Client Center, IBM Germany R&D
Frank Heimes – Senior IT Architect - <frank.heimes@de.ibm.com>

Linux on System z – Update

Current & Future Technology

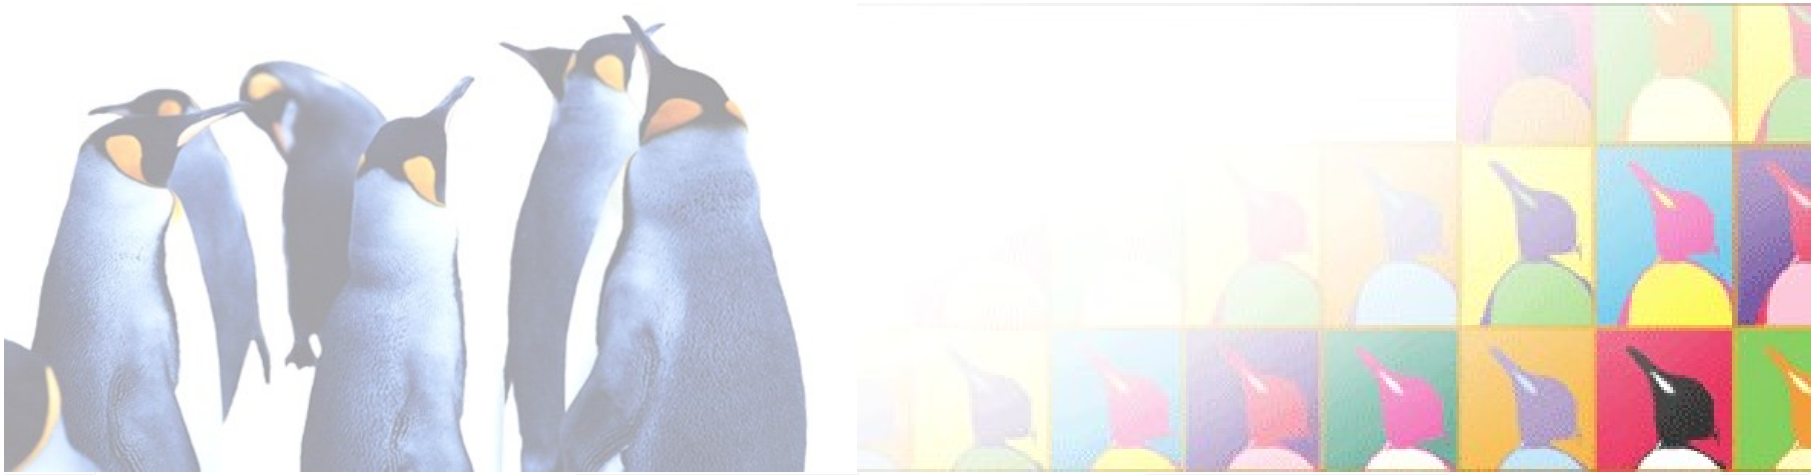

Thanks to Martin Schwidefsky

Linux Kernel Source

- The Linux kernel is **highly portable** (even if initially not developed to be) and supports the following computer architectures:
- One **single set of source code** available from “The Linux Kernel Archives”: http://kernel.org – http://git.kernel.org
- Linux for different achitectures is all **build from** the **same source**!
- All **architecture dependent code** is included:
  - Linux Kernel Sources: /linux/arch/s390 /linux/drivers/s390
  - Linux Kernel Headers (include/asm)
- Linux is always an **ASCII** operating system, even Linux on System z, running on the Extended Binary Coded Decimal Interchange Code (EBCDIC) encoded z/Architecture

[linux/kernel/git/torvalds/linux-2.6.git] / arch /

- Alpha architecture (alpha)
- Analog Devices (blackfin)
- ARM architecture (arm)
- Atmel AVR32 (avr32)
- Axis Communications' ETRAX CRIS (cris)
- Freescale's (formerly Motorola's) 68k (m68k)
- Fujitsu FR-V (frv)
- HP PA-RISC (parisc)
- H8 Renesas Technology, formerly Hitachi (H8300)
- **IBM System/390 (31-bit), z/Architecture (64-bit) (s390)**
- Intel IA-64 Itanium, Itanium II (ia64)
- **x86 architecture: 80386 ... x86_64 (x86)**
- M32R from Mitsubishi (m32r)
- Microblaze from Xilinx (microblaze))
- MIPS architecture (mips)
- MN103 from Panasonic Corporation (mn10300)
- OpenRISC (openrisc)
- **IBM POWER architecture (powerpc)**
- PowerPC architecture (powerpc)
- IBM's Cell (powerpc)
- SPARC, UltraSPARC (sparc)
- SuperH (sh)
- S+core (score)
- Tilera (tile)
- Xtensa from Tensilica (xtensa)
- UniCore32 (unicore32)
- User Mode Linux, UML (um)

Linux on IBM System z in 2Q2012

Installed Linux MIPS at 39% CAGR*

- 21% of Total installed MIPS run Linux as of 2Q12
- Installed IFL MIPS increased 26% YoY (2Q12)
- 36% of System z Customers have IFL's installed as of 2Q12
- 67 of the top 100 System z Customers are running Linux on the mainframe as of 2Q12
- 31% of all System z servers have IFLs

Installed Capacity Over Time

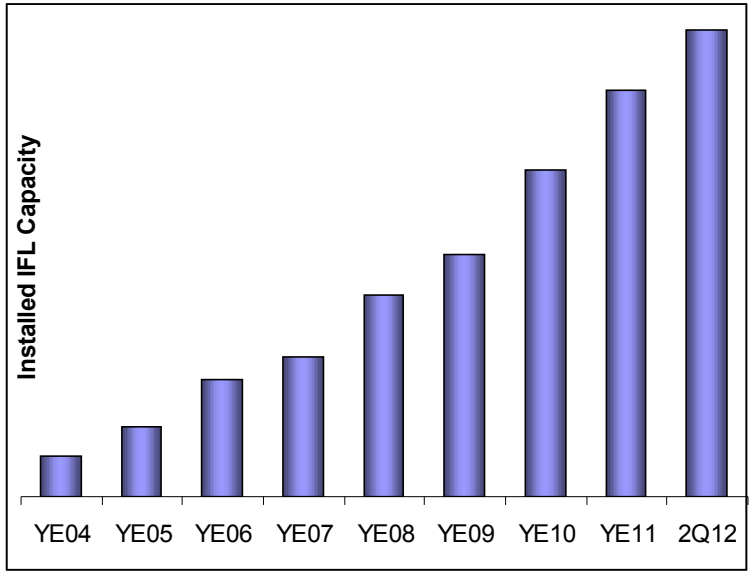

* Based on YE 2004 to 2H 2012

Linux on System z Deployment Trends

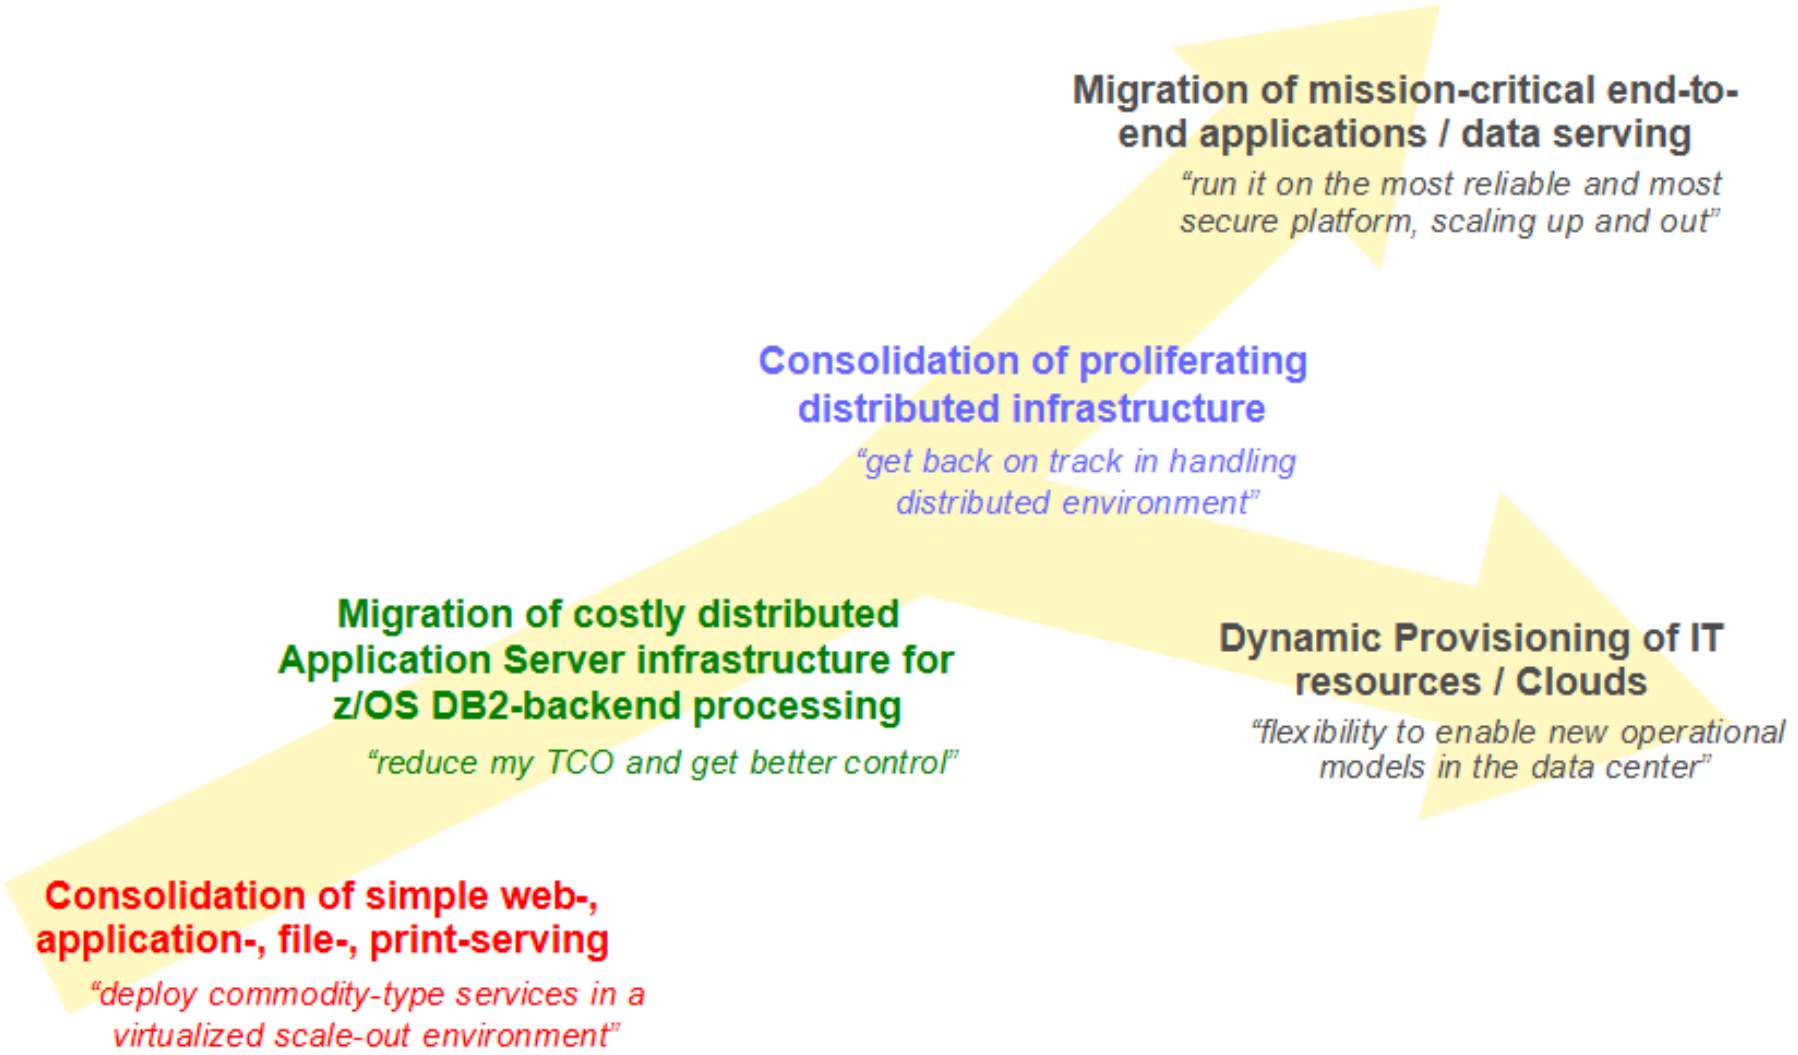

How Linux on System z is developed

How does the “community” work?

IBM Linux on System z Development

IBM Linux on System z Development contributes in the following areas: Kernel, s390-tools, Open Source Tools (e.g. eclipse, ooprofile), GCC, GLIBC, Binutils

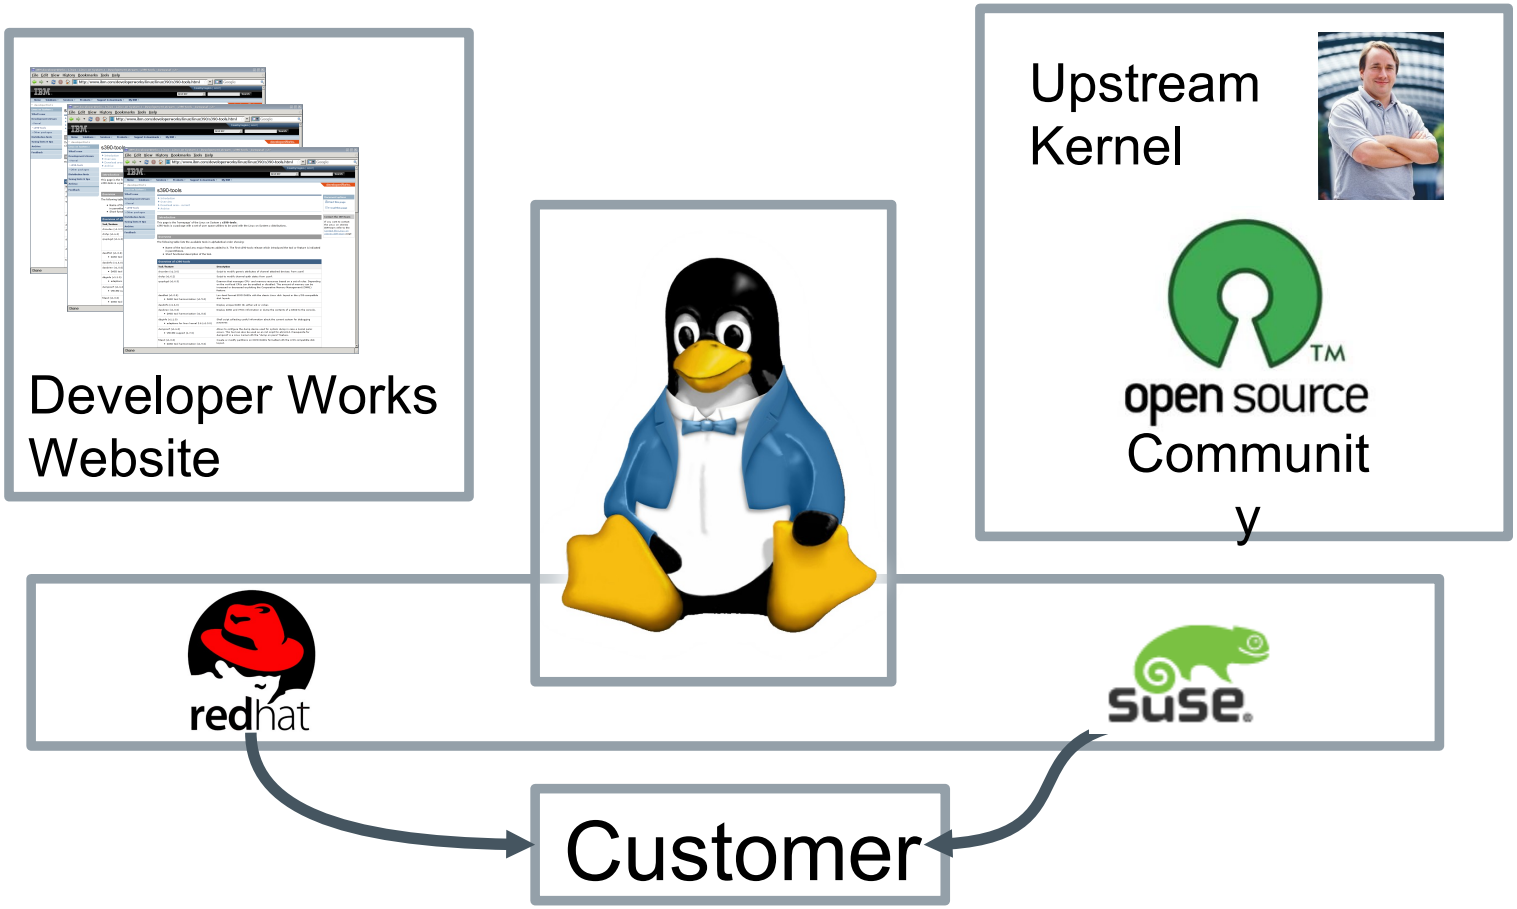

Linux kernel development: rate of change

Average for the last 7 years (without renames): 102 days per release, 5897 lines added, 2586 lines removed and 1221 lines modified **per day**

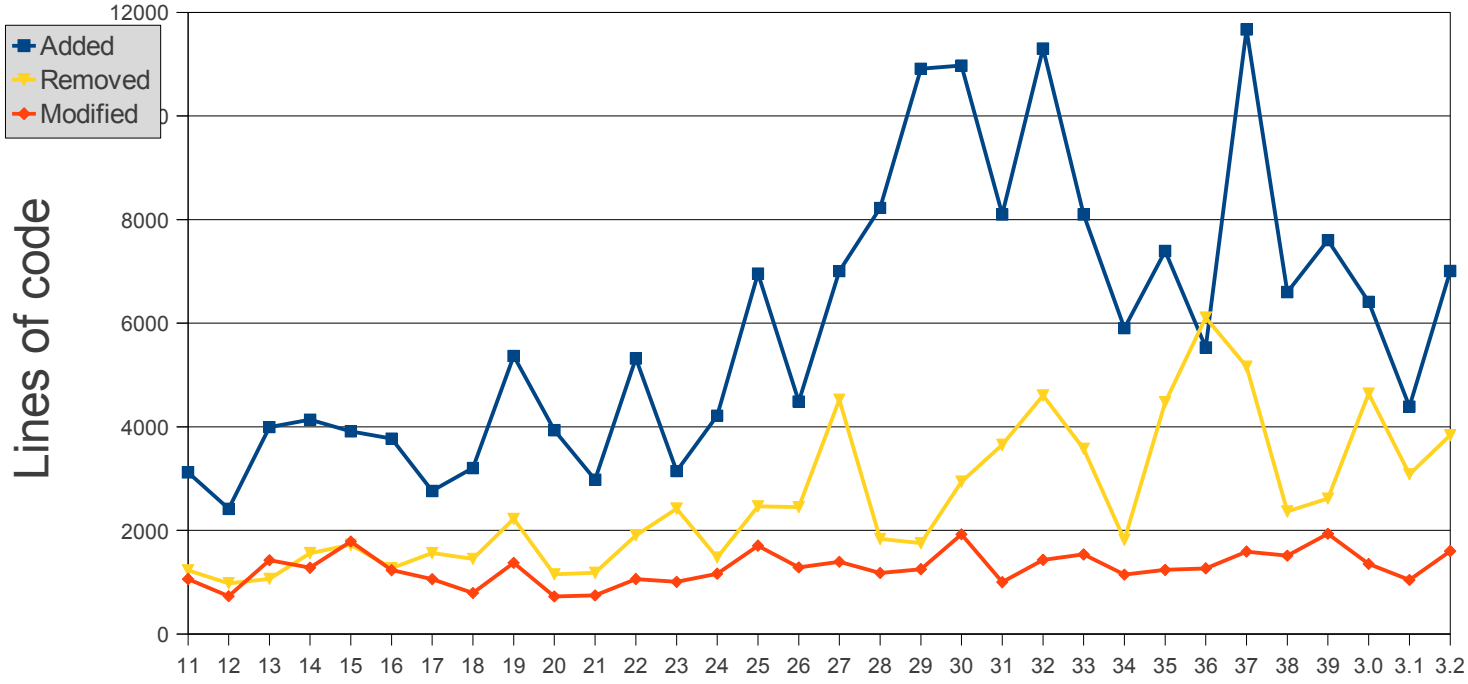

Linux kernel development: System z contributions

Changesets per 2.6.x / 3.x kernel release

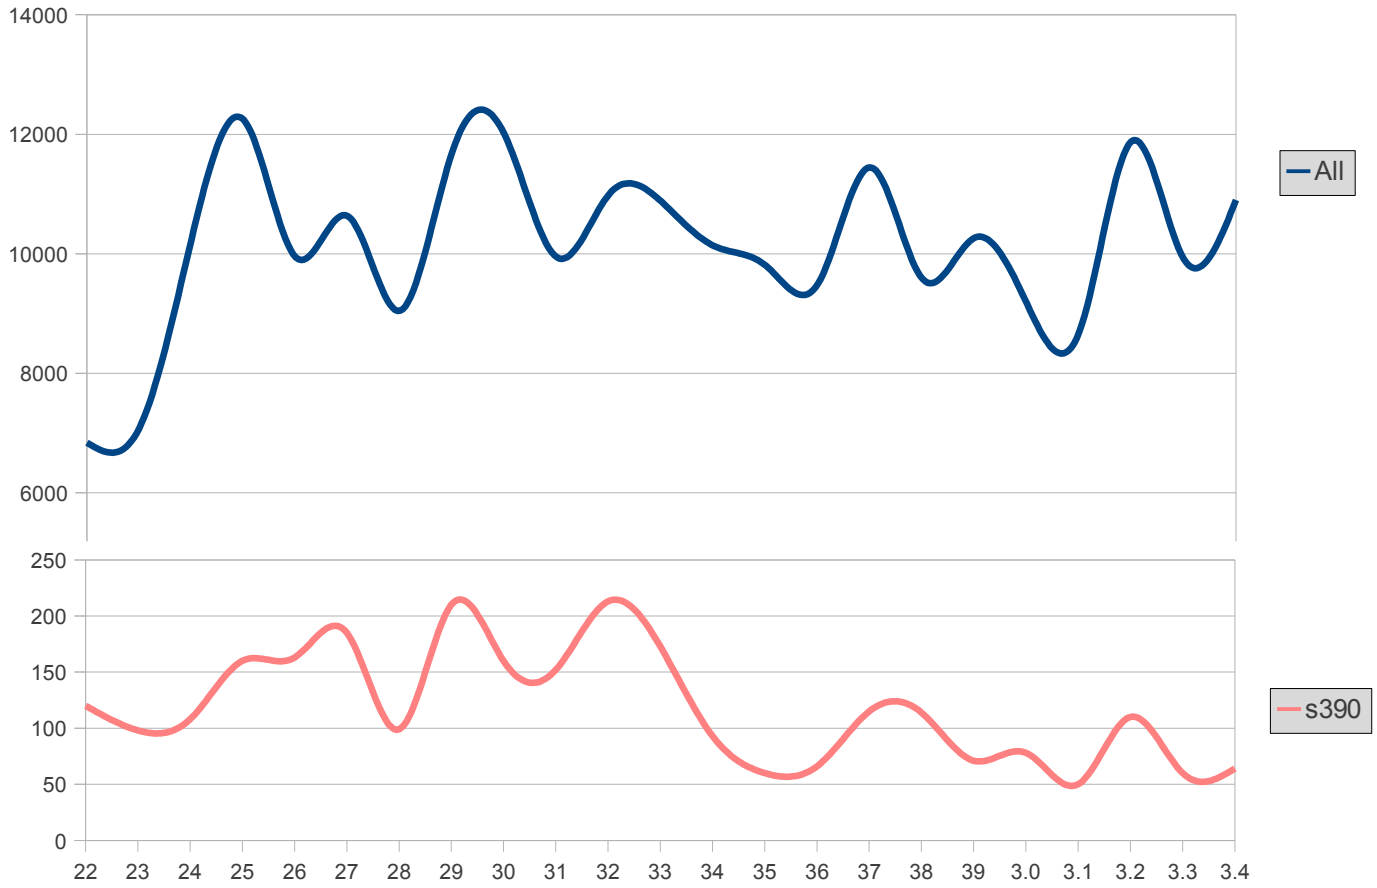

History of Enterprise Linux

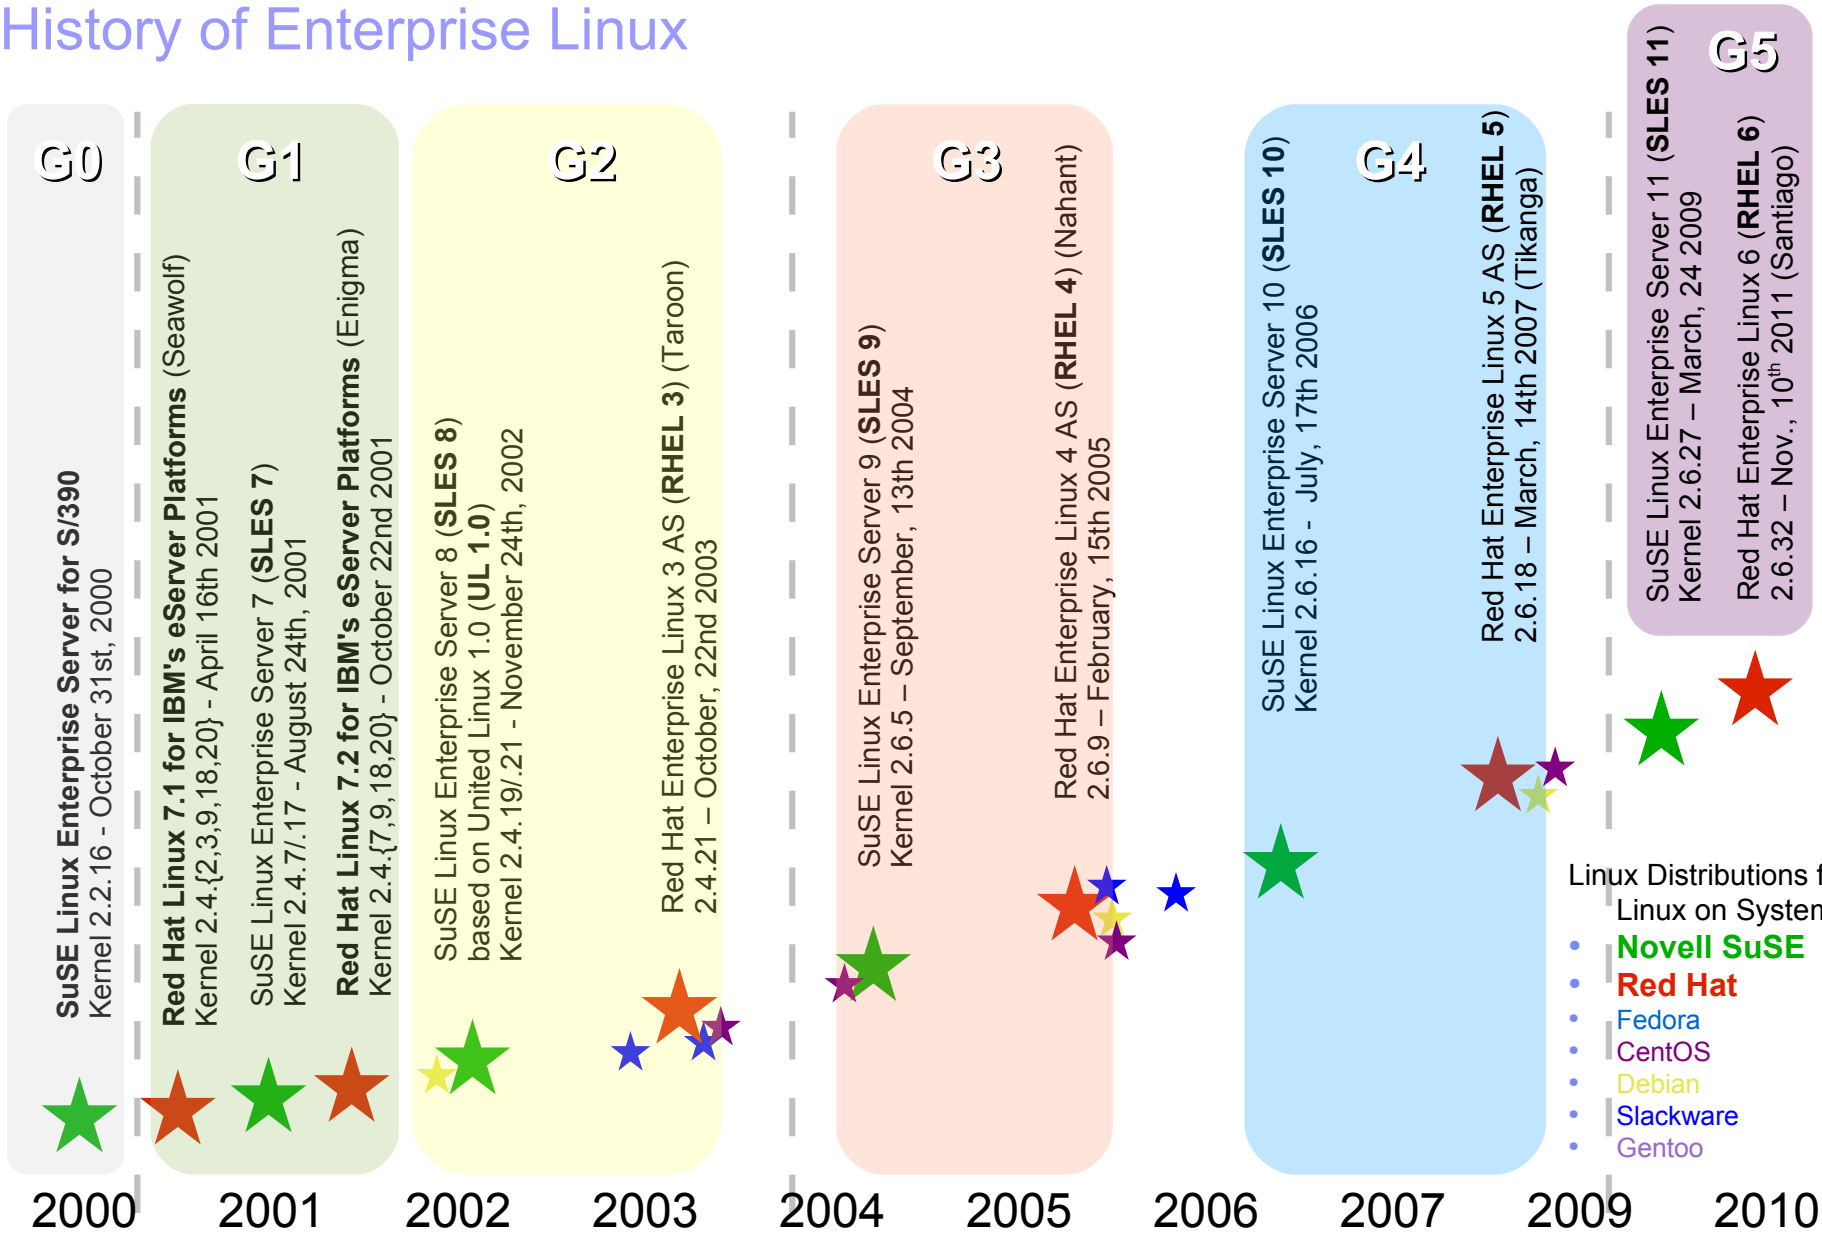

Enterprise Linux Distributions – Tested & Supported

The table below shows IBM tested Linux environments. IBM remote technical support for these environments is provided when you obtain a Support Line contract.
You may also find support for these environments by contracting with a third party provider.

Hardware Platform and Operating System Software Compatibility (64-bit)					
Generation	N – 4	N – 3	N – 2	N – 1	N
Release	zSeries	System z9	System z10	z196/z114	zEC12
SLES 9 (*)	✓	✓	✓	✓ (6)	✗
SLES 10	✓	✓	✓	✓	✓ (4)
SLES 11	✗	✓	✓	✓	✓ (3)
RHEL 4 (*)	✓	✓	✓	✓ (5)	✗
RHEL 5	✓	✓	✓	✓	✓ (2)
RHEL 6	✗	✓	✓	✓	✓ (1)

✓ Indicates that the Linux 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.

(1) Recommended level: RHEL 6.3
(2) Recommended level: RHEL 5.8
(3) Recommended level: SLES 11 SP2
(4) Recommended level: SLES 10 SP4 with latest maintenance updates
(5) RHEL 4.8 only. Some functions have changed or are not available with the z196, e.g. the Dual-port OSA cards support to name one of several. Please check with your service provider regarding the end of service.
(6) SLES 9 SP4 with latest maintenance updates only. Some functions have changed or are not available with the z196, e.g. the Dual-port OSA cards support to name one of several. Please check with your service provider regarding the end of service.

✗ Indicates that the distribution is not supported by IBM on this server

(*) The distribution is out of service, extended support is required.

To retrieve interoperability support information for Enterprise Storage products when used in a supported host server environment see the **System Storage Interoperation Center**: http://www.ibm.com/systems/support/storage/config/ssic/index.jsp

zEC12 – Toleration Tests and Certifications

- Certifications tests executed and published by SUSE and Red Hat:
  - SLES 10 SP4:
    https://www.suse.com/nbswebapp/yesBulletin.jsp?BulletinNumber=138473
  - SLES11 SP2:
    https://www.suse.com/nbswebapp/yesBulletin.jsp?bulletinNumber=138472
  - RHEL 5 Update 8:
    https://hardware.redhat.com/show.cgi?id=852529
  - RHEL 6 Update 3:
    https://hardware.redhat.com/show.cgi?id=852525
- Be sure to install the z/VM PTFs for zEC12 before you upgrade a system:
  http://www-01.ibm.com/support/docview.wss?uid=isg1_2827DEVICE_2827-ZVM
- It's possible to upgrade the hardware without changing the Linux stack!
  - No compatibility package needed!
- Matrix of Tested & Supported Linux Distributions:
  http://www.ibm.com/systems/z/os/linux/support_testedplatforms.html
- **RHEL4 and SLES9 are only supported with special agreements – plan for migration now!**

Current Linux on System z Technology

Features & Functionality contained in the Novell & Red Hat Distributions

Selected New Features for Linux on System z

- Linux Open Source Development is happening “upstream”
- Features are integrated later in distributions
  - Customer requests can help speed up this process
- Complete listing is available here:
  http://www-05.ibm.com/de/events/linux-on-z/pdf/whats-new-mar13.pdf

System z kernel features – Core

- **Improved QDIO performance statistics (kernel 2.6.33)**
  - Converts global statistics to per-device statistics and adds new counter for the input queue full condition

- **Breaking event address for user space programs (kernel 2.6.35)**
  - Remember the last break in the sequential flow of instructions
  - Valuable aid in the analysis of wild branches

- **z196 enhanced node affinity support (kernel 2.6.37)**
  - Allows the Linux scheduler to optimize its decisions based on the z196 topology

- **Performance indicator bytes (kernel 2.6.37)**
  - Display capacity adjustment indicator introduced with z196 via /proc/sysinfo

CMSFS user space file system support

- Allows to mount a z/VM minidisk to a Linux mount point
- z/VM minidisk needs to be in the enhanced disk format (EDF)
- The cmsfs fuse file system transparently integrates the files on the minidisk into the Linux VFS, no special command required

```
# cmsfs-fuse /dev/dasde /mnt/cms
# ls -la /mnt/fuse/PROFILE.EXEC
-r--r----- 1 root root 3360 Jun 26  2009 /mnt/fuse/PROFILE.EXEC
```

- By default no conversion is performed
  - Mount with '-t' to get automatic EBCDIC to ASCII conversion

```
# cmsfs-fuse -t /dev/dasde /mnt/cms
```

- Write support is work in progress, almost completed
  - use “vi” to edit PROFILE.EXEC anyone ?
- Use fusermount to unmount the file system again

```
# fusermount -u /mnt/cms
```

Deliver z/VM CP special messages as uevent

6.1 11.2

- Allows to forward SMSG messages to user space programs
  - Message needs to start with “APP”
- The special messages cause uevents to be generated
- See “Writing udev rules for handling CP special messages” in the Device Drivers Book

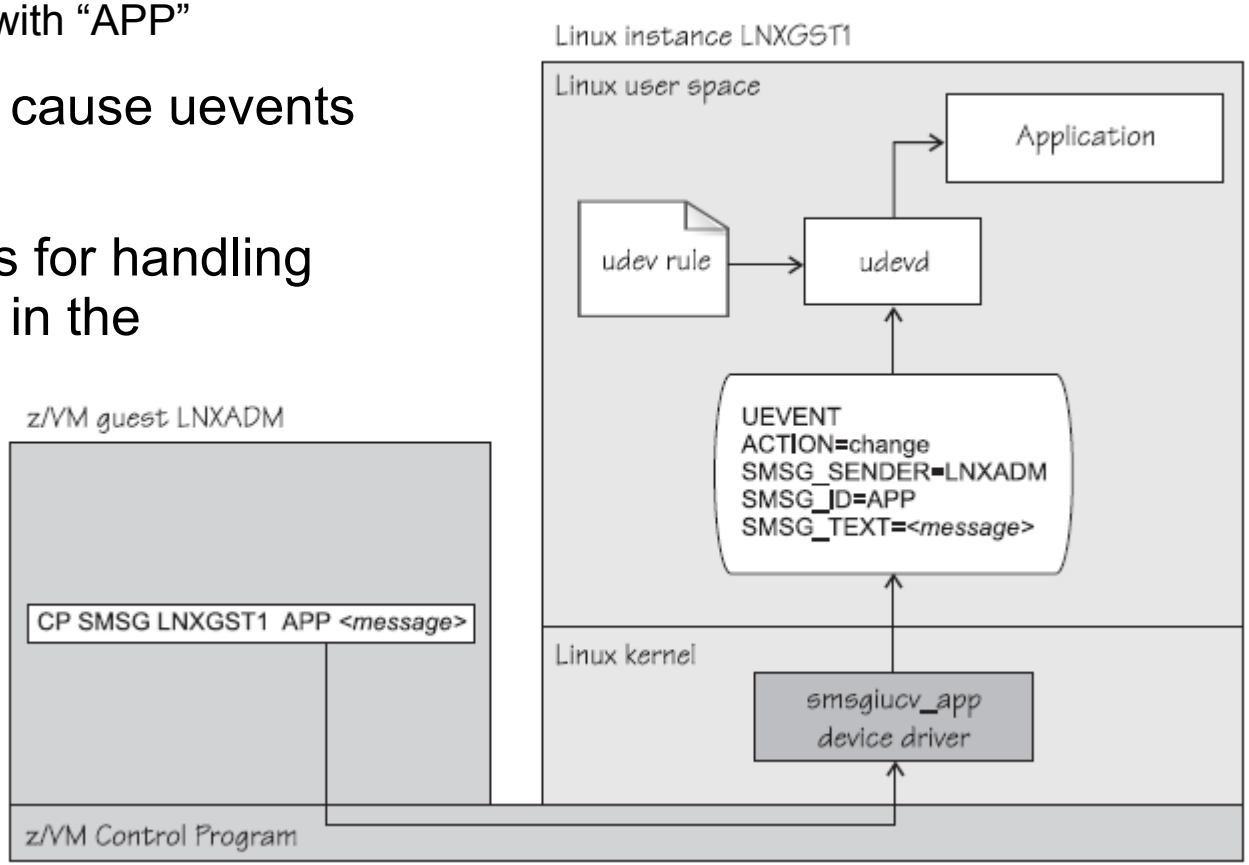

System z kernel features – Usability / RAS

- **Dump on panic – prevent reipl loop (s390-tools 1.8.4)**
  - Delay arming of automatic reipl after dump.
  - Avoids dumps loops where the restarted system crashes immediately.

- **Add support for makedumpfile tool (kernel 2.6.34, s390-tools 1.9.0)**
  - Convert Linux dumps to the ELF file format
  - Use the makedumpfile tool to remove user data from the dump.
  - Multi-volume dump will be removed.

- **Address space randomization (kernel 2.6.38)**
  - Enable flexible mmap layout for 64 bit to randomize start address for the runtime stack and the mmap area

- **Get CPC name (kernel 2.6.39)**
  - Useful to identify a particular hardware system in a cluster
  - The CPC name and the HMC network name are provided

System z kernel features – Networking

- **Offload outbound checksumming (kernel 2.6.35)**
  - Move calculation of checksum for non-TSO packets from the driver to the OSA network card

- **OSX/OSM CHPIDs for hybrid data network (kernel 2.6.35)**
  - The OSA cards for the zBX Blade Center Extension will have a new CHPID type
  - Allows communication between zBX and Linux on System z

- **Toleration of optimized latency mode (kernel 2.6.35)**
  - OSA devices in optimized latency mode can only serve a small number of stacks / users. Print a helpful error message if the user limit is reached.
  - Linux does not exploit the optimized latency mode

- **NAPI support for QDIO and QETH (kernel 2.6.36)**
  - Convert QETH to the NAPI interface, the “new” Linux networking API
  - NAPI allows for transparent GRO (generic receive offload)

LNXHC – Linux Health Checker

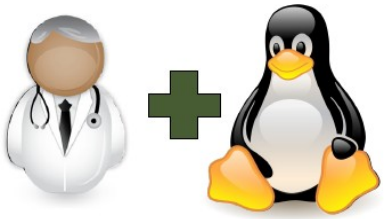

- The Linux Health Checker is a command line tool for Linux.
- Its purpose is to identify potential problems before they impact your system's availability or cause outages.
- It collects and compares the active Linux settings and system status for a system with the values provided by health-check authors or defined by you. It produces output in the form of detailed messages, which provide information about potential problems and the suggested actions to take.
- The Linux Health Checker will run on any Linux platform which meets the software requirements. It can be easily extended by writing new health check plug-ins.
- The Linux Health Checker is an open source project sponsored by IBM. It is released under the Eclipse Public License v1.0
  - Homepage: http://lnxhc.sourceforge.net/
  - IBM Linux Information Center: http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/topic/com.ibm.trouble.doc/lnxhc/lnxhc_linuxonz.htm
  - Live Virtual Class (LVC): http://www.vm.ibm.com/education/lvc/
  - User's Guide: http://lnxhc.sourceforge.net/documentation.html

zEC12 Out of Order Detail

Out of order yields significant performance benefit through

1) Re-ordering instruction execution
   - Instructions stall in a pipeline because they are waiting for results from a previous instruction or the execution resource they require is busy
   - In an in-order core, this stalled instruction stalls all later instructions in the code stream
   - In an out-of-order core, later instructions are allowed to execute ahead of the stalled instruction
2) Re-ordering storage accesses
   - Instructions which access storage can stall because they are waiting on results needed to compute storage address
   - In an in-order core, later instructions are stalled
   - In an out-of-order core, later storage-accessing instructions which can compute their storage address are allowed to execute
3) Hiding storage access latency
   - Many instructions access data from storage
   - Storage accesses can miss the L1 and require 10 to 500 additional cycles to retrieve the storage data
   - In an in-order core, later instructions in the code stream are stalled
   - In an out-of-order core, later instructions which are not dependent on this storage data are allowed to execute

System z toolchain (1)

- zEnterprise 196 exploitation (gcc 4.6)
  - Use option -march=z196 to utilize the new instructions added with z196
  - Use -mtune=z196 to schedule the instruction appropriate for the new out-of-order pipeline of z196
  - Re-compiled code/apps get further performance gains through 110+ new instructions

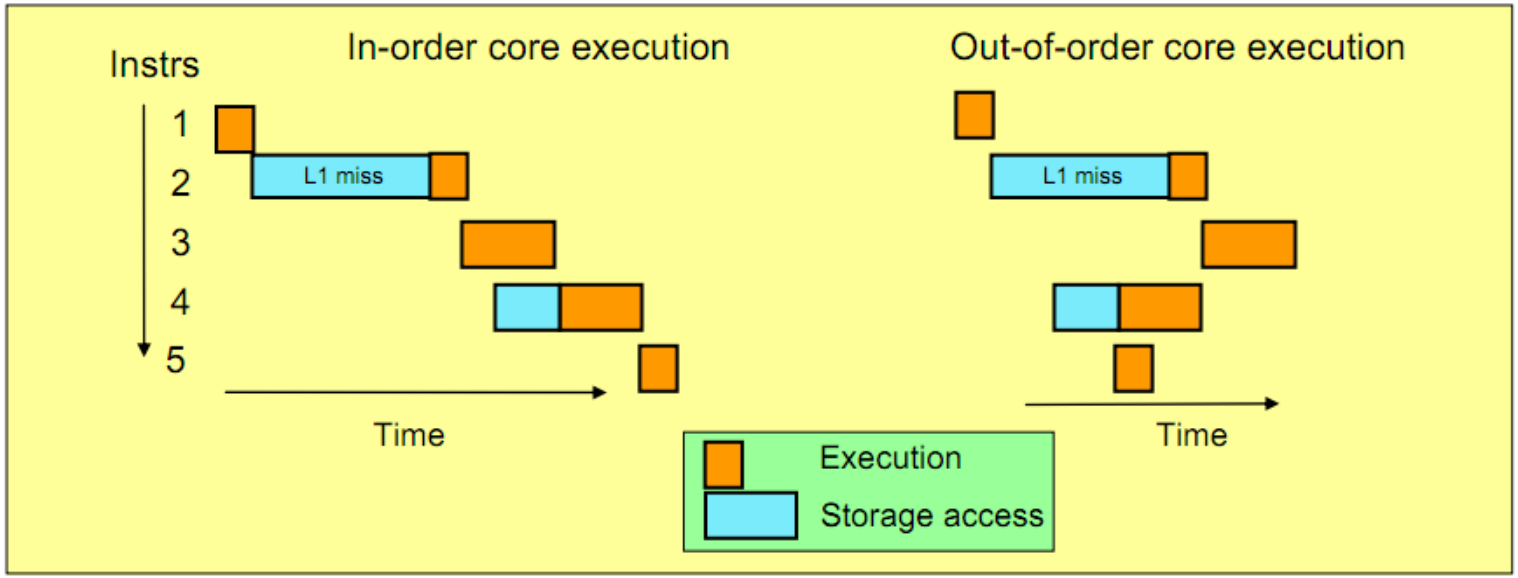

System z toolchain (2)

- zEC12 exploitation (gcc 4.x)
  - Use option -march=zEC12 to utilize the new instructions added with zEC12
  - Use -mtune=zEC12 to schedule the instruction appropriate for the improved out-of-order pipeline of zEC12
  - http://gcc.gnu.org/ml/gcc-patches/2012-10/msg00943.html
  - http://sourceware.org/ml/binutils/2012-10/msg00040.html

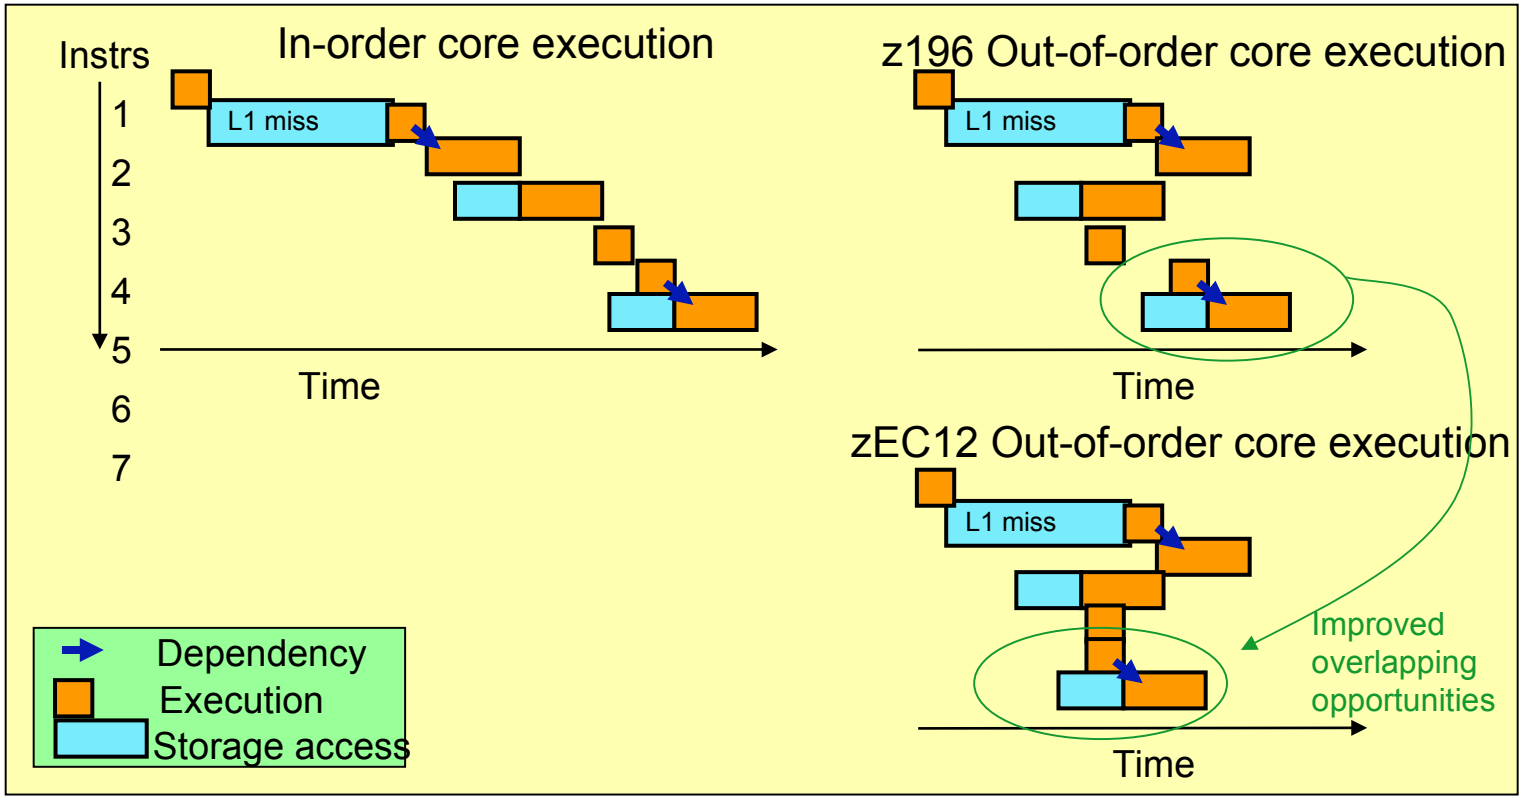

Future Linux on System z Technology

*Software which has already been developed and integrated into the Linux Kernel – but is **not** yet available in any Enterprise Linux Distribution*

System z kernel features – Storage FICON

- **DASD sanity check to detect path connection errors (kernel 3.3)**
  - An incorrect physical connection between host and storage server which is not detected by hardware or microcode can lead to data corruption
  - Add a check in the DASD driver to make sure that each available channel path leads to the same storage server
- **Extended DASD statistics (kernel 3.1)**
  - Add detailed per-device debugging of DASD I/Os via debugfs
  - Useful to analyze problems in particular for PAV and HPF

Extended DASD statistics

- Start data collection

```
# dasdstat -e dasda 0.0.1234
```

- Reset statistics counters

```
# dasdstat -r dasda
```

- Read summary statistics

```
# dasdstat
statistics data for statistic: 0.0.6527
start time of data collection: Fri Feb 24 16:00:19 CET 2012

1472 dasd I/O requests
with 14896 sectors(512B each)
0 requests used a PAV alias device
0 requests used HPF
 __<4 ___8 __16 __32 __64 _128 _256 _512 __1k __2k __4k __8k _16k _32k _64k 128k
 _256 _512 __1M __2M __4M __8M _16M _32M _64M 128M 256M 512M __1G __2G __4G _>4G
Histogram of sizes (512B secs)
    0    0 1441    8   13    5    2    2    0    1    0    0    0    0    0    0
    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
Histogram of I/O times (microseconds)
    0    0    0    0    0    0    1 1160   49   52   61  142    7    0    0    0
    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0    0
```

System z toolchain

- **64 bit register in 31 bit compat mode**
  - Make use of 64 bit registers in 31 bit application running in z/Architecture mode.
  - Allows to use instruction operating on 64 bits, e.g. 64 bit multiplication
  - Needs kernel support for asynchronous signals
- **ATLAS support (libatlas 3.9.52)**
  - Add support for System z to the “Automatically Tuned Linear Algebra Software”.
  - Improve performance of the library functions for System z.

System z toolchain

- **Oprofile support for hardware sampling introduced with z10 (2.6.39)**
  - Provide CPU measurement data to applications for performance tuning
  - Based on hardware counters and samples built into the CPU
  - Use oprofile to communicate the information to user space programs
- **Oprofile z196 hardware customer mode sampling (kernel 3.3)**
  - Extend the hardware sampling to support z196.
- **Valgrind System z support**
  - Valgrind is a generic framework for creating dynamic analysis tools and can be used for memory debugging, memory leak detection and profiling (e.g. cachegrind)
  - Valgrind is in essence a virtual machine using just-in-time (JIT) compilation techniques
  - Memory debugging is available with Valgrind version 3.7.0

Valgrind System z support

- valgrind –tool=memcheck [--leak-check=full] [--track-origins] <program>
  - Detects if your program accesses memory it shouldn't
  - Detects dangerous uses of uninitialized values on a per-bit basis
  - Detects leaked memory, double frees and mismatched frees
- valgrind –tool=cachegrind
  - Profile cache usage, simulates instruction and data cache of the cpu
  - Identifies the number of cache misses
  - Needs cache line size, Extract Cache Attributes (ECAG) instruction introduced with z10
- valgrind –tool=massif
  - Profile heap usage, takes regular snapshots of program's heap
  - Produces a graph showing heap usage over time

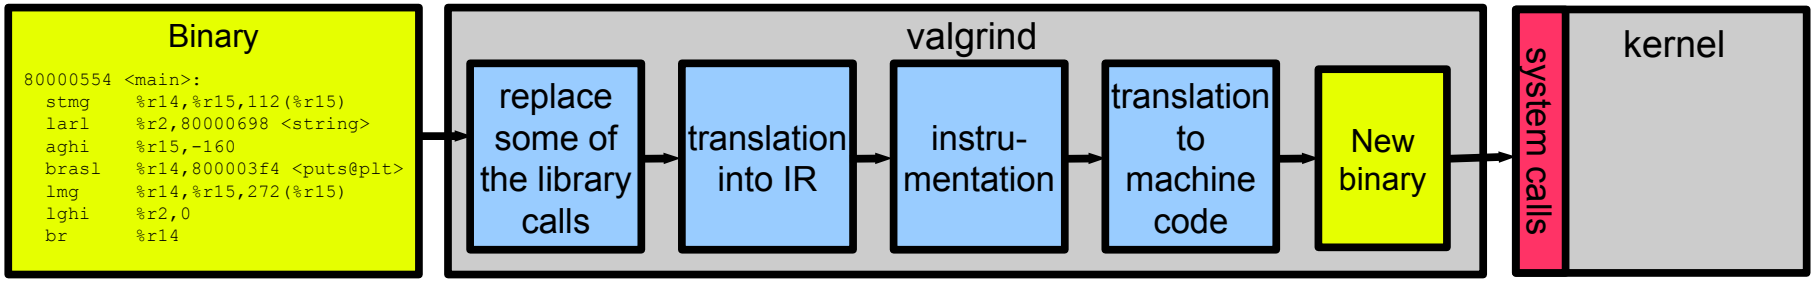

s390-tools package: what is it?

- s390-tools is a package with a set of user space utilities to be used with the Linux on System z distributions.
  - It is **the** essential tool chain for Linux on System z
  - It contains everything from the boot loader to dump related tools for a system crash analysis .
- This software package is contained in all major (and IBM supported) enterprise Linux distributions which support s390
  - RedHat Enterprise Linux 4
  - RedHat Enterprise Linux 5
  - RedHat Enterprise Linux 6
  - SuSE Linux Enterprise Server 9
  - SuSE Linux Enterprise Server 10
  - SuSE Linux Enterprise Server 11
- Website: http://www.ibm.com/developerworks/linux/linux390/s390-tools.html
- Feedback: linux390@de.ibm.com

s390-tools package: the content

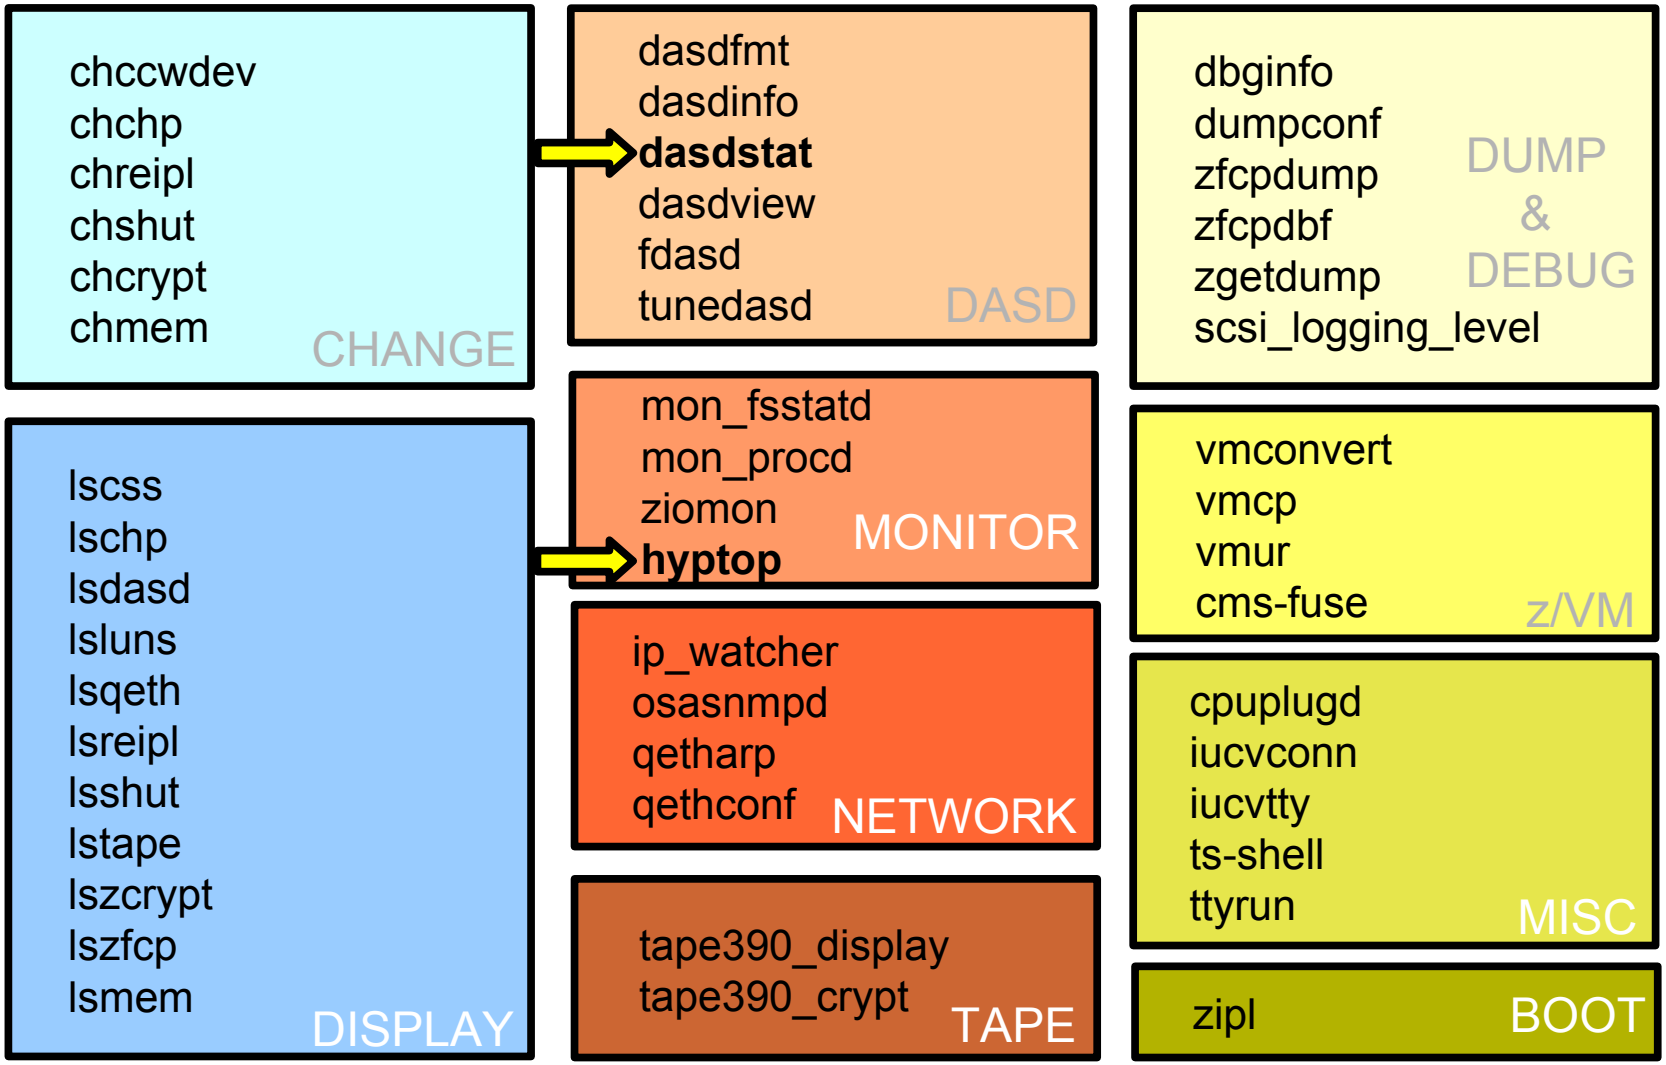

zEC12 Technology

Planned exploitation of features, introduced with zEnterprise EC12 (zEC12).

zEC12 – Optimization and scale improvements Starting at the Core

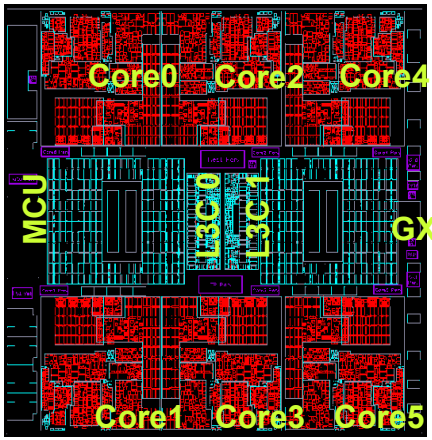

- **New 5.5 GHz 6-Core Processor Chip continues our leadership in microprocessor design with a boost in performance for all workloads**
  - Second generation out of order execution design
- **Larger caches to optimize data serving environments**
- **New hardware functions optimized for software performance**
  - **Transactional Execution Facility** for parallelism and scalability
  - **Runtime Instrumentation Facility** is intended to help reduce Java overhead
  - **2 GB page frames** are intended to offer performance improvements for DB2 buffer pools and Java heaps
  - New IBM Enterprise PL/I compiler is planned to exploit and get a performance boost from **decimal format conversions facility**
- **Integrated cryptographic function available on each core characterized as a CP or IFL with enabling microcode**

Hardware Transactional Memory (HTM)

- **Allow lockless interlocked execution of a block of code called a 'transaction'**
  - **Transaction:** Segment of code that appears to execute 'atomically' to other CPUs
    - Other processors in the system will either see **all-or-none** of the storage up-dates of transaction
- **How it works:**
  - TBEGIN instruction starts speculative execution of 'transaction'
  - Storage conflict is detected by hardware if another CPU writes to storage used by the transaction
  - Conflict triggers hardware to roll-back state (storage and registers)
    - transaction can be re-tried, or
    - a fall-back software path that performs locking can be used to guarantee forward progress
  - Changes made by transaction become visible to other CPUs after TEND

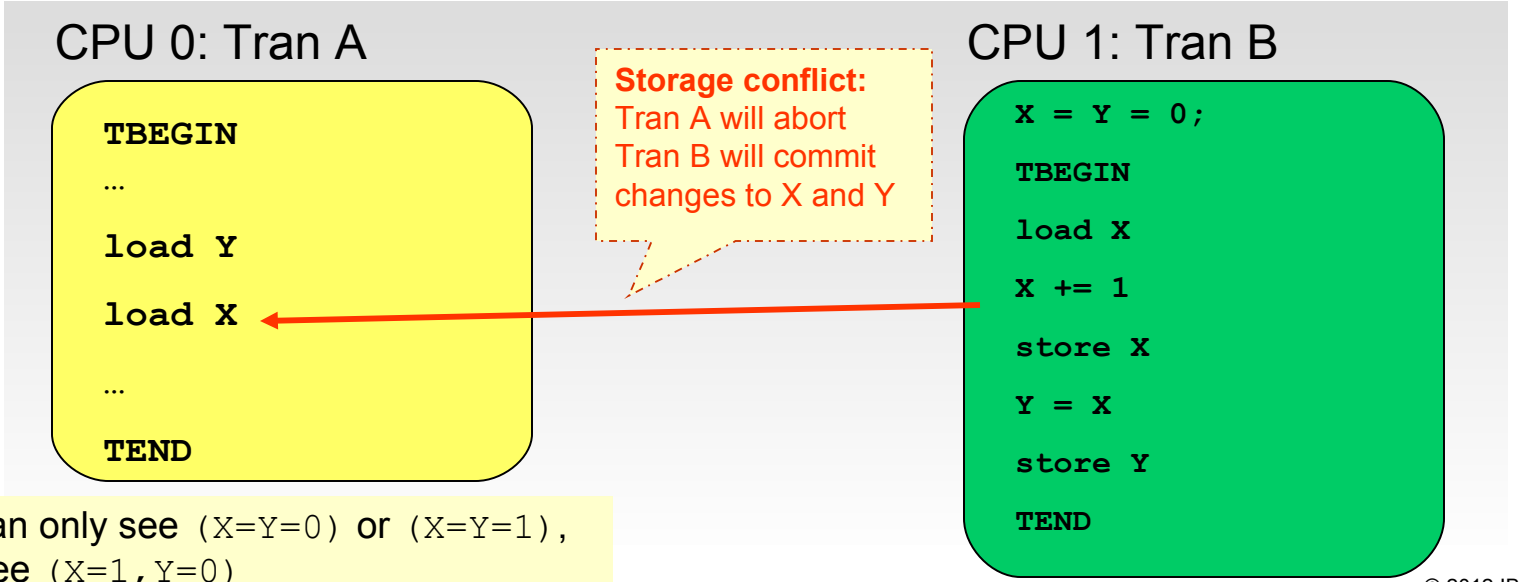

CPU 0 can only see `(X=Y=0)` or `(X=Y=1)`, cannot see `(X=1,Y=0)`

New zEC12 features help to improve availability and maximize service levels

- Companies competing for the highest quality of service want **outstanding availability** to deliver superior SLAs
- **Flash Express** is a new adapter card that **fits in the PCIe I/O drawer** and is designed to help improve availability and performance
  - Automatically improve availability for key workloads at critical processing times
  - Drive performance for critical business workloads that cannot tolerate paging spikes
  - Slashes latency for critical application processing such as diagnostics collection
- **Secure, resilient, easily configurable** for immediate use solution
- IBM is working with its Linux distribution partners to include support in future **Linux on System z** distribution releases,
- Flash Express will be first exploited by z/OS
  - **DB2 for z/OS V10 and Java SDK7 SR3** will support pageable large pages and be able to take advantage of Flash Express when the z/OS support is available
  - **IMS V13 beta program** (announced 4Q12) announced the intent to support pageable large pages

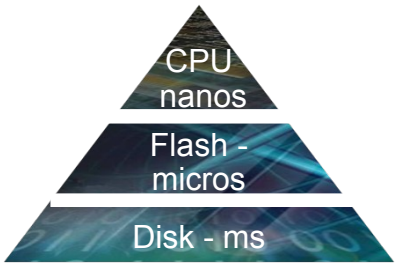

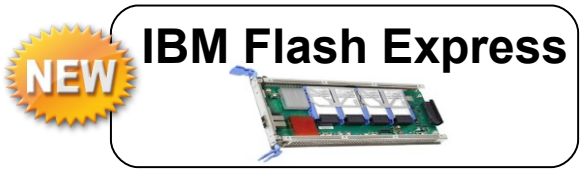

Flash Express – What is it?

FLASH Express

- Physically comprised of internal storage on Flash SSDs
- Used to deliver a new tier of memory - storage class memory (SCM)
- Supported on z/OS V1.13 plus web deliverable (PTF)
- Uses **PCIe I/O drawer**

One Flash Express Card

- Sized to accommodate *all LPAR paging*
  - Each **card pair** provides **1.6 TB** usable storage (3.2 TB total)
  - Maximum 4 card pairs (4 X1.6 = 6.4 TB)

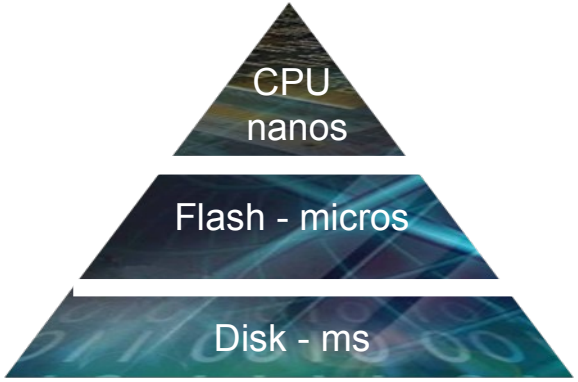

- Immediately usable
  - No capacity planning needed
  - No intelligent data placement needed
- Robust design
  - Delivered as a **RAID10** mirrored pair
  - Designed for long life
  - Designed for concurrent firmware upgrade
- Secured
  - Flash Express adapter is protected with 128-bit AES encryption.
  - Key Management provided based on a Smart Card

Linux on System z Support

	z10 EC WdfM	z10 BC WdfM	z196	z114	zEC12	Availability Date
RHEL 5	**x**	**x**	**x**	**x**	**x**	**03/2007**
RHEL 6	**x**	**x**	**x**	**x**	**x**	**11/2010**
SLES 10	**x**	**x**	**x**	**x**	**x**	**08/2006**
SLES 11	**x**	**x**	**x**	**x**	**x**	**03/2009**

	End of Production Ph 1	End of Production Ph 2	End of Production Ph 3
RHEL 5 support*	**4Q 2011**	**4Q 2012**	**03/31/2014**
RHEL 6 support*	**4Q 2014**	**4Q 2015**	**11/30/2017**

	General support	Extended support	Self support
SLES 10 support*	**07/31/2013**	**07/31/2016**	**07/31/2016**
SLES 11 support*	**03/31/2016**	**03/31/2019**	**03/31/2019**

- **For latest information and details contact your Linux distributor**
- Recommendation: use RHEL 6 or SLES 11 for new projects
- RHEL 4 and SLES 9 are not supported on zEC12
- For latest information about supported Linux distributions on System z refer to: http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

* SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux
Support dates may be changed by Linux distributors

z/VM zEC12 Compatibility Support

- Compatibility → z/VM support for host / guests on zEC12 at the z196 functional level with limited exploitation of new functions (some transparent)
  - Support available as PTFs concurrently with the availability of zEC12.
    Be sure to install these z/VM PTFs for zEC12 before you upgrade a system:
    http://www-01.ibm.com/support/docview.wss?uid=isg1_2827DEVICE_2827-ZVM
  - Includes PTFs for EREP, IOCP, HCD, HLASM, OSASF, and Performance Toolkit
  - For z/VM 5.4, 6.1, and 6.2

- Compatibility Support for several new architectural facilities

- Support for the PCIe I/O Drawer cage and new OSA-Express4S 1000BASE-T and Crypto Express4S features

- Supported releases:
  - z/VM 5.4 EoS 9/2013
  - z/VM 6.1 EoS 4/2013, requires a z10 ALS (Advanced Lifecycle Support)
  - z/VM 6.2 EoS 4/2015, requires a z10 ALS

Operating System Support for zEC12

- Currency is key to operating system support and exploitation of future servers
- The following releases of operating systems are supported on zEC12
  (refer to PSP buckets for any required maintenance):

Operating System	Supported levels
z/OS	• z/OS V1R13 with PTFs (Exploitation) • z/OS V1R12 with PTFs (Exploitation) • z/OS V1R11 with PTFS (Toleration, Lifecycle extension required after 09/12) • z/OS V1R10 (Toleration, Lifecycle Extension Required)

Linux on System z	• Linux distributions:
	  • SUSE SLES 10 and SLES 11
	  • Red Hat RHEL 5 and RHEL 6

z/VM	• z/VM 5.4, 6.1, and 6.2 will support zEC12 with PTFs

z/VSE	• z/VSE compatibility with PTFs, support for:
	  • z/VSE 4.3
	  • z/VSE 5.1

z/TPF	• V1.1

- Support for p Blades in zBX Model 003
  - AIX 5.3 Technology Level 12 or higher, AIX 6.1 Technology Level 5 or higher, AIX 7.. All with PowerVM™ Enterprise Edition
- Support for Linux and Windows* environments on select System x blades in zBX Model 003
  - 64 bit version support only
  - Red Hat RHEL 5.5 and higher, 6.0 and higher and SLES 10 (SP4), 11 SP1 and higher
  - Microsoft Windows Server 2008 R2 and Microsoft Windows Server 2008 (SP2) (Datacenter Edition recommended

Thank you – Questions ?

Obrigado

Portuguese

Merci

French

Thank You

English

Gracias

Spanish

Danke

German

ibm.com/education/students

Thanks to Martin Schwidefsky ...

Trademarks & Disclaimer

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml: AS/400, DB2, e-business logo, ESCON, eServer, FICON, IBM, IBM Logo, iSeries, MVS, OS/390, pSeries, RS/6000, S/390, System Storage, System z9, VM/ESA, VSE/ESA, WebSphere, xSeries, z/OS, zSeries, z/VM.

The following are trademarks or registered trademarks of other companies

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries. LINUX is a registered trademark of Linux Torvalds in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC. Intel is a registered trademark of Intel Corporation. * All other products may be trademarks or registered trademarks of their respective companies.

NOTES: Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured 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. References in this document to IBM products or services do not imply that IBM intends to make them available in every country. Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use. The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.