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

IBM

Experiences, People and Ideas Converge to Power Business Outcomes
Celebrating 60 Years of SHARE

March 1-6
Sheraton Seattle
Seattle, WA

SHARE in Seattle 2015

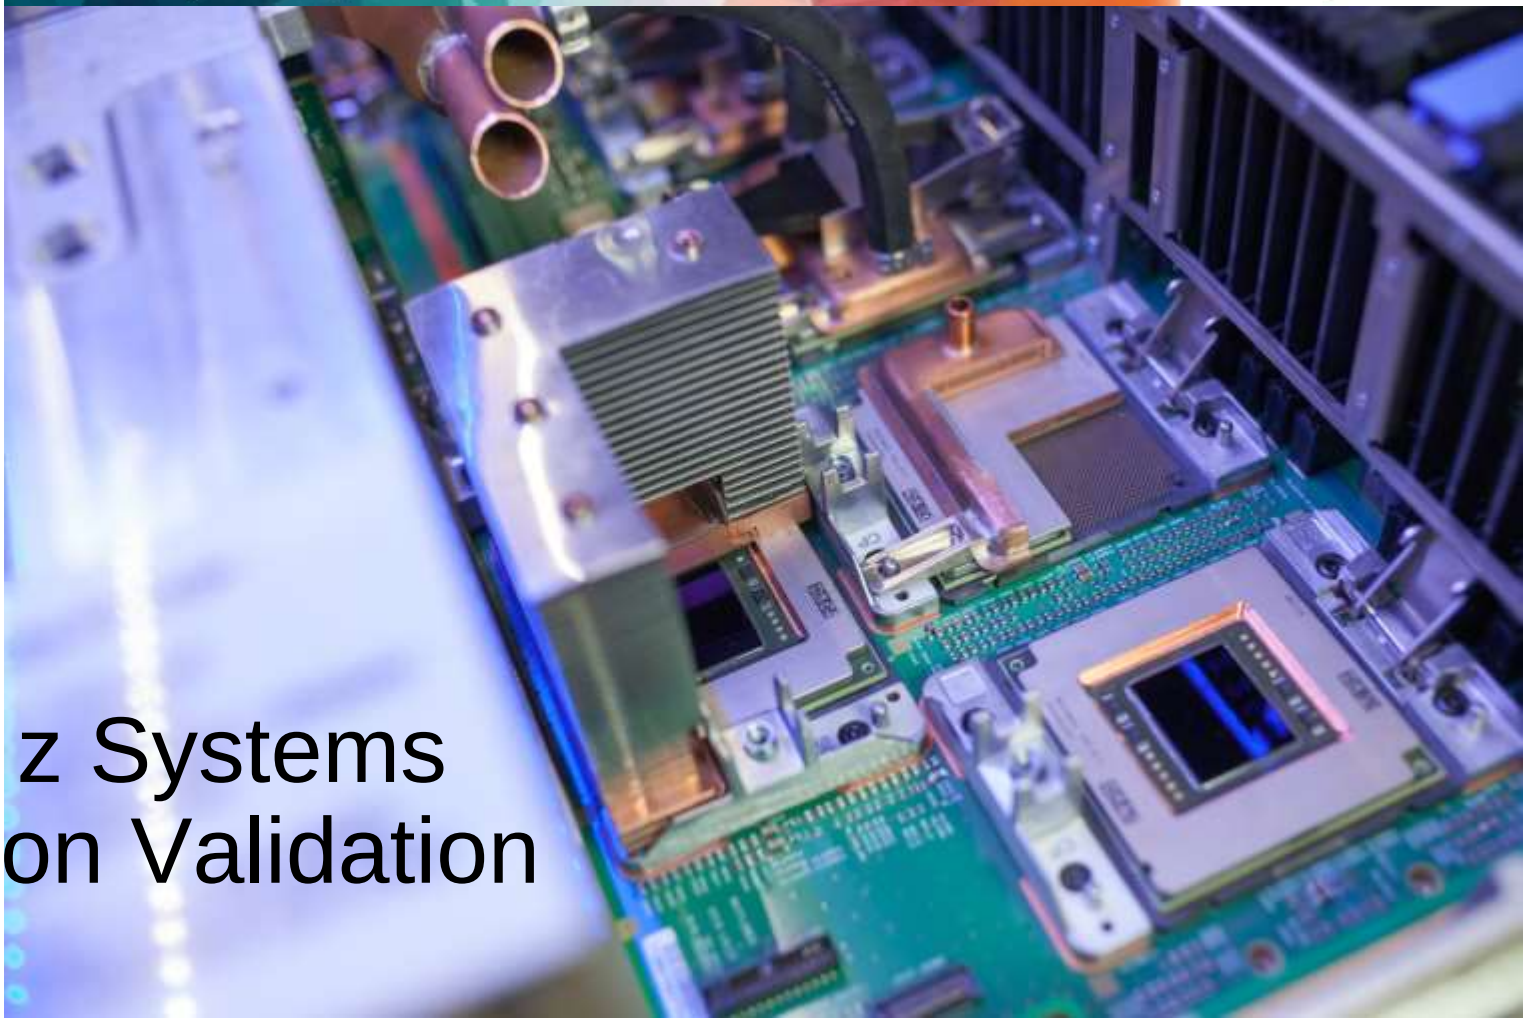

Linux on z Systems Distribution Validation for z13

Trademarks

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

BlueMix
BigInsights
Cognos*
DB2*
DB2 Connect
Domino*
DS8000*
ECKD
FICON*
FileNet*
FlashSystem
GDPS*
GPFS
IBM*
Ibm.com
IBM (logo)*
IMS
Informix*
InfoSphere
Maximo*
MQSeries*
Performance Toolkit for VM
POWER*
Quickr*
Rational*
Sametime*
Smarter Cities*
Smarter Analytics
SPSS*
Storwize*
System Storage*
Tivoli*
WebSphere*
XIV*
z13
zEnterprise*
z/OS*
z Systems
z/VSE*
z/VM*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

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.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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.

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.

* 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. **All z13 numbers have been measured on pre GA hardware with pre GA software.**

Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here. All z13 numbers have been measured on pre GA hardware.

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

- **z13 (Linux) overview**
- Distribution support timelines
- z13 Linux testing
- IBM tested platforms
- Distribution certifications
- Outlook and recommendations

Linux on IBM z13 (z13)

- **Data center simplicity inside one server**
- **Trusted operations**
- **Unrivaled economics**

LPAR	**Logical Partition** = subset of hardware resources, virtualized as a separate computer; up to 85 LPARs can be configured
IFL	**Integrated Facility for Linux = core**; up to 141 cores (IFLs) on IBM z13™ (z13)
Virtual. Mgmt.	Hypervisor providing efficiency at scale and virtualization management for easy administration, provisioning, automation
Linux Guest	virtual Linux guests running workloads such as mobile, analytics, databases, Java™ apps, etc. – in a cloud; up to thousands Linux guests can be hosted on a single z13

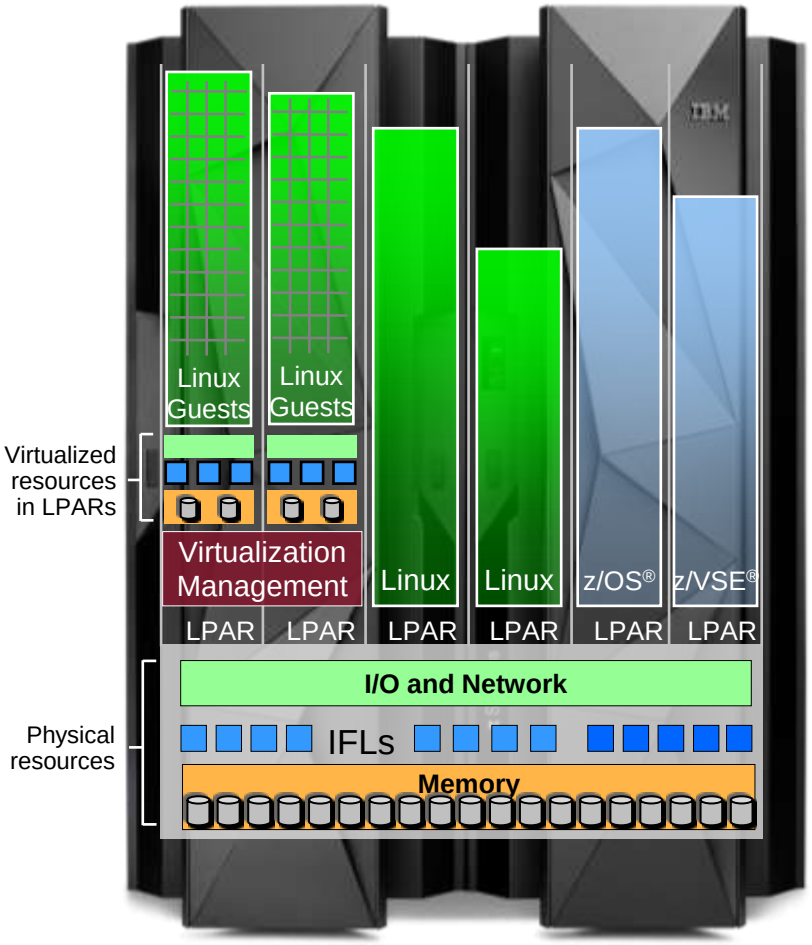

z13 System Design Changes

- 22nm Processor with SIMD, SMT
- Integrated I/O with PCIe Direct Attach
- Single Chip Modules
- Drawer-Based CPC Design
- Cable-Based SMP Fabric
- Oscillator Backplane
- Flexible Service Processor (FSP2)
- Integrated Sparing
- On-chip power/thermal monitor / control

- New Memory Controller
- Crypto Express5S
- FICON Express16S
- 1U Support Element
- Standalone zBX Node Hybrid Computing
- 2.7M lines of firmware changed
- Radiator Design improvements
- Expanded operating environment (Rear Doors)

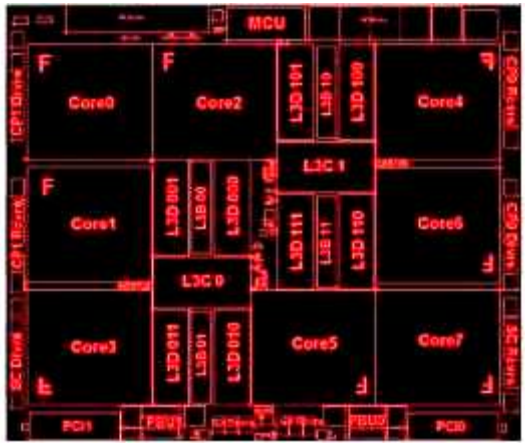

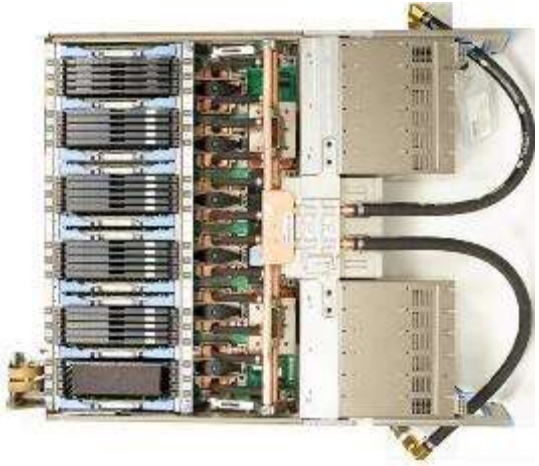

Agenda

- z13 (Linux) overview
- **Distribution support timelines**
- z13 Linux testing
- IBM tested platforms
- Distribution certifications
- Outlook and recommendations

Distribution support timeline (Red Hat)

* The life cycle time spans and dates are subject to adjustment.

Release	GA	End of production 3
RHEL5	3/15/2007	3/31/2017
RHEL6	10/11/2010	11/30/2020
RHEL7	6/10/2014	6/30/2024

General policy is that customers need to upgrade with the release of a new minor version

Red Hat offers Extended Update Support for certain distributions:

- 5.9 (ends March 31, 2015)
- 6.5 (ends November 30, 2015)
- 6.6 (ends October 31, 2016)
- 7.0 (N/A)

http://www.redhat.com/security/updates/errata/

Distribution support timeline (SUSE)

- SLES 10 – general suppport end 31 Jul 2013

Service Pack Release	FCS Date	General Ends
SUSE Linux Enterprise Server 10	17 Jul 2006	31 Dec 2007
SUSE Linux Enterprise Server 10 SP1	18 Jun 2007	31 Nov 2008
SUSE Linux Enterprise Server 10 SP2	19 May 2008	11 Apr 2010
SUSE Linux Enterprise Server 10 SP3	12 Oct 2009	11 Oct 2011
SUSE Linux Enterprise Server 10 SP4	12 Apr 2011	31 Jul 2013

- SLES 11 – general suppport end 31 Mar 2019

Service Pack Release	FCS Date	General Ends
SUSE Linux Enterprise Server 11	24 Mar 2009	31 Dec 2010
SUSE Linux Enterprise Server 11 SP1	02 Jun 2010	31 Aug 2012
SUSE Linux Enterprise Server 11 SP2	29 Feb 2012	31 Jan 2014
SUSE Linux Enterprise Server 11 SP3	01 Jul 2013	TBD

- SLES 12 – general suppport end 31 Oct 2024

- https://www.suse.com/lifecycle/

Agenda

- z13 (Linux) overview
- Distribution support timelines
- **z13 Linux testing**
- IBM tested platforms
- Distribution certifications
- Outlook and recommendations

Test Areas for a new hardware

- **Goal:** find bugs that are being introduced with the new hardware
- Side effect: find additional distribution bugs
- Manual installation tests
- Regression tests including error injects for
  - Kernel
  - Memory
  - Crypto
  - ECKD
  - FCP
  - Networking
- Stress test
- Limit test
- Certification support

Test Challenges

- z13 hardware is all new
- 5 distros (at defined kernel levels): RHEL 5.11, RHEL 6.6, RHEL 7.1, SLES 11.3, SLES 12.0
- Many Regression and manual tests
- z/VM and LPAR
- moving microcode
- interrupts for updates, repros etc
- fixing cycles
- ....

→ Testing everything is impossible

→ Need automation (home grown)
and intelligent testing (IBM Functional Coverage Unified Solution)
http://researcher.watson.ibm.com/researcher/view_group.php?id=1871

Types of bugs found – some examples on the next pages

- Hardware bugs (usually in early stages)
  - none known left ☺
- Firmware bugs
  - pay close attention to the first set of MCL bundles after GA
  - mostly in the areas not widely covered by z/OS
- Limits of Linux code exposed with the larger machine
- Race conditions in Linux code due to new timing
- Toleration enablement needed to even access the hardware

Toleration Enablement – CEX5

- toleration code needed
  - RHEL 5.11: next zStream update
  - RHEL 6.6: next zStream update
  - RHEL 7.1: next zStream update
  - SLES 11 SP3: next maintenance update
  - SLES12: 3.12.36-38.1

- if needed earlier please contact your distribution partner

Limit Example

- /proc/sysinfo used only one page. With more CPUs (IFL, CP, ZIIP, ZAAP, ...) in the system this is no longer enough.
- All systems with 4 drawers are exposed (it somewhat depends on the naming and if z/VM is in the picture or not)
- Reading /proc/sysinfo on exposed systems leads to a memory overwrite
- Fix is already released and included in the following kernel levels:
  - RHEL 5.11: 2.6.18-400.el5 (https://rhn.redhat.com/errata/RHSA-2014-1959.html)
  - RHEL 6.6: GA level
  - RHEL 7.0: System z GA level
  - SLES 11.3: 3.0.101-0.40.1
  - SLES 12.0: GA level

- **It is mandatory to apply the patches before upgrading the hardware!**
- **If you have a LTS contract for an older version, then please contact your service provider**

Example - New Race Condition

- IPL / reboot failure with a kernel panic due to failure in module loading
- a few percent of reboot fail, for small guests some say it can be neglected
- with larger the #of CPU the probability increases and it‘s definitely a real problem
- Implication:
  - HA failover / restart may fail

- For me it always failed when I could tolerate it the least
- RHEL 7.1 and SLES12 problem
- Fix available upstream:
  - http://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=4982223e51e8ea9d09bb33c8323b5ec1877b2b51
- Fix coming for distributions asap

Example – common code problems

- Stress test was targeted at 3 different file systems
  - XFS
  - btrfs
  - ext3/4

- The file systems are in various stages of support in the distributions under test

- Only ext3/4 survived all

- Opened distro bugs for the others, where supported

Agenda

- z13 (Linux) overview
- Distribution support timelines
- z13 Linux testing
- **IBM tested platforms**
- Distribution certifications
- Outlook and recommendations

IBM tested platform web page (1)

- IBM information is spublished on
  - http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html
- Watch for Updates there

Distribution	z13	zEnterprise - zBC12 and zEC12	zEnterprise - z114 and z196	System z10 and System z9
RHEL 7	✔ (1,3)	✔ (4)	✔ (4)	✘
RHEL 6	✔ (1,3)	✔ (5)	✔	✔
RHEL 5	✔ (1,3)	✔ (6)	✔	✔
RHEL 4 (*)	✘	✘	✔ (9)	✔
SLES 12	✔ (2,3)	✔	✔	✘
SLES 11	✔ (2,3)	✔ (7)	✔	✔
SLES 10 (*)	✘	✔ (8)	✔	✔
SLES 9 (*)	✘	✘	✔ (10)	✔

snapshot as of 2/3/2015

IBM tested platform web page (2) – foot notes

✔ Indicates that the distribution (version) has been tested by IBM on the hardware platform, will run on the system, and is an IBM supported environment. Updates or service packs applied to the distribution are also supported. Please check with your service provider which kernel-levels are currently in support.

(1) Red Hat Hardware Certification statements are available for RHEL 7.0, RHEL 6.6, and RHEL5.11 at:
https://hardware.redhat.com/&quicksearch=z13

The following kernel-levels are the currently known required minimum-levels for z13:

RHEL 7.0: 3.10.0-123.4.2.el7
RHEL 6.6: 2.6.32-504.8.1.el6
RHEL 5.11: 2.6.18-400.el5

These minimum required levels do not include Crypto Express5S (CEX5S) toleration support

(2) SUSE YES CERTIFIED Bulletins are available for SLES 12 and SLES 11 SP3 at:
https://www.suse.com/yessearch/SResults.jsp?bulletinNumber=&keywords=z13

The following kernel-levels are the currently known required minimum-levels for z13:

SLES 12: 3.12.28-4.6
SLES 11 SP3: 3.0.101-0.40.1

These minimum required levels do not include Crypto Express5S (CEX5S) toleration support

(3) When the Crypto Express5S toleration support is tested, the recommended kernel-levels will be published on this page.

snapshot as of 2/3/2015

Agenda

- z13 (Linux) overview
- Distribution support timelines
- z13 Linux testing
- IBM tested platforms
- **Distribution certifications**
- Outlook and recommendations

Hardware Certifications (Red Hat)

- Red Hat has a new portal for searching supported hardware:
  - https://access.redhat.com/ecosystem/search/#/category/Server
  - enter z13 or zEC12 there and you find the certs for all Red Hat distributions
    - z13: https://access.redhat.com/ecosystem/hardware/1350433
    - zEC12: https://access.redhat.com/ecosystem/hardware/952163

IBM - IBM z Systems - IBM z13 (2964)

Vendor: IBMIBM z Systems
IBM z13
(2964)

CERTIFICATIONS

Product	Versions	Level
Red Hat Enterprise Linux (s390x)	6.6 - 6.x	Certified
Red Hat Enterprise Linux (s390x)	7.0 - 7.x	Certified
Red Hat Enterprise Linux (s390x)	5.11 - 5.x	Certified

Hardware certification (SUSE)

- SUSE portal:
  - https://www.suse.com/yessearch/Search.jsp
  - search for z13 or zEC12
    - z13
    - zEC12

YES CERTIFIED Bulletin Search

Keywords:	Company:	Product:	Hardware/Software:
z13	Any	Any	Any

SEARCH AGAIN

search result hits: 2

Bulletin	Company	Product	Category	OS	Date
142729	IBM	**IBM z Systems z13 (2964)**	Network Server	SUSE® Linux Enterprise Server for System z 12	2015-02-10
142728	IBM	**IBM z Systems z13 (2964)**	Network Server	SUSE® Linux Enterprise Server 11 for IBM zSeries® Service Pack 3 for SUSE® SLES 11	2015-02-10

SEARCH AGAIN

Agenda

- z13 (Linux) overview
- Distribution support timelines
- z13 Linux testing
- IBM tested platforms
- Distribution certifications
- **Outlook and recommendations**

Recommendations

- Update your z/VM to the required level
  - install all released service as well
- If not already done
  - Update your Linux guests to something that‘s supported and covers what you need
- Update the hardware to z13
  - upgrade the firmware before bringing up the system
- If anything doesn‘t work – don‘t be shy to open a PMH / PMR
- Plan for another update cycle (firmware, z/VM, Linux) later this year

The journey continues

- So far we only worked on toleration i.e. making sure upgrades are running smoothly
- Next is exploitation of all the nice new features, offloads etc
- Focus for exploitation will be on the new distributions (SLES12 and RHEL7)
- See SHARE session from Matin Schwidefsky **16450: What's New in Linux on System z?**

Toleration is required to ensure that existing JVMs in the field can exploit z9, z10, z196, zEC12 optimizations

Java Release	SR or FP	Aavailability
Java6	SR16 FP3	Jan 2015
Java7	SR8 FP3	Jan 2015
Java7.1	SR2 FP10	Jan 2015

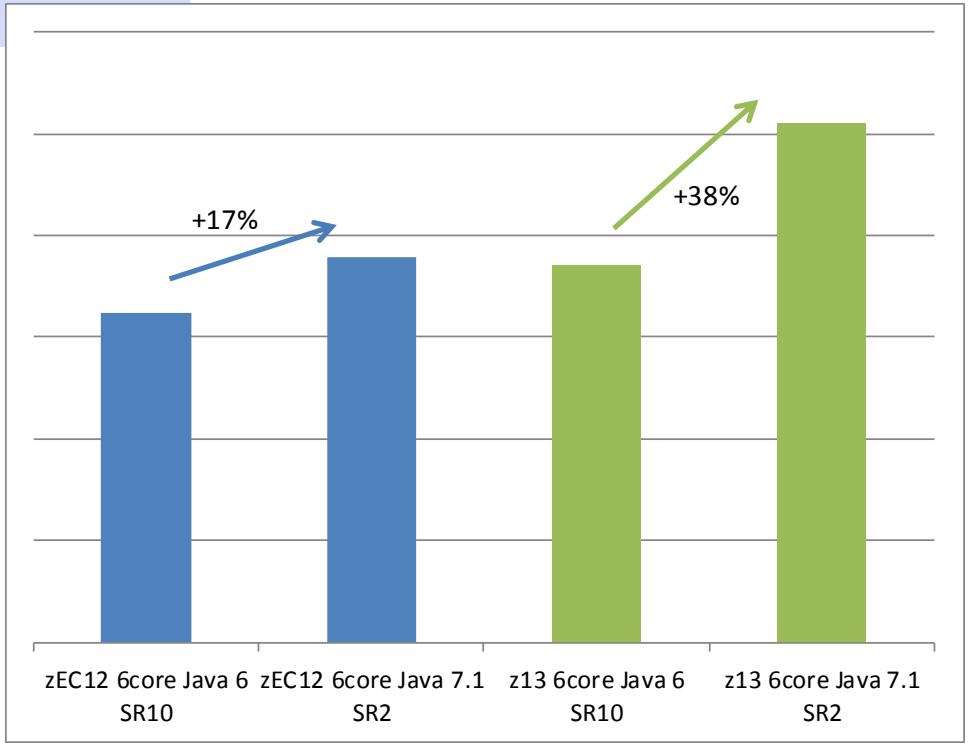

Ensure that you update all the middleware that comes with an embedded Java version

Dr. Eberhard Pasch
epasch@de.ibm.com

Linux on System z – Tuning hints and tips:
http://www.ibm.com/developerworks/linux/linux390/perf/index.html

Mainframe Linux blog: http://linuxmain.blogspot.com