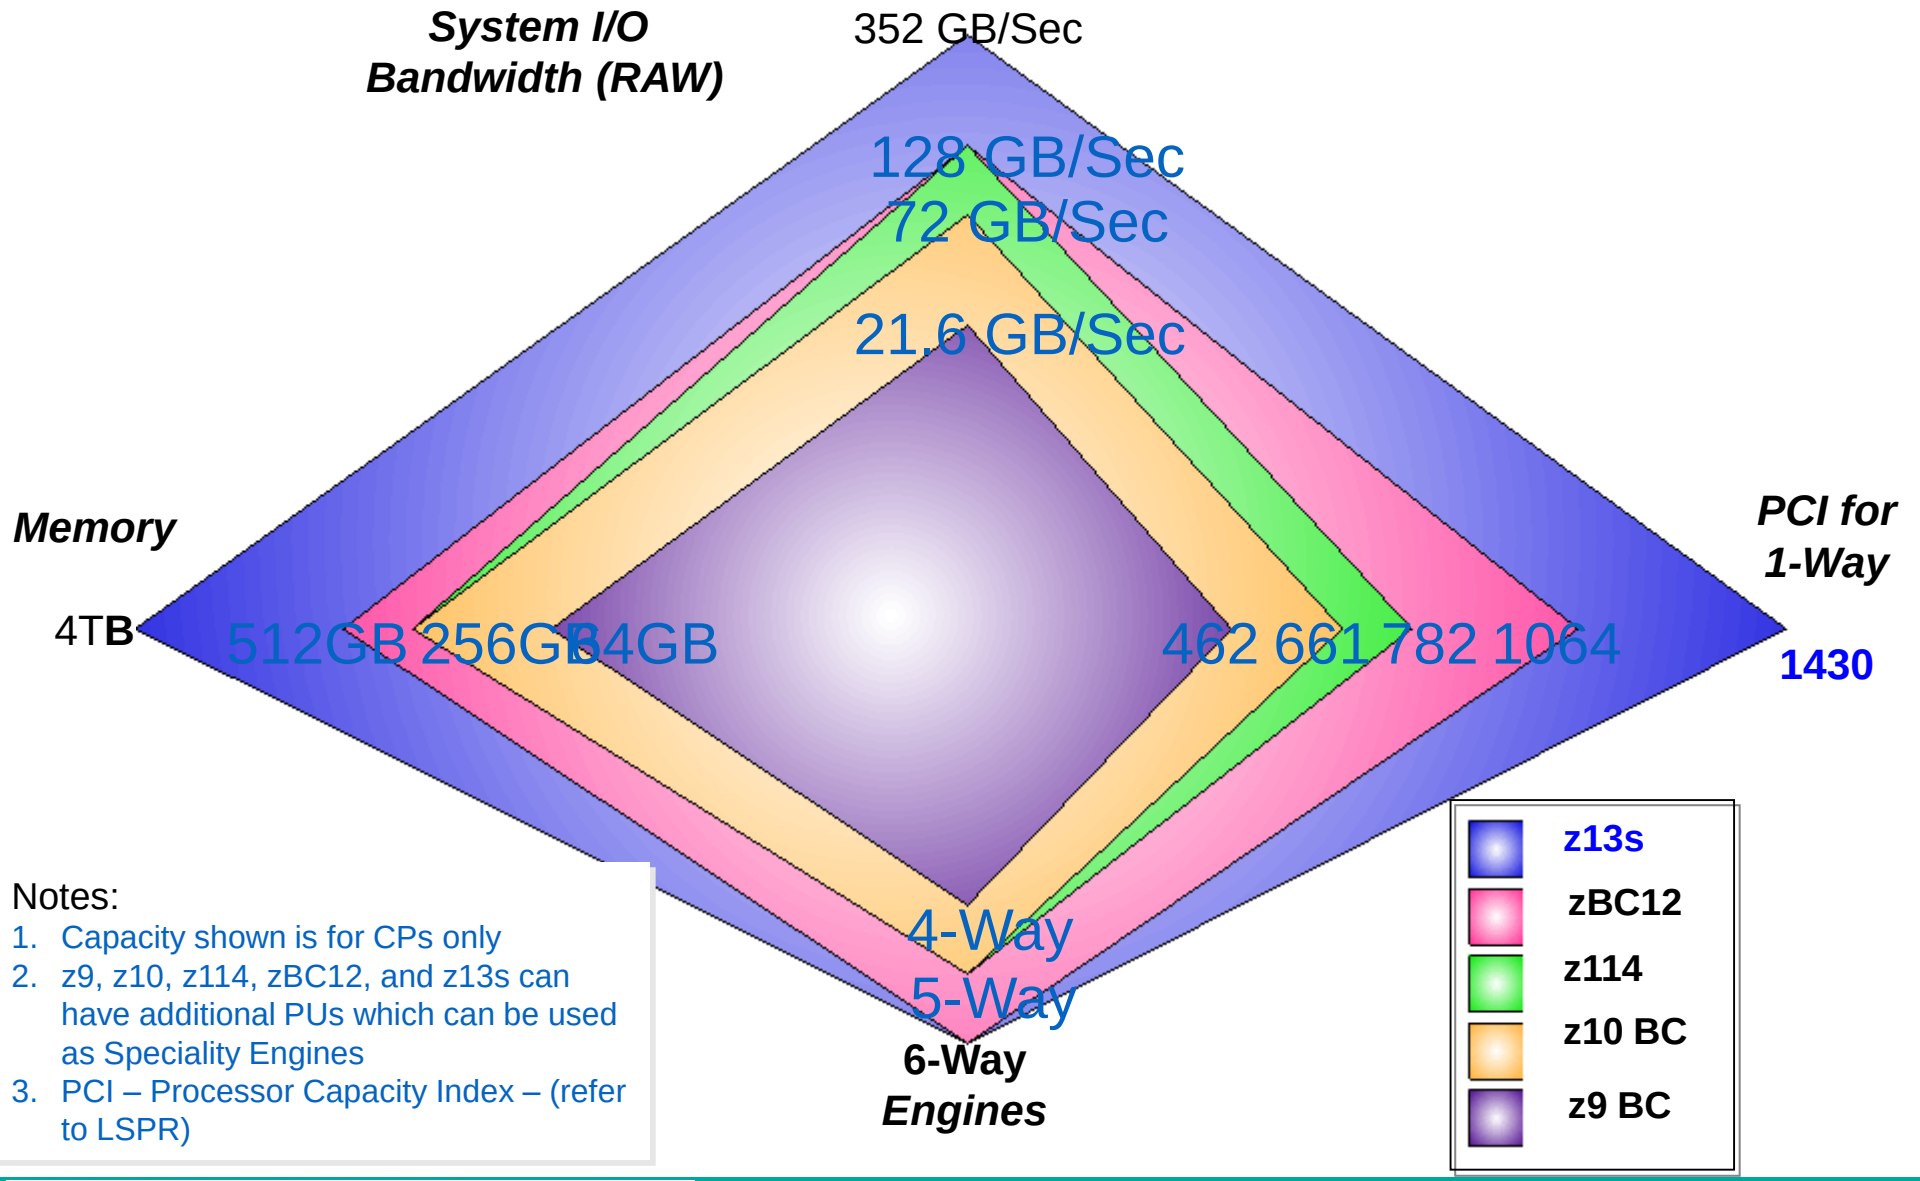

Notes:

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

z13s Processor Unit Allocation and Usage

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

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

z13s Sub-capacity Processor Granularity

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

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

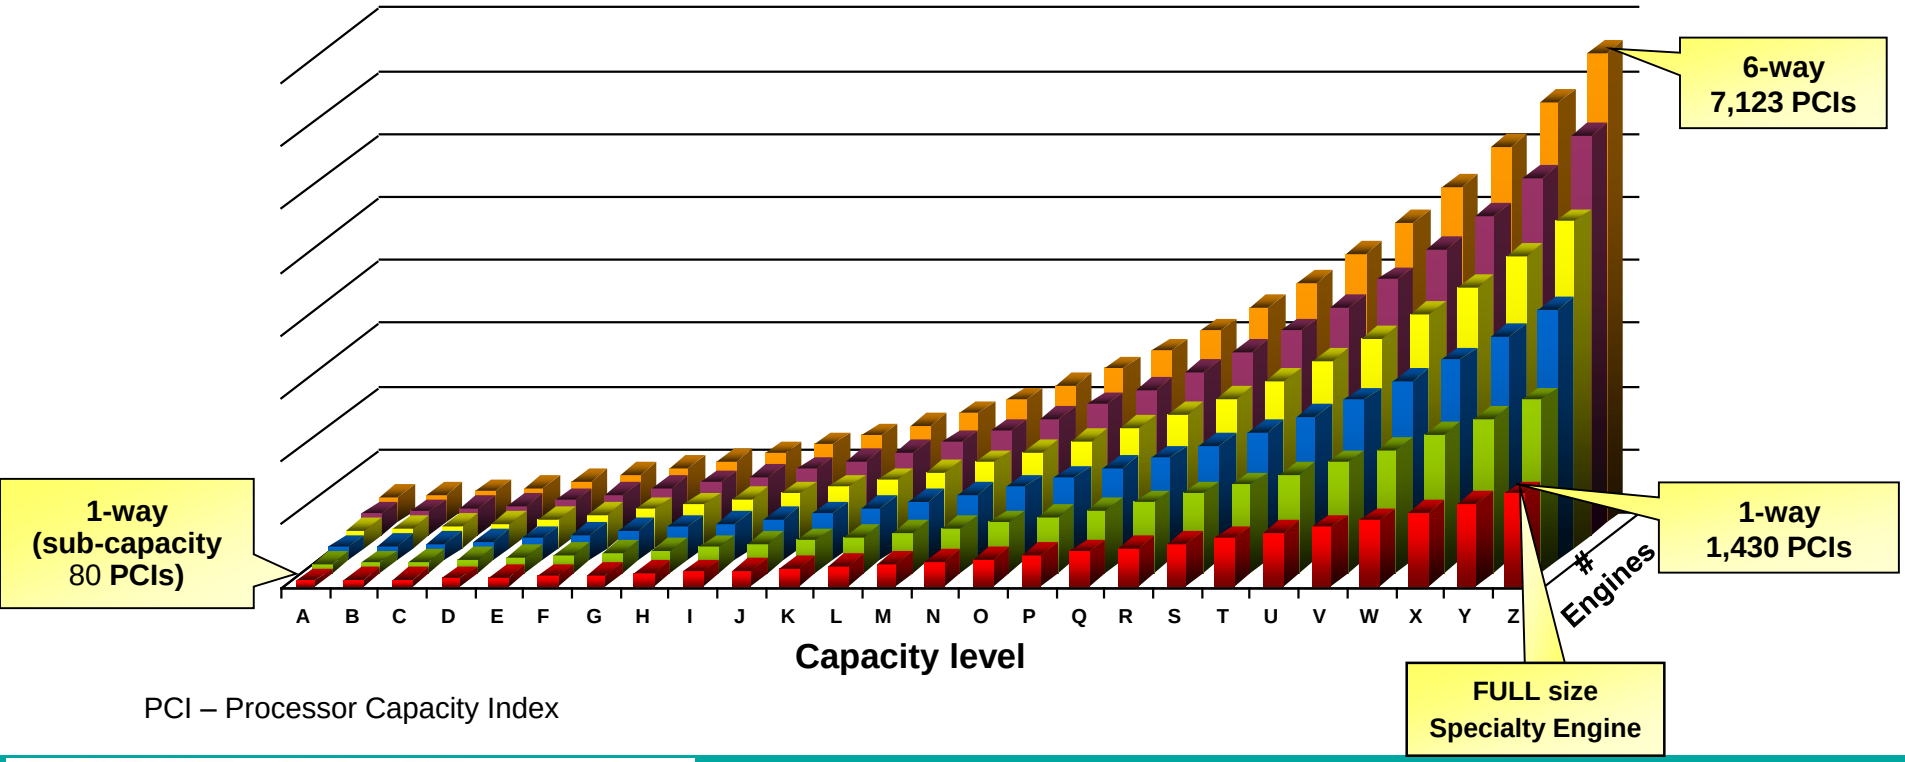

PCI – Processor Capacity Index

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

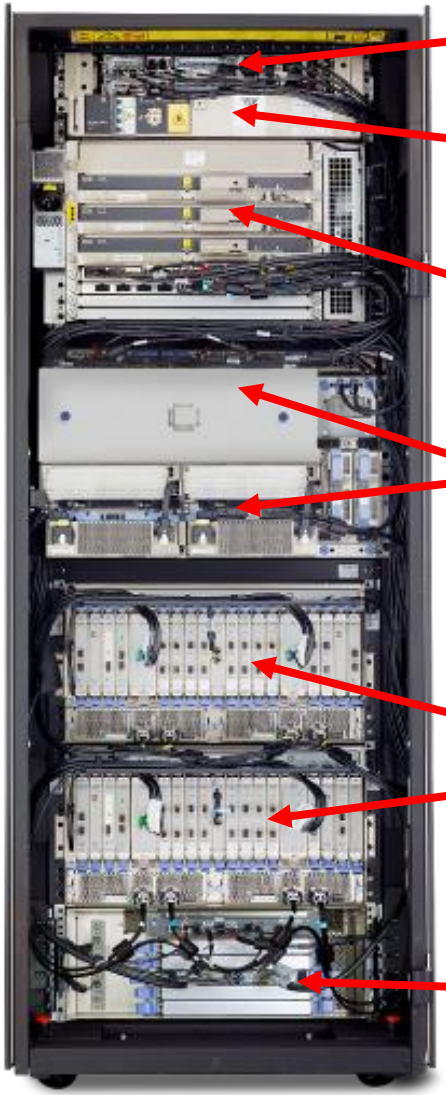

Rear View

Support Elements

Internal Batteries (optional)

Power Supplies

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

PCIe I/O Drawers

I/O Drawer

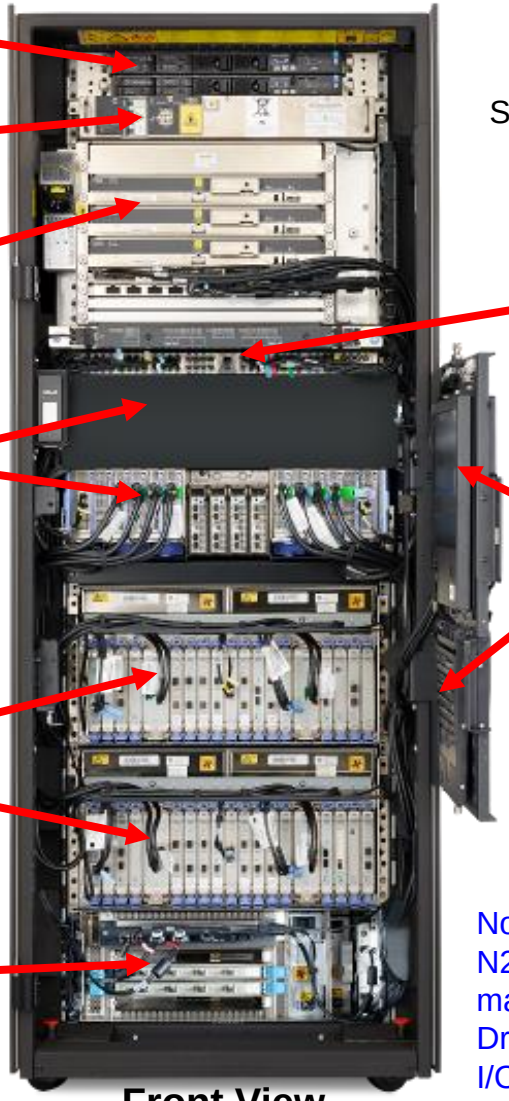

Front View

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

SCH (x2)

Monitors + keyboards for SEs

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

LinuxOne

Sharpen your competitive edge

2016 IMS Technical Symposium

March 7 – 10, 2016

Wiesbaden, Germany

www.ims-symposium.com

IBM LinuxONE Portfolio

IBM LinuxONE *Systems*
The most trusted, efficient and high performance enterprise-grade Linux platform

IBM LinuxONE Emperor™

with 6-141 processors
96GB+ memory

IBM LinuxONE Rockhopper™

with 2-20 processors
40GB+ memory

IBM LinuxONE *Solutions*
Designed for the Digital Economy

Mobile

Analytics

Cloud

DevOps

New IBM LinuxONE Elastic *Pricing*

IBM LinuxONE *Services*

What is Linux on IBM z Systems

Linux is Linux

- Pure Linux®, no emulation
- Not a unique version of Linux
- No changes in Look & Feel

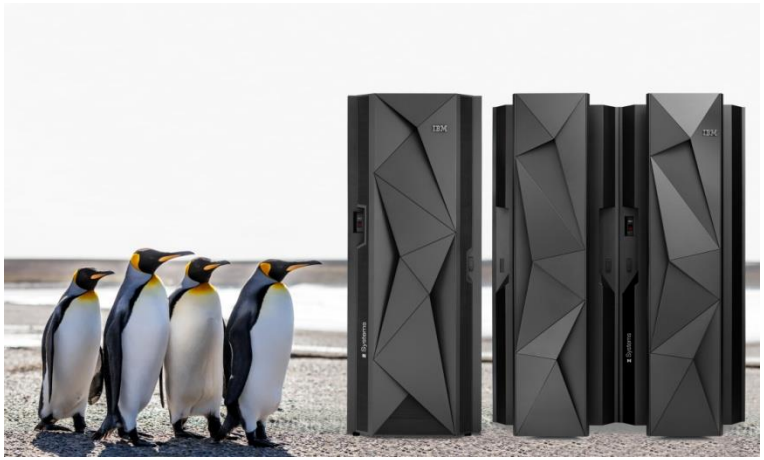

Supported Linux distributions

- redhat
- SUSE
- ubuntu Supported by Canonical *(planned)*

See "Tested Platforms"

Supported Virtualization

- IBM z/VM® + IBM Wave for z/VM
- KVM for IBM z Systems™
- Logical Partitions (LPAR)

See "z Systems Virtual Servers"

2000 — In the market since 2000, well accepted and growing — 2016

IBM z/VM and KVM for IBM z

z/VM

- World class quality, security, reliability - powerful and versatile
- Extreme scalability creates cost savings opportunities
- Exploitation of advanced technologies, such as:
  - Shared memory (Linux kernel, executables, communications)
- Highly granular control over resource pool
- Valuable tool for resiliency and Disaster Recovery
- Provides virtualization for all z Systems operating systems

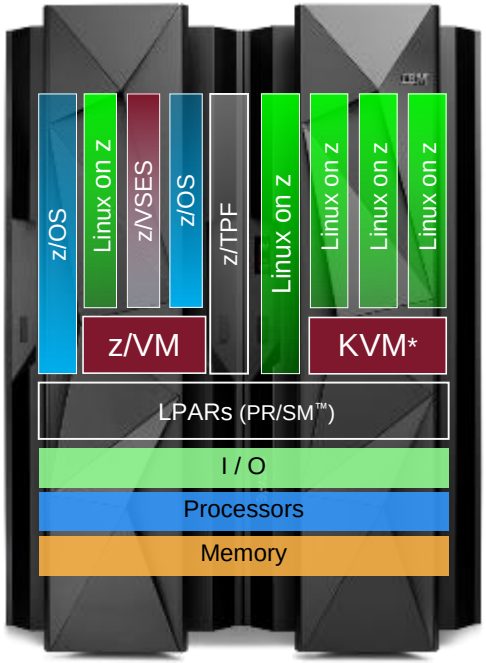

KVM

- Simplifies configuration and operation of server virtualization
- Leverage common Linux administration skills to administer virtualization
- Flexibility and agility leveraging the Open Source community
- Provides an Open Source virtualization choice
- Easily integration into Cloud/OpenStack environments

Robust solutions from IBM, ISVs & Open Source Community

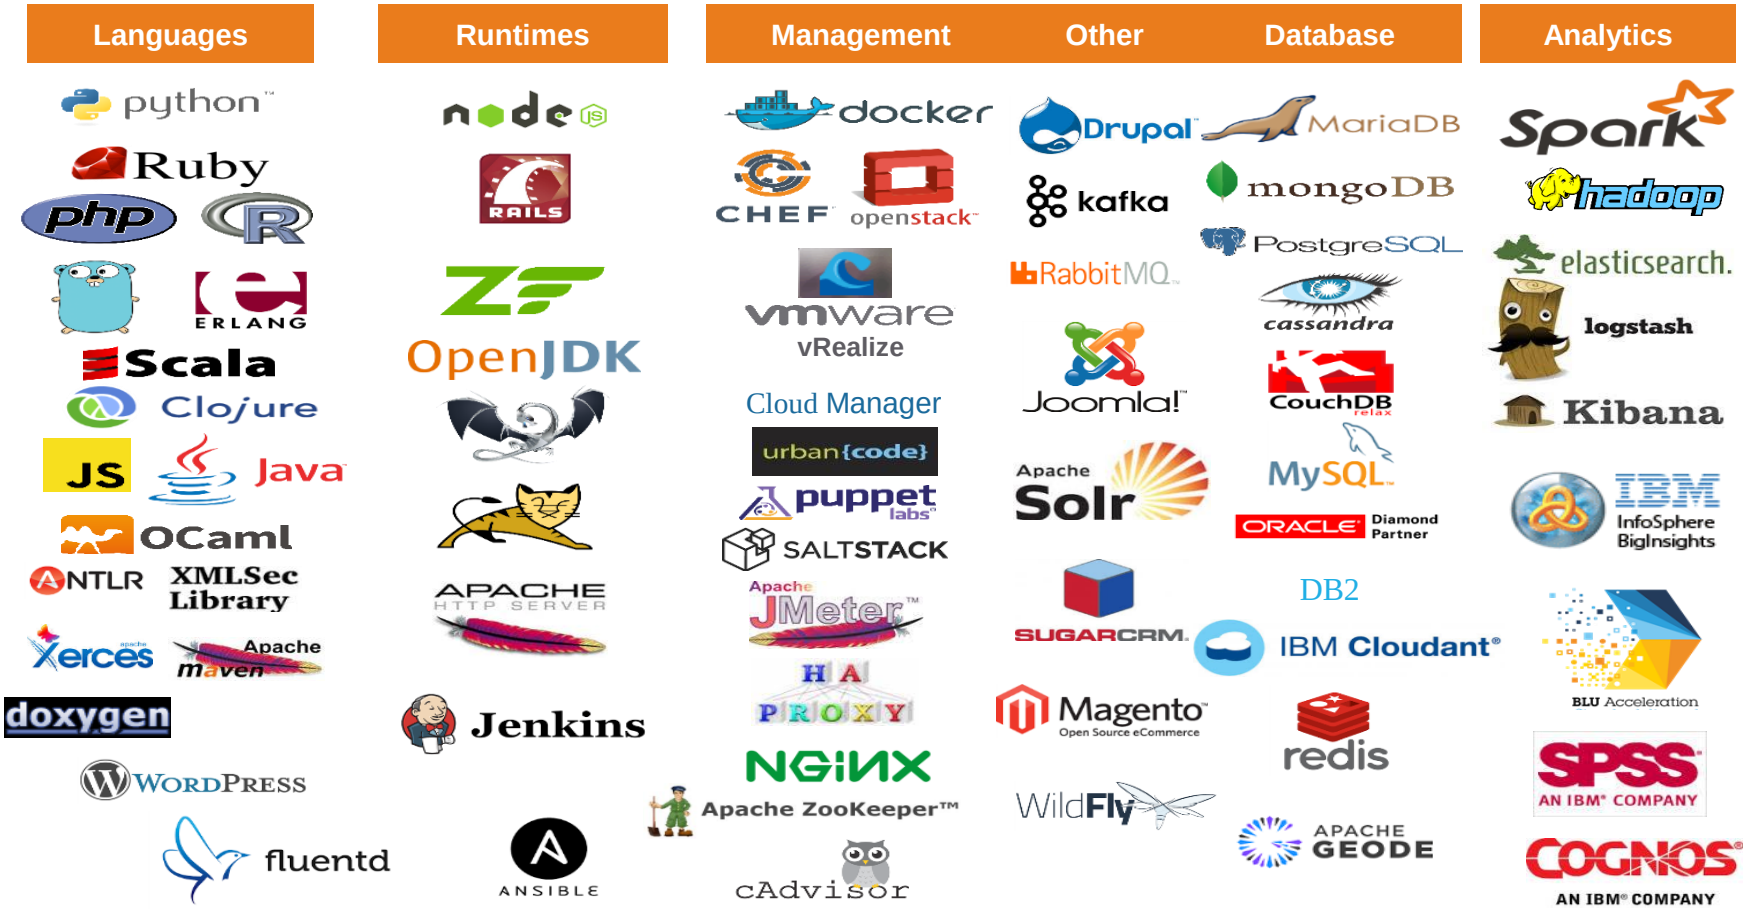

z13 – Redesigned for efficient and trusted cloud services

Enterprise-grade Linux cloud services at **half** the cost, **half** the energy, and **half** the floor space of alternatives

Up to 10 TB Memory on z13
Improves consolidation ratios

GDPS for Linux on z Systems
Disaster Recovery solution for mission-critical workloads

SMT-2 technology on z13
Improves performance and throughput of workloads

Increase in # of LPARs on z13
Improves TCO

KVM
New industry-standard hypervisor (SOD)

Cloud Manager w/ OpenStack V4.2
Heterogeneous platform management from z Systems

Elastic Storage for Linux on z Systems
Enables new class of workloads

8,000 virtual servers in a single system
Reduce cost and administration overhead

Crypto Express 5S
Security & performance

"Smaller enterprises often choose public cloud services, but encounter issues with cost and complexity when they expand. With … the Enterprise Cloud System – which can accommodate more than 6,000 VMs – we can offer clients the cost effective scalability they need to take their business to the next level." - Steve Groom, CEO of Vissensa

Mainframe vs Distributed Terminology

Mainframe

- System programmer
- POR / IML
- IPL
- 4-way
- Dispatcher
- Main storage
- **DASD** ← 'external' disk storage
- *OSA*
- *CP / **IFL*** ← *Specialty Engines / processors*

Distributed

- System administrator
- Coldstart / Boot
- Warmstart / (Re-)Boot
- 4-processor machine
- Scheduler
- Main memory
- Disk
- *NIC*
- *CPU*

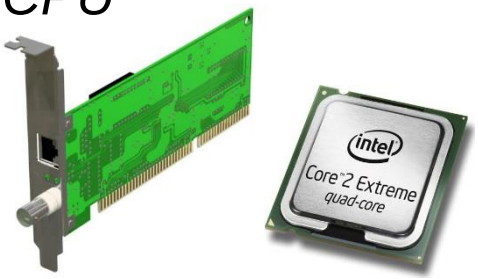

IBM z Systems – Reliable, Scalable, Secure and Virtualized

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

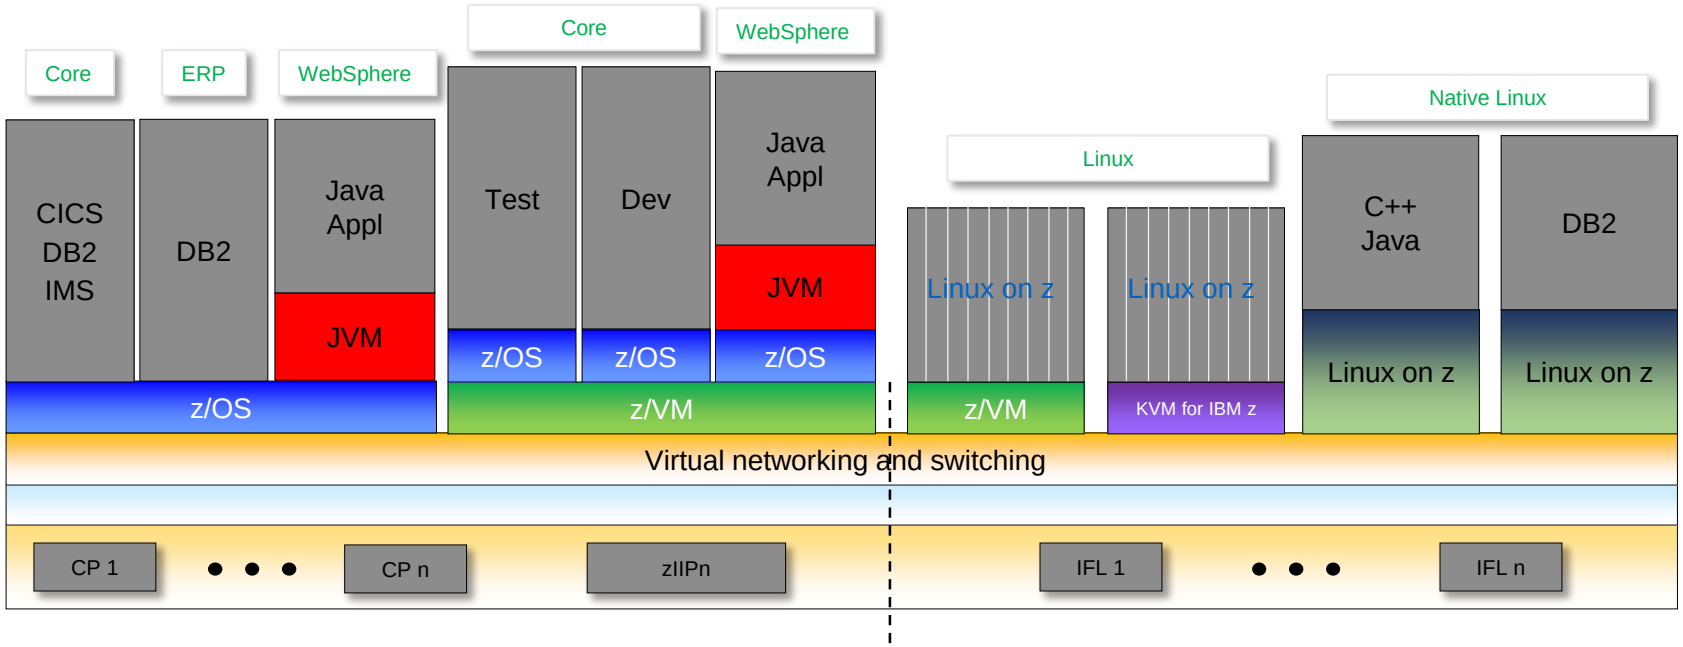

- Large scale, robust consolidation platform
- Built-in Virtualization
- 100s to 1000s of virtual servers on z/VM
- Intelligent and autonomic management of diverse workloads and system resources

Some technical details

Performance increase forever?

- **Moore's Law is a computing term which originated around 1970; the simplified version of this law states that processor speeds, or overall processing power for computers will double every two years. A quick check among technicians in different computer companies shows that the term is not very popular but the rule is still accepted.**
- **Future challenges:**
  - Density
  - Heat
- **Gordon Moore stated in … law cannot be sustained indefinitely: "It … The nature of exponentials is that you push th… disaster happens." He also noted that transi… tually reach the limits of miniaturization at atomic levels.**

Moore's Law is no longer valid in terms of processor speed

z System Servers Continue to Scale with z13

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

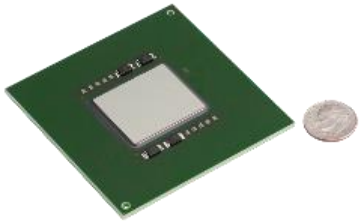

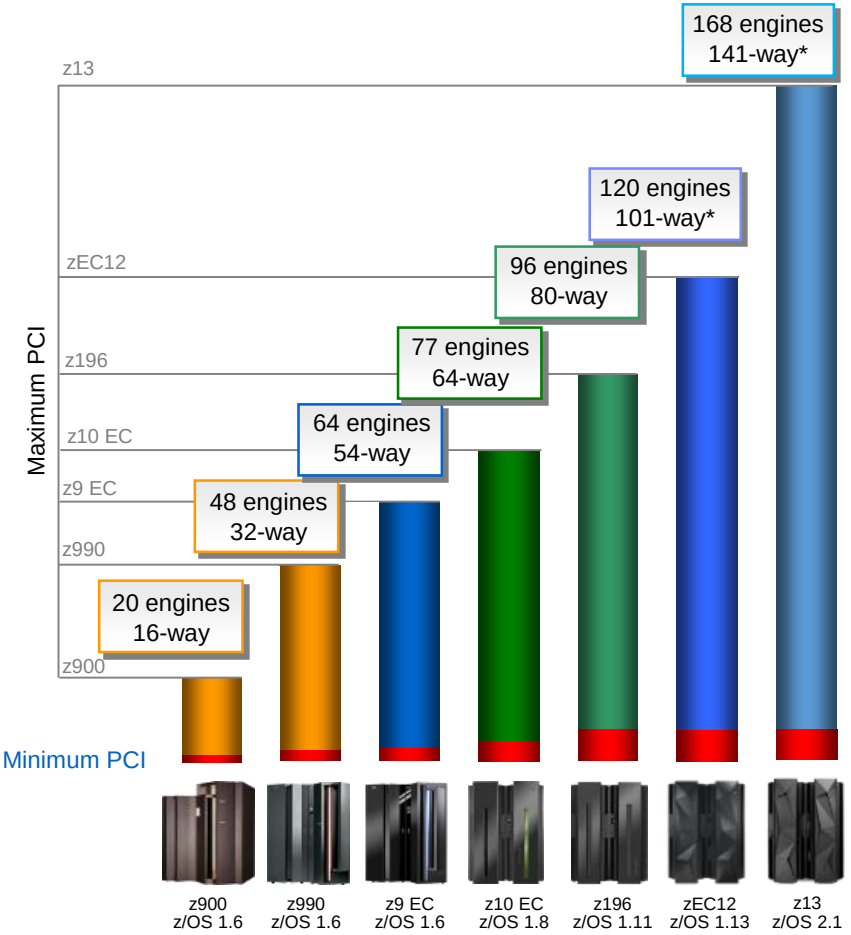

PCI - Processor Capacity Index
*z/OS supports up to a 100-way only

z Systems - Processor Roadmap

z10 2/2008

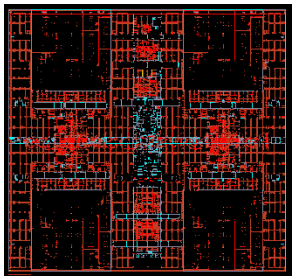

Workload Consolidation and Integration Engine for CPU Intensive Workloads

Decimal FP

Infiniband

64-CP Image

Large Pages

Shared Memory

z196 9/2010

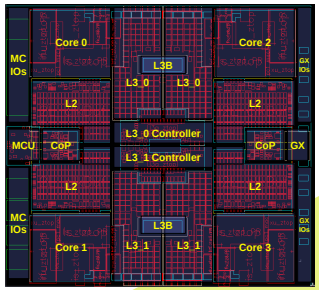

Top Tier Single Thread Performance,System Capacity

Accelerator Integration

Out of Order Execution

Water Cooling

PCIe I/O Fabric

RAIM

Enhanced Energy Management

zEC12 8/2012

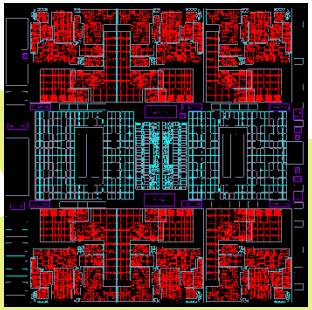

Leadership Single Thread, Enhanced Throughput

Improved out-of-order

Transactional Memory

Dynamic Optimization

2 GB page support

Step Function in System Capacity

z13 1/2015

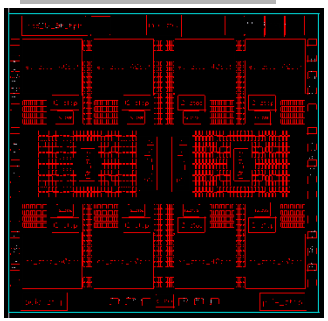

Leadership System Capacity and Performance

Modularity & Scalability

Dynamic SMT

Supports two instruction threads

SIMD

Business Analytics Optimized

Accelerate Key Workloads with Special-Purpose Hardware

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

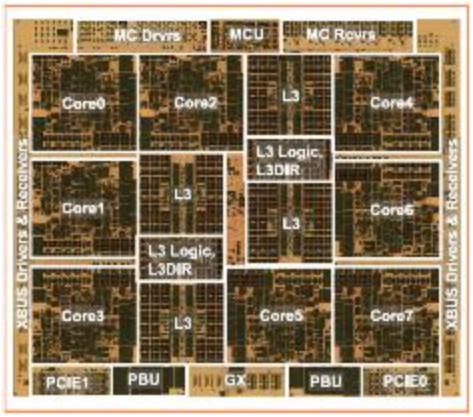

IBM DB2 Analytics Accelerator built on Netezza Technology

Cache Latency

- **Why needs the CPU access Memory? Instruction / Data**

- **Cache latency for z196 (1, 4, 12, 32 & 77 are relative access times)**
  Ratios are still accurate

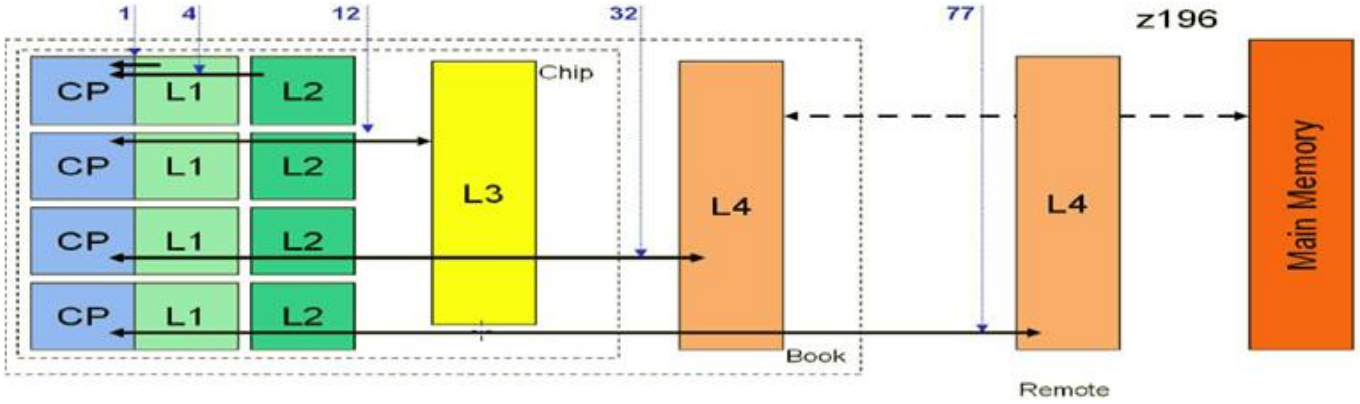

z13 CPC Drawer Cache Hierarchy Detail

Single CPC Drawer View (N30 Model) – 2 Nodes

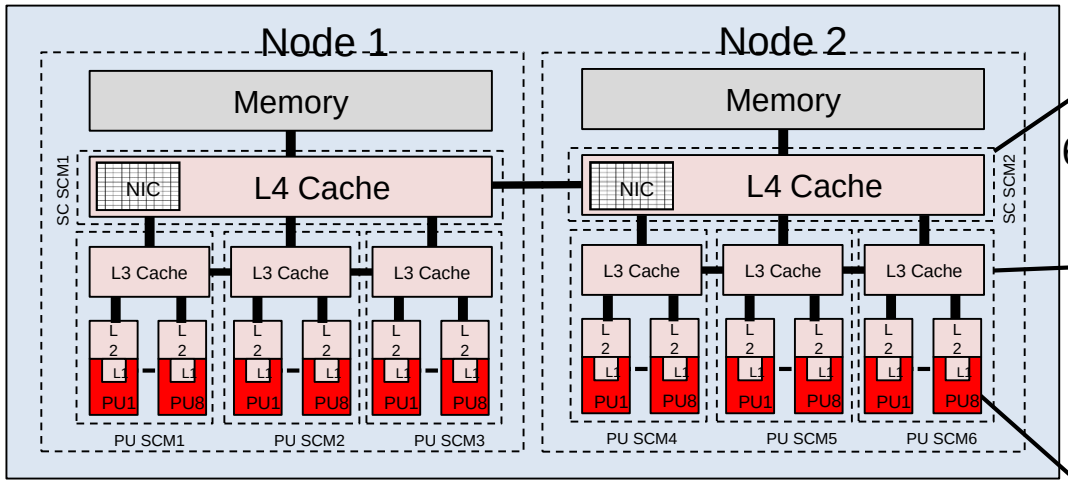

2 SC SCM Chips
(1 per node)

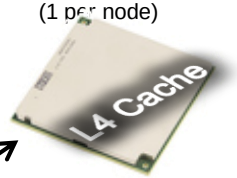

6 PU SCM Chips
(3 per node)

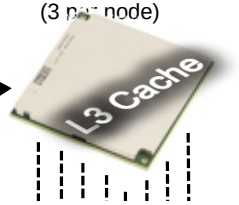

* Up to 8 PU cores per chip

Single PU core

* Not all PU's active

Node 1 - Caches

- L1 private 96k i, 128k d
- L2 private 2 MB i + 2 MB d
- L3 shared 64 MB / chip
- L4 shared 480 MB / node
  - plus 224 MB NIC

Node 2 - Caches

- L1 private 96k i, 128k d
- L2 private 2 MB i + 2 MB d
- L3 shared 64 MB / chip
- L4 shared 480 MB / node
  - plus 224 MB NIC

z13: Simultaneous Multi-Threading

- **Today**
  - –Each CPU support a single instruction stream
    - System z workload tend to experience non-trivial number of cache misses
    - CPU generally unproductive while resolving cache misses

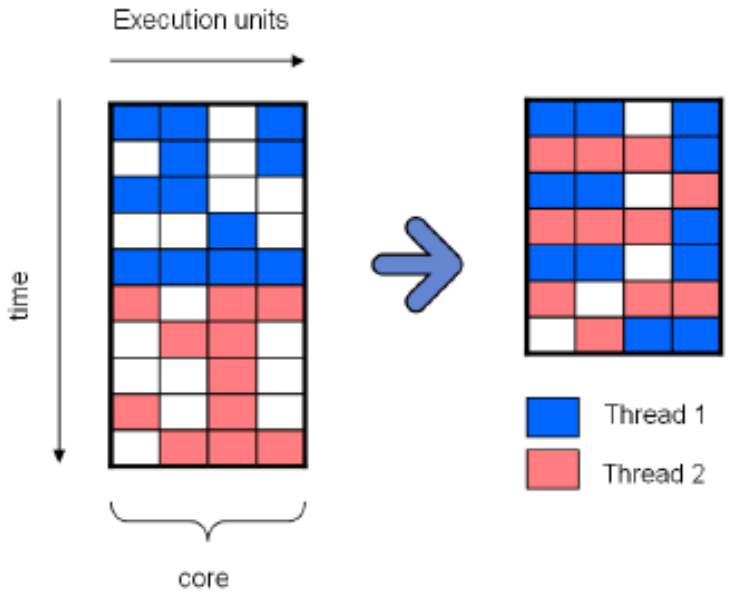

Simultaneous Multithreading (SMT) on z13

- Simultaneous multithreading allows instructions from one or two threads to execute on a zIIP or IFL processor core.
- SMT helps to address memory latency, resulting in an overall **capacity*** (throughput) improvement per core
- Capacity improvement is variable depending on workload. We see in the field about 20-40% capacity increase

- 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

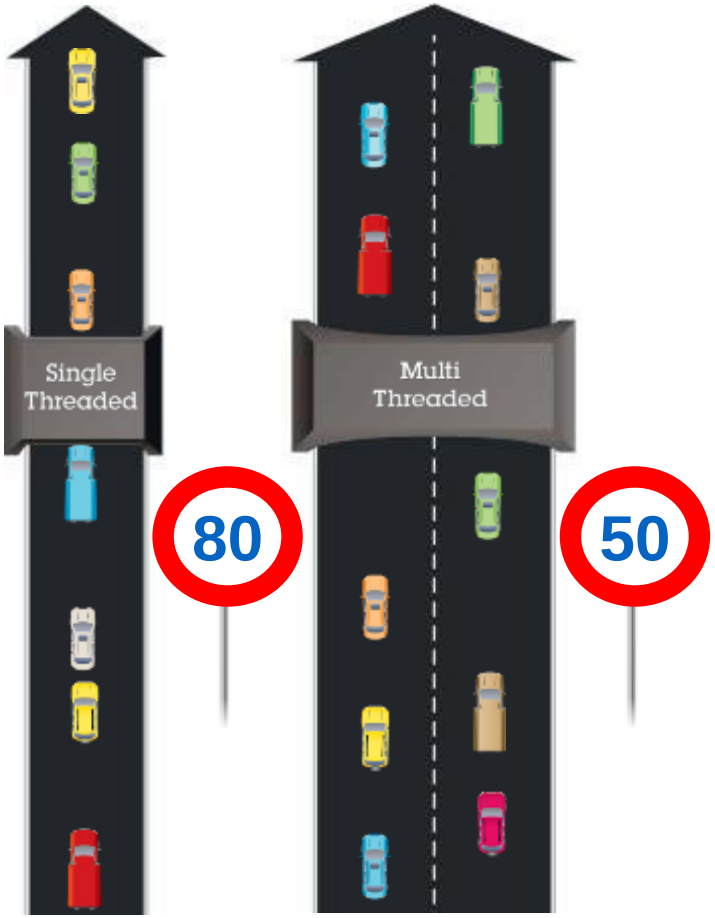

*Which approach is designed for the highest volume** of traffic? Which road is faster?*

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

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

Why Single Instruction Multiple Data (SIMD) on z Systems

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

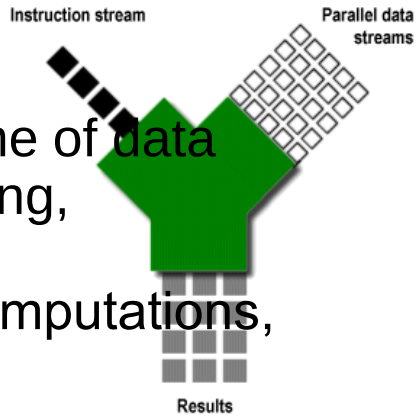

- **Customer *perception* of Analytics and z Systems**
  - z Systems handle OLTP and Batch jobs types of workload
  - Mathematical and data intensive operations can lead to unaffordable MIPS usage

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

SIMD – Single Instruction Multiple Data

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

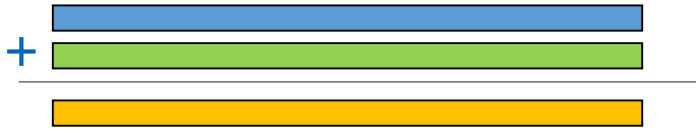

- **New: Single instruction operates on multiple data in parallel**

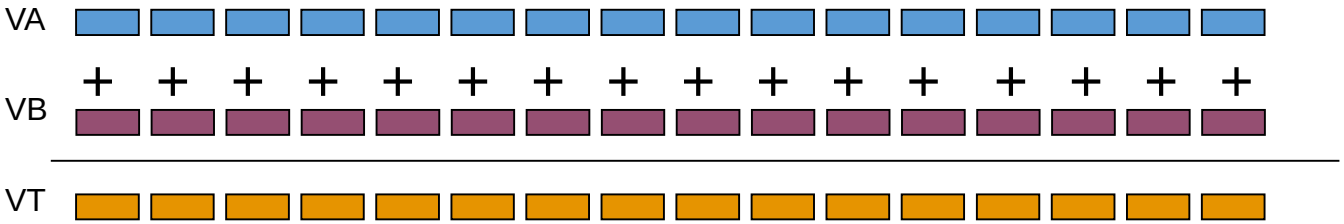

- **Each register contains multiple data elements of a fixed size**
  - Byte, Halfword, Word, Doubleword, Quadword
  - The collection of elements in a register is also called a **vector**
  - Field in the instruction word specifies data format type

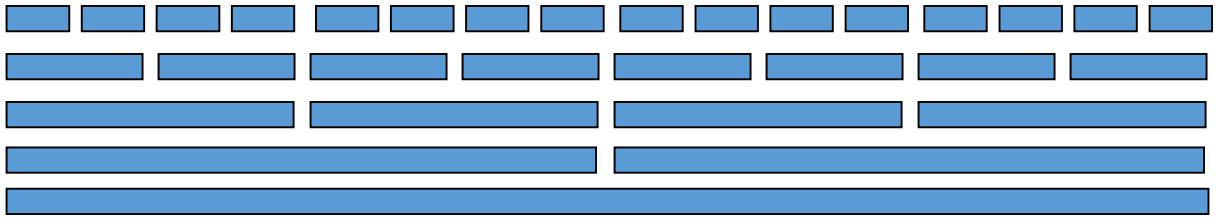

128b wide vector:
16xB, 8xHW, 4xW,
2xDW, 1xQW

SIMD Hardware Accelerator

New SIMD FXU

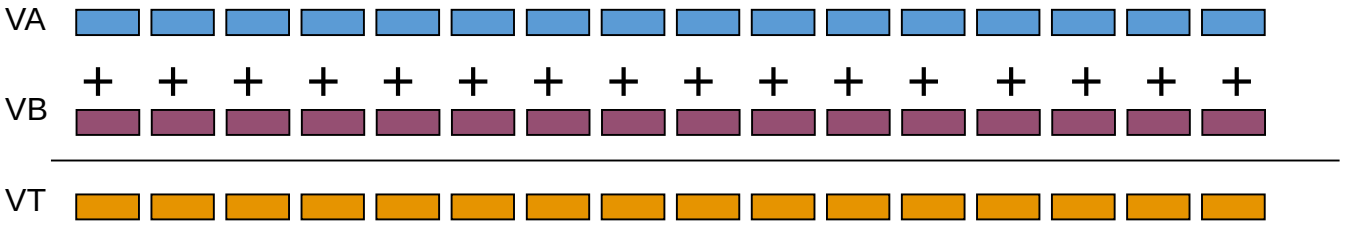

Three distinct data types

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

Implementation

New string engine

Enhancing BFU for SIMD

Exploitation

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

Analytics workloads like ILOG, SPSS, ALGO

z13 String: Vector Find Any Element Equal

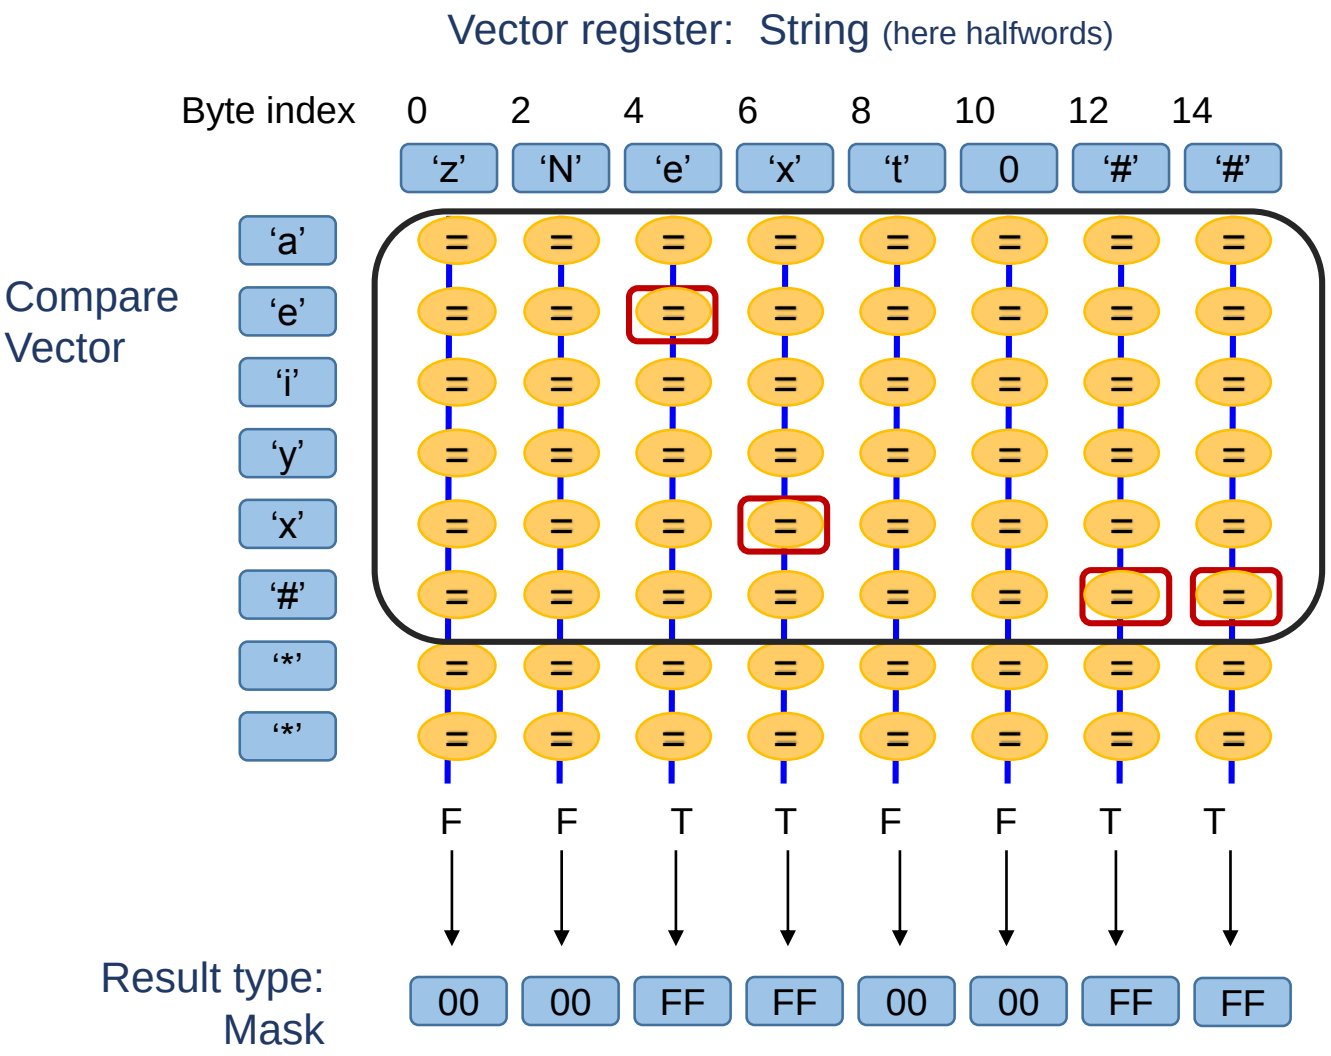

OR-ing all range checks per element

zEC12:
Depending on coding
Up to 8x6 HW compares
→ 48 instructions

z13:
Single instruction
High parallelism

z13: SIMD String Support

- **Big comparator array**
  - Supporting strings with 16x8b, 8x16b, 4x32b
  - Comparators dynamically re-arranged to match required width

- **Very high parallelism for small data types**

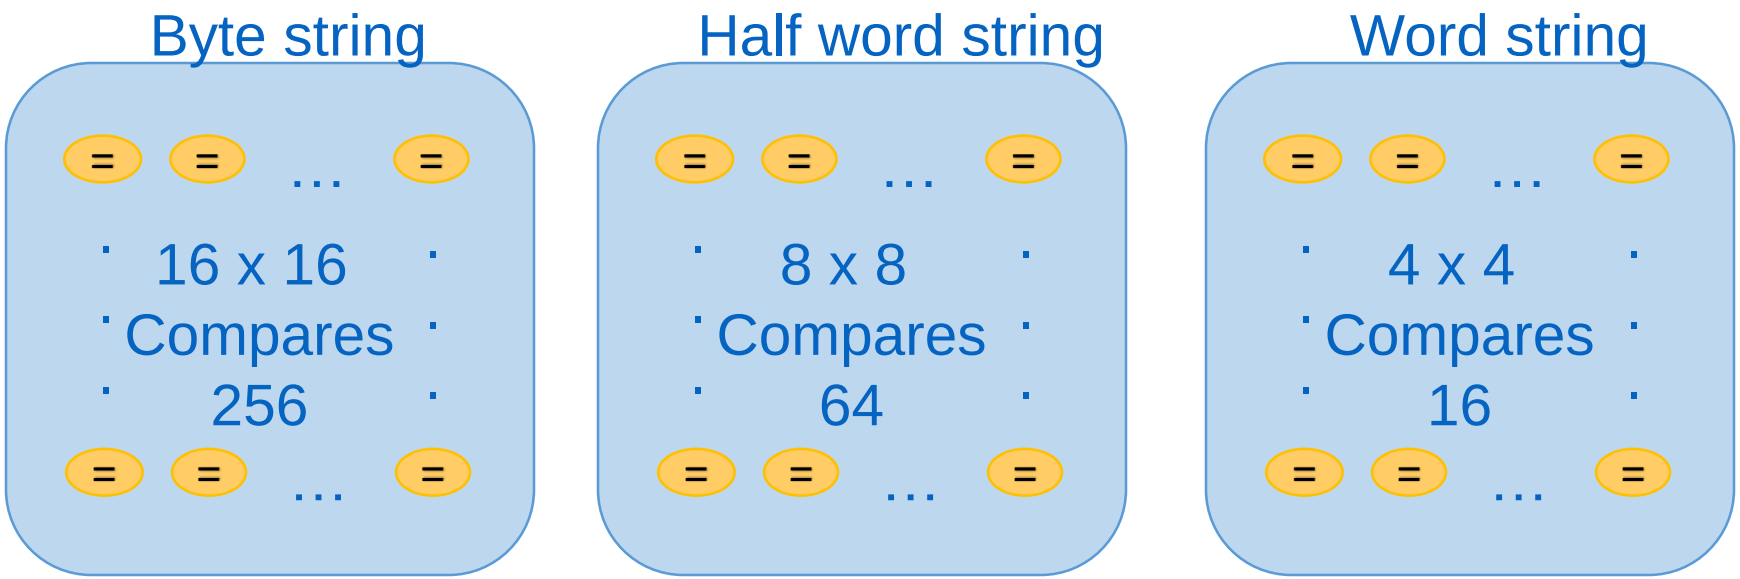

Loop Optimization to Increase Instruction Level Parallelism

```
For (i=1, i<1024, i++)
{  s += a[i] * b[i]; }
```

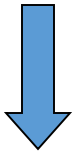

```
For (i=1, i<1024, i+=8)
{  s0 += a[i] * b[i];
   s1 += a[i+1] * b[i+1];
   s2 += a[i+2] * b[i+2];
         ...
   s7 += a[i+7] * b[i+7];
}
s = s1 + s2 + s3 + ... + s7;
```

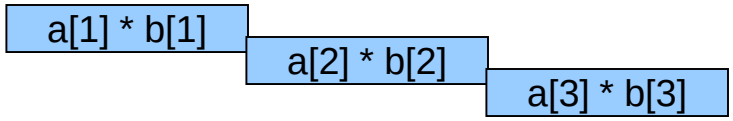

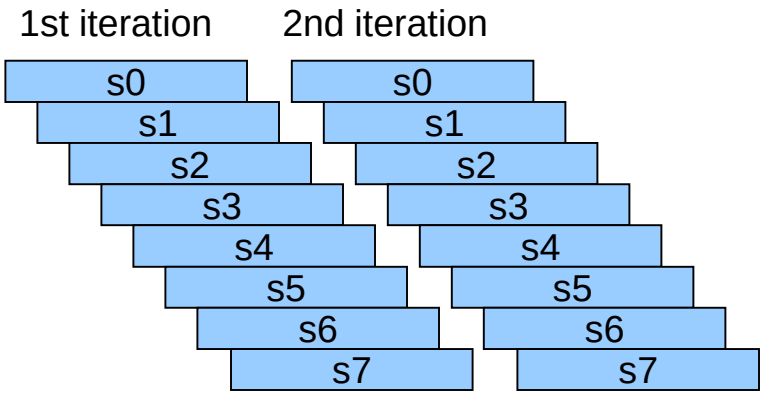

SIMD Exploitation and Enablement – Things IBM is doing for you

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

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

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

Architecture Exploitation
0 to 11 In One Step

All Pre-V5 COBOL compiler releases generated only base ESA/390 level code

ABO, like COBOL V5, generates code up to z13 ARCH(11)

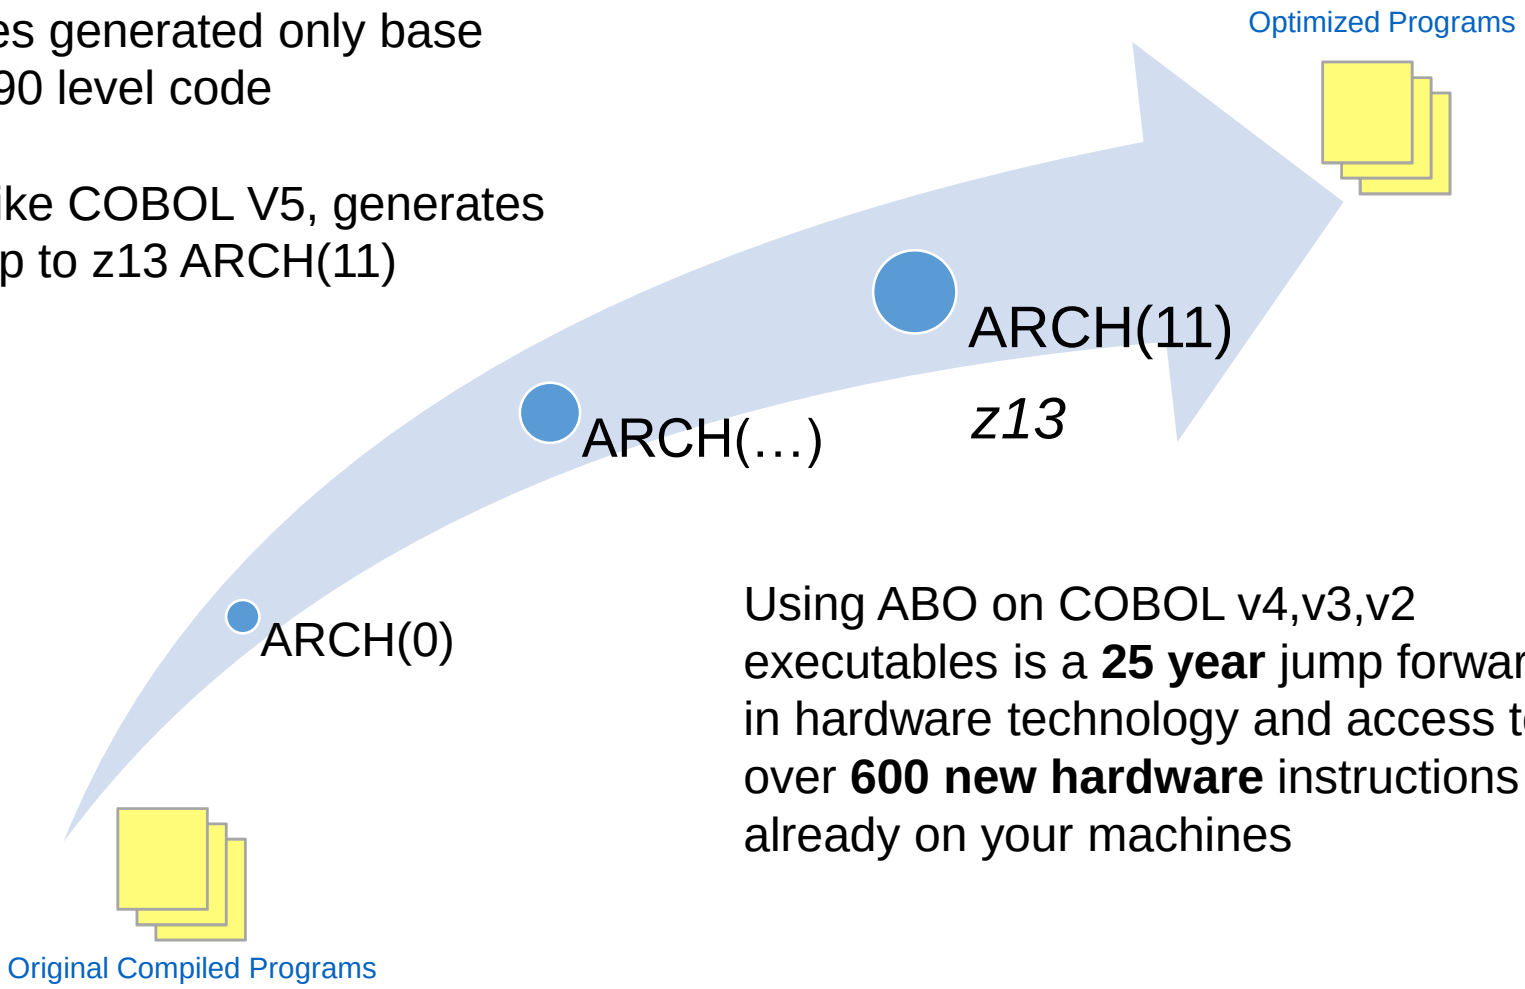

Using ABO on COBOL v4,v3,v2 executables is a **25 year** jump forward in hardware technology and access to over **600 new hardware** instructions already on your machines

IBM Automatic Binary Optimizer (ABO) for z/OS
Overview

http://www-03.ibm.com/software/products/en/z-compilers-optimizer
Available November 6th 2015 – z/OS 2.2

- ABO improves performance of already compiled **COBOL v3 & v4 programs**
  - Optimize directly from the compiled program
    - No source level migration or recompilation or options tuning required
  - Leverage latest advanced COBOL optimization technology
  - Generate code to target latest z Systems (e.g. zEC12,zBC12 & z13)
  - Support in z/OS® 2.2 to automatically load optimized modules to target latest z Systems

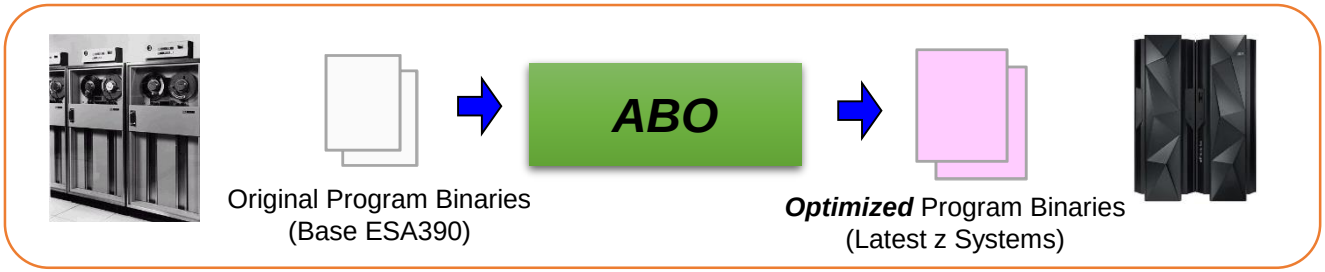

Version 1.1 and Trial Version Available Now – Requires z/OS 2.2
***z/OS 2.1 Support** Planned Availability in 1Q2016*

ABO and COBOL Compiler Positioning

- They serve different but complementary purposes

Use Case	ABO	Compiler
Significant Performance Improvement* **No* Source, Migration or Options Tuning Required	✓	
Interoperability/Legacy Compatibility PDS supported, pre-Enterprise COBOL etc.	✓	
Built in Support for Targeting Multiple Hardware Levels At Deployment	✓	
No need to downgrade ARCH setting to match DR* machine Original compiled program always available for DR	✓	
New COBOL development and new features		✓
Maintenance on existing COBOL programs		✓
Maximum Performance Improvement* *Source, Migration and Options Tuning Required		✓

*DR → Disaster Recovery Machine : Down level machine used for emergency situations. Usually 1 or 2 revisions old so puts limits on Compiler ARCH setting (and performance improvements possible) based on this older level

COBOL Compiler Releases Eligible for Optimization

The compiler releases potentially eligible in the future is currently being reviewed

- *Please provide feedback on which releases should be made eligible*

Program Produced by Compiler Release	Eligible in first release	Potentially Eligible In Future
OS/VS COBOL	✗	?
VS COBOL II	✗	?
COBOL/370 1.1 and COBOL for MVS & VM V1R2	✗	?
COBOL for OS/390 & VM V2R1→V2R2	✗	?
Enterprise COBOL V3R1 → V3R4	✓	✓
Enterprise COBOL V4R1 → V4R2	✓	✓
Enterprise COBOL V5 →	✗	✓*

*eligible via possible future 'Smart Binary' technology in the Compiler and ABO

Performance

Internal Benchmark Suite and Early Customer Results
Higher is better

- **Early customer results show performance gains of 5% → 21% for a mix of v3 and v4 compiled input programs**
- **Performance gains will vary by application but expected to average 15%**

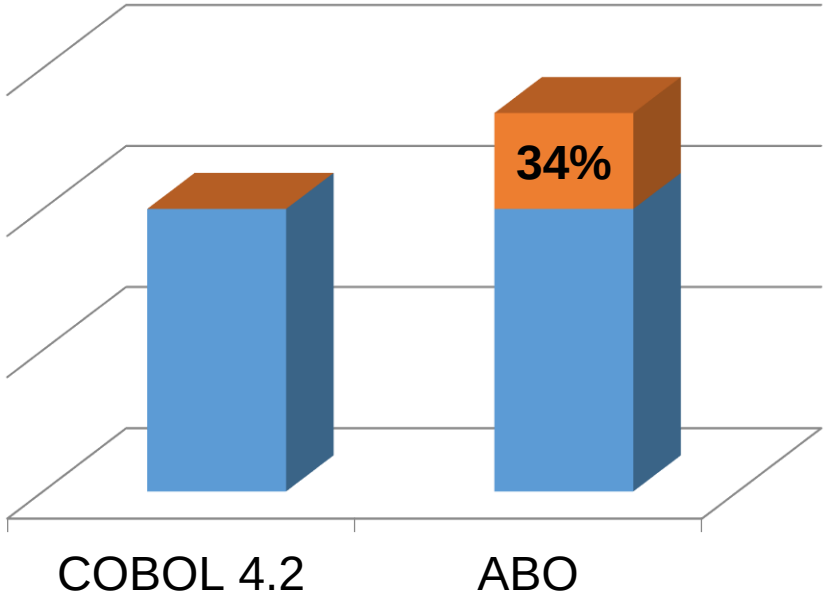

Internal Benchmarks : Mix of Compute and I/O Bound Applications – z13

Higher is better : ABO gives a 34% Improvement

*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.

Tooling Support For The Optimized Modules
Status

- **IBM Problem Determination Tools for z/OS support includes:**
  - Fault Analyzer (FA) for z/OS
  - Debug Tool (DT) for z/OS
  - Application Performance Analyzer (APA) for z/OS
- **Several 3rd party tooling vendors were involved in our beta program this year**
  - Please contact your tools vendor directly to ask about support for ABO

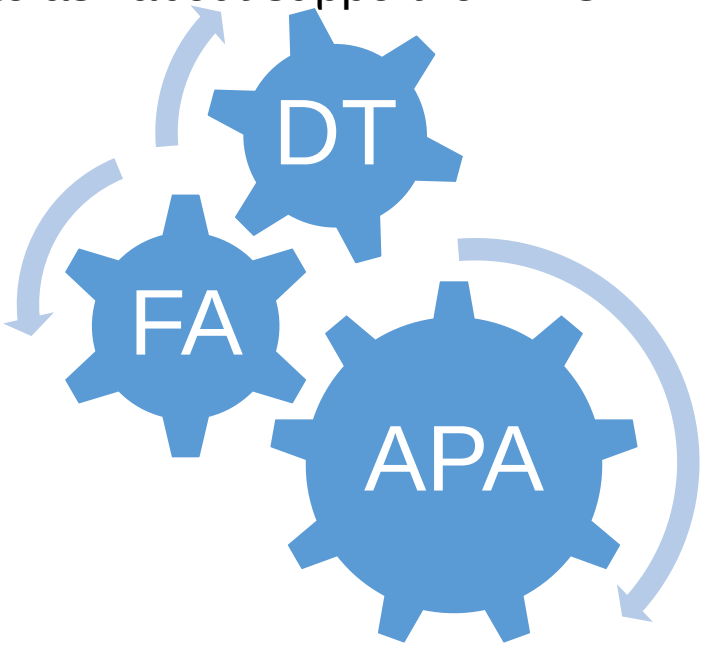

IBM zEnterprise Data Compression (zEDC) capability

- The cost of storing and handling data in CP consumption and DASD space is growing.
- Data compression using software can address this today. There is substantial benefit, but it comes with a cost: CP time.
- Simplistically, data can be classified two ways:
  - Not compressed for <u>frequent access</u> – CPU time used to compress/decompress would be wasted to compress/decompress each time data is accessed. Examples: BSAM/QSAM writing data sequentially and reading it back; DB2 using sequential write to create a report; and SMF logger.
  - Compressed for <u>infrequent access</u> - Historical data that is written out to tape and archived for a few years. Very little need to access this data.
- Goal of IBM zEnterprise Data Compression (zEDC) is to save storage (create storage "white space") and improve wall clock time for compression.

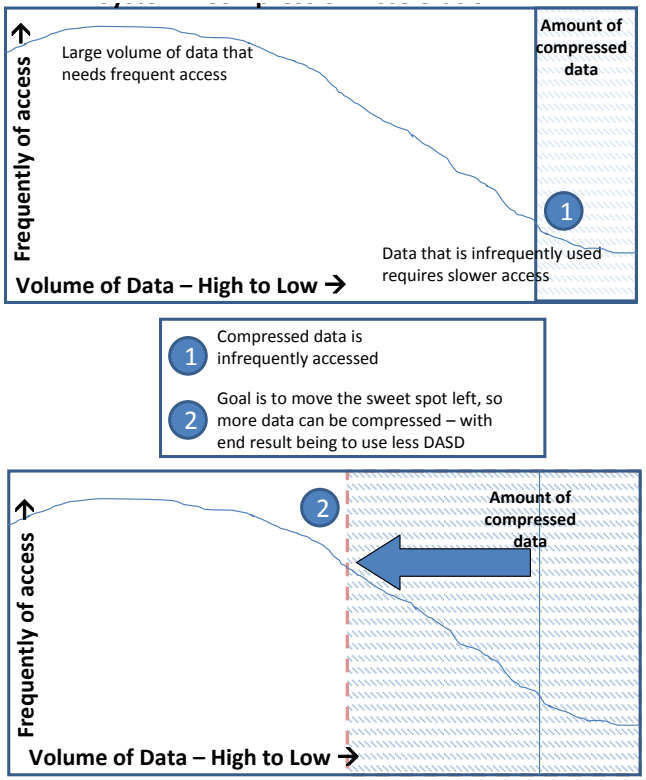

zEDC Express feature

- Configuration:
  - One coprocessor per PCIe I/O feature
  - Supports concurrent requests from up to 15 LPARs
  - Up to 8 features supported by zEC12/zBC12
  - Minimum two feature configuration recommended
- Exploitation and Compatibility
  - Exclusive to zEC12 GA2 and z/OS support in V2R1
  - z/OS Support Planned:
    - z/OS V2.1 – Hardware exploitation for SMF, September 2013, and BSAM/QSAM, 1Q2014*
    - z/OS V1.13 and V1.12 with PTFs - Software decompression support only
    - Authorized APIs for ISV use are planned
    - Includes new PCIE activity report in RMF
- Great results for archived logs (DB2)
- IMS SLDS are also good candidates for zEDC

zEDC Express
FC 0420

Note: Full performance benefits are not achieved unless all systems sharing data are enabled

BSAM/QSAM zEDC Compression Results

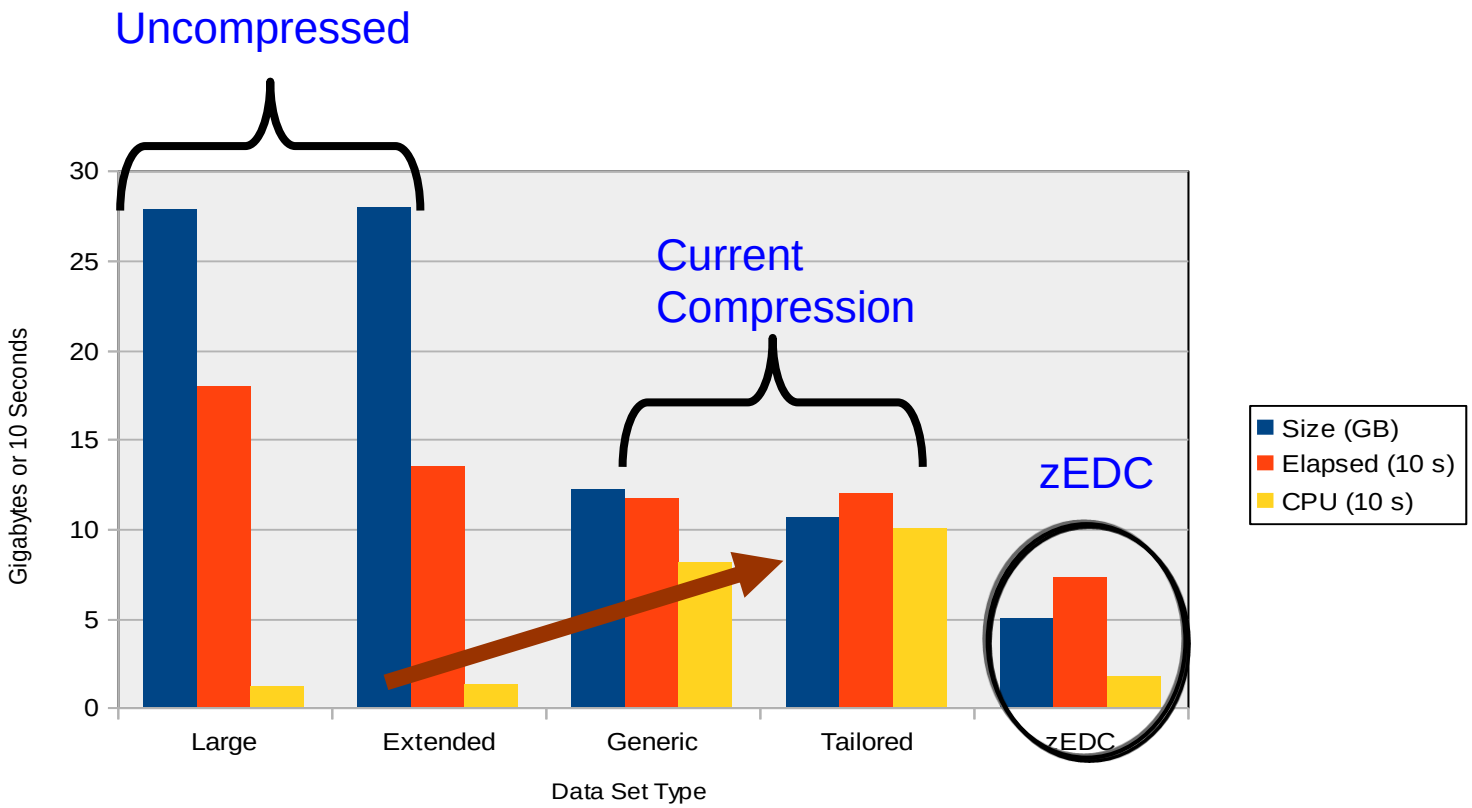

*Measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

zEDC Product usage Overview

- Value buy-in from IBM products across the product stack
- z Systems ISVs have been very active with zEDC

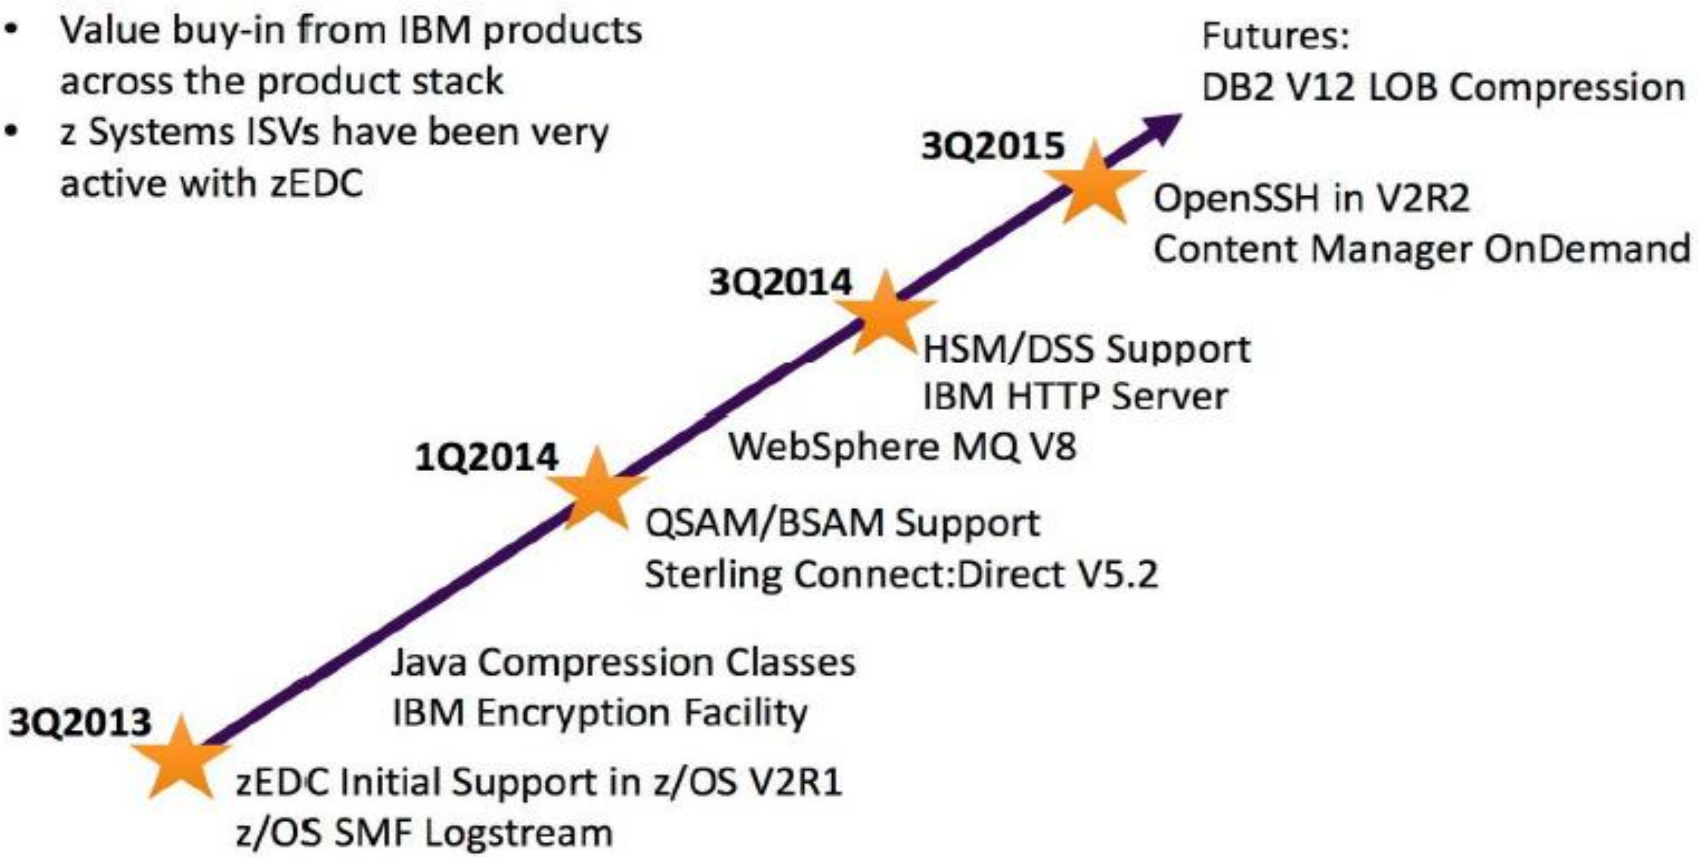

Some performance measurements

IMS Full Function Workload

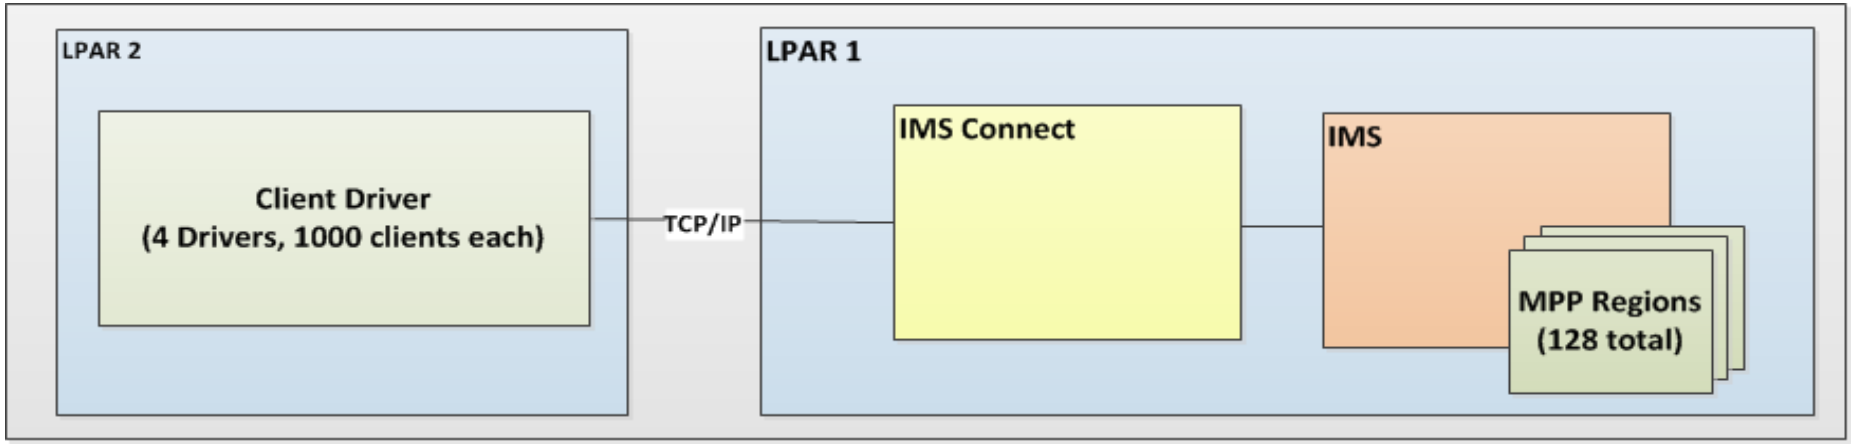

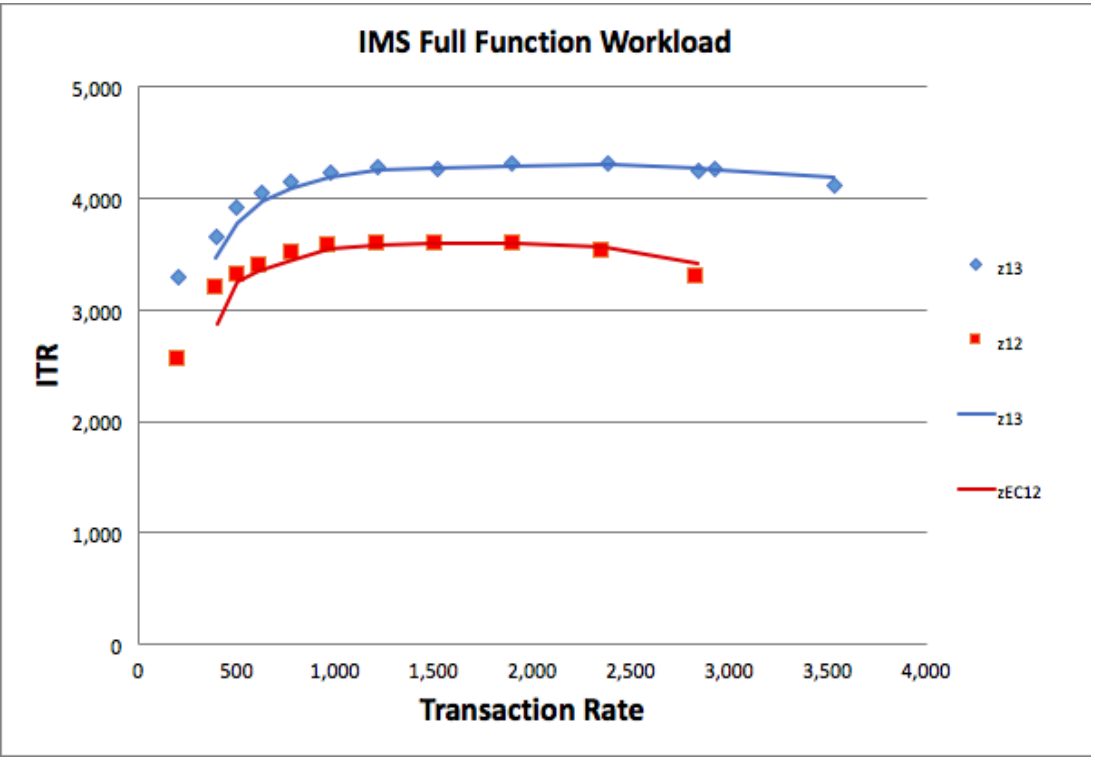

*IMS 13, when running the IMS Full Function workload (1-way IMS, non-data sharing) on IBM z13, showed as much as a **20%** increase in throughput at equivalent CPU as compared to zEC12*

(Controlled measurement environment, results may vary)

IMS Shared Queues Workload

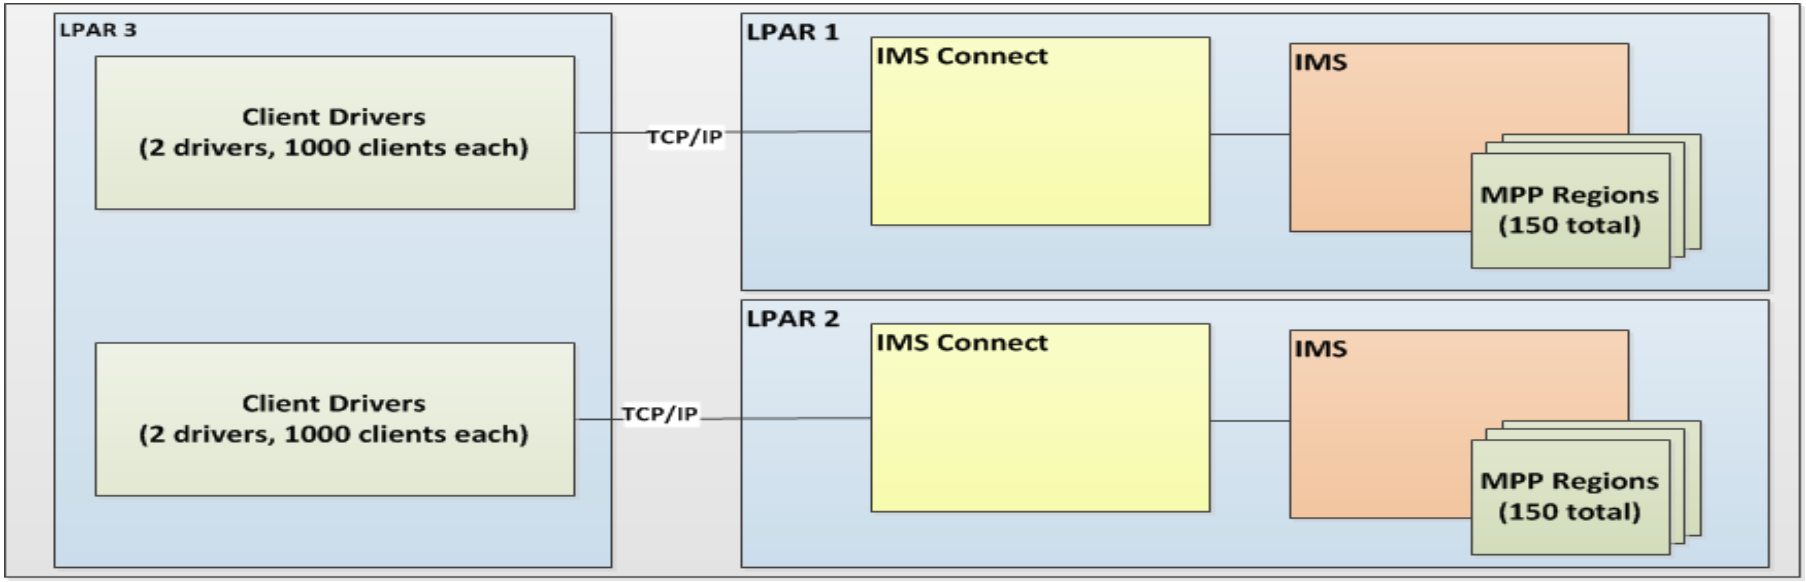

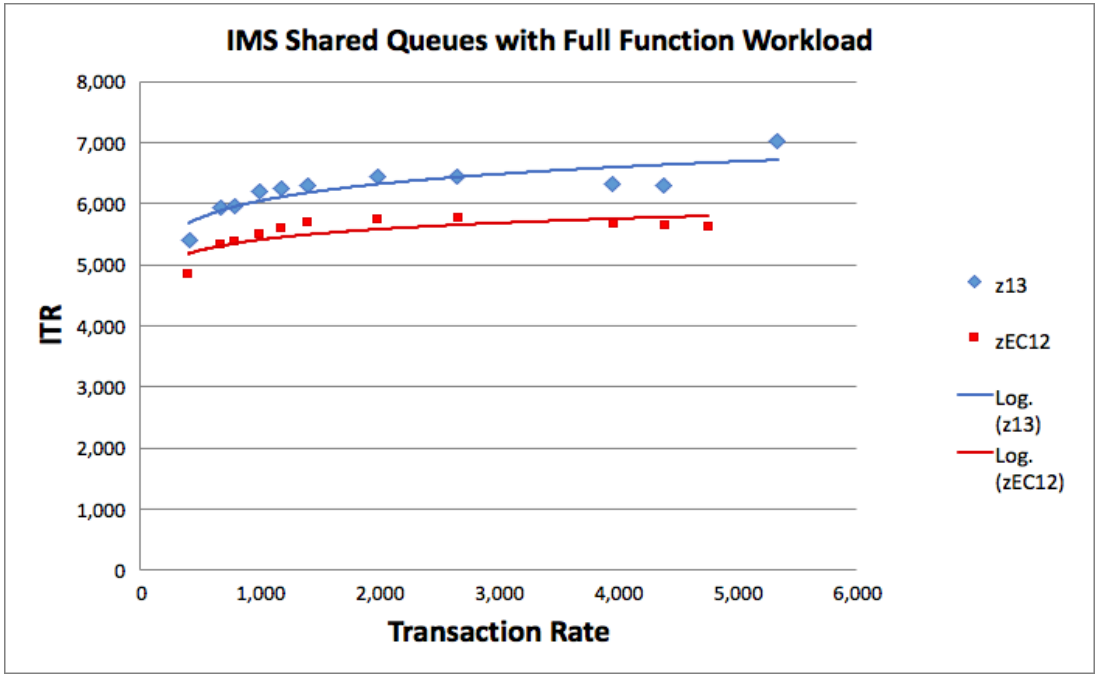

*IMS 13, when running the IMS Shared Message Queues workload (2-way IMS, data sharing) on IBM z13, showed as much as **11%** increase in throughput at equivalent CPU as compared to zEC12*

(Controlled measurement environment, results may vary)

IMS Open Database DRDA Workload

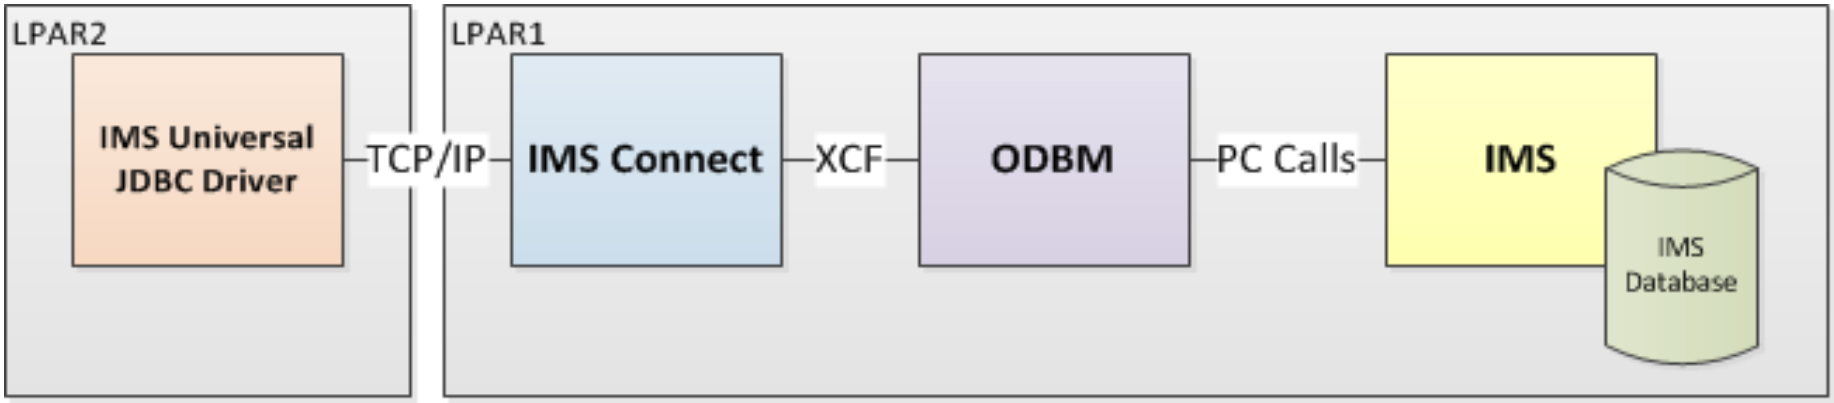

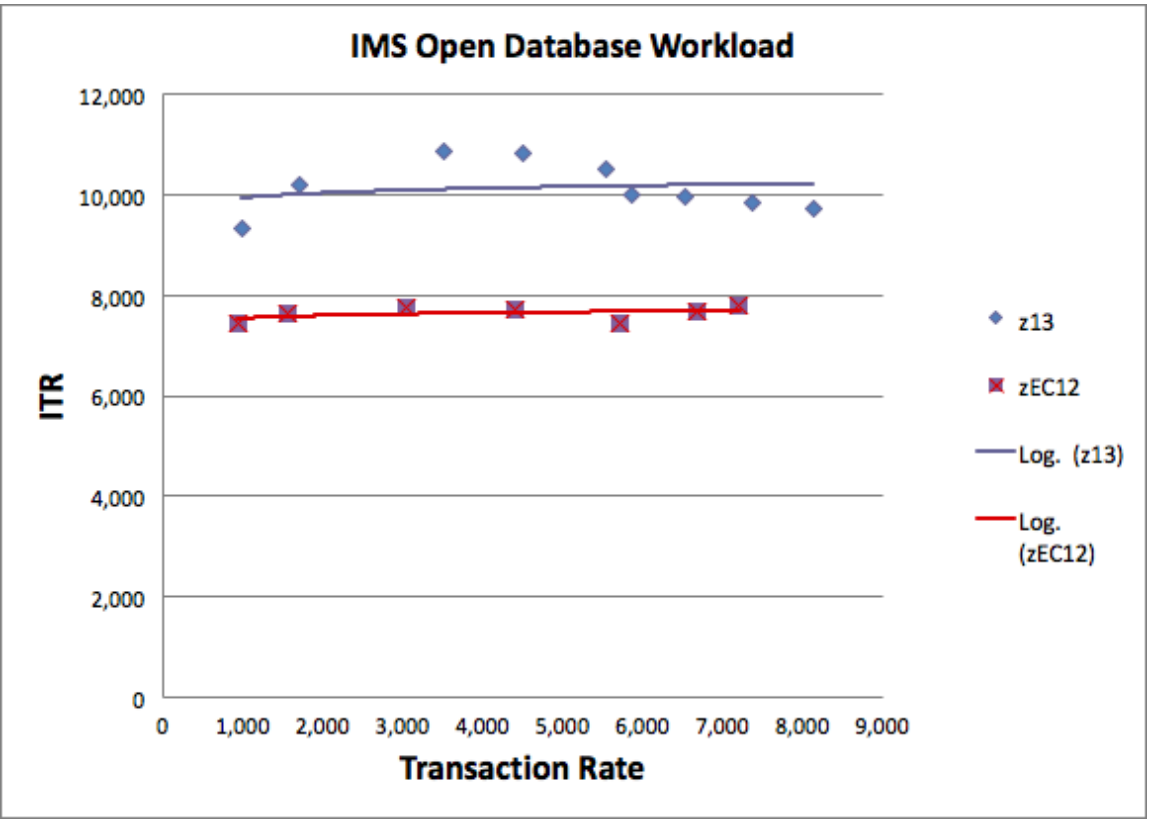

*Open Database DRDA workload on IBM z13, showed as much as **29%** increase in throughput at equivalent CPU as compared to zEC12*

(Controlled measurement environment, results may vary)

IMS Java Transaction Processing Workload

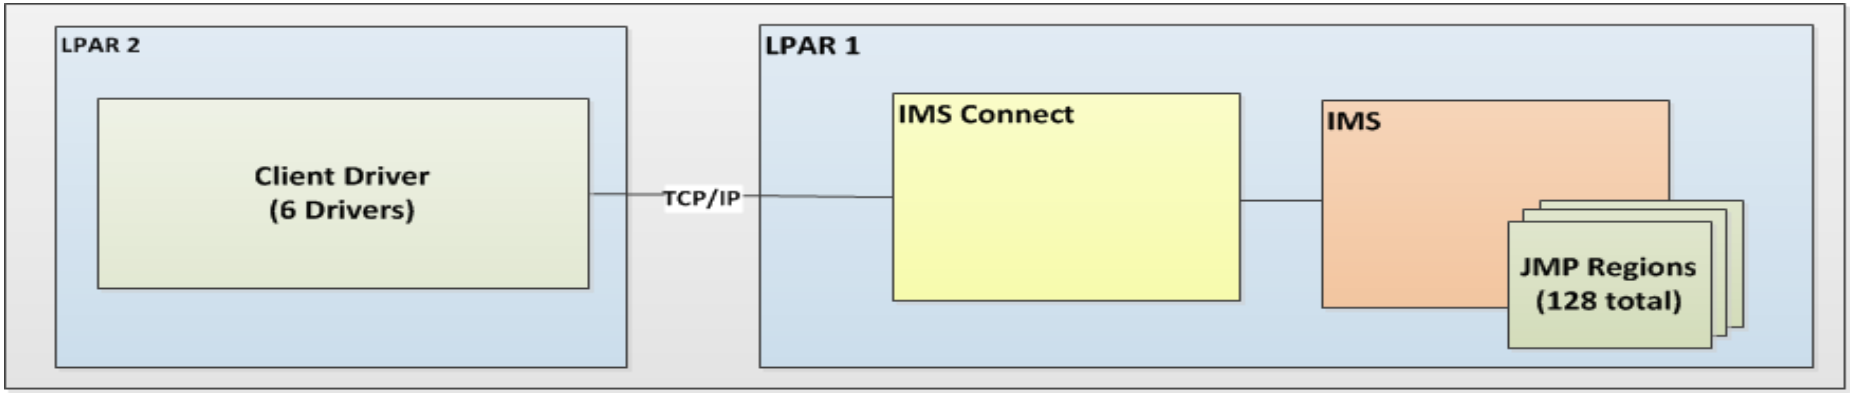

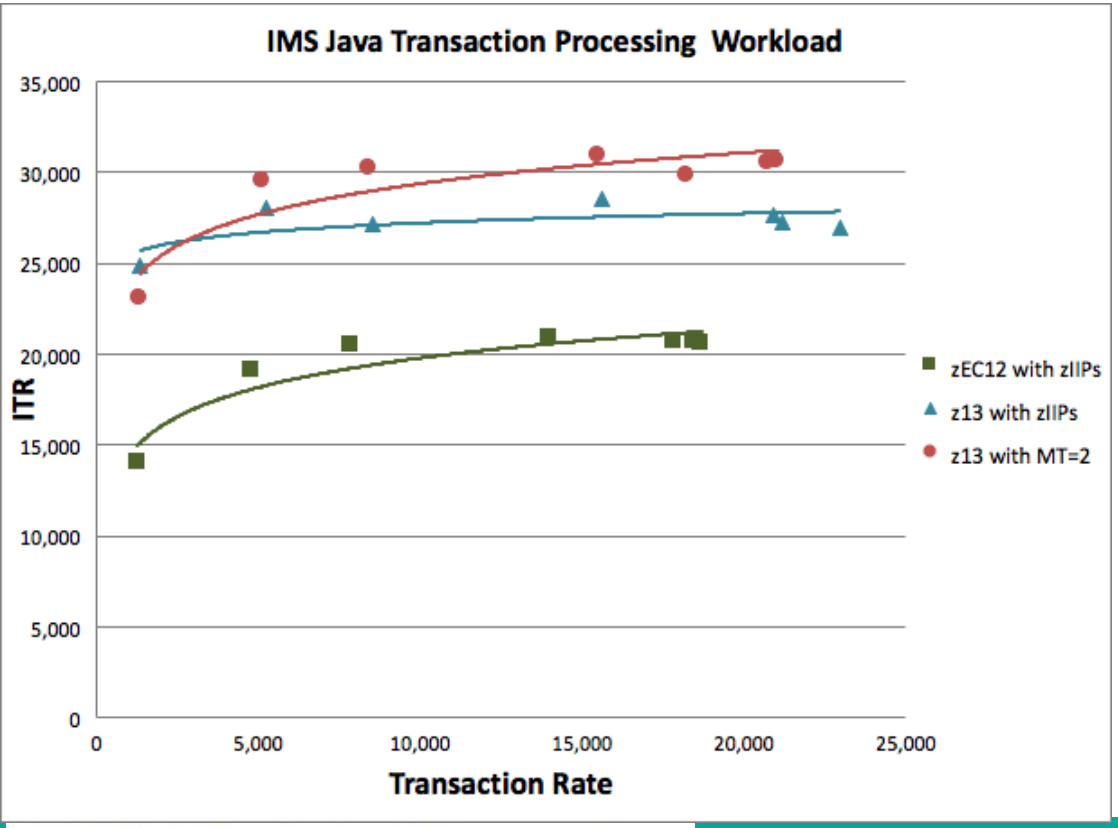

*IMS 13, when running the IMS Java transaction processing workload on IBM z13, showed as much as a **38%** increase in throughput at equivalent CPU as compared to zEC12*

*IMS 13, when running the IMS Java transaction processing workload on IBM z13 using zIIPs with multithreading (**SMT**) enabled, showed as much as an **8%** increase in throughput at equivalent CPU as compared to z13 using zIIPs without multithreading enabled.*

(Controlled measurement environment, results may vary)

z/OS Connect IMS Mobile Feature Pack Workload

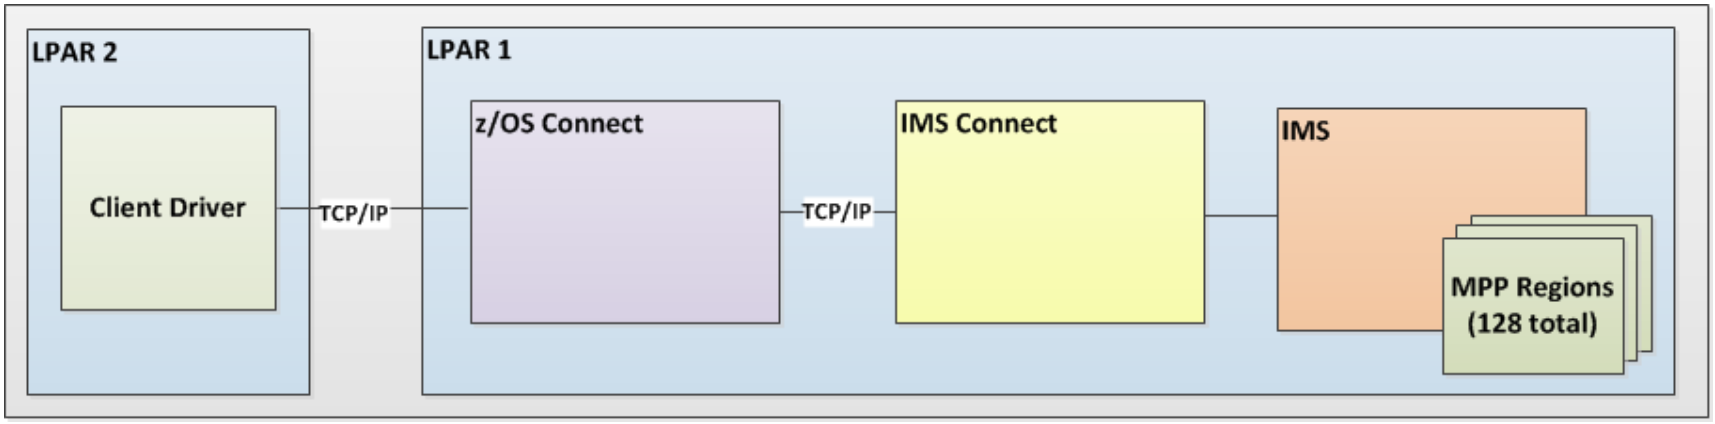

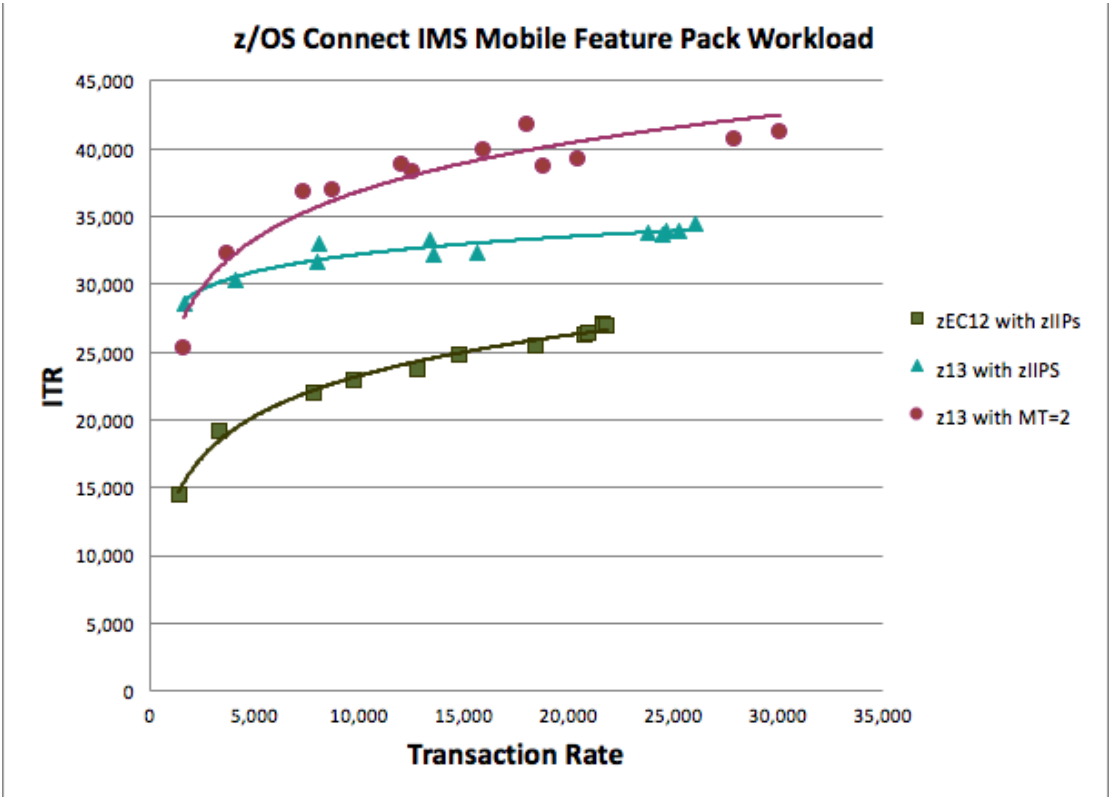

*IMS 13, when running the z/OS Connect IMS Mobile Feature Pack workload on IBM z13, showed as much as **27%** increase in throughput at equivalent CPU as compared to zEC12*

*IMS 13, when running the z/OS Connect IMS Mobile Feature Pack workload on IBM z13 using zIIPs with multi-threading (**SMT**) enabled, showed as much as **22%** increase in throughput at equivalent CPU as compared to z13 using zIIPs without multi-threading enabled*

(Controlled measurement environment, results may vary)

Areas to look into for using large memory

- What about paging rates – be serious – “some” right
- Page fixing is a great idea to save CPU cycles
- DB2 Buffer-Pools
- 1 MB pages (all z/OS) 2 2 GB Pages (Java and DB2) help to reduce CPU cycles.
- MQ Series Version 8 can nicely exploit large memory (allocation above the bar)
- Sort can benefit from large memory (no sortwork on DASD)
- Linux is always memory hungry. To give Linux 20% more memory as it actually needs is good idea. 50% more memory is not any better than 20% for Linux.
- Java Heap size can benefit greatly from additional memory. But – be careful – this is like candies, too much is not good
- Application redesign using much more memories as today (ask your friends in the Intel world how this works
- Have a look at: IBM Redpaper

Summary: The all new IBM z13 and z13s: Pushing the boundaries of system innovations

Up to **10TB** RAIM Memory delivers up to 50% better response time

Accelerated Analytics for Numeric-Intensive Workloads with Single Instruction Multiple Dataset (**SIMD**)

30% Better Capacity for Linux and Java with Simultaneous Multi-Threading (**SMT**)

Specialty Engines: **zIIPs, IFLs, and ICFs** to optimize performance across diverse workloads

Crypto Express5S providing dedicated cryptographic processing for security of transactions and data, 2x faster

Up to **141 Processor** Cores with 5GHz performance and unprecedented scales for data and transaction growth

Up to **320 Separate** Channels of Dedicated I/O for massive data and transaction throughput

Up to **17x Faster Analytics** than the Competition with IBM DB2 Analytics Accelerator

Up to **8,000 Virtual** Machines in one System with new open-standards based KVM hypervisor

zEDC accelerated data compression to reduce data transfer volumes & storage costs by up to 75%

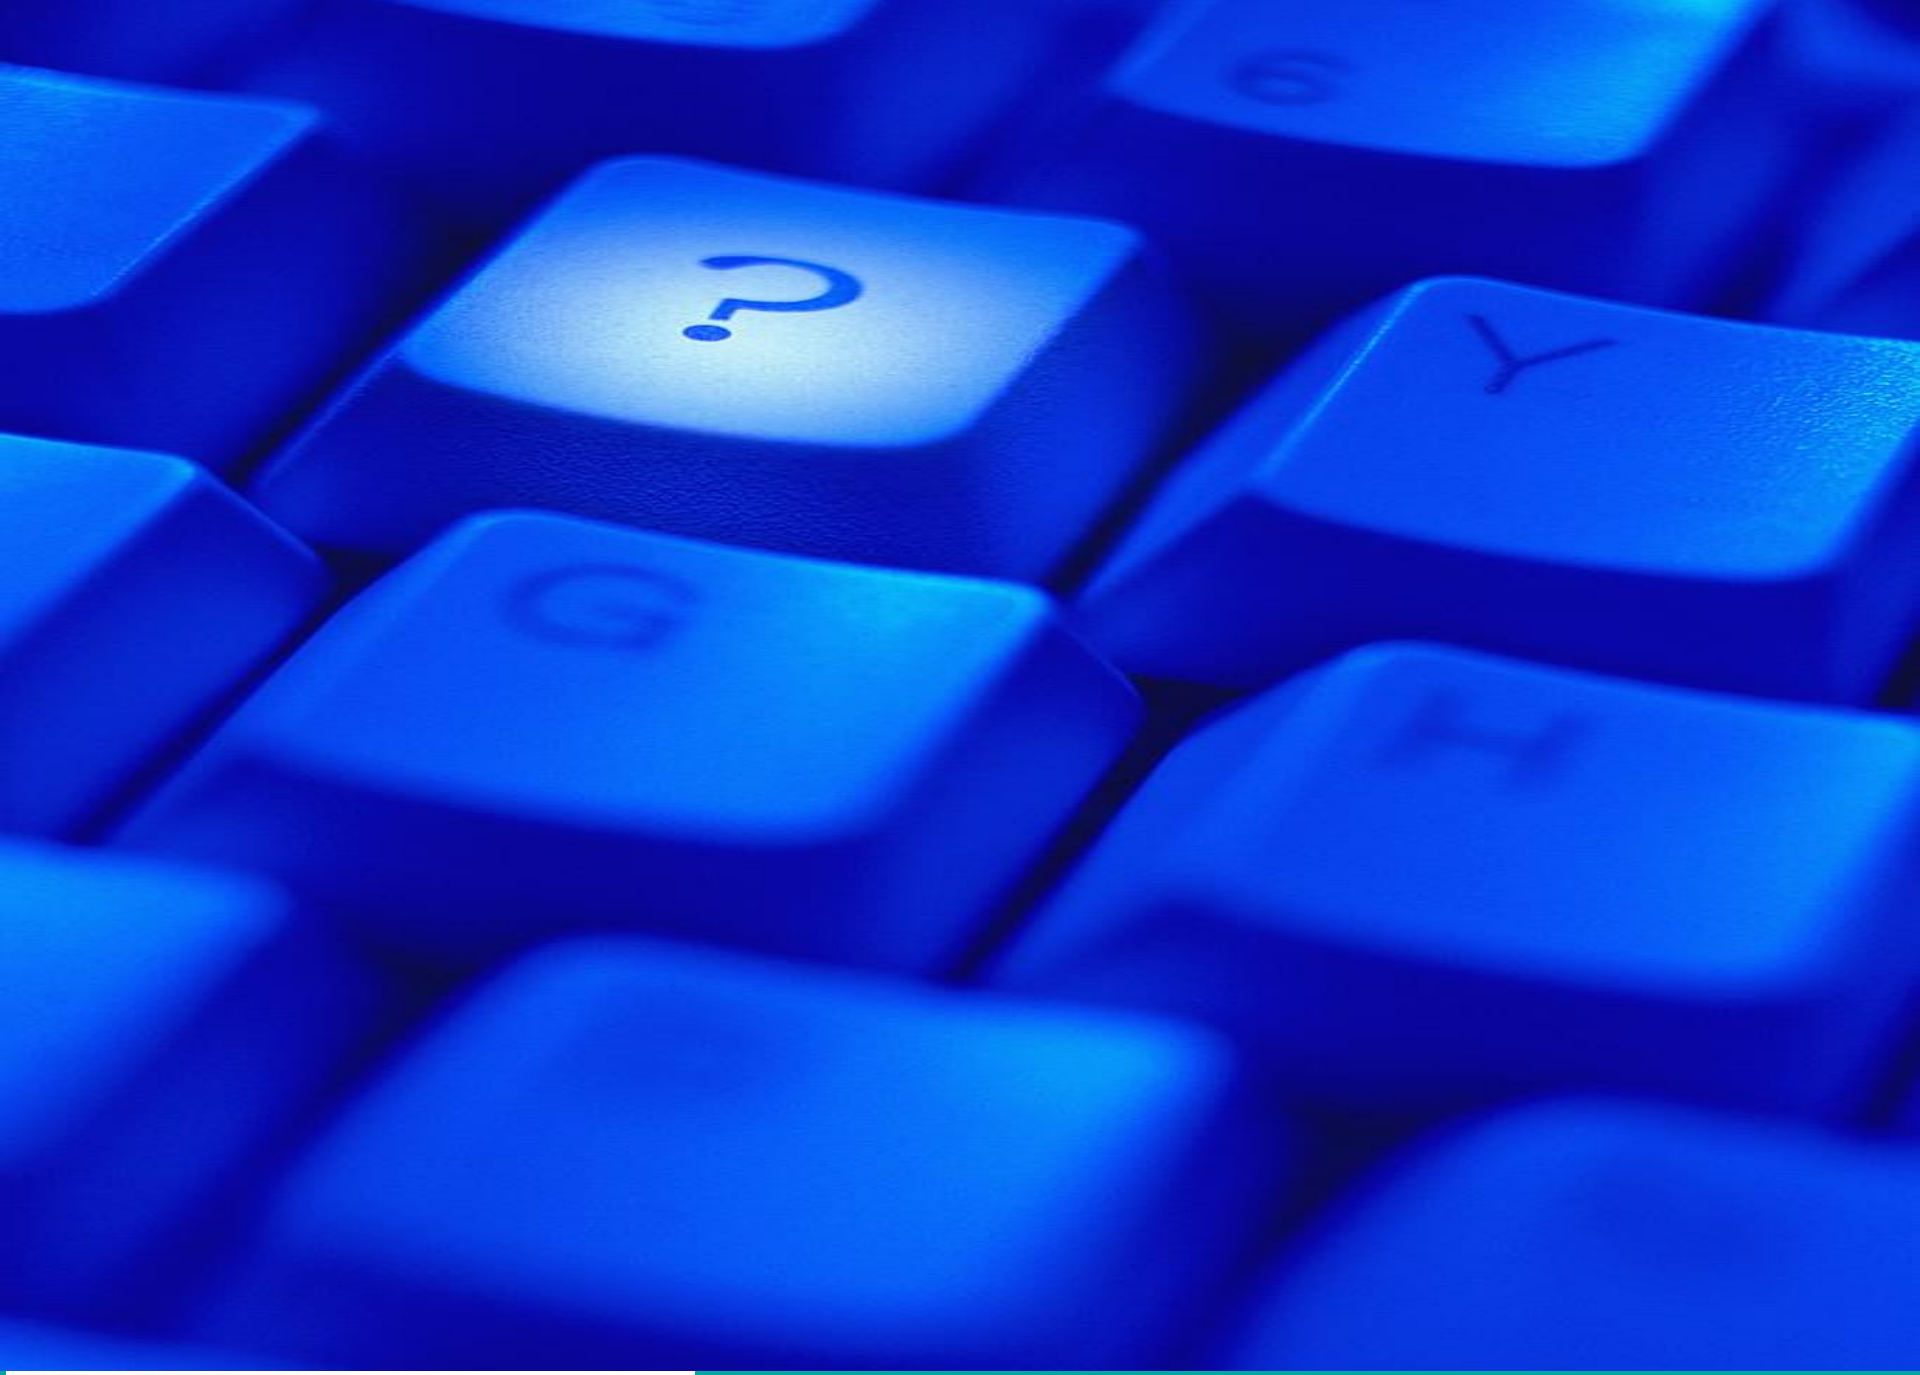

IBM z13 Redbooks

- New – IBM z13 and IBM z13s Technical Introduction, SG24-8250-01
- Updated - IBM z13 Technical Guide, SG24-8251-01
- New - IBM z13s Technical Guide, SG24-8294
- IBM z13 Configuration Setup, SG24-8260
- IBM z Systems Connectivity Handbook, SG24-5444
- Updated – IBM z Systems Functional Matrix, REDP-5157-01

- The z13 IBM Redbooks launch page will be:

  http://www.redbooks.ibm.com/redbooks.nsf/pages/z13?Open

Get Connected with IMS!

Find us on Facebook

facebook.com/IMSFans

slideshare
present yourself

slideshare.net/ibmims

twitter

twitter.com/IBM_IMS

instagram.com/ibm_ims

LinkedIn

linkd.in/IBM_IMS

YouTube

bit.ly/IMS_YouTube