Linux on z14
Testing and Performance experiences

—

Dr. Eberhard Pasch
(epasch@de.ibm.com)

My blog:
https://linuxmain.blogspot.com

Trademarks

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

BladeCenter*	DFSMShsm	HiperSockets	Ibm.com	LinuxONE Emperor	z13*	z/OS*
CICS*	ECKD	HyperSwap*	IBMZ*	PR/SM	z14	z/VM*
DB2*	FICON*	IBM*	Infiniband*	System Storage*	zEnterprise*	z/VSE*
DFSMSdss	GDPS*	IBM (logo)*	LinuxONE	WebSphere*	zHyperLink	

* Registered trademarks of IBM Corporation

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.
ITIL is a Registered Trade Mark of AXELOS Limited.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
UNIX is a registered trademark of The Open Group in the United States and other countries.
VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.
Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Agenda

- **z14 hardware overview**
- Linux z14 support and certification
- Linux z14 exploitation (future)
- z14 hardware features & performance
- z14 performance comparisons

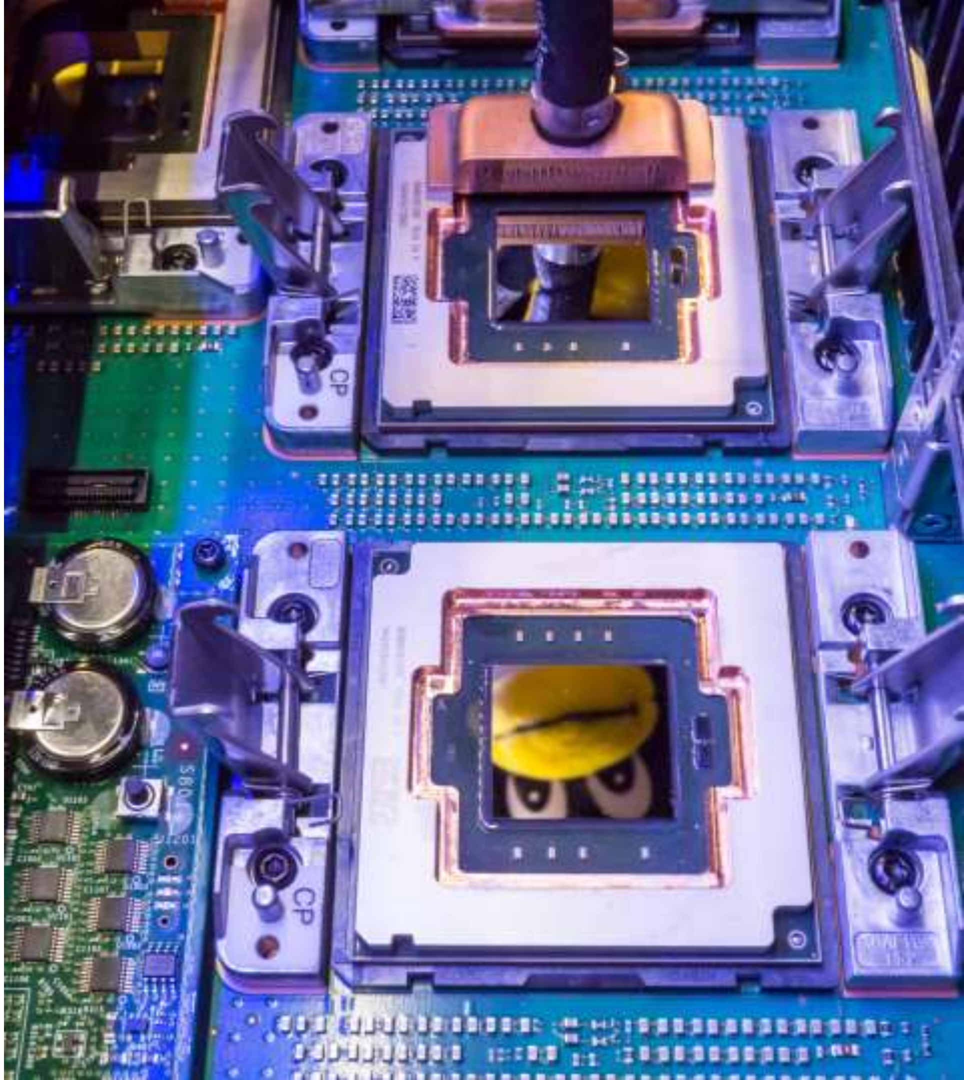

IBM Z – Processor Roadmap

65 nm

z10
2/2008

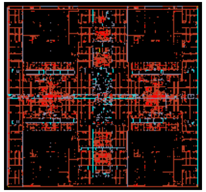

Workload Consolidation and Integration Engine for CPU Intensive Workloads

Decimal FP

Infiniband

64-CP Image

Large Pages

Shared Memory

45 nm

z196
9/2010

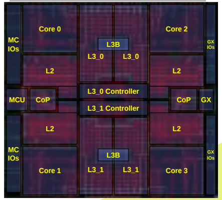

Top Tier Single Thread Performance,System Capacity

Accelerator Integration

Out of Order Execution

Water Cooling

PCIe I/O Fabric

RAIM

Enhanced Energy Management

32 nm

zEC12
8/2012

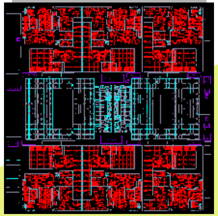

Leadership Single Thread, Enhanced Throughput

Improved out-of-order

Transactional Memory

Dynamic Optimization

2 GB page support

Step Function in System Capacity

22 nm

z13
1/2015

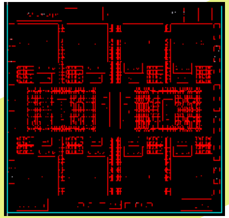

Leadership System Capacity and Performance

Modularity & Scalability

Dynamic SMT

Supports two instruction threads

SIMD

PCIe attached accelerators

Business Analytics Optimized

14 nm

z14
7/2017

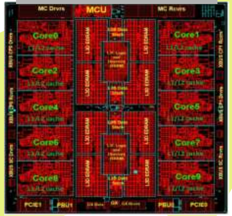

Pervasive encryption

Low latency I/O for acceleration of transaction processing for DB2 on z/OS

Pause-less garbage collection for enterprise scale JAVA applications

New SIMD instructions

Optimized pipeline and enhanced SMT

Virtual Flash Memory

z14 Continues the CMOS Mainframe Heritage

~ 10% for equal z13 n-way
Up to 35% max capacity 170-way vs 141-way (z13)
SMT vs Single Thread ~ 10–40% (average 25%)
- both zIIP & IFL
SMT z14 vs z13 ~ 15% (z/VM Guests)

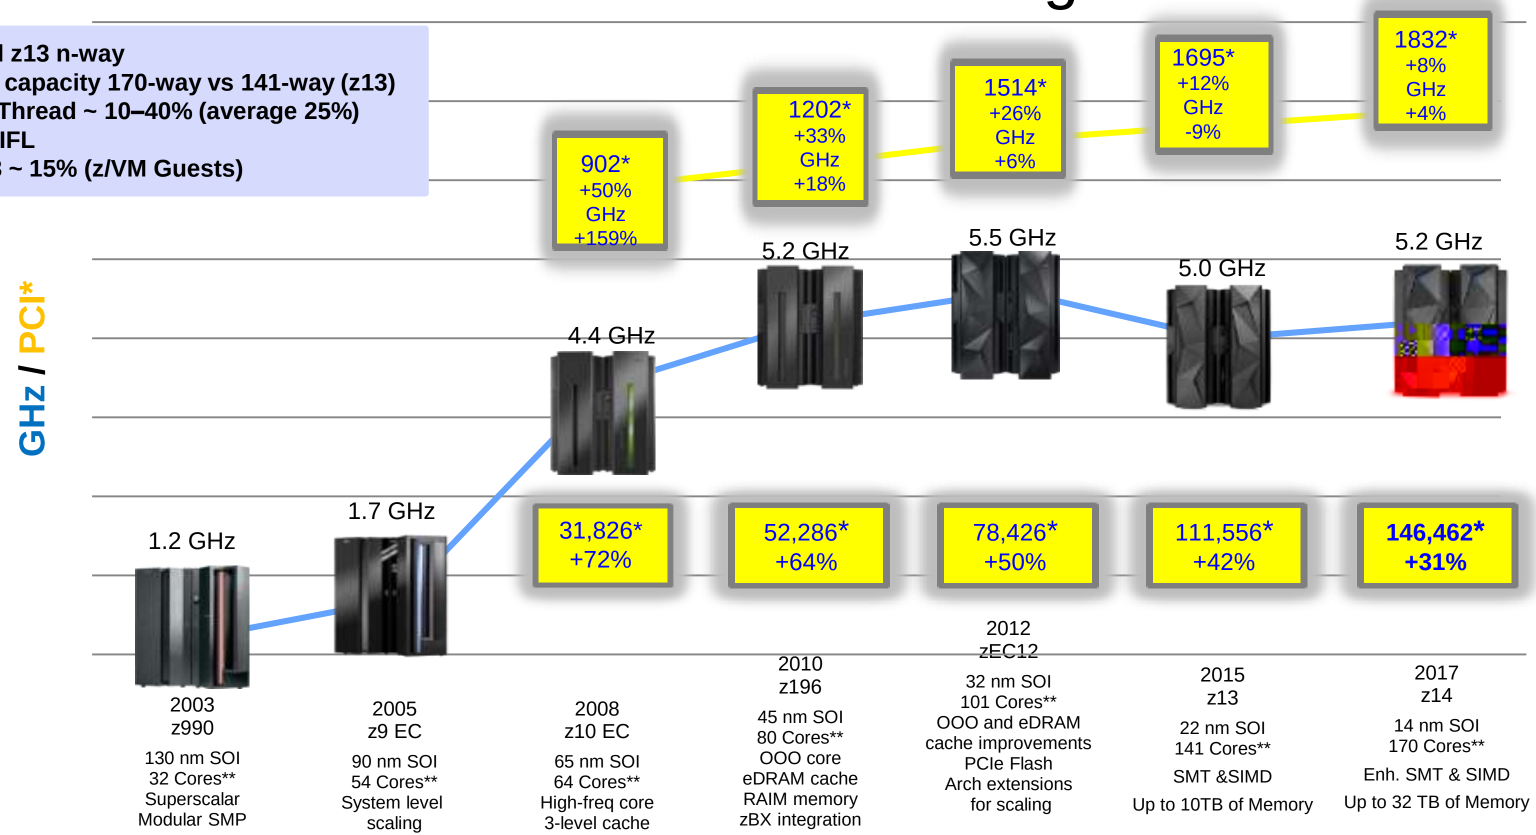

* MIPS Tables are NOT adequate for making comparisons of IBM Z processors. Additional capacity planning required
** Number of PU cores for customer use

z14 System Design Changes

- 14 nm Processor with improved SIMD, SMT, CPACF
- 10 Cores per CP SCM design
- 5 or 6 CP SCMs per Drawer
- Integrated I/O with PCIe Direct Attach
- Single System Controller Chip
- Simplified CPC Drawer SMP Fabric

- Crypto Express6S
- OSA-Express6S
- FICON Express16S+
- RoCE Express2
- IBM zHyperLink Express
- Coupling Express Long Reach
- Radiator Design improvements
- Expanded operating environment (ASHRAE Class A3)
- Thin doors (optional)

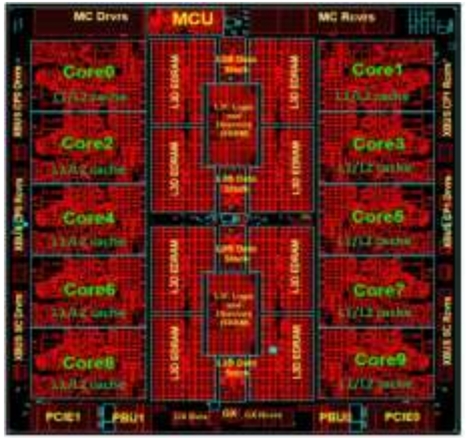

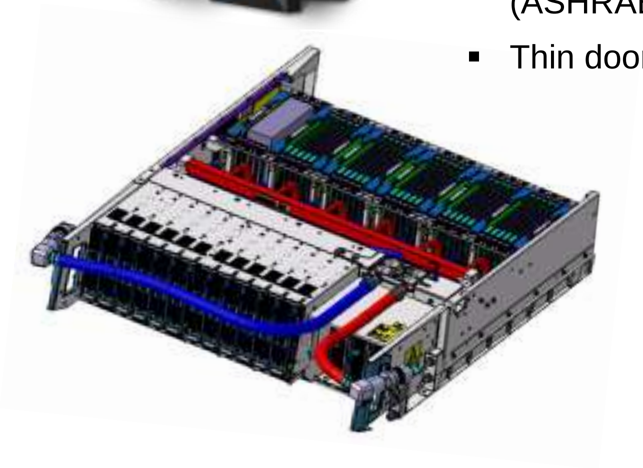

Balanced system design

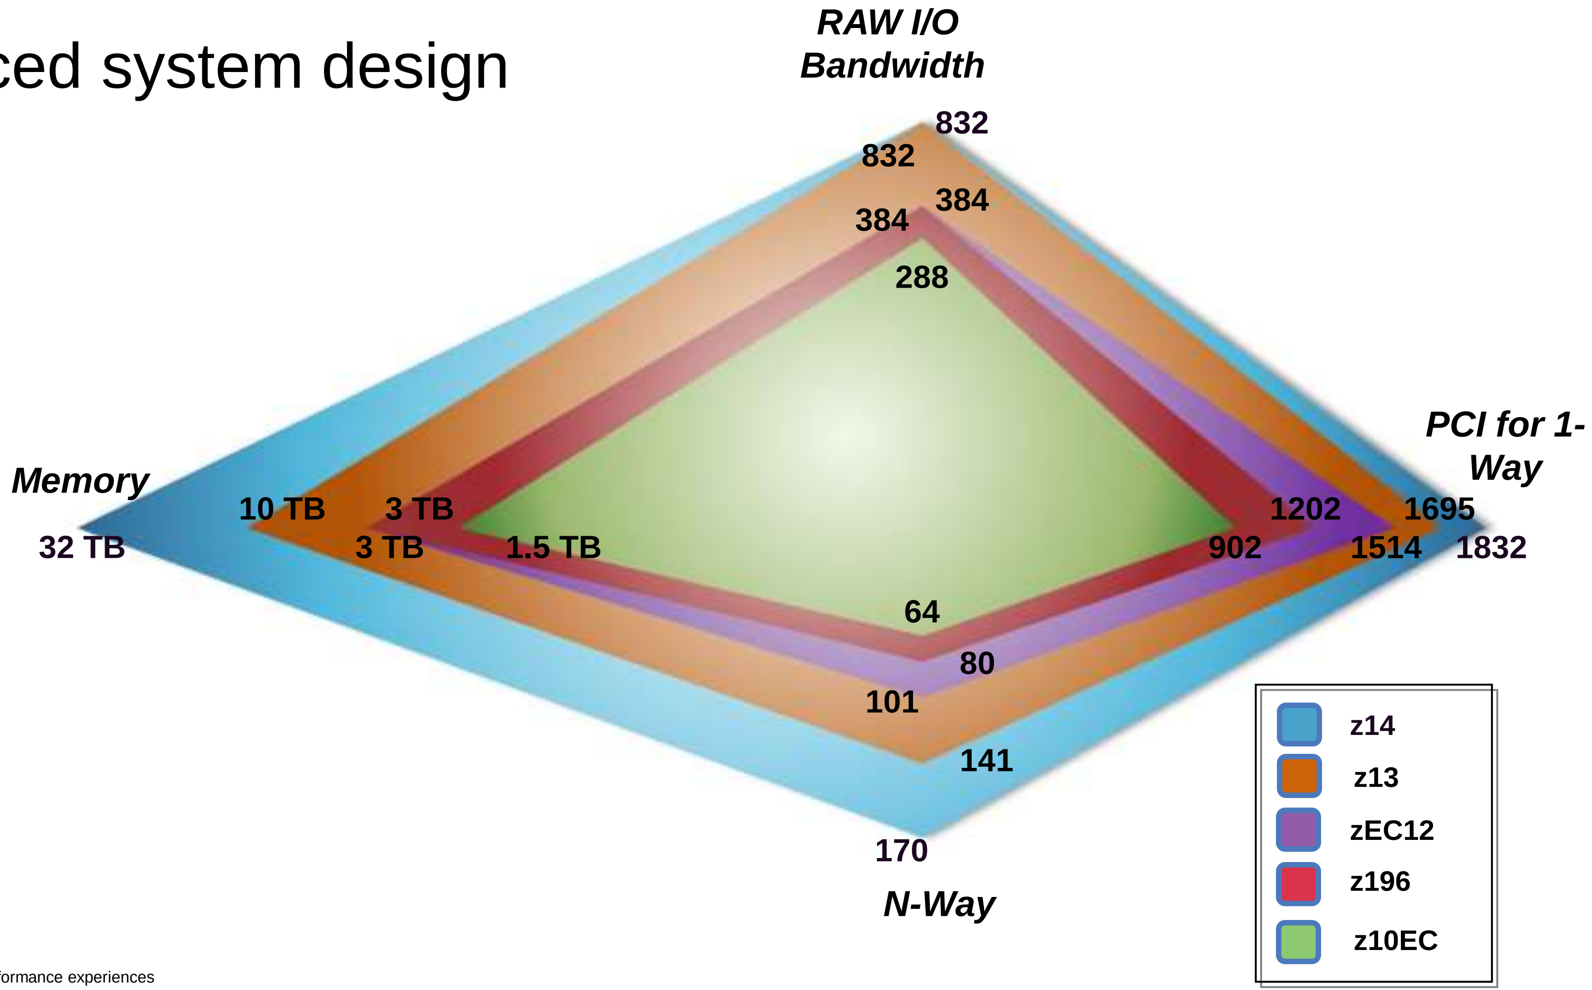

z14 vs. z13 HW SMP Topology

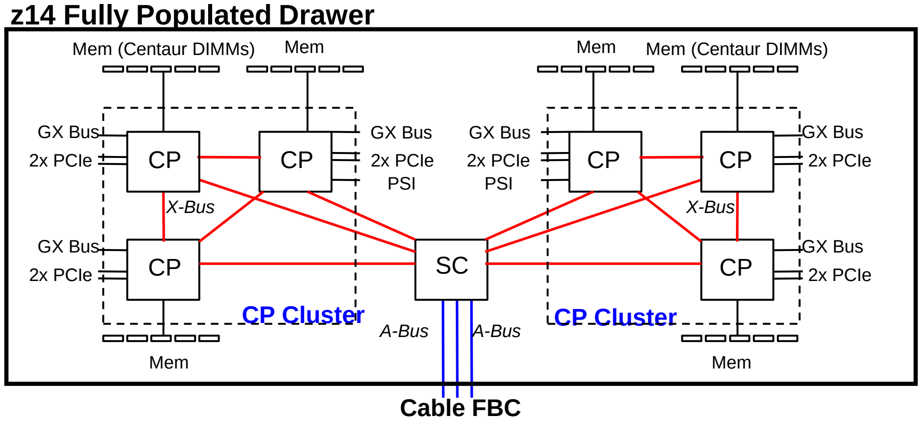

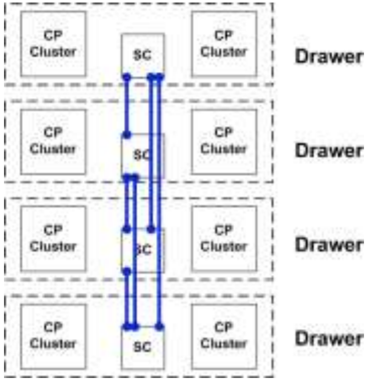

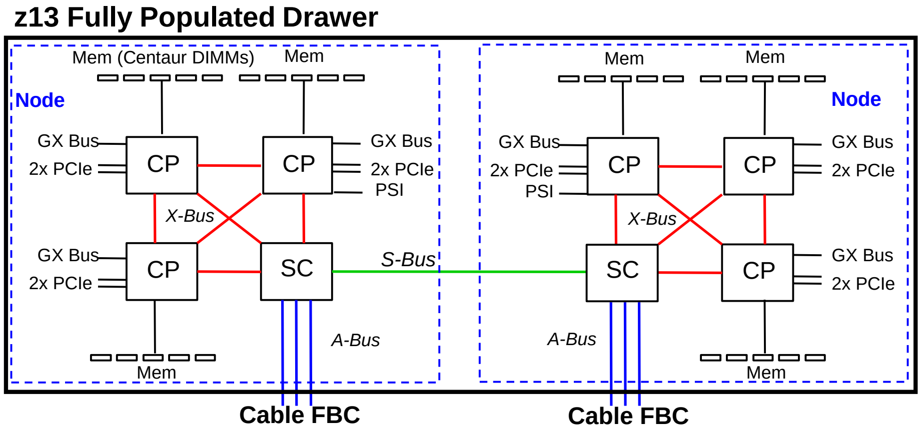

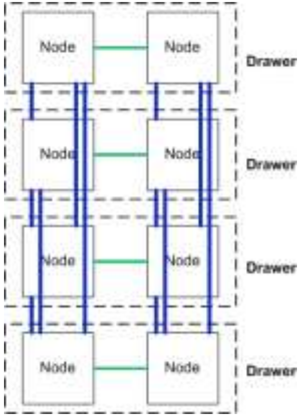

Agenda

- z14 hardware overview
- **Linux z14 support and certification**
- Linux z14 exploitation (future)
- z14 hardware features & performance
- z14 performance comparisons

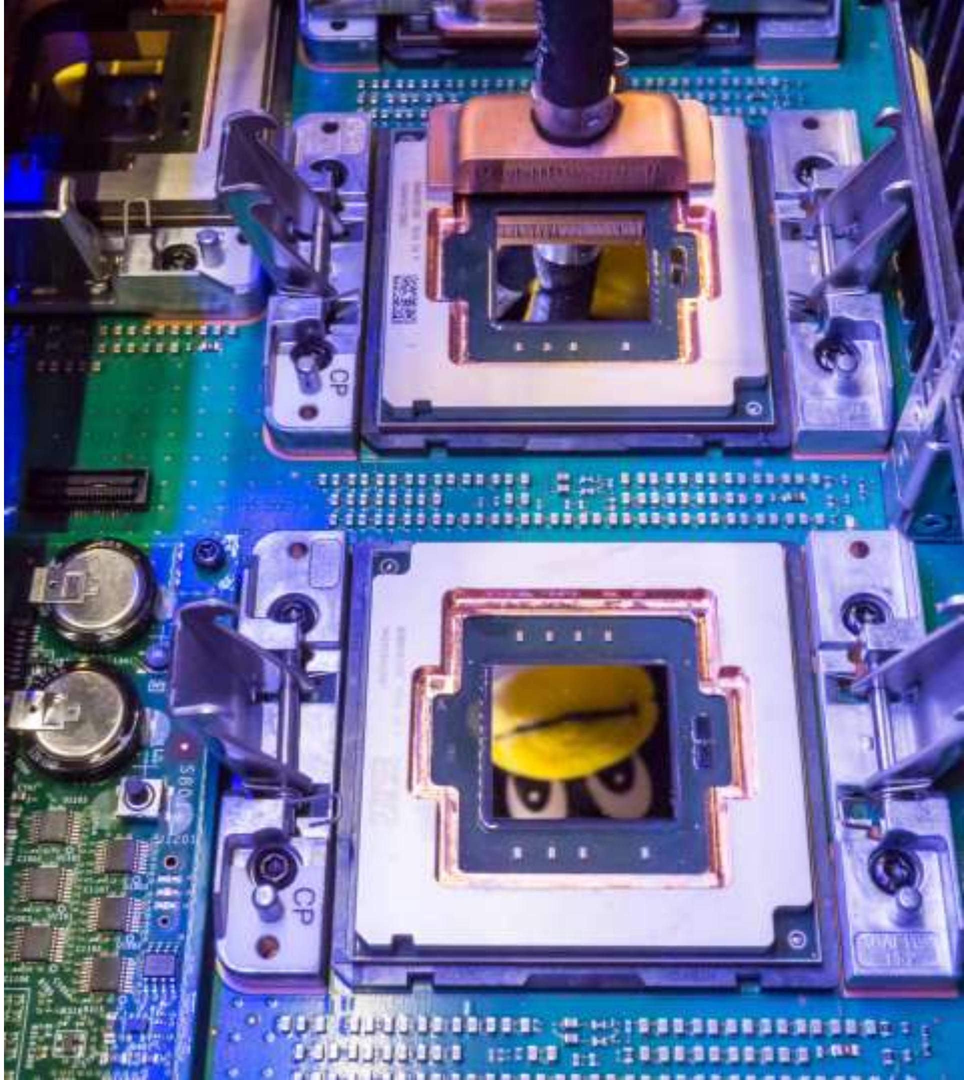

Best practice: Update Operating Systems before HW upgrade

- For z14 some enabling PTFs are required
  - check all guests before upgrading the hardware
  - install software updates before the hardware upgrade
  - No support below the minimal levels

- Linux
  - Tested platforms website with minimal kernel / distribution levels: http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

- z/VM
  - z14 requirement site: http://www.vm.ibm.com/service/vmreqz14.html

Best practice: Update Firmware after hardware upgrade

- System Z firmware is continuously improved
- System Z firmware can applied concurrently

- After hardware update
  - Do a firmware update
  - Check restriction letter

- Plan for a few regular firmware updates outside of freeze periods

IBM tested platform web page

- IBM information is published on
  - http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html
- Screen shot as of 11/20/2017
  - RHEL 6 update is pending depending on fix release by Red Hat

Distribution	LinuxONE Emperor II	LinuxONE Emperor	LinuxONE Rockhopper
	z14	z13	z13s
RHEL 7	✓ (1)	✓ (4)	✓ (4)
RHEL 6	✓ (**)	✓ (4)	✓ (4)
RHEL 5	✗	✓ (4)	✗
RHEL 4 (*)	✗	✗	✗
SLES 12	✓ (2)	✓ (5)	✓ (5)
SLES 11	✓ (2)	✓ (5)	✓ (5)
SLES 10 (*)	✗	✗	✗
SLES 9 (*)	✗	✗	✗
Ubuntu 16.04	✓ (3)	✓ (6)	✓ (6)

IBM tested platform web page – foot notes (2)

- Screen shots as of 11/20/2017
  - RHEL 6 update is pending depending on fix release by Red Hat

[**]IBM is working with the Linux partner to support selected levels of the distribution on z14.

RHEL6 support is planned to be based on a service update of RHEL 6.9

Note: the required patch levels and additional details will be provided soon.

[1] Red Hat Hardware Certification statement for RHEL 7.3 is available at: https://access.redhat.com/ecosystem/hardware/3014651

The following kernel-level and cryptography libraries are the currently known required minimum-levels for z14:

RHEL 7.3: 3.10.0-514.el7 ; Crypto libs: TKE9.0, csulcca-5.2.23-12, ep11-host-1.3.0-3, ep11-host-devel-1.3.0-3

[2] SUSE Hardware Certification statement for SLES 12 SP2 is available at: https://www.suse.com/nbswebapp/yesBulletin.jsp?bulletinNumber=145823

The following kernel-level and cryptography libraries are the currently known required minimum-levels for z14:

SLES 12 SP2: kernel-default-4.4.74-92.35.1 ; Crypto libs: TKE9.0, csulcca-5.2.23-12, ep11-host-1.3.0-3, ep11-host-devel-1.3.0-3

SUSE Hardware Certification statement for SLES 11 SP4 is available at: https://www.suse.com/nbswebapp/yesBulletin.jsp?bulletinNumber=145817

The following kernel - level and cryptography libraries are the currently known required minimum - levels for z14: SLES 11 SP4: kernel - default - 3.0.101 - 108.10 ; Crypto libs: TKE9.0, csulcca - 5.2.23 - 12, ep11 - host - 1.3.0 - 3, ep11 - host - devel - 1.3.0 - 3

RoCE Express2 is not supported

[3] Canonical Hardware Certification statement for Ubuntu 16.04.3 LTS is available at: https://certification.ubuntu.com/server/models/?query=z14&vendors=IBM

The following kernel-level and cryptography libraries are the currently known required minimum-levels for z14:

Distribution certifications

- **Red Hat**
  - RHEL 7.3 / 6.9
    - https://access.redhat.com/ecosystem/hardware/3014651
- **SUSE**
  - SLES 12.3
    - https://www.suse.com/nbswebapp/yesBulletin.jsp?bulletinNumber=145823
  - SLES 11.4
    - https://www.suse.com/nbswebapp/yesBulletin.jsp?bulletinNumber=145817
- **Canonical**
  - Ubuntu 16.04 LTS
    - https://certification.ubuntu.com/server/models/?query=z14&vendors=IBM

Caveats

- Timing changes in IO
  - Faster IO cards – 3x IOPS
  - Race conditions possible (after upgrade double check your logfiles)

- STCK instruction is slower
  - Documented in the Principles of Operations since years – now really visible
  - Alternative instructions available
    - STCKF
    - STCKE

Agenda

- z14 hardware overview
- Linux z14 support and certification
- **Linux z14 exploitation (future)**
- z14 hardware features & performance
- z14 performance comparisons

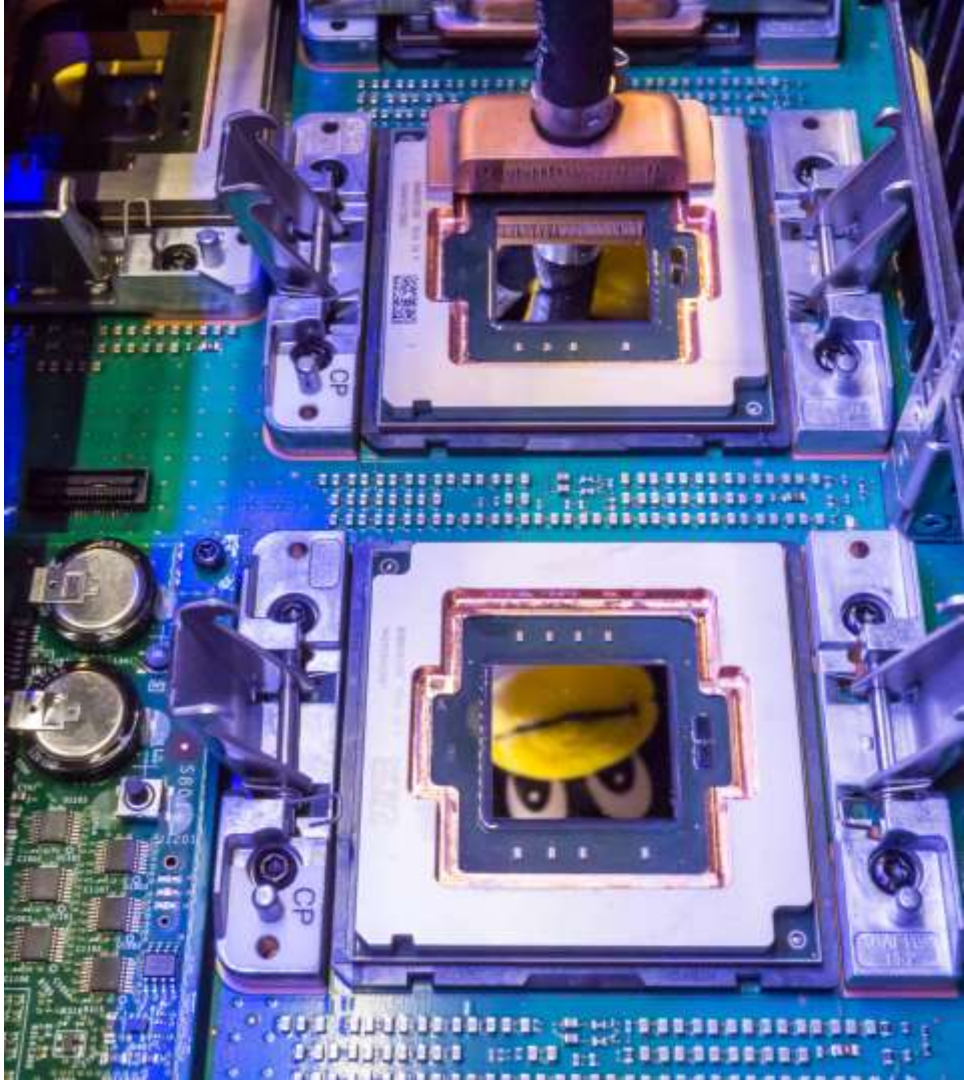

z14 exploitation

- Next major releases from Red Hat, SUSE, Canonical
- Selected features in upcoming updates
  - RHEL 7.5 + 7.6
  - SLES 12 SP4
  - Ubuntu 16.04.5
- New compiler and toolchain is available for two out of three new distributions
  - Red Hat DTS 7.0
    - https://access.redhat.com/documentation/en-us/red_hat_developer_toolset/7/html/7.0_release_notes/
  - SUSE toolchain module update released
    - https://download.suse.com/Download?buildid=RMVZVdUyf60~
  - Ubuntu – next major release

Agenda

- z14 hardware overview
- Linux z14 support and certification
- Linux z14 exploitation (future)
- **z14 hardware features & performance**
- z14 performance comparisons

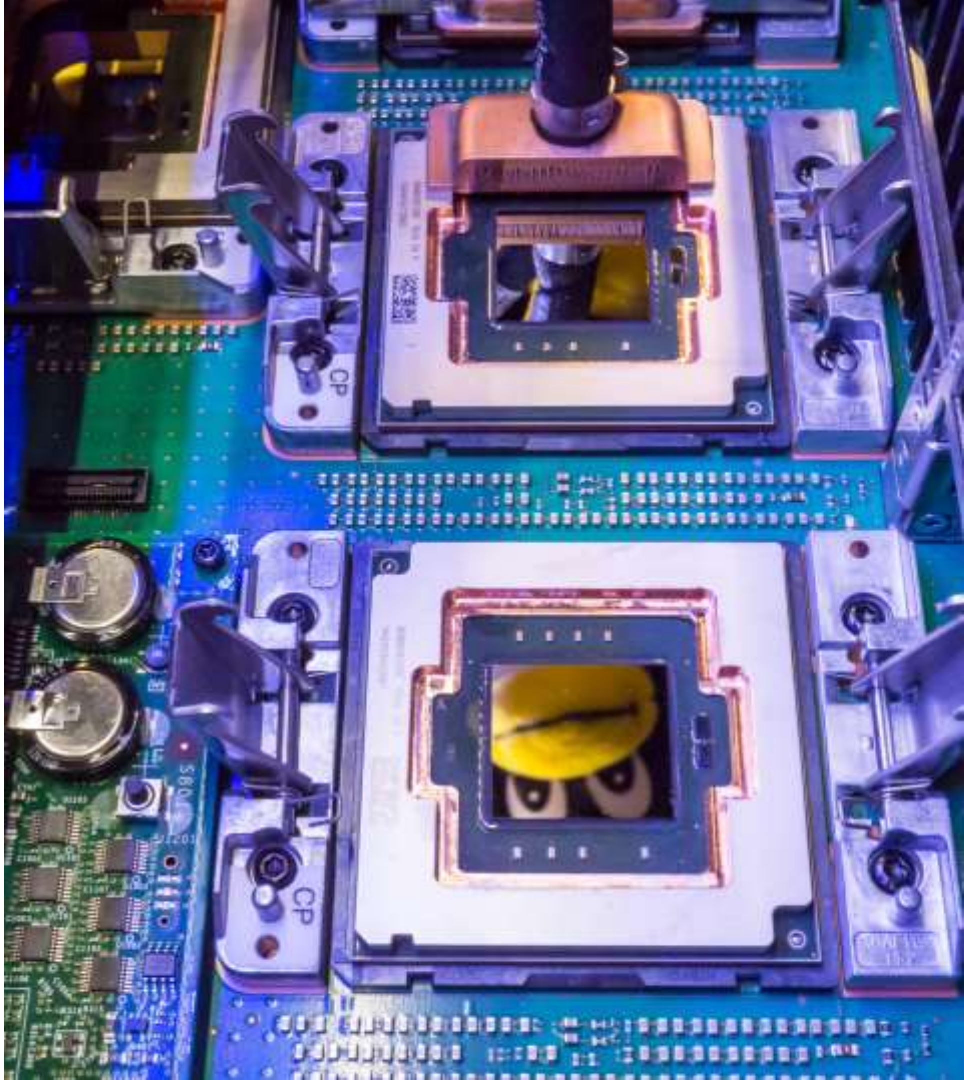

z14 processor design summary

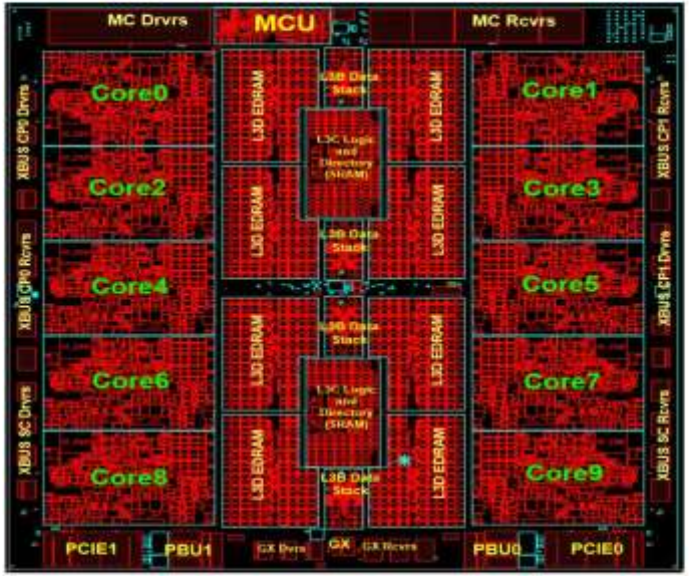

- 6.1 Billion transistors
- 25.3 x 27.5 mm chip area
- 14nm SOI technology,
- 17 layers of metal
- 10 cores per CP-chip,
- 5.2GHz

- Cache Improvements:
  - New power efficient logical directory design
  - 33% larger L1 I$ (128K)
  - 2x larger L2 D$ (4MB)
  - 2x larger L3 Cache with symbol ECC
- New Translation/TLB2 design
  - 4 concurrent translations
  - Reduced latency
  - Lookup integrated into L2 access pipe
  - 2x CRSTE growth
  - 1.5X PTE growth
  - New 64 entry 2gig TLB2
- Pipeline Optimizations
  - Improved instruction delivery
  - Faster branch wakeup
  - Reduced execution latency
  - Improved OSC* avoidance
  - Optimized 2nd generation SMT2
- Better Branch Prediction
  - 33% Larger BTB1 & BTB2
  - New Perceptron Predictor
  - New Simple Call Return Stack

CPU benchmarks – relative improvements single thread

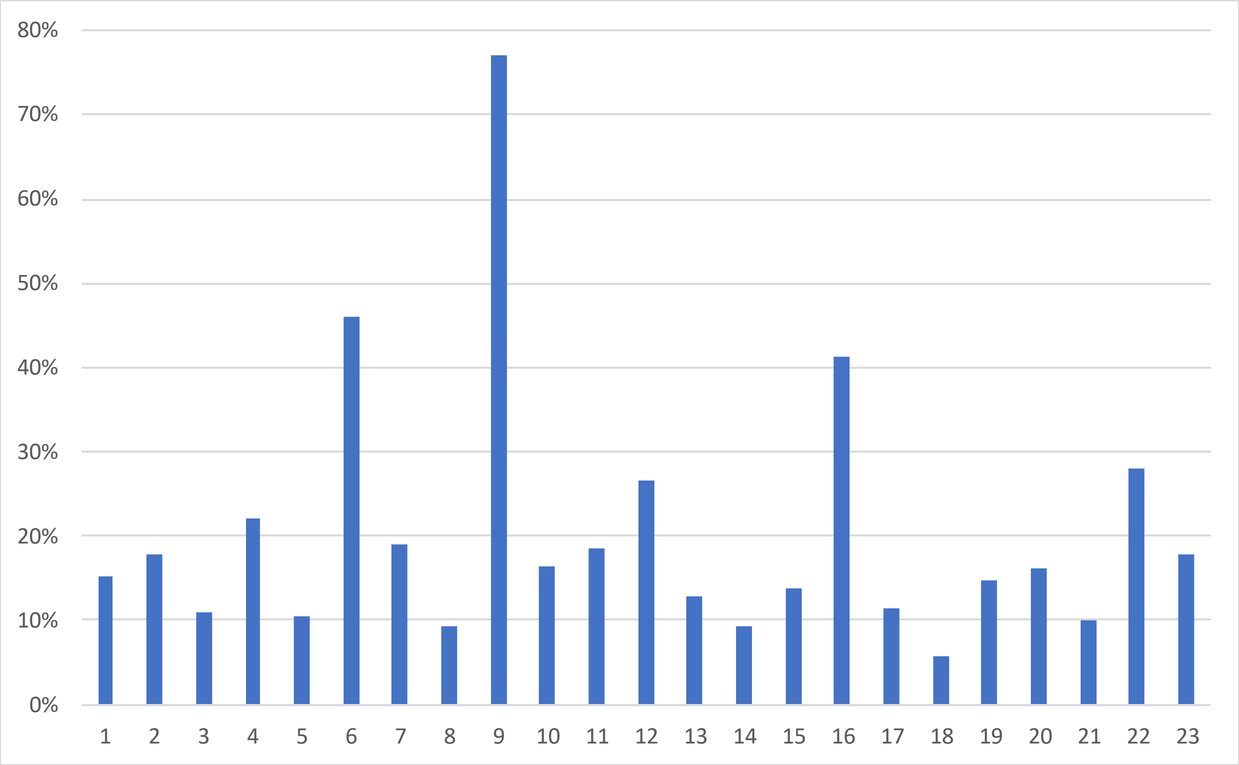

- new compiler (SLES 12.2 tool chain module)
- Optimized for use of z13 instructions
- Hardware is optimized for future workloads
- All results on the positive side
  - If you have z13 optimized code that degrades on z14 please contact me!
- If you have the opportunity: recompile!

CPU benchmarks – relative improvements single thread (2)

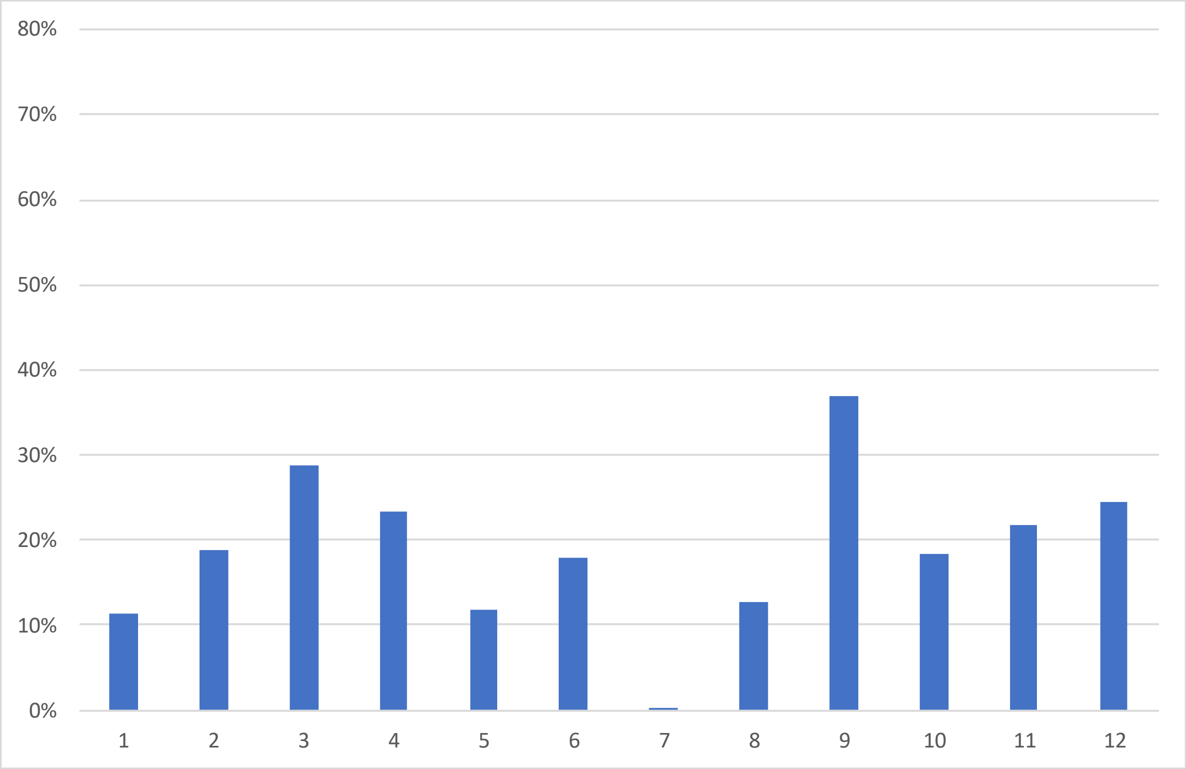

- SLES 12.2 standard compiler
- -march=z196, -mtune=zEC12
- Your mileage will vary

Java 8 – performance example

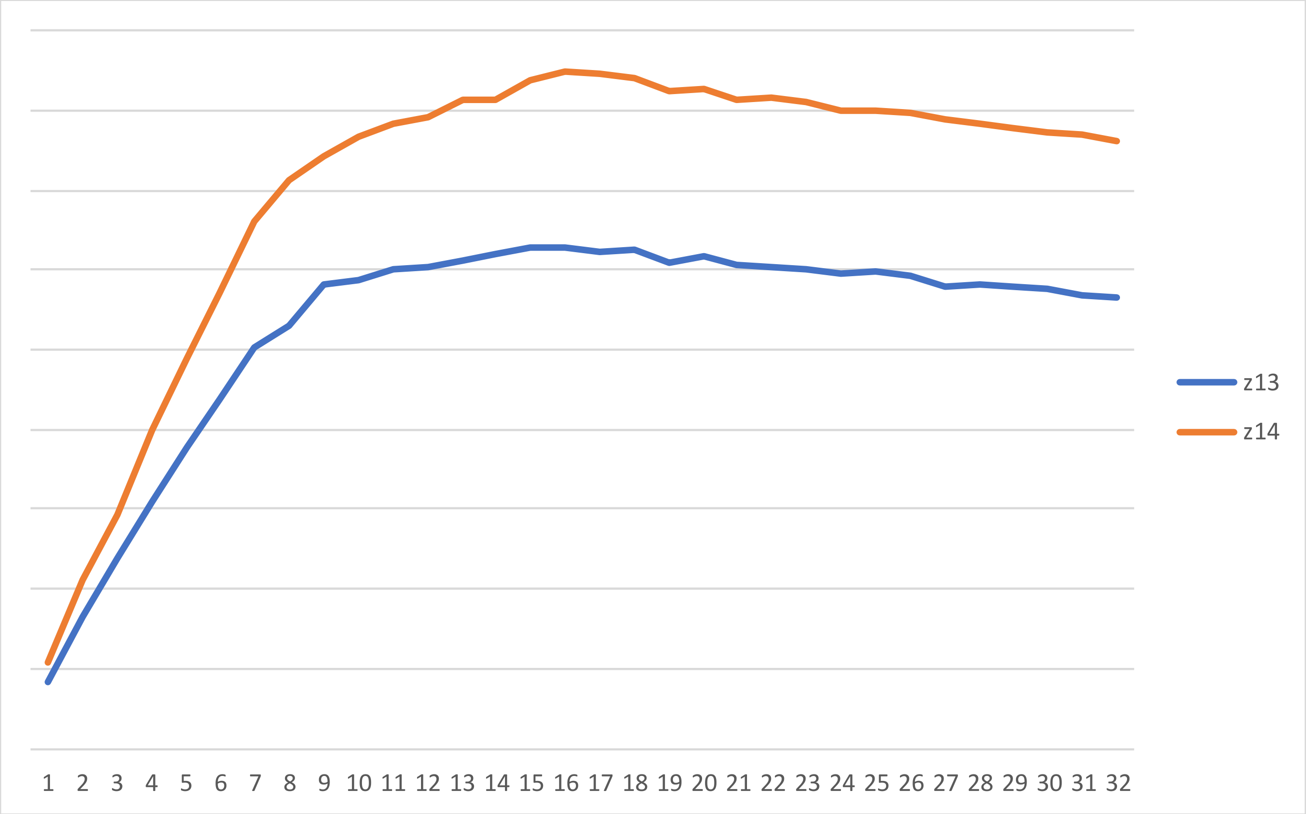

- CPU bound benchmark
- “old” Java code (JVM not exploiting z14)
- SLES 12 SP2+
- 8 cores – SMT enabled
- Overall: +35% on average

IBM z14: Performance that Changes the Game for Security

Performance with Integrated Cryptographic Hardware

– *6x* faster encryption for like modes and data sizes with enhanced on-chip (CPACF) cryptographic performance compared to z13[1]

– *2X* the SSL handshake performance on z14 with Crypto Express6S compared to z13 with Crypto Express5S[1]

Datasets automatically protected with z/OS Dataset Encryption

– Protect z/OS data sets[2] automatically throughout their life cycle

– Enforce consistent policy over access to encrypted content

Protection in the sysplex

– Data is encrypted/decrypted at a host and is protected in flight and at rest inside the Coupling Facility (CF)[3]

[1] Based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13. Results may vary.
[2] On October 4th, 2016 IBM announced a Statement of Direction to deliver z/OS dataset encryption capability in z/OS V2.2 (Announcement Letter)
[3] IBM z/OS 2.3 Preview Announcement Letter

CPACF relative performance (OpenSSL AES-256 encryption)

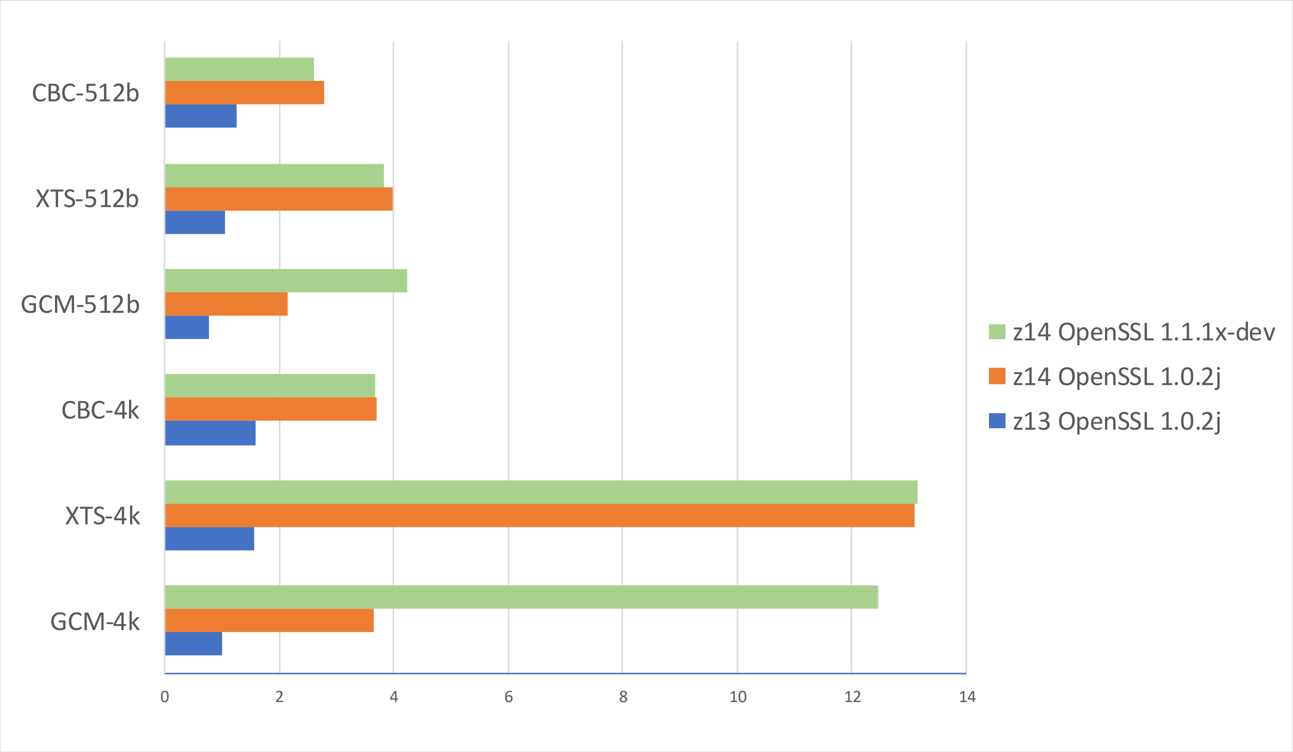

- SLES 12.2+, single core
- OpenSSL speed with 4k and 512b payloads
- All numbers relative to GCM-4K
- OpenSSL 1.0.2j as included with SLES 12.2+
- OpenSSL 1.1.1x development version
- Larger payload helps throughput

Network throughput improvements with Hipersockets (MTU 8k)

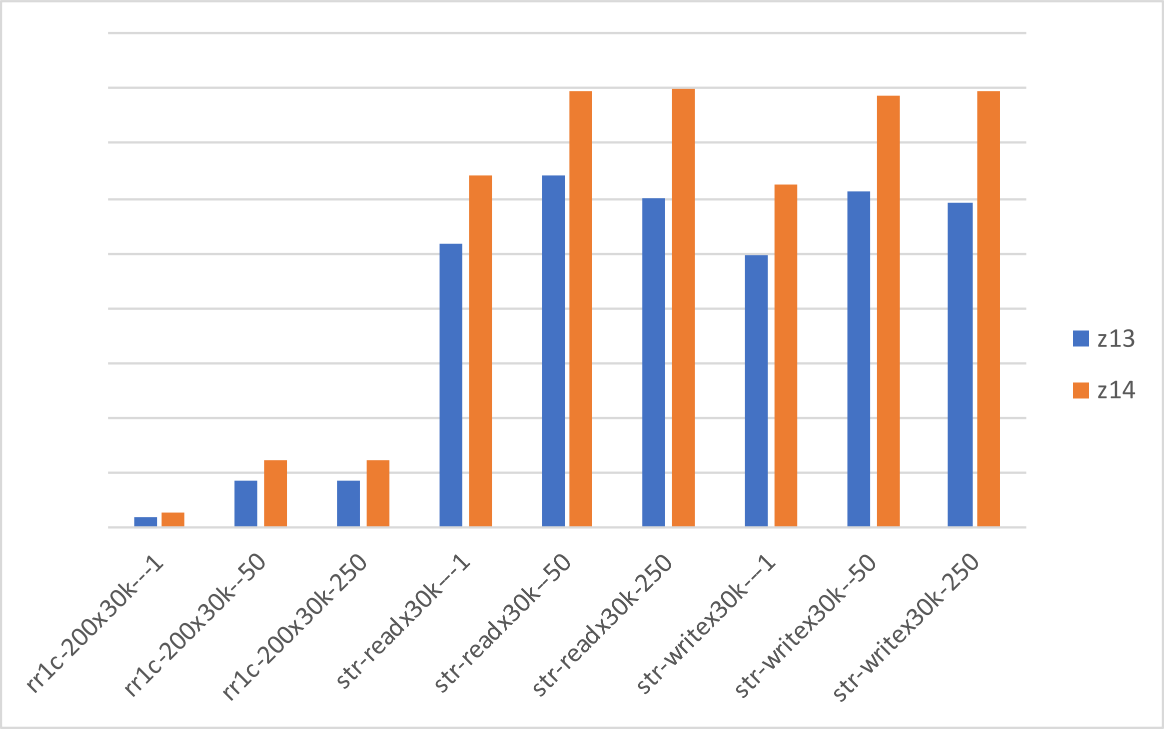

- SLES 12 SP2+, SMT2
- Linux to Linux connection cross drawer
- On average ~20% improvement
  - Faster CPU
  - Improved microcode

OSA-Express6S 1000BASE-T Ethernet feature

NEW

- PCIe form factor feature supported by PCIe I/O drawer
  - One two-port CHPID per feature
  - Half the density of the OSA-Express3 version
- Small form factor pluggable (SFP+) transceivers
  - Concurrent repair/replace action for each SFP
- Exclusively Supports: Auto-negotiation to 100* or 1000 Mbps and full duplex only on Category 5 or better copper
  - No 10Mbps
  - RJ-45 connector
  - Operates at “line speed”

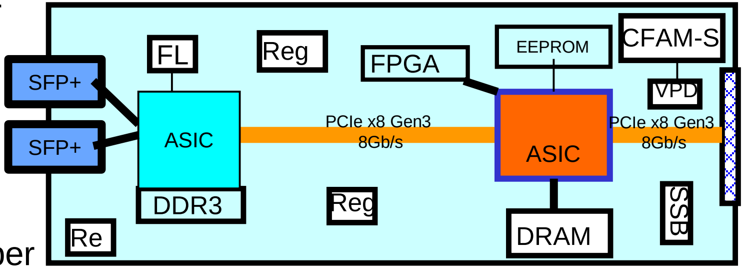

FC #0426

Connector = RJ-45

CHPID TYPE Support:

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

* OSA-Express6S 1000BASE-T adapters (#0426) will be the last generation of OSA 1000BASE-T adapters to support connections operating at 100 Mb/second link speed. Future OSA-Express 1000BASE-T adapter generations will support operation only at 1000 Mb/second (1Gb/s) link speed.

Network throughput improvements with OSA (MTU 1500)

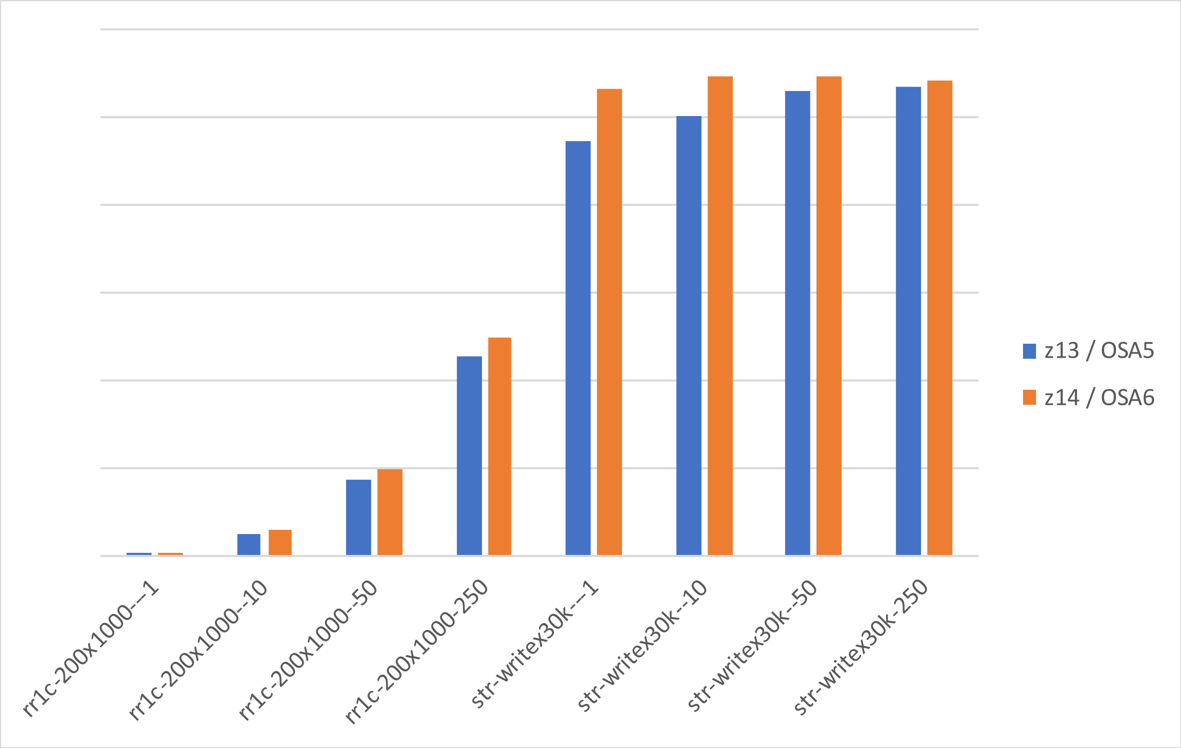

- SLES 12 SP2+, SMT2
- Improvements mainly due to improved processor speed
- Low double digit for transactional workloads
- Single digit improvement for streaming (reaching bandwidth limit of 10 Gb card earlier)

RoCE - What is it and Why we do it?

What?

- IBM Z introduces the next generation of RoCE technology with the IBM 10GbE RoCE Express2
- The 10GbE RoCE Express2 provides a technology refresh for RoCE on IBM Z. Most of the technology updates are related to internal aspects of the RoCE (RNIC) architecture (e.g. Queue Pair related technology).

Why?

- RoCE is an evolving technology. It is critical to keep the IBM Z RoCE technology current within the industry.
- Technology currency provides many improvements in the base technology that will provide benefits for application workloads.

Benefits

- Technology currency allows Z customers to benefit from the latest advancements in RoCE architecture, technology and specifications.
- The 10GbE RoCE Express2 provides two physical 10GbE ports (no change)
- Key difference: RoCE Express2 provides increased virtualization (sharing) capabilities allowing RoCE to be extended to more workloads:
  - RoCE Express2 supports 63 Virtual Functions (VFs) per physical port for a total of 126 VFs per PCHID.
  - RoCE Express supports 31 Virtual Functions (VFs) per PCHID.

Network throughput improvements with RoCE

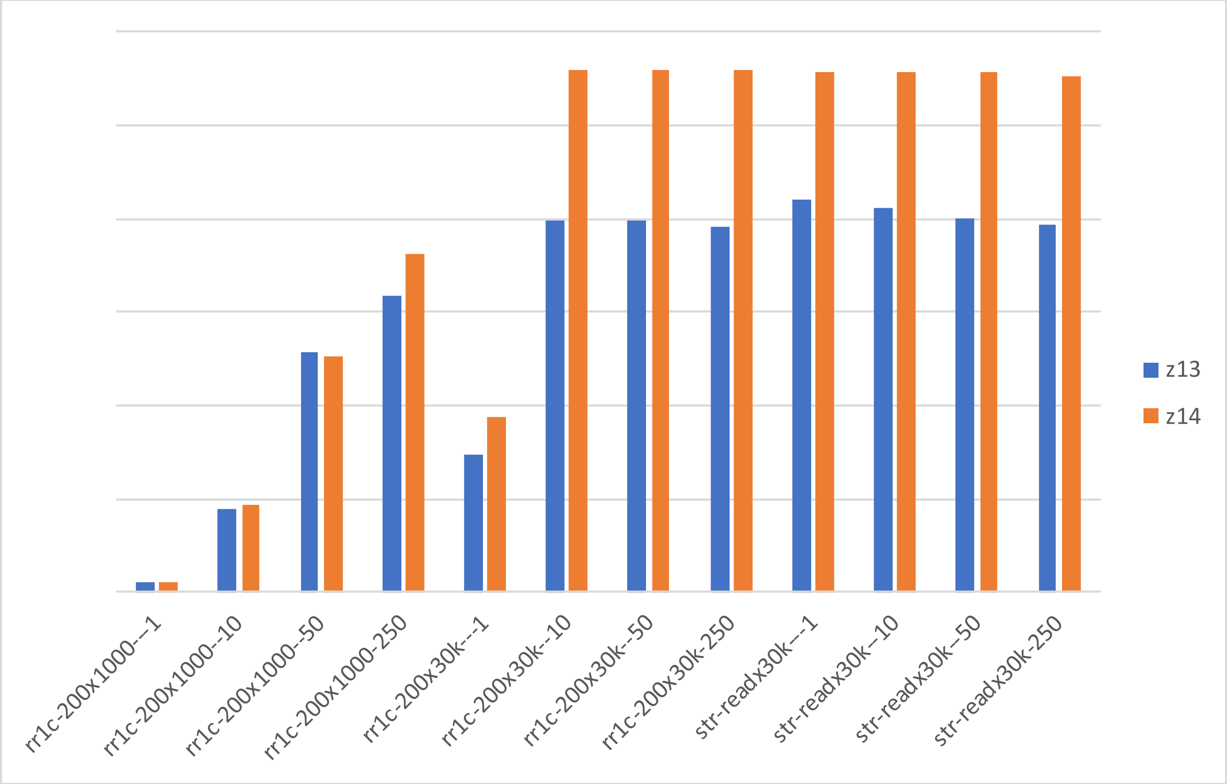

- SLES 12 SP2+, SMT2
- z13 RoCE1 vs z14 RoCE2
- Linux to Linux connection with cards cross drawer
- Bandwidth improvement ~1/3 if data is being moved
  – Faster CPU
  – Better cards
  – Better PCI

FICON Express16S+

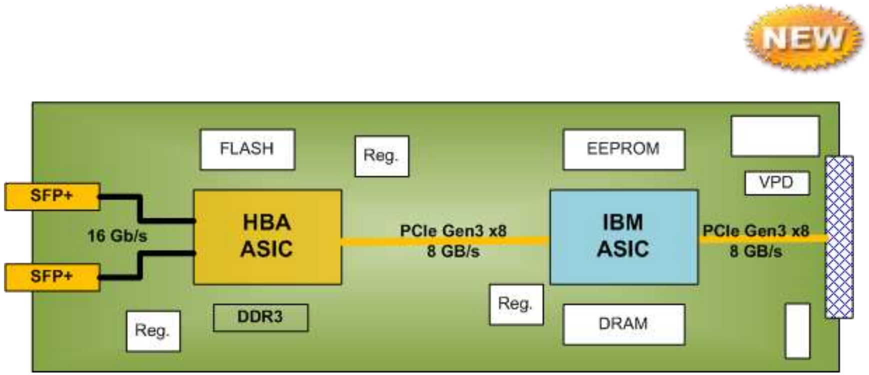

- **For FICON, zHPF, and FCP**
  - CHPID types: FC and FCP
  - *Both ports must be same CHPID type*
    - 2 PCHIDs / CHPIDs
- **Auto-negotiates to 4, 8, or 16 Gbps**
  - 2 Gbps connectivity not supported
  - FICON Express8S will be available for 2Gbps (carry forward only)
- **Increased performance compared to FICON Express16S**
- **Small form factor pluggable (SFP) optics**
  - Concurrent repair/replace action for each SFP
  - 10KM LX - 9 micron single mode fiber
  - SX - 50 or 62.5 micron multimode fiber
- **2 channels of LX or SX (no mix)**

FC #0427 – 10KM LX, FC #0428 – SX

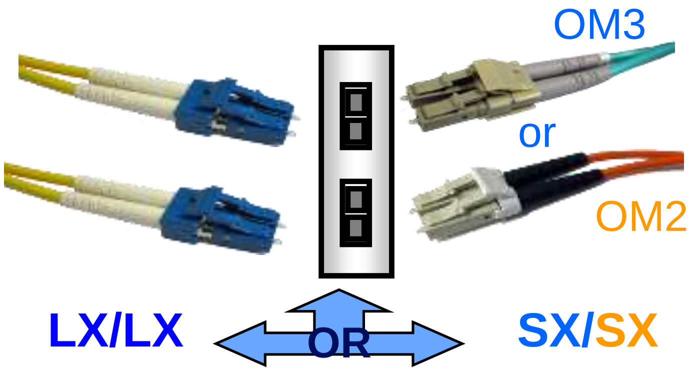

zHPF and FICON Performance* z14

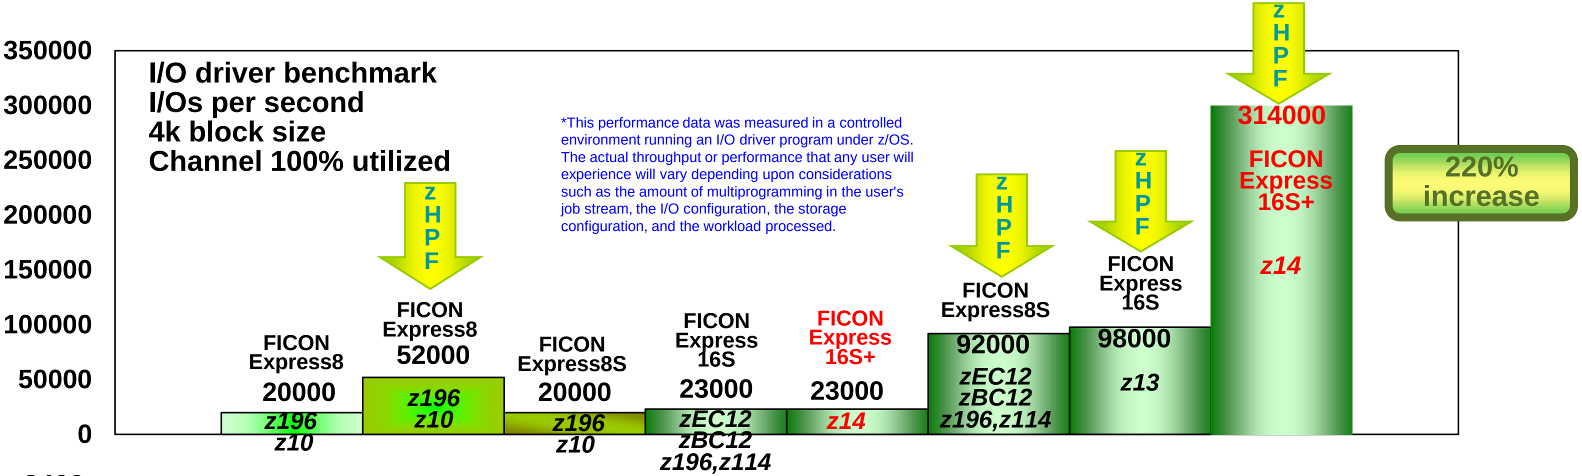

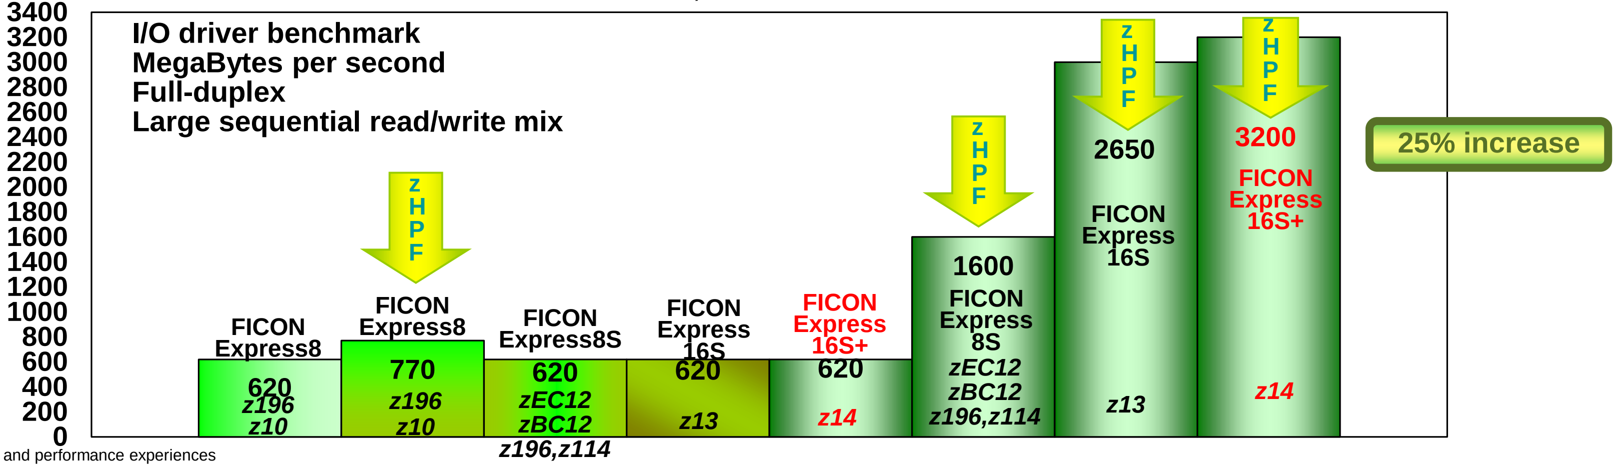

FCP Performance* for z14

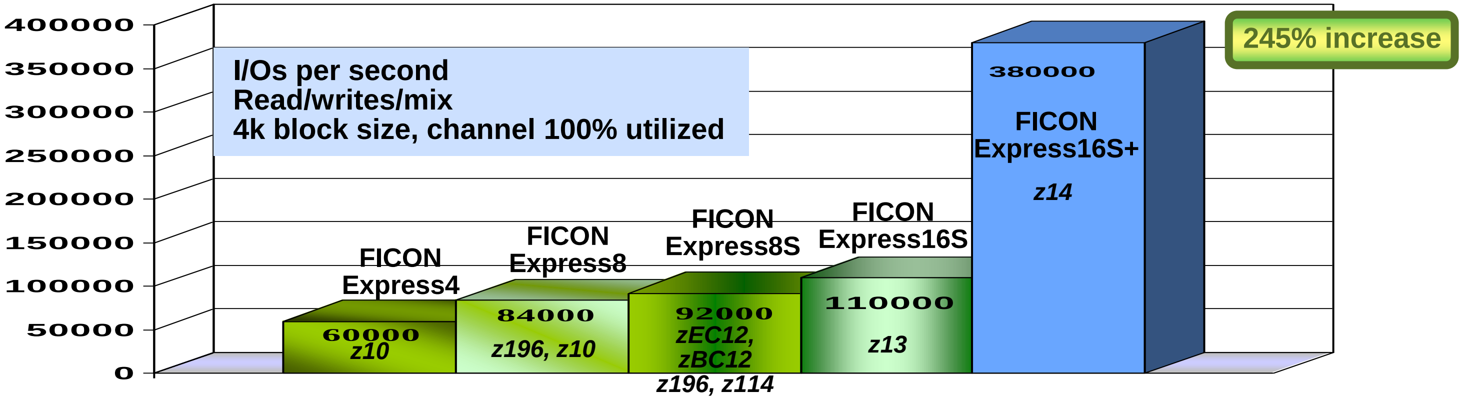

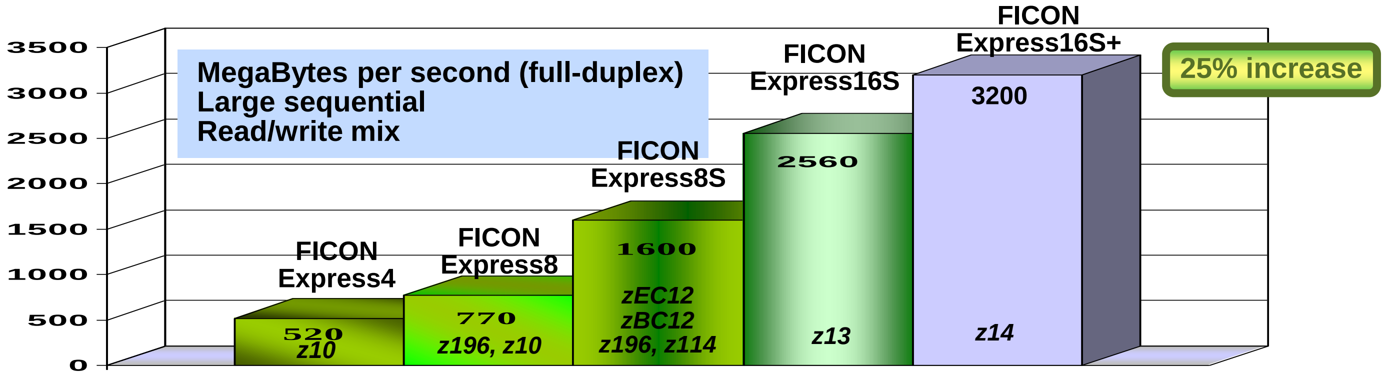

*This performance data was measured in a controlled environment running an I/O driver program under z/OS. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration,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

Improved adapter IOPS - implications

- Revisit your SAN and IOCDS
  – One host adapter per CHPID is not enough!

- Spreading your workload is even more important now
  – IOPS limit lifted – what’s the next bottleneck in your environment?

- Less adapters for same throughput
  – If your limit has been IOPS
  – There is no bandwidth increase!

Agenda

- z14 hardware overview
- Linux z14 support and certification
- Linux z14 exploitation (future)
- z14 hardware features & performance
- **z14 performance comparisons**

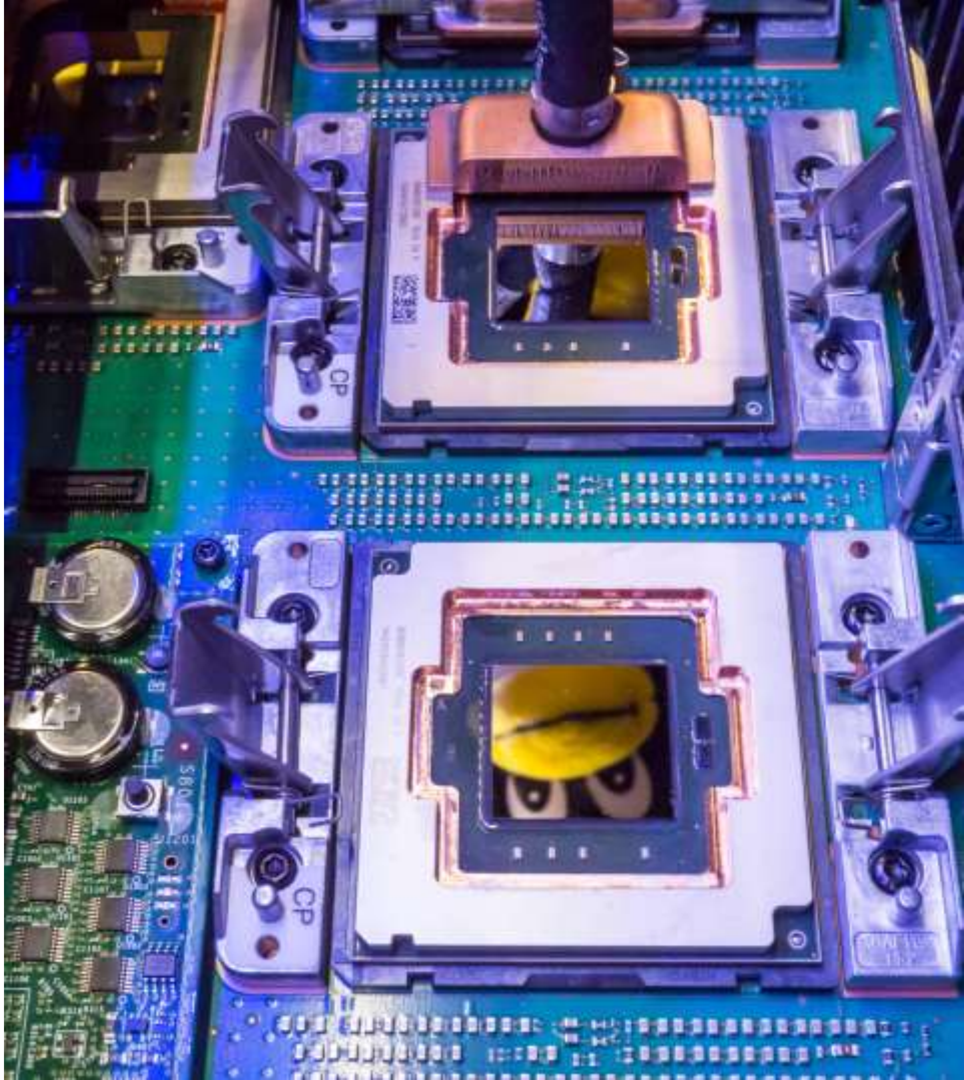

PostgreSQL Performance on z14 vs x86 Broadwell – Benchmark Configuration

- **Benchmark Setup**
  - Ran pgBench workload driver locally with 32 concurrent threads
    - read-only (100% read)
    - write-only (100% write)
  - Database size 20 GB
- **System Stack**
  - z14
    - LPAR with 2-16 dedicated IFL, 64 GB memory, and 40 GB DASD storage running SLES 12 SP2 with SMT enabled
    - PostgreSQL 9.6.1, pgBench 9.6
  - x86
    - 2-16 Intel E5-2697 v4 cores @ 2.30GHz w/ Hyperthreading turned on, 64 GB of memory, and 500 GB local RAID-5 HDD storage running SLES 12 SP2
    - PostgreSQL 9.6.1, pgBench 9.6

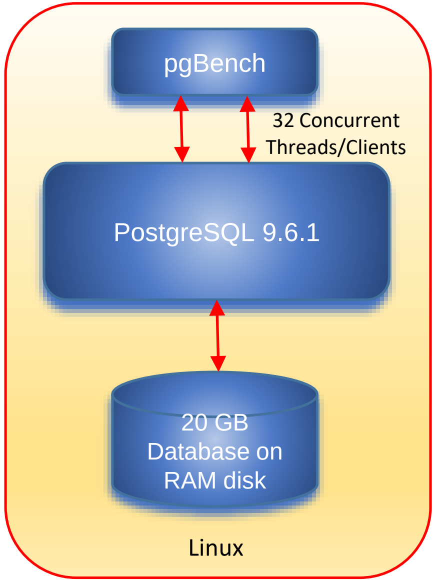

PostgreSQL Performance on z14 vs x86 Broadwell

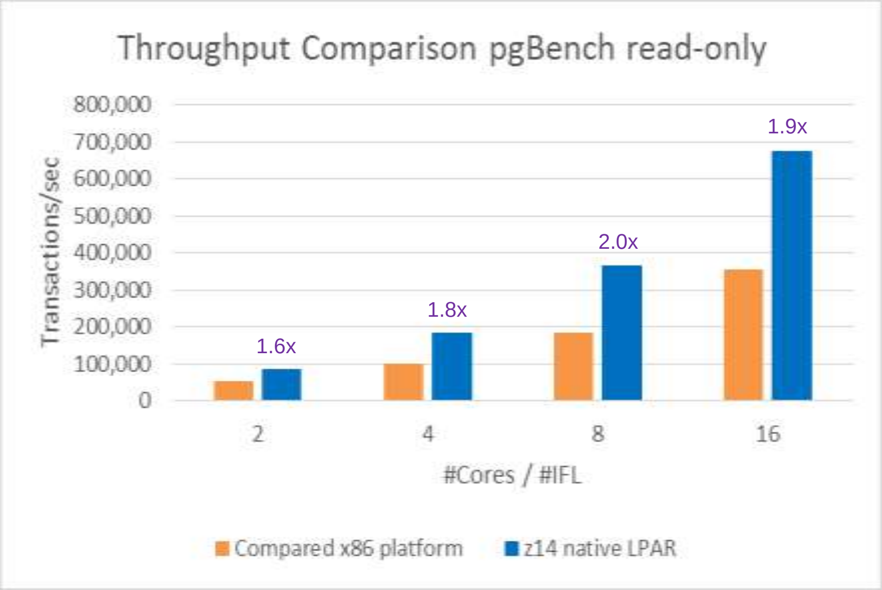

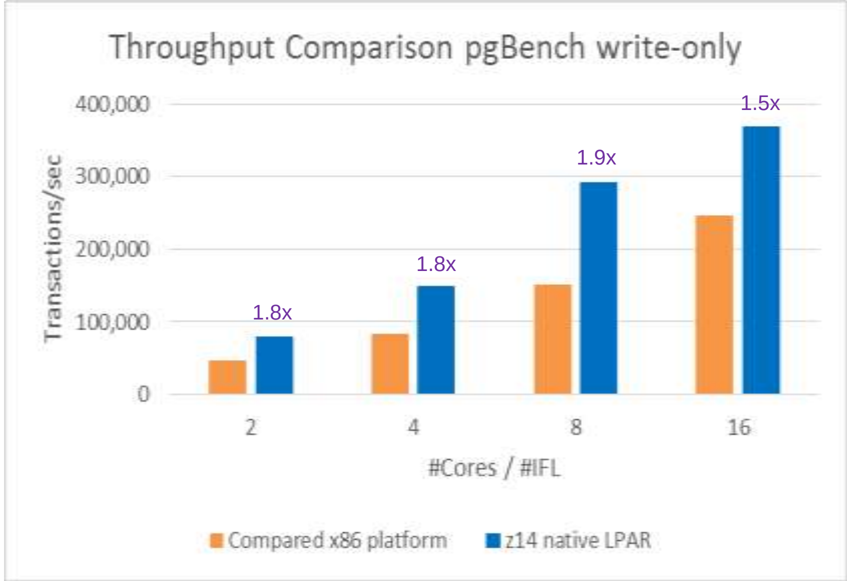

Run the pgBench benchmark on PostgreSQL 9.6.1 with up to 2x more throughput per core on a z14 LPAR versus a compared x86 platform

Disclaimer: Performance result is extrapolated from IBM internal tests running pgbench 9.6 benchmark on PostgreSQL 9.6.1 (20 GB database in RAM disk). Results may vary. x86 configuration: 2-16 Intel E5-2697 v4 cores @ 2.30GHz with Hyperthreading turned on, 64GB memory, and 500 GB local RAID-5 HDD storage, SLES12 SP2. z14 configuration: LPAR with 2-16 dedicated IFLs, 64GB memory, and 40 GB DASD storage, SLES12 SP2 (SMT mode).

MongoDB Performance on z14 vs x86 Broadwell – Benchmark Configuration

- **Benchmark Setup**
  - Ran YCSB workload driver locally
    - read-only (100% read)
    - write-heavy (50% write)
  - Database size 5 GB
- **System Stack**
  - z14
    - LPAR with 36 dedicated IFLs (2-8 IFLs dedicated to MongoDB, 20-28 IFLs dedicated to YCSB), 64 GB memory, and 120 GB DASD storage running SLES 12 SP2 with SMT enabled
    - MongoDB 3.4.1, YCSB 0.11.0
  - x86
    - 36 Intel E5-2697 v4 cores @ 2.30GHz w/ Hyperthreading turned on on (2-8 cores dedicated to MongoDB, 20-28 cores dedicated to YCSB), 64 GB of memory, and 480 GB local RAID-5 HDD storage running SLES 12 SP2
    - MongoDB 3.4.1, YCSB 0.11.0

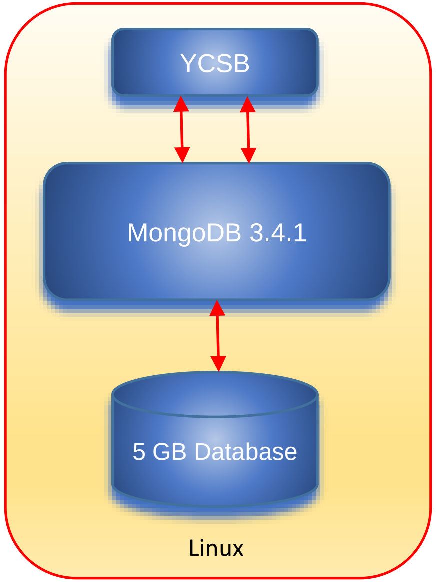

MongoDB Performance on z14 vs x86 Broadwell

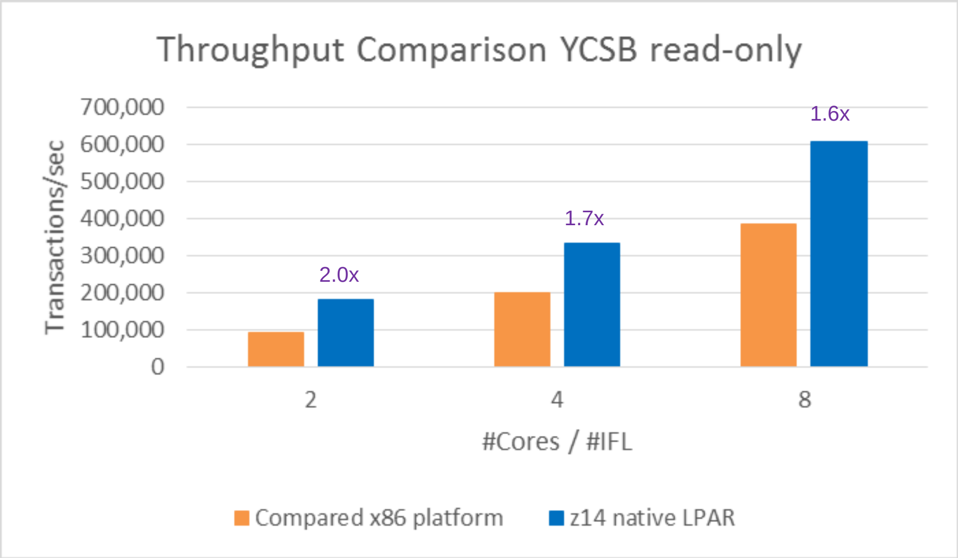

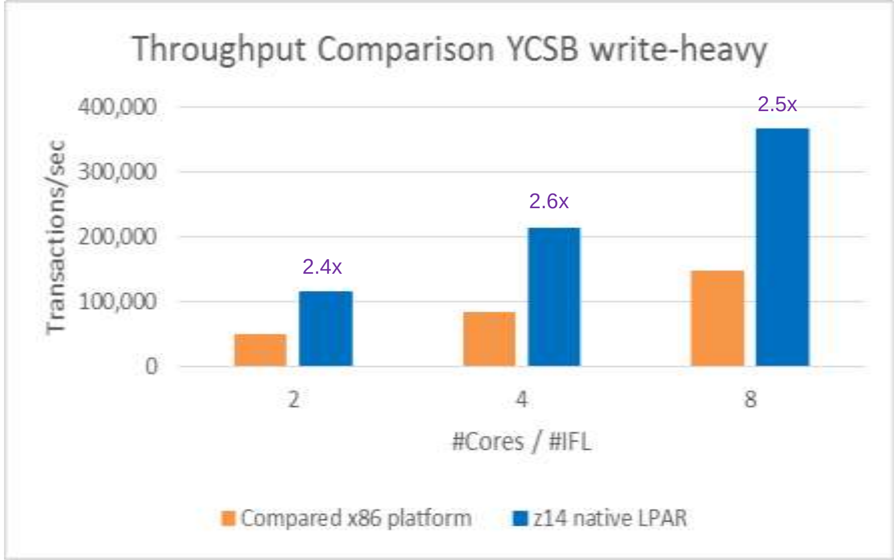

Run the YCSB benchmark on MongoDB 3.4.1 with up to 2.6x more throughput per core on a z14 LPAR versus a compared x86 platform

Disclaimer: Performance results based on IBM internal tests running YCSB 0.11.0 (write-heavy, read-only) on local MongoDB Enterprise Release 3.4.1 (Database size 5GB). Results may vary. x86 configuration: 36 Intel E5-2697 v4 cores @ 2.30GHz with Hyperthreading turned on (2-8 cores dedicated to MongoDB, 20 or 28 cores dedicated to YCSB), 64GB memory, and 480 GB local RAID-5 HDD storage, SLES12 SP2. z14 configuration: LPAR with 36 dedicated IFLs (2-8 cores dedicated to MongoDB, 20 and 28 cores dedicated to YCSB), 64GB memory, and 120 GB DASD storage, SLES12 SP2 (SMT mode).

Node.js Performance on z14 vs x86 Broadwell – Benchmark Setup

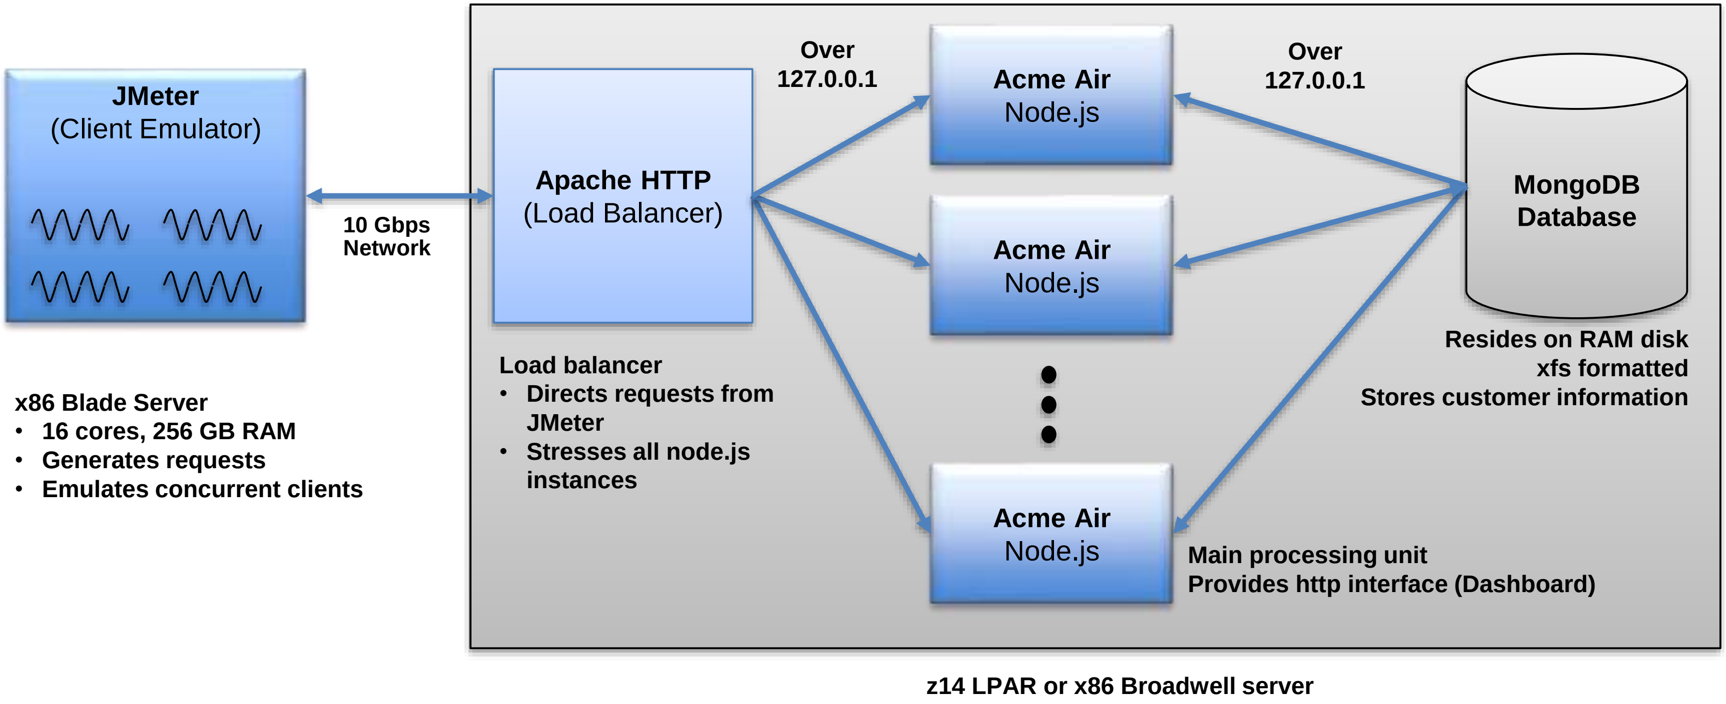

Node.js Performance on z14 vs x86 Broadwell

Run the Acme Air benchmark on node.js 6.10 with up to 2.5x more throughput per core on a z14 LPAR versus a compared x86 platform

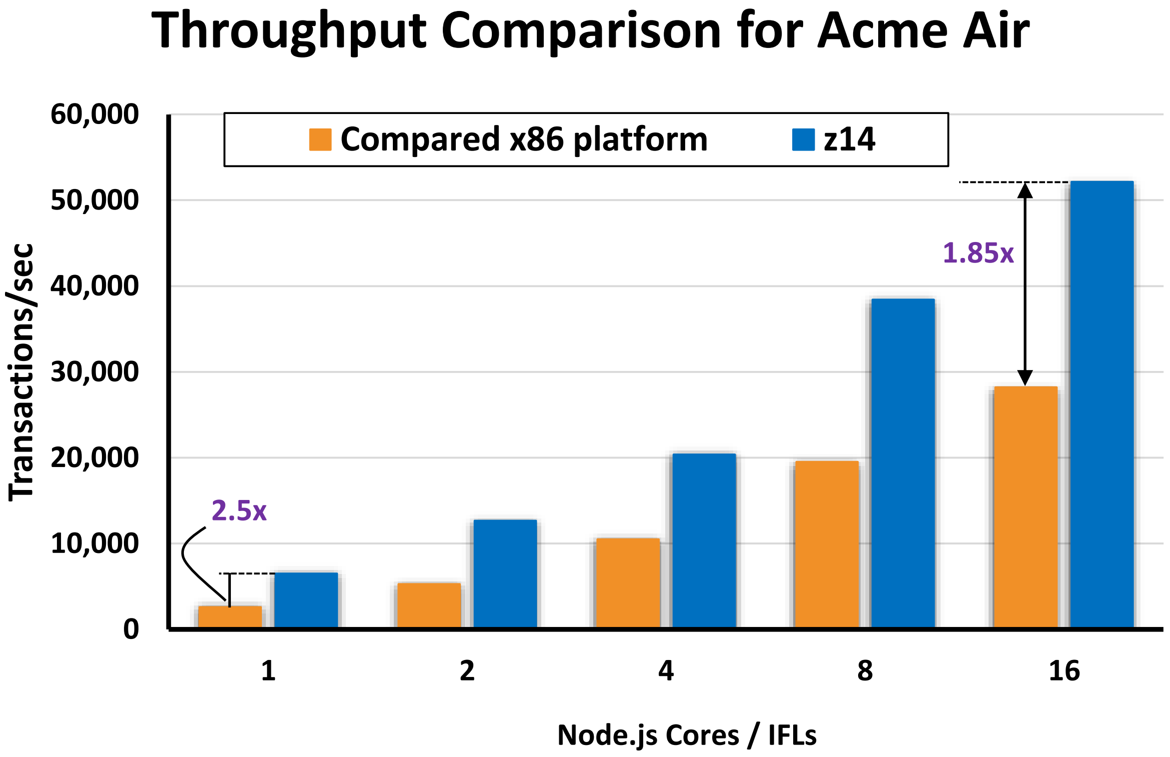

Disclaimer: Performance results based on IBM internal tests running Acme Air with 10,000 customers on Node.js v6.10.0 against MongoDB Enterprise 3.4.2 driven remotely by 250 JMeter 2.13 threads. Apache HTTP server 2.4.23 was used as load balancer. Results may vary. x86 configuration: 36 Intel E5-2697 v4 cores @ 2.30GHz, Apache HTTP server pinned to 1 core, Node.js pinned to 1-16 cores, MongoDB pinned to 2-4 cores, 768GB memory, SLES12-SP2 with Hyperthreading, application logs and database on the RAM disk. z14 configuration: LPAR with 32 dedicated IFLs, Apache HTTP server pinned to 1 IFL, Node.js pinned to 1-16 IFLs, MongoDB pinned to 2-4 IFLs, 768GB memory, 40 GB DASD storage, SLES12-SP2 with SMT, application logs and database on the RAM disk.

WebSphere Application Server Performance on z14 vs x86 Broadwell – Benchmark Configuration

- **Benchmark Setup**
  - DayTrader Benchmark (15000 Users, 10000 Stocks) (ftp://public.dhe.ibm.com/software/webservers/appserv/was/DayTrader3Install.zip)
  - ibm-java-x86_64-sdk-8.0-3.22
  - Two driving x86 server, each trading for 7500 users
  - 2-6 driver threads (channels) per WAS compute thread
- **System Stack**
  - z14
    - LPAR with 2-16 IFL, 64 GB memory running SLES12 SP2 with SMT enabled, DS8K DASD storage
    - WAS 8.5.5.9 with Java 8.0-3.22 pinned to half of the IFLs
    - DB2 11.1.1.1 pinned to half of the IFLs
  - x86
    - 2-16 Intel Xeon CPU E5-2697 v4 @ 2.30GHz, 1.5TB memory running SLES 12 SP2 with Hyperthreading enabled, local HDD storage
    - WAS 8.5.5.9 with Java 8.0-3.22 pinned to half of the cores
    - DB2 11.1.1.1 pinned to half of the cores

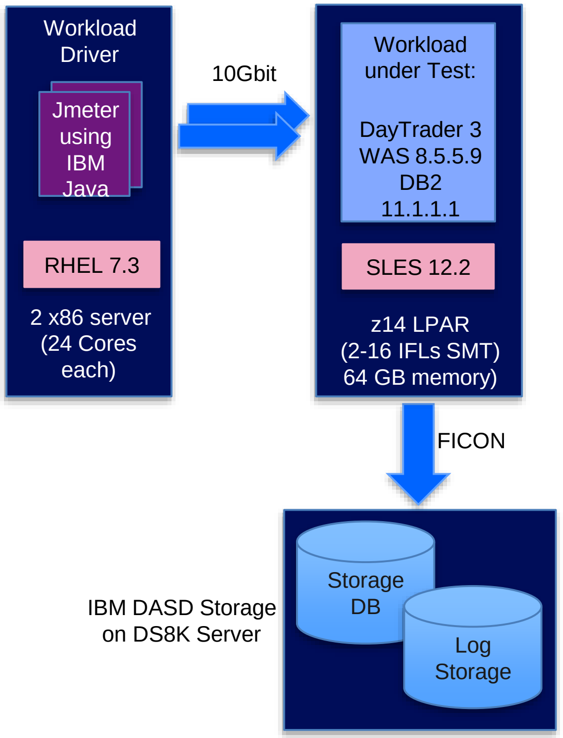

WebSphere Application Server Performance on z14 vs x86 Broadwell

Run the DayTrader benchmark on WebSphere Application Server 8.5.5.9 with up to 1.9x more throughput per core on a z14 LPAR versus a compared x86 platform

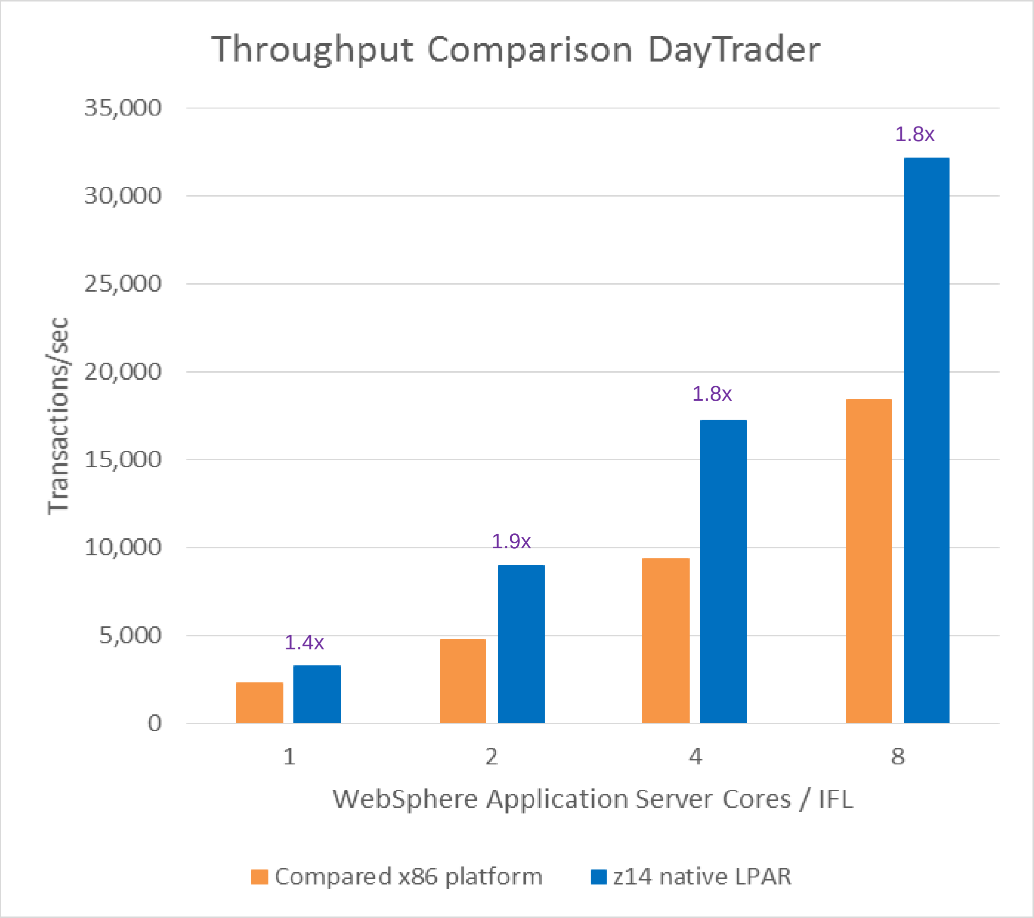

Disclaimer: Performance results based on IBM internal tests running Daytrader 3 web application benchmark on Websphere Application Server WAS 8.5.5.9 with IBM Java 1.8.0 (SR3). Database DB2 LUW 11.1.1.1 located on the same system was used to persist application data. Half of the compute cores for each system variation under test were bound to DB2, the other half to WAS. The workload was driven remotely by Apache JMeter to trade 10000 stocks among 15000 users. The utilization of the workload was adjusted by the number of driver threads. Results may vary. x86 configuration: 2-16 Intel(R) Xeon(R) CPU E5-2697 v4 @ 2.30GHz, 1.5TB fast TruDDR4 2400MHz Memory, and 400GB local HDD storage, SLES12 SP2 with Hyperthreading enabled. z14 configuration: LPAR with 2-16 IFLs, running under SLES12 SP2 (SMT mode), 64GB memory, 80GB DASD storage, HyperPAV=8.

DB2 LUW Performance with FICON Express16S+ Cards – Benchmark Configuration

- **Benchmark Setup**
  – BDI workload driver based on TPC-DS
  – 8 parallel users performing predefined SQL Queries
  – DB2 database size 500 GB
- **System Stack**
  – z13
    - LPAR with 8 dedicated IFLs and 64GB memory running RHEL 7.3 with SMT enabled
    - 11 TB LUN on IBM FlashSystem 900 attached via FICON Express16S card
    - DB2 LUW 11.1.1, IBM Java 1.8
  – z14
    - LPAR with 8 dedicated IFLs and 64GB memory running RHEL 7.3 with SMT enabled
    - 11 TB LUN on IBM FlashSystem 900 attached via FICON Express16S+ card
    - DB2 LUW 11.1.1, IBM Java 1.8

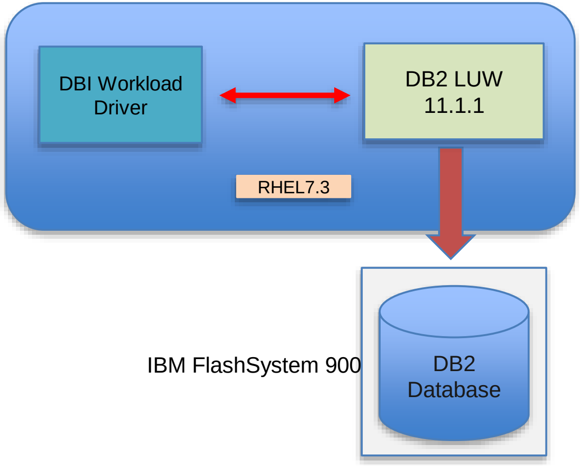

DB2 LUW Performance with FICON Express16S+ Cards

Run the BDI benchmark on DB2 LUW 11.1.1 with up to 20% more throughput using FICON Express16S+ cards on z14 compared to using FICON Express16S cards on z13

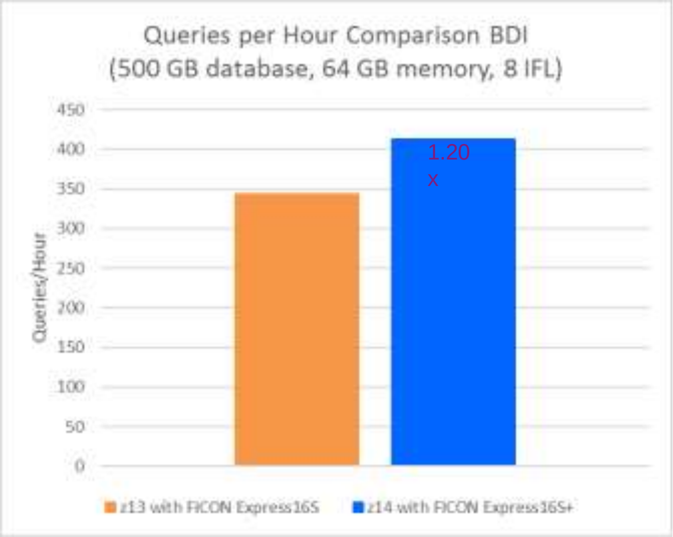

<u>Disclaimer:</u> Performance results based on IBM internal tests running the BDI benchmark, which is based on TPC-DS, on DB2 LUW with BLU Acceleration. The BDI benchmark was configured to run a fixed sequence of queries. DB2 database size was 500 GB. Results may vary. z13 configuration: LPAR with 8 dedicated IFLs, 64GB memory, and 11 TB LUN on IBM FlashSystem 900 attached via FICON Express16S cards, RHEL 7.3 (SMT mode) running DB2 LUW 11.1.1, IBM Java 1.8, and BDI. z14 configuration: LPAR with 8 dedicated IFLs, 64GB memory, and 11 TB LUN on IBM FlashSystem 900 attached via FICON Express16S+ cards, RHEL 7.3 (SMT mode) running DB2 LUW 11.1.1, IBM Java 1.8, and BDI.

MongoDB Consolidation under z/VM on z14 versus z13

Run 25% more MongoDB guests with the same throughput under z/VM 6.4 on z14 compared to z13

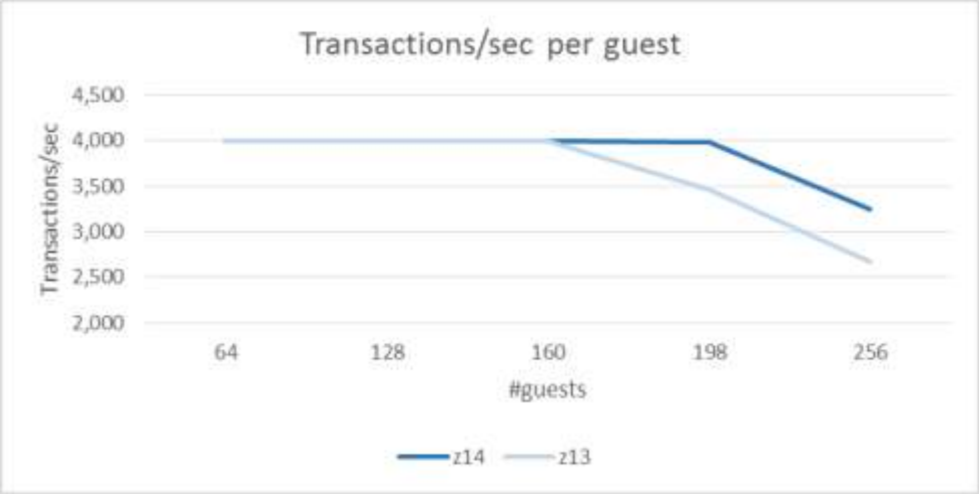

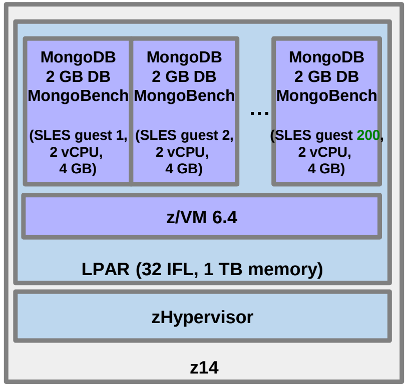

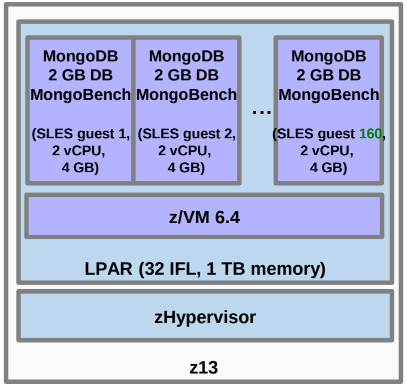

<u>Disclaimer</u>: Performance result based on IBM internal tests comparing MongoDB performance under z/VM 6.4 with the PTF for APAR VM65942 on z14 versus z13 driven locally by MongoBench (https://github.com/linux-on-ibm-z/mongo-bench) issuing 90% read and 10% write operations. Results may vary. z14 configuration: LPAR with 32 dedicated IFLs and 1 TB memory running a z/VM 6.4 with the PTF for APAR VM65942 instance in SMT mode with 200 guests. Each guest was configured with 2 vCPUs and 4 GB memory and ran a MongoDB Enterprise Server 3.4.1 instance (no sharding, no replication) with a 2 GB database. The databases were located on a FCP-attached DS8700 LUN with multi-pathing enabled. z13 configuration: LPAR with 32 dedicated IFLs and 1 TB memory running a z/VM 6.4 with the PTF for APAR VM65942 instance in SMT mode with 160 guests. Each guest was configured with 2 vCPUs and 4 GB memory and ran a MongoDB Enterprise Server 3.4.1 instance (no sharding, no replication) with a 2 GB database. The databases were located on a FCP-attached DS8700 LUN with multi-pathing enabled.

Summary

- Toleration for Linux is here, exploitation is coming
- Good performance improvements with z14 compared to last hardware generation
- z14 performance is more predictable than z13

One year hard work!