Hans-Joachim Picht, Linux on System z Initiatives & Linux Champion EMEA <hans@de.ibm.com>

Linux on System z Update:
Current & Future Linux on System z Technology

Frühjahrstagung 2011 der Arbeitsgruppen z/VSE mit CICS und DB2 z/VM mit Linux on System z
Hilton Düsseldorf, 02.-04. Mai 2011

IBM collaborates with the Linux community

- ...has been an active participant since 1999
- ...is one of the leading commercial contributors to Linux
- ...has over 600 full-time developers working with Linux and open source

Linux Kernel & Subsystem Development

Kernel Base

Security

Systems Mgmt

Virtualization

Filesystems,

and more...

Expanding the Open Source Ecosystem

Apache

Eclipse

Mozilla Firefox

OpenOffice.org,

and more...

Promoting Open Standards & Community Collaboration

The Linux Foundation

Linux Standards Base

Common Criteria certification,

and more...

Foster and Protect the Ecosystem

Software Freedom Law Center

Free Software Foundation (FSF),

and more...

Facts on Linux

- Last year, **75%** of the Linux code was developed by **programmers working for corporations**.
- **$7.37 billion:** projected cost to produce the 283 million lines of code which are contained in Linux Distribution **in a commercial environment.**
- IDC forecasts show that **Linux server revenue will grow by 85.5%** between 2008 and 2012 **in the non-x86 server space** equalling a four year compound annual growth rate of 16.7%.
- **Linux is Linux**, but …features, properties and quality differ dependent on your platform

Source: Intelligence Slideshow: 40 Fast Facts on Linux http://www.baselinemag.com/c/a/Intelligence/40-Fast-Facts-on-Linux-727574/
http://www.internetnews.com/dev-news/article.php/3659961
http://public.dhe.ibm.com/software/au/downloads/IBM_zLinux_DAG_FINAL.pdf

Linux kernel development: Rate of Change

Average: 6683 lines added, 3774 lines removed, 1797 lines changed every day for the last 5 1/2 years.

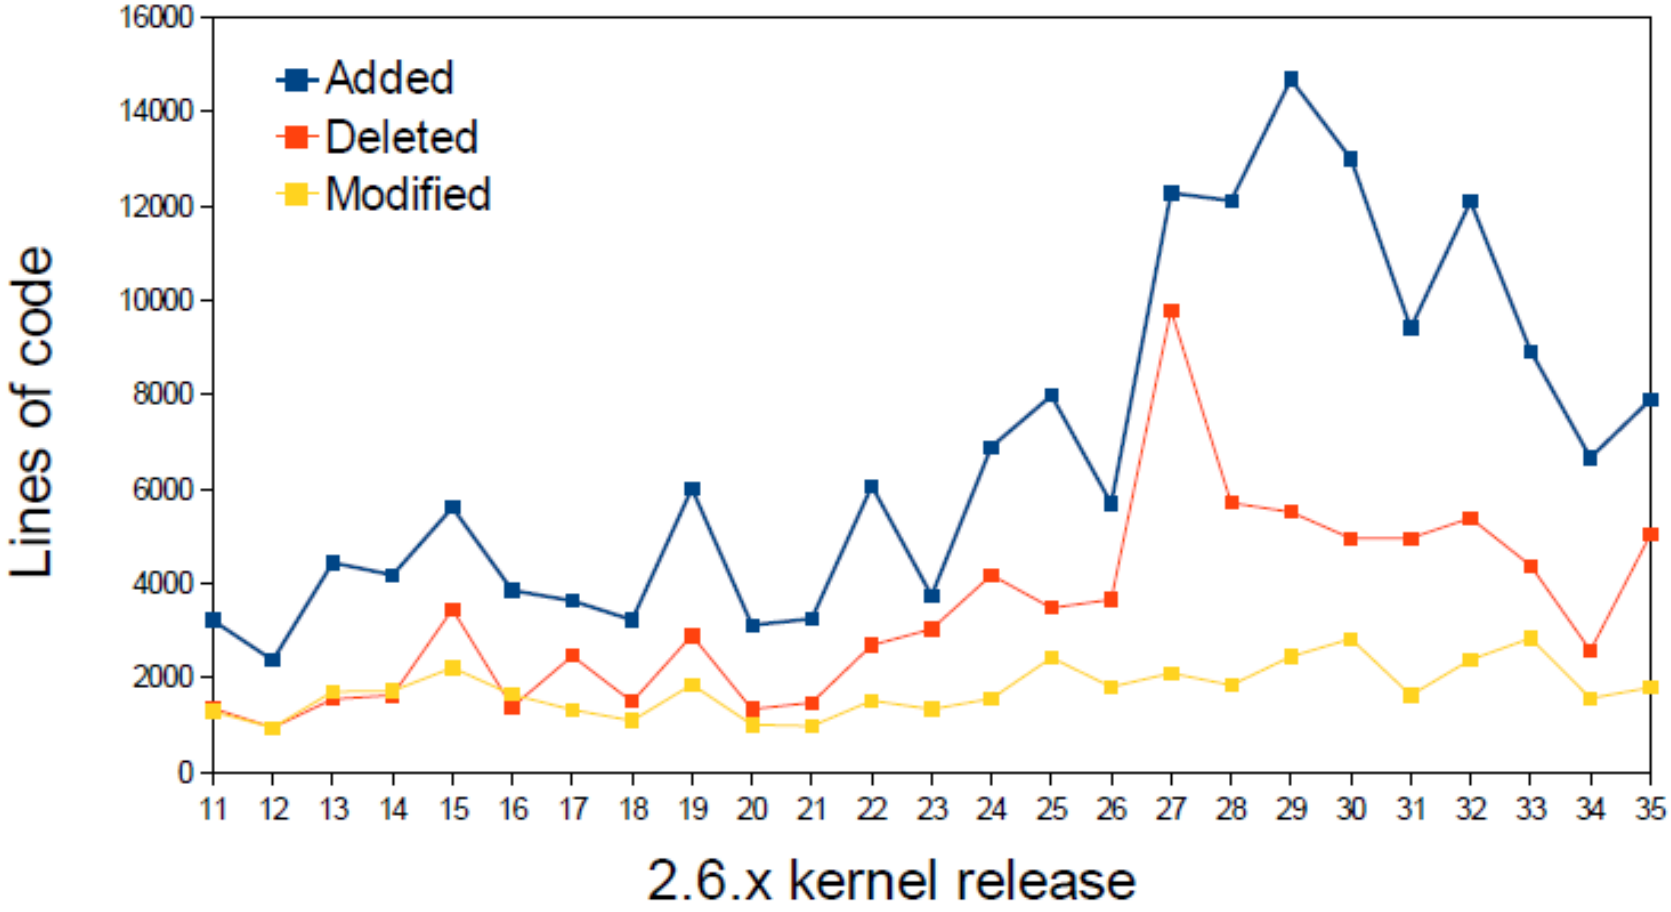

Source: http://www.linuxfoundation.org/docs/lf_linux_kernel_development_2010.pdf

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

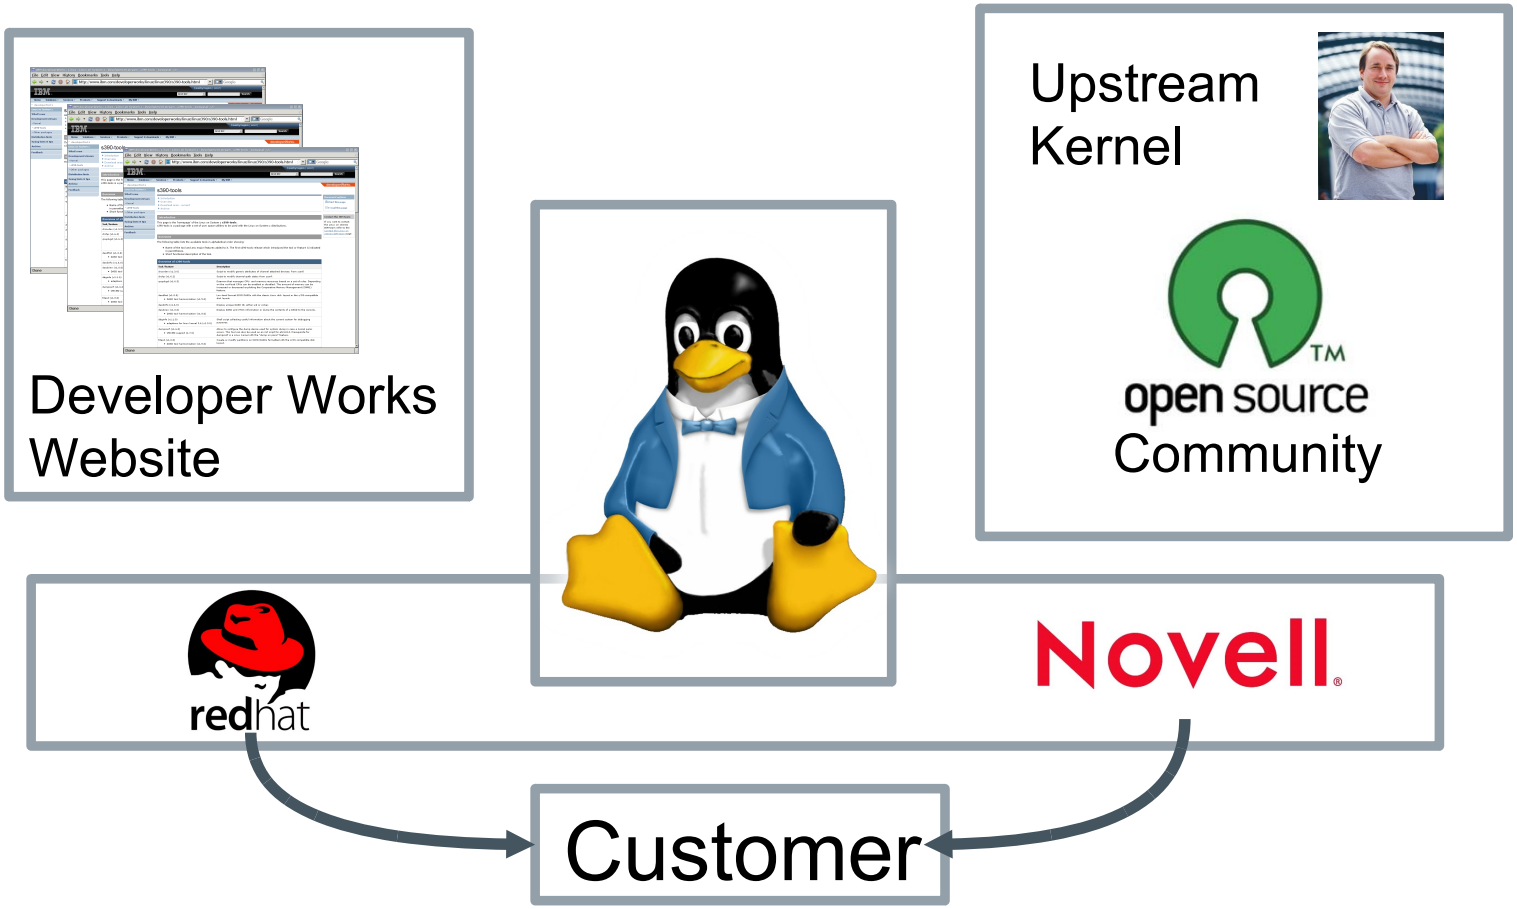

....the code you use is the result of the efforts of an anonymous army of blue penguins involved in developing, testing, documenting, ...

IBM Supported Linux Distributions for System z

	z196	z10	z9	zSeries
RHEL 6	✓	✓	✓	✗
RHEL 5	✓	✓	✓	✓
SLES 10	✓	✓	✓	✓
SLES 11	✓	✓	✓	✗

http://www-03.ibm.com/systems/z/os/linux/resources/testedplatforms.html

Structure of Linux on System z

Many Linux software packages did not require any code change to run on Linux on System z

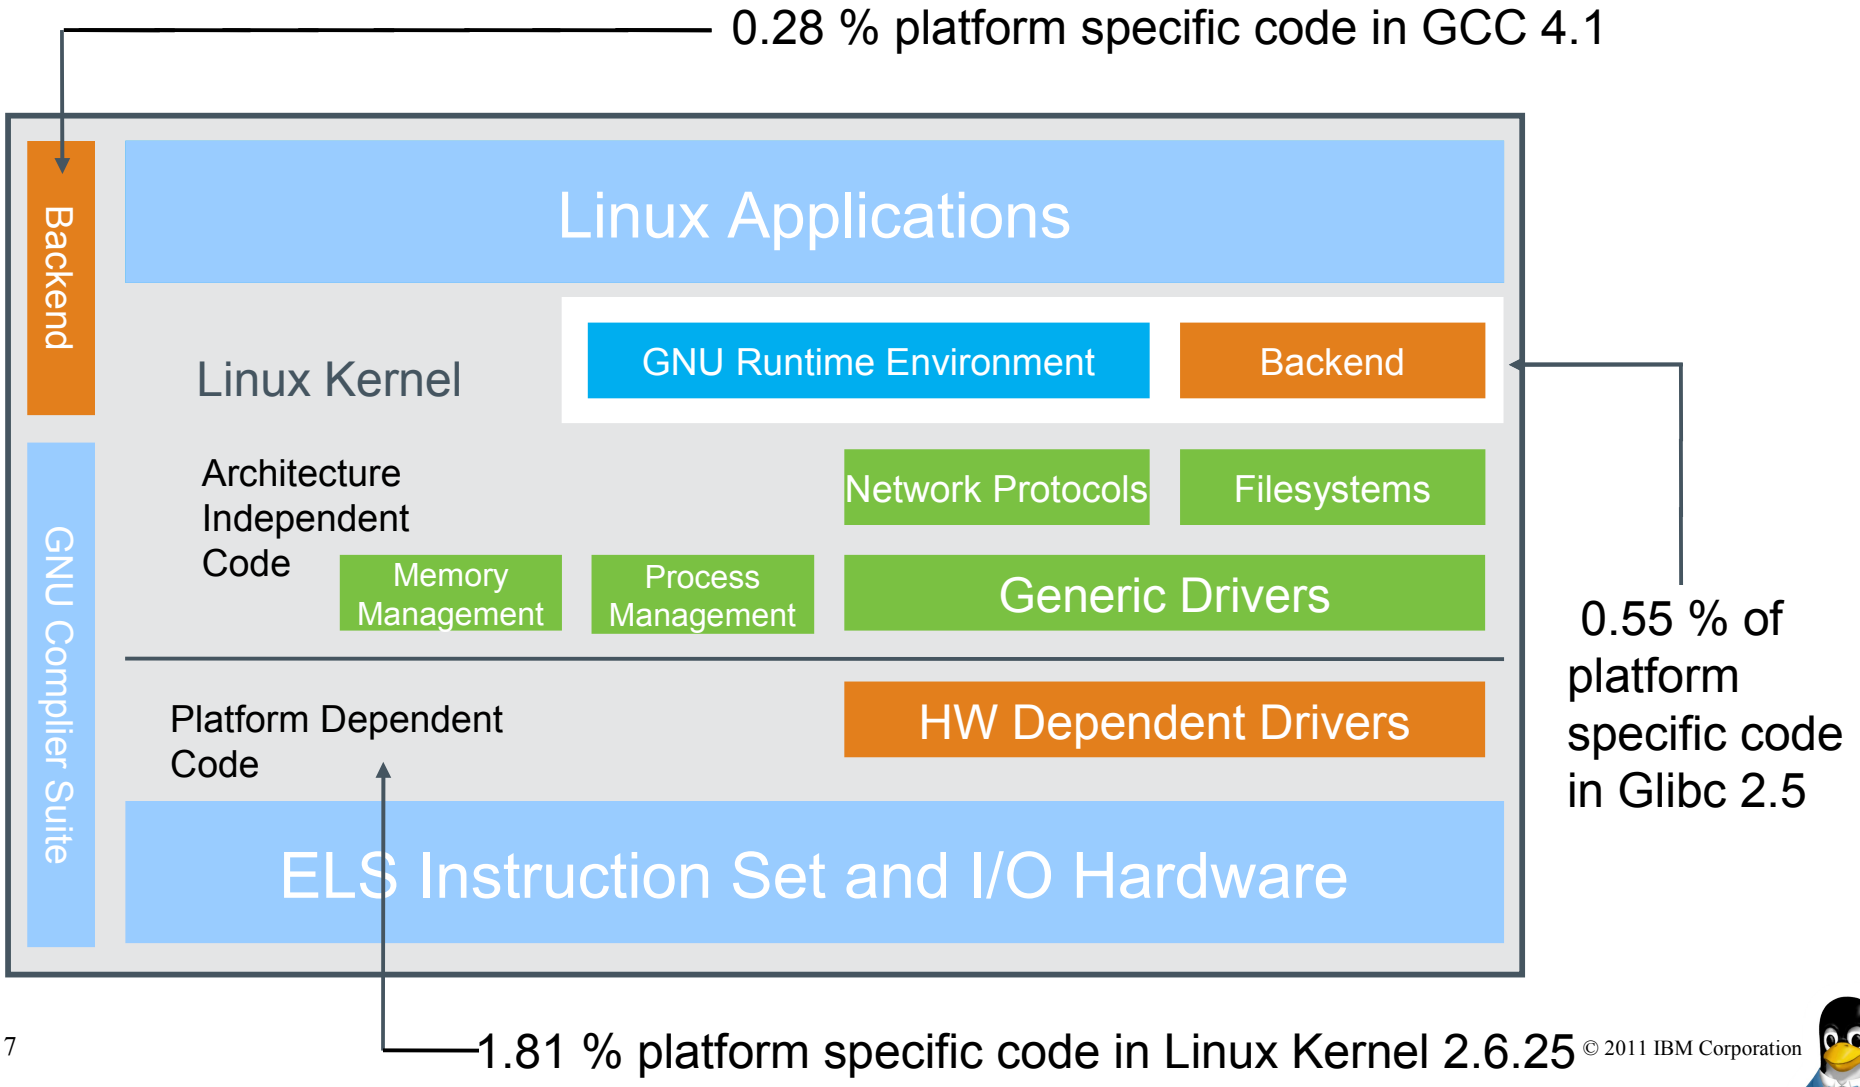

Linux kernel development: System z contributions

Changesets per 2.6.x kernel release

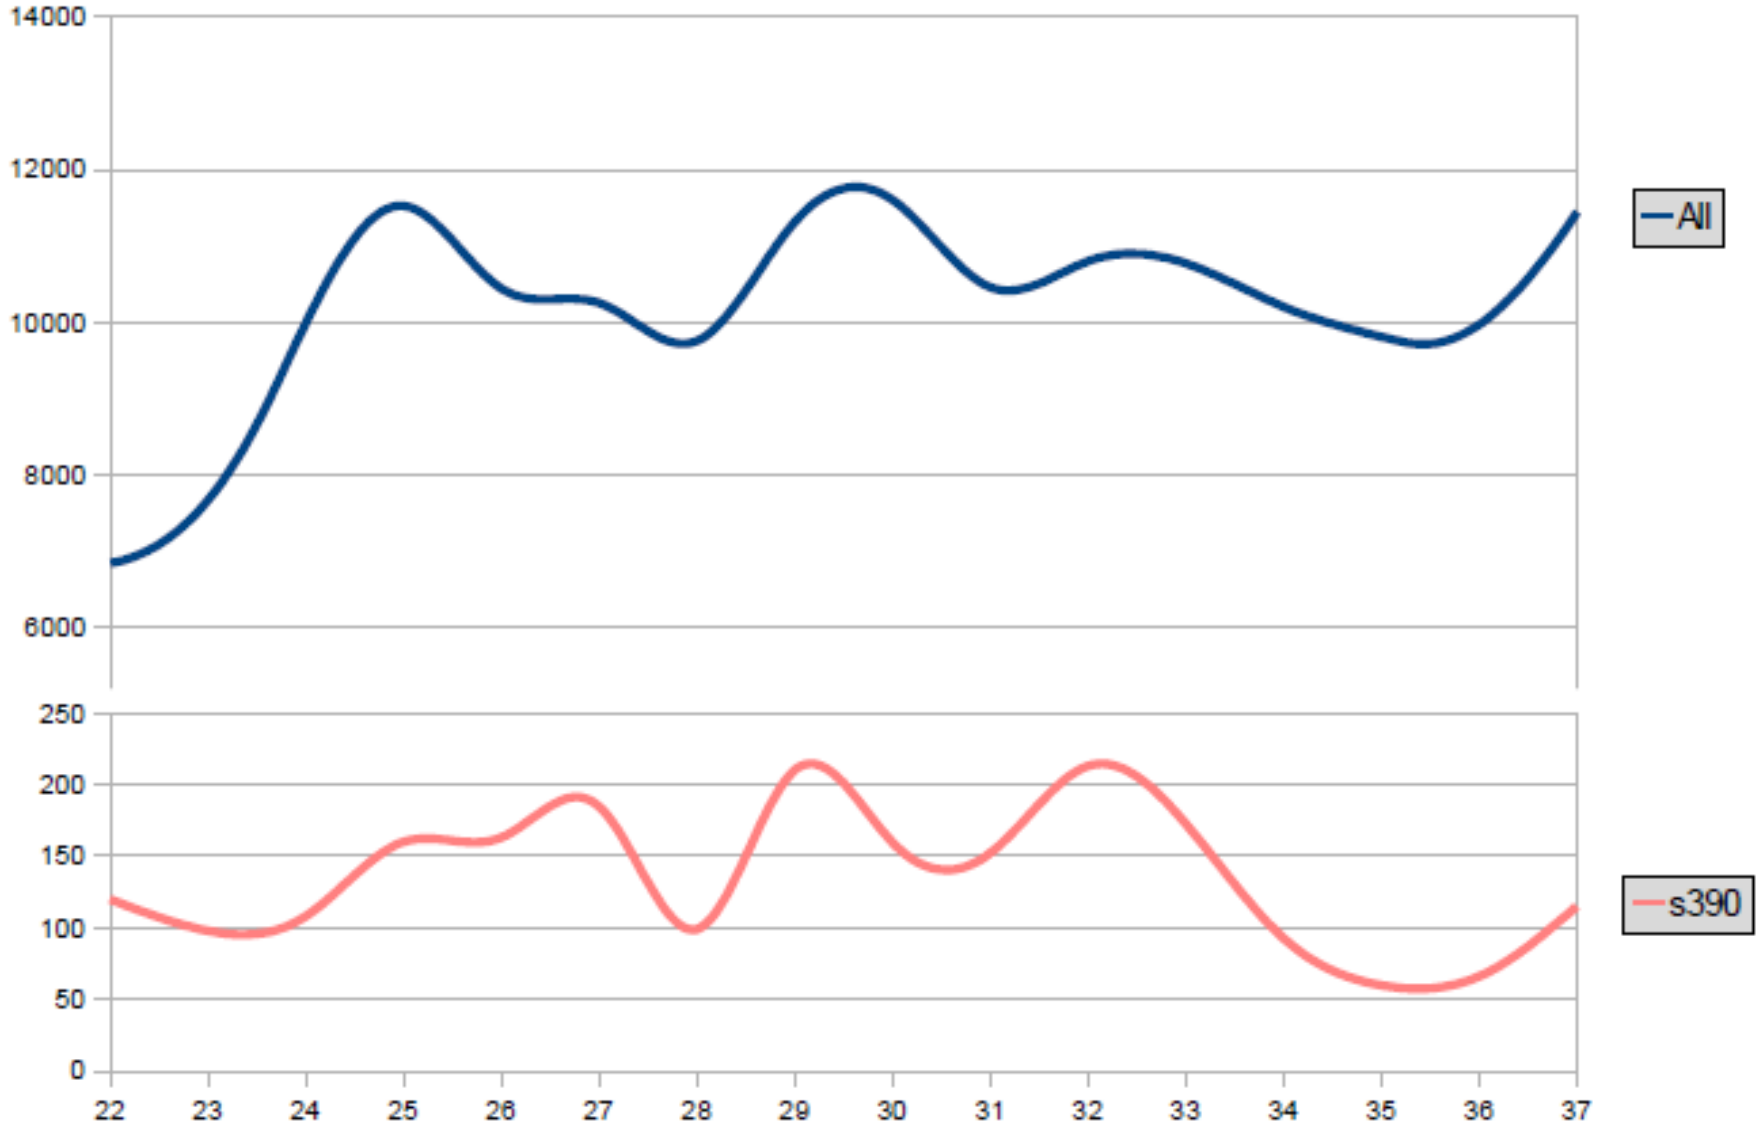

Linux on System z Development Focus

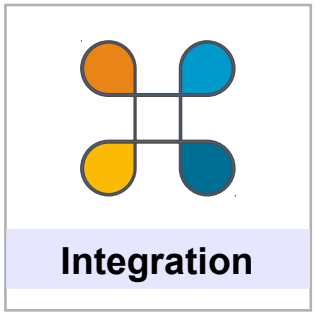

Integration

Application Serving

- z/OS & z/VSE integration

Data Hub

- Database Consolidation

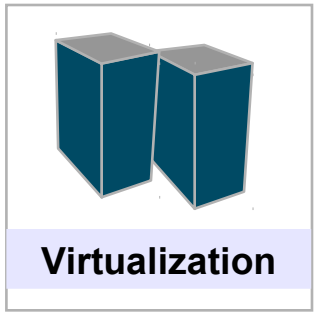

Virtualization

Virtualization &

Virtualization Management

- Ease of Use
- Serviceability
- Hosting capacity

Security

Security

- Certifications
- Data security & privacy

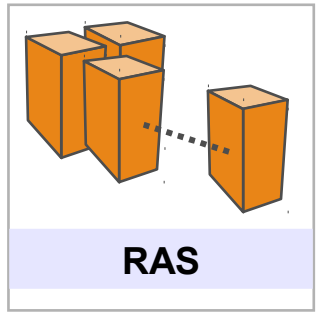

RAS

Continuous Availability

& Data Replication

- RAS
- Differentiation for mission critical workloads

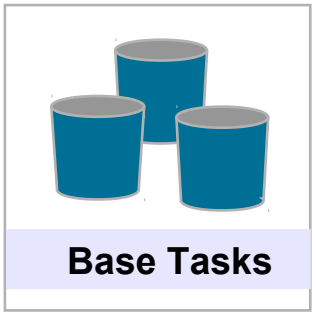

Base Tasks

Customer Requirements

- Address customer observed deficiencies

Competitiveness

- Close competitive gaps
- Differentiation / innovation that matters

Hardware Support

- Exploitation of new System z HW
- Storage exploitation

Linux

- Maintainership & code currency

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*

Kernel news – Common code

Linux version 2.6.35 (2010-08-01)

- Filesystems: btrfs improvements, XFS delayed logging
- Support for multiple multicast route tables
- Support for Layer 2 Tunneling Protocol L2TP Version 3
- Memory compaction

Linux version 2.6.36 (2010-10-20)

- Tilera architecture support
- Concurrency-managed workqueues
- Improve VM-related desktop responsiveness
- Integration of AppArmor
- New out-of-memory killer (OOM)

Linux version 2.6.37 (2011-01-04)

- Filesystems: better SMP scalability for ext4, XFS scalability improvements
- Removal of the BKL: Big Kernel Lock
- I/O throttling support for process groups
- Jump labels: performance optimization for disabled tracepoints

Linux version 2.6.38-rc4 (2011-02-07)

- Automatic process grouping (SCHED_AUTOGROUP)
- RCU-based path name lookup (dcache scalability)
- Transparent huge pages
- Transmit packet steering (XPS) for multiqueue devices

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 fill condition.
- **Breaking event address for user space programs (kernel 2.6.35)**
  Store the breaking-event-address for user space programs
  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.
- **QDIO outbound scan algorithm (kernel 2.6.38)**
  Improved scheduling of QDIO tasklets, OSA / HiperSockets / zfcp need different thresholds.
- **Precise process accounting (> kernel 2.6.38)**
  Extend the taskstats interface to provide better process accounting **values**
  Quality goal is a resolution of 10ths of microseconds in a snapshot over all tasks

System z Kernel Features - z/VM

- **Deliver z/VM CP special messages as uevent (kernel 2.6.34)**
  Allows to forward SMSG messages starting with “APP” to user space.udev rules can be used to trigger application specific actions
- **Automatic detection of read-only devices (2.6.34)**
  Improve usability by automatically detection of read-only dasd devices with diagnose 210
- **CMSFS user space file system support**
  (s390-tools 1.9.0 for the read-only cmsfs support) Implement a FUSE file system that allows to read from and write to CMSFS minidisks.
  Writing is difficult, the record based CMSFS does not fit well into the byte steam oriented Linux VFS
- **CMSFS configurable code page conversion (s390-tools 1.12.0)**
  Adds a configuration file to CMSFS that defines which CMS files are automatically converted from EBCDIC to ASCII.

How can you read files on a CMS disk with Linux?

About the CMS user space file system (fuse) support

11.1

- 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

```
root@larsson:~> cmsfs-fuse /dev/dasde /mnt/cms
root@larsson:~> ls -la /mnt/cms/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

```
root@larsson:~> 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

```
root@larsson:~> fusermount -u /mnt/cms
```

The future version will have write support

hyptop - Display hypervisor performance data

The hyptop command provides a dynamic real-time view of a hypervisor environment on System z.

- It works with both the z/VM and the LPAR PR/SM hypervisor.
- Depending on the available data it shows, for example, CPU and memory information about running LPARs or z/VM guest operating systems.

The following things are required to run hyptop:

- The debugfs file system must be mounted.
- The hyptop user must have read permission for the required debugfs files:
  - z/VM: <debugfs mount point>/s390_hypfs/diag_2fc
  - LPAR: <debugfs mount point>/s390_hypfs/diag_204
- To monitor all LPARs or z/VM guest operating systems of the hypervisor, your system must have additional permissions:
  - For z/VM: The guest must be privilege class B.
  - For LPAR: On the HMC or SE security menu of the LPAR activation profile, select the Global performance data control checkbox.

hyptop – Displaying hypervisor performance data

Displaying performance data for the z/VM hypervisor

```
10:11:56 CPU-T: UN(16)                                                ?=help
system    #cpu      cpu      Cpu+    online memuse memmax wcur
(str)      (#)      (%)      (hm)     (dhm)  (GiB)  (GiB)  (#)
T6360003     6   506.92  3404:17  44:20:53   7.99   8.00  100
T6360017     2   199.58     8:37  29:23:50   0.75   0.75  100
T6360004     6    99.84   989:37  62:00:00   1.33   2.00  100
T6360005     2     0.77     0:16   5:23:06   0.55   2.00  100
T6360015     4     0.15     9:42  18:23:04   0.34   0.75  100
T6360035     2     0.11     0:26   7:18:15   0.77   1.00  100
T6360027     2     0.07     2:53  62:21:46   0.75   0.75  100
T6360049     2     0.06     1:27  61:17:35   0.65   1.00  100
T6360010     6     0.06     5:55  61:20:56   0.83   1.00  100
T6360021     2     0.06     1:04  48:19:08   0.34   4.00  100
T6360048     2     0.04     0:27  49:00:51   0.29   1.00  100
T6360016     2     0.04     6:09  34:19:37   0.30   0.75  100
T6360008     2     0.04     3:49  47:23:10   0.35   0.75  100
T6360006     2     0.03     0:57  25:20:37   0.54   1.00  100
NSLCF1       1     0.01     0:02  62:21:46   0.03   0.25  500
VTAM         1     0.00     0:01  62:21:46   0.01   0.03  100
T6360023     2     0.00     0:04   6:21:20   0.46   0.75  100
PERFSVM      1     0.00     2:12   7:18:04   0.05   0.06    0
AUTOVM       1     0.00     0:03  62:21:46   0.00   0.03  100
FTPSERVE     1     0.00     0:00  62:21:47   0.01   0.03  100
TCPIP        1     0.00     0:01  62:21:47   0.01   0.12 3000
DATAMOVE     1     0.00     0:06  62:21:47   0.00   0.03  100
VMSERVU      1     0.00     0:00  62:21:47   0.00   0.03 1500
OPERSYMP     1     0.00     0:00  62:21:47   0.00   0.03  100
```

hyptop – Displaying hypervisor performance data

Displaying performance data for a single LPAR

```
10:16:59 H05LP30 CPU-T: IFL(18) CP(3) UN(2)                    ?=help
cpuid  type    cpu  mgm visual
(#)   (str)    (%)  (%) (vis)
0      IFL   29.34 0.72 |##############                              |
1      IFL   28.17 0.70 |##############                              |
2      IFL   32.86 0.74 |################                            |
3      IFL   31.29 0.75 |###############                             |
4      IFL   32.86 0.72 |################                            |
5      IFL   30.94 0.68 |###############                             |
6      IFL    0.00 0.00 |                                            |
7      IFL    0.00 0.00 |                                            |
8      IFL    0.00 0.00 |                                            |
9      IFL    0.00 0.00 |                                            |
=:V:N       185.46 4.30
```

Networking

- **OSA QDIO Data Connection Isolation (kernel 2.6.33)**
  Isolate data traffic from Linux on System z guests sharing an OSA card
  Communication between guests needs to go over via external entity

- **HiperSockets Network Traffic Analyser (kernel 2.6.34)**
  Trace HiperSockets network traffic for problem isolation and resolution.
  Supported for layer 2 and layer 3

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

- **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

Note: this list comprises selected items only

Networking (cont)

- **NAPI support for QDIO and QETH (> kernel 2.6.35)**
  Convert QETH to the NAPI interface, the “new” Linux networking API
  NAPI allows for transparent GRO (generic receive offload)

- **QETH debugging per single card (> kernel 2.6.35)**
  Split some of the global QETH debug areas into separate per-device areas
  Simplifies debugging for complex multi-homed configurations

- **Configuration tool for System z network devices (s390-tools 1.8.4)**
  Provide a shell script to ease configuration of System z network devices

- **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

Note: this list comprises selected items only

znetconf network device configuration tool

- Allows to list, add, remove & configure System z network devices
- For example: list all potential network devices:

```
root@larsson:~> znetconf -u
Device Ids                       Type Card Type CHPID Drv.
------------------------------------------------------------
0.0.f500,0.0.f501,0.0.f502 1731/01 OSA (QDIO) 00 qeth
0.0.f503,0.0.f504,0.0.f505 1731/01 OSA (QDIO) 01 qeth
```

- Configure device 0.0.f503

```
root@larsson:~> znetconf -a 0.0.f503
```

- Configure device 0.0.f503 in layer2 mode and portname “myport”

```
root@larsson:~> znetconf -a 0.0.f503 -o layer2=1 -o
portname=myport
```

- Remove network device 0.0.f503

```
root@larsson:~> znetconf -r 0.0.f503
```


New Linux on System z Storage Features (FICON)

- **Resume reordered devices (kernel 2.6.34)**
  Allow resume of a guest with different subchannels for individual devices
  Allow suspend of a system with devices in the disconnected state
- **Unit check handling (kernel 2.6.35)**
  Improve handling of unit checks for internal I/O started by the common-I/O layer.
  After a unit check certain setup steps need to be repeated, e.g. for PAV
- **Store I/O status and initiate logging (SIOSL) (kernel 2.6.36)**
  Enhance debug capability for FCP attached devices
  Enables operating system to detect unusual conditions on a device of channel path
- **Tunable default grace period for missing interrupts in DASD (kernel 2.6.36)**
  Provide a user interface to specify the timeout for missings interrupts for standard I/O operations on DASD

Note: this list comprises selected items only

New Linux on System z Storage Features (FICON) (cont)

- **Dynamic PAV toleration (kernel 2.6.35)**
  Tolerate dynamic Parallel Access Volume changes for base PAV
  System management tools can reassign PAV alias device to different base devices.
- **CHPID reconfiguration handling (kernel 2.6.37)**
  Update data structures after channel-path related information change
  Inform device drivers about relevant changes
- **Query DASD reservation status (kernel 2.6.37)**
  New DASd ioctl to read the 'Sense Path Group ID' data
  Allows to determine the reservation status of a DASD in relation to the current Linux
- **Multi-track extension for HPF (kernel 2.6.38)**
  Allows to read from and write to multiple tracks with a single CCW
- **Improve handling of stolen DASD reservation (kernel 2.6.38)**
  Provide alternatives to handle unit checks that indicate stolen reservations
  Fail any request to a device until it is set offline
  Queue I/O until reservation is release again

Note: this list comprises selected items only

New Linux on System z Storage Features (FICON) (cont)

- **Access to raw ECKD data from Linux (kernel 2.6.38)**
  This item allows to access ECKD disks in raw mode
  Use the 'dd' command to copy the disk level content of an ECKD disk to a Linux file, and vice versa.
  Storage array needs to support read-track and write-full-track command.

- **Automatic menu support in zipl (s390-tools 1.11.0)**
  Zipl option that will create a boot menu for all eligible non-menu sections in the zipl configuration file

- **reIPL from device-mapper devices (s390-tools 1.12.0)**
  The automatic re-IPL function only works with a physical device
  Enhance the zipl support for device-mapper devices to provide the name of the physical device if the zipl target is located on a logical device

Usability / RAS

- **Dump on panic – prevent reipl loop (s390-tools 1.8.4)**
  Delay arming of automatic reipl after dump
  Avoids dump 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 tape dump will be removed

- **Breaking event address for user space (kernel 2.6.35)**
  Store the breaking-event-address for user space programs
  Valuable aid in the analysis of wild branches

- **Precise process accounting (> kernel 2.6.36)**
  Extend the taskstats interface to provide better process accounting values
  Quality goal is a resolution of 10ths of microseconds

System z toolchain

- **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-orderpipeline of z196

- **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

- **Oprofile hardware customer mode sampling**
  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

- **Valgrind System z support**
  Valgrind is a generic framework for creating dynamic analysis tools
  Valgrind is in essence a virtual machine using just-in-time (JIT) compilation techniques
  Valgrind can be used for memory debugging, memory leak detection, and profiling (e.g. cachegrind)

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

- **valgrind –tool=massif**
  Profile heap usage, takes regular snapshots of program's heap
  Produces a graph showing heap usage over time

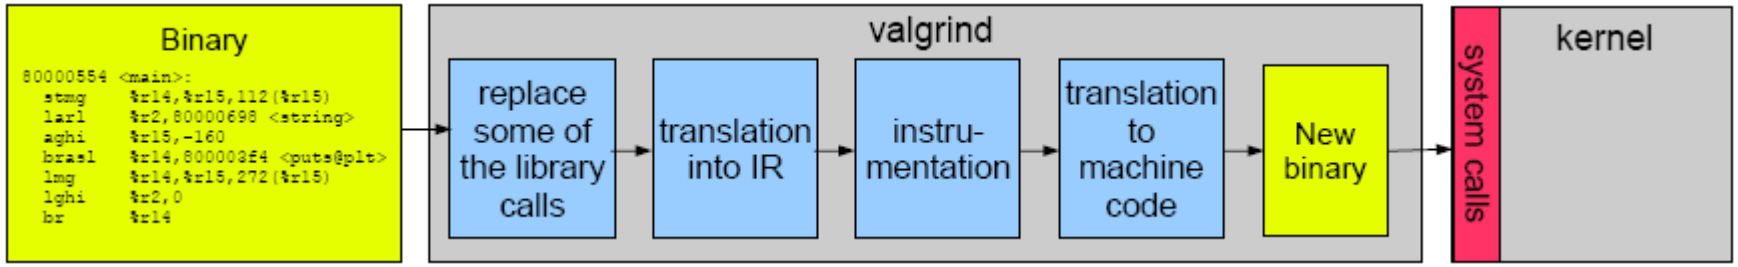

System z kernel features - Security

- **4096 bit RSA fast path (kernel 2.6.38)**
  Make use of 4096 bit RSA acceleration available with Crypto Express 3 GA2 cards.

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

- **New libica APIs for supported crypto modes**
  Provide a programmatic way to query for supported crypto ciphers, modes and key sizes.
  Deliver information whether the cryptographic features are implemented in hardware or in software

Current Linux on System z Technology

Features & Functionality contained in the Novell & Red Hat Distributions

Integration

- **AF_IUCV SOCK_SEQPACKET support**
  Introduce AF_IUCV sockets of type SOCK_SEQPACKET that map read/write operations to a single IUCV operation.
  The socket data is not fragmented.
  The intention is to help application developers who write applications using the native IUCV interface, e.g. Linux to z/VSE.

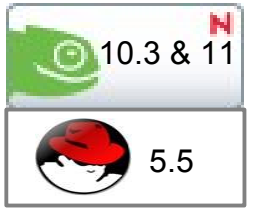

- **HiperSockets Layer3 support for Ipv6**
  Providing Layer3 IPv6 communication, for communication to z/OS

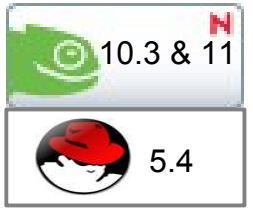

- **Linux to add Call Home data if running in LPAR**
  Also referred to as Control Program Identification (CPI) or SCLP_CPI
  Allows the user to set information about the LPAR which will be displayed on the HMC/SE

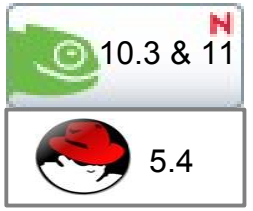

Note: this list comprises selected items only

Virtualization

- **TTY terminal server over IUCV**

  Provide central access to the Linux console for the different guests of a z/VM.

  Fullscreen applications like *vi* are usable on the console.

  Access Linux instances with no external network because IUCV is independent from TCP/IP

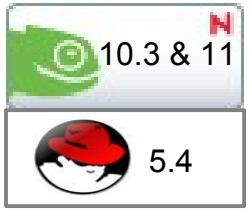

- **Dynamic memory attach/detach**

  Allows to attach/detach memory for Linux as a guest without needing to reipl.

- **Extra kernel parameter via VMPARM**

  Allows to use z/VM VMPARM variable to add or substitute the kernel command line.

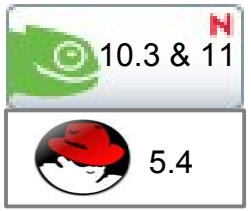

- **Provide CMS script for initial IPL**

  Avoids having to create an script to start a new installation under z/VM.

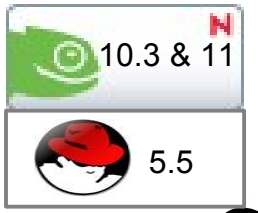

Note: this list comprises selected items only

IUCV terminal environment

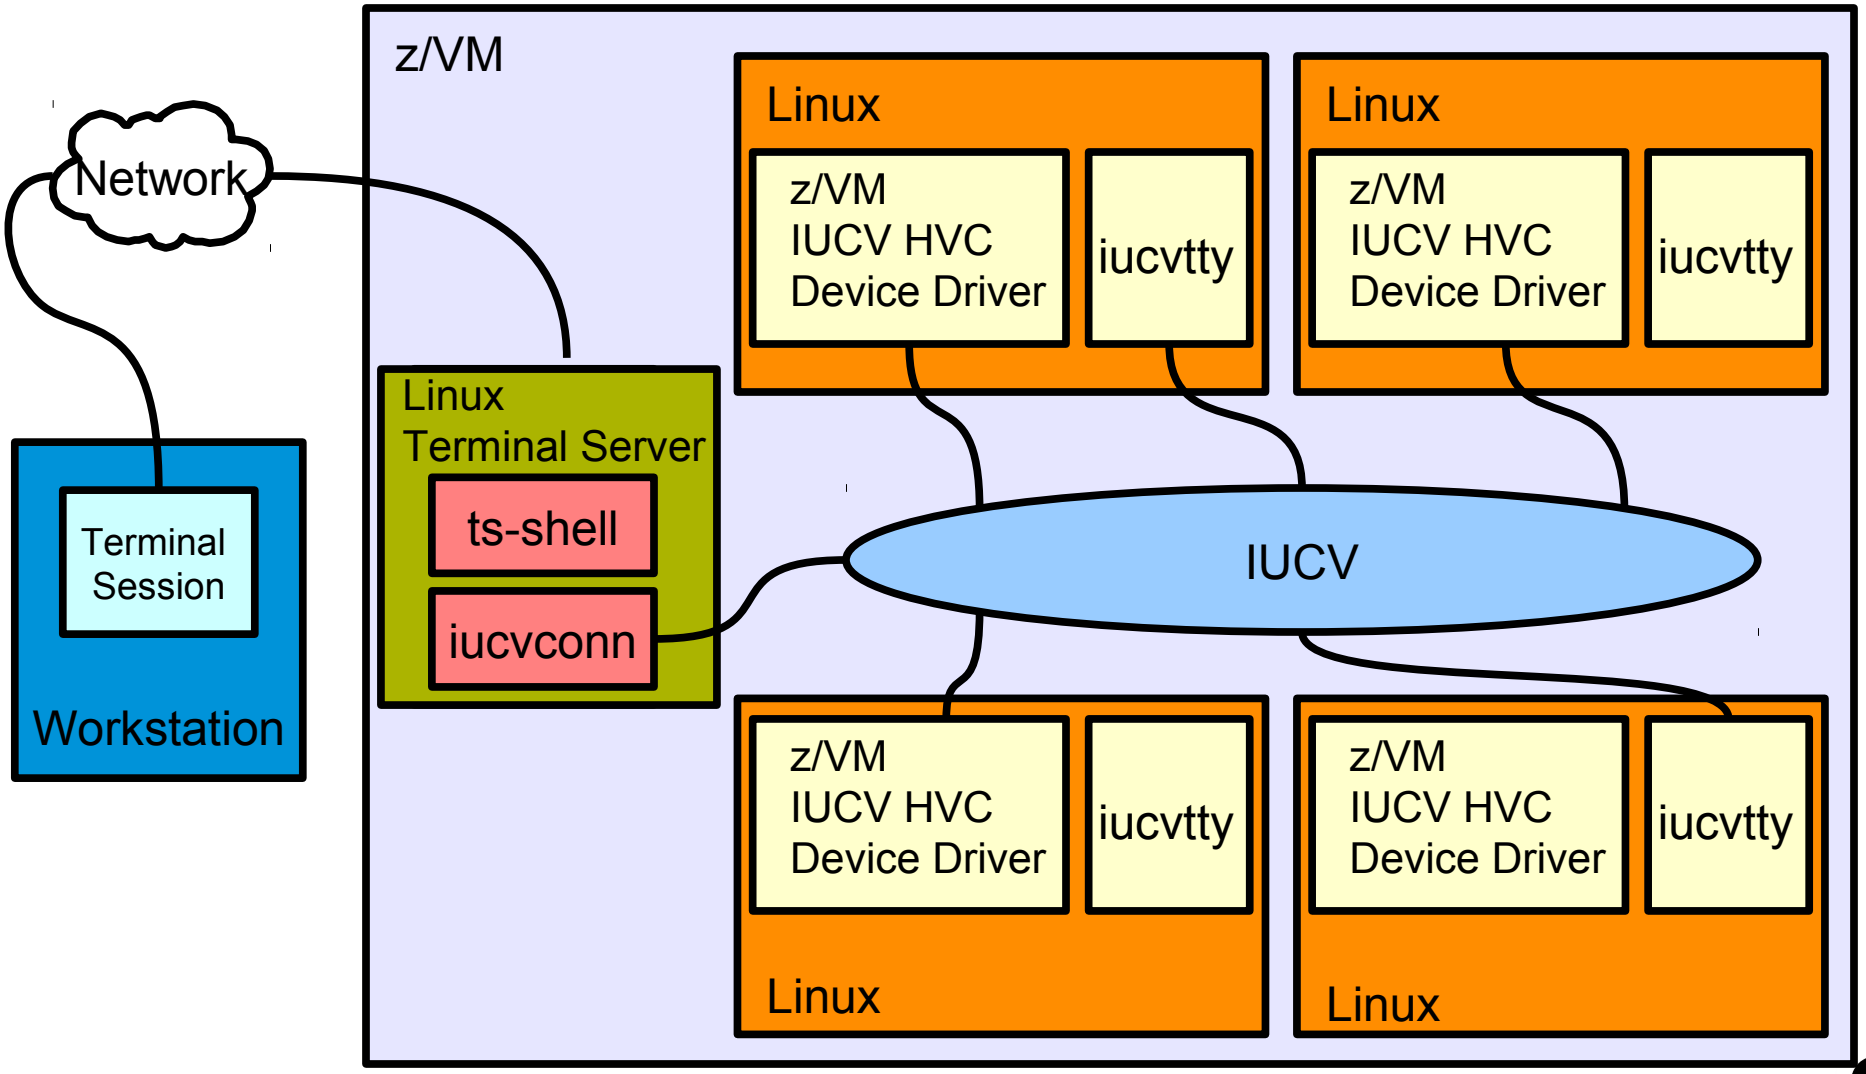

Virtualization (cont.)

- **Exploitation of DCSSs above 2G**
  Solves restriction to use DCSS above or greater than 2GB.

- **Provide service levels of HW & Hypervisor in Linux**
  Improves serviceability by providing uCode and z/VM levels via /proc interface

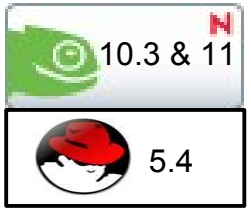


```
root@larsson:~> cat /proc/service_levels
VM: z/VM Version 5 Release 2.0
service level 0801(64-bit)
qeth: 0.0.f5f0 firmware level 087d
```


Note: this list comprises selected items only

Security

- **Long Random Numbers Generation**
  Provide access to the random number generator feature on the Crypto card (high volume random number generation, compared to a CPU based solution)

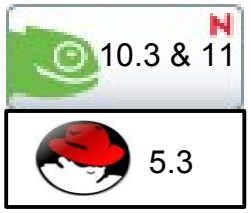

- **Crypto Express3 cards enablement**
  Support for Crypto Express3 Accelerator (CEX3A) and Crypto Express3 Coprocessor (CEX3C)

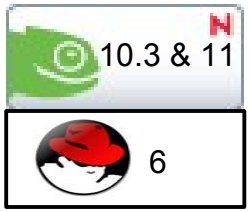

- **Crypto device driver use of thin interrupts**
  Provides better performance and lower CPU consumption.

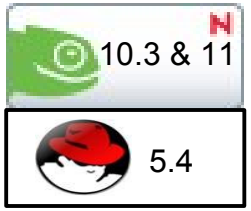

Note: this list comprises selected items only

RAS

- **Shutdown Actions Interface**
  The shutdown actions interface allows the specification of a certain shutdown action (stop, ipl, reipl, dump, vmcmd) for each shutdown trigger (halt, power off, reboot, panic)
  Possible use cases are e.g. to specify that a vmdump should be automatically triggered in case of a kernel panic or the z/VM logoff command should be executed on halt.

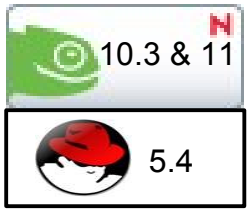

- **Automatic IPL after dump**
  The new shutdown action dump_reipl introduces a system configurations which allows to create a dump in case of a Linux panic, followed by a re-ipl of the system, once the dump was successfully created.
  Allows to configure system to re-ipl after a dump is taken.

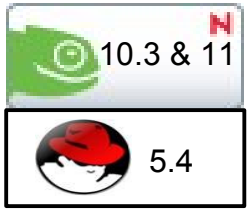

Note: this list comprises selected items only

RAS

- **Suspend / resume support (kernel 2.6.31)**

  Add the ability to stop a running Linux system and resume operations later on. The image is stored on the swap device and does not use any system resource while suspended.

  Only suspend to disk is implemented, suspend to RAM is not supported.

- **Add Call Home data on halt and panic if running in LPAR (kernel 2.6.32)**

  Report system failures (kernel panic) via the service element to the IBM service organization. Improves service for customers with a corresponding service contract. (by default this features is deactivated)

- **Large image dump on DASD**

  Solves restriction to dump only 48GB of memory to DASD. Now up to 32 ECKD DASDs can be used in a multiple volume configuration

Note: this list comprises selected items only

Suspend / resume support

- Ability to stop a running Linux on System z instance and later continue operations
- Memory image is stored on the swap device specified with a kernel parameter: **`resume=/dev/dasd<x>`**
- Lower the swap device priority for the resume partition

```
root@larsson:~> grep swap /etc/fstab
/dev/dasdb1 swap swap pri=-1 0 0
/dev/dasdc1 swap swap pri=-2 0 0
```

- Suspend operation is started with a simple echo:

```
root@larsson:~> echo disk > /sys/power/state
```

- Resume is done automatically on next IPL
- Use signal quiesce to automatically suspend a guest

```
ca::ctrlaltdel:/bin/sh -c "/bin/echo disk > \
/sys/power/state || /sbin/shutdown -t3 -h now"
```


I/O

- **High Performance FICON**

  Adds support for the zHPF protocol to the DASD driver. zHPF provides a much simpler link protocol than FICON: Promises increased I/O bandwidth due to better channel utilization
  This features is available with DS8000 R4.1

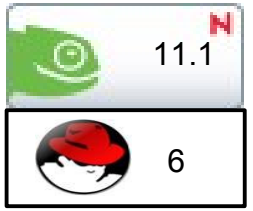

- **FCP - SCSI error recovery hardening**

  Improve error recovery cooperation between SCSI-mid-layer and zFCP by allowing the SCSI error recovery to wait for completion of the zFCP error recovery. This Increases the stability and availability in scenarios of error recovery, for example firmware/uCode upgrades.

- **FICON DS8000 Large Volume (EAV) Support**

  Large Volume Support is a feature that allows to use ECKD devices with more than 65520 cylinders (>50GB).
  This features is available with DS8000 R4.0 Allows to exploit

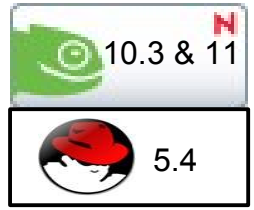

Note: this list comprises selected items only

Where to Find More Information

More Information

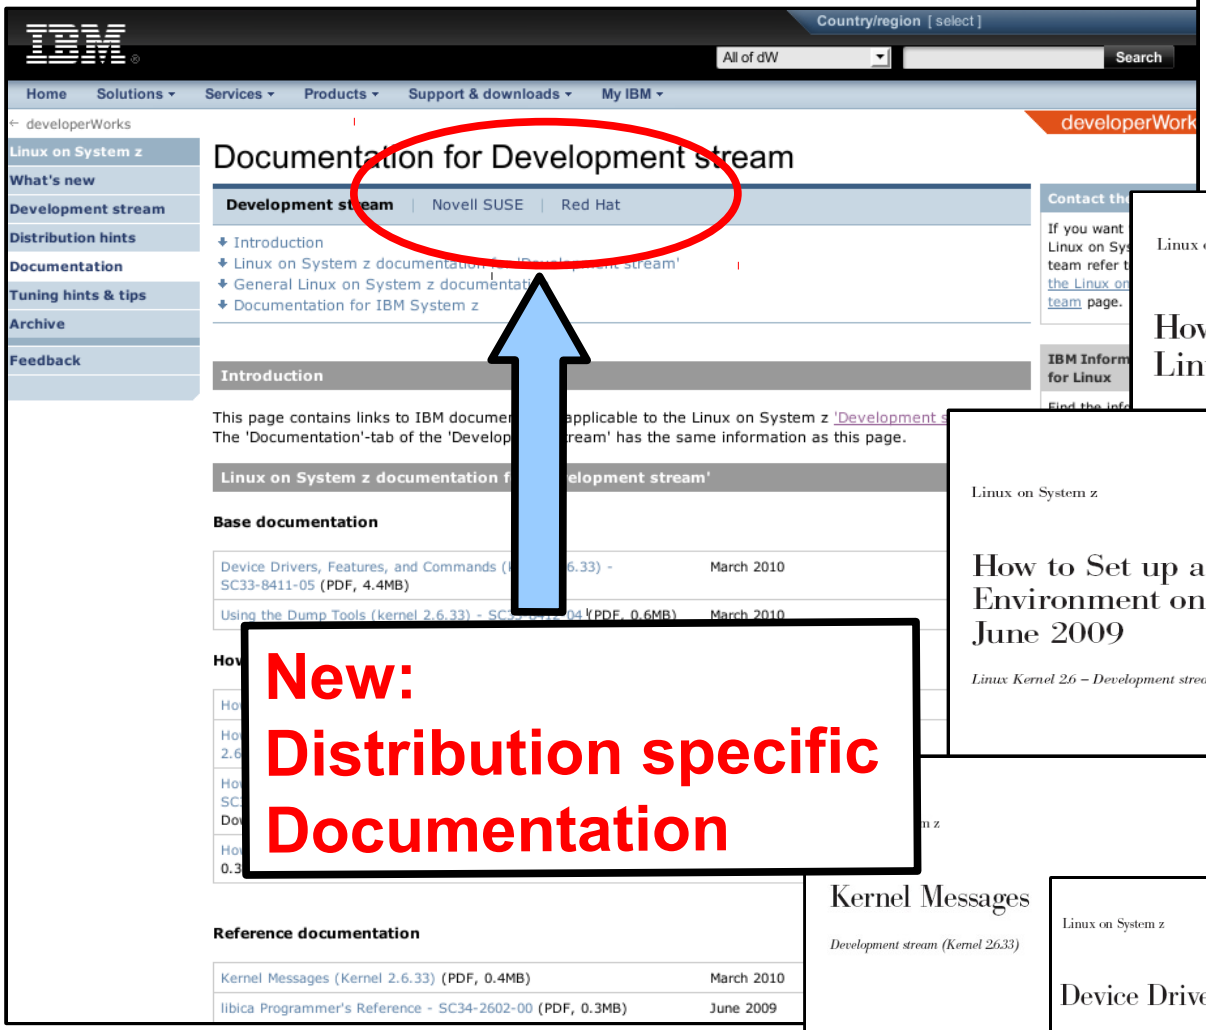

New:
Distribution specific
Documentation

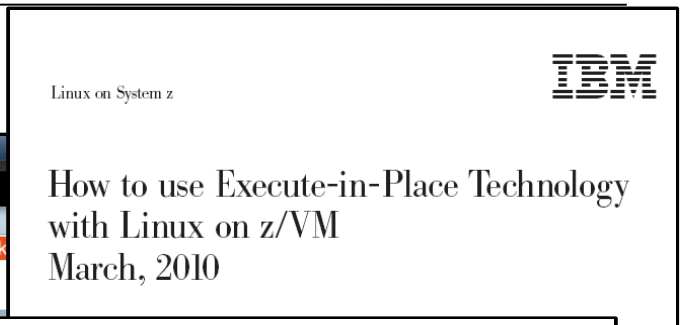

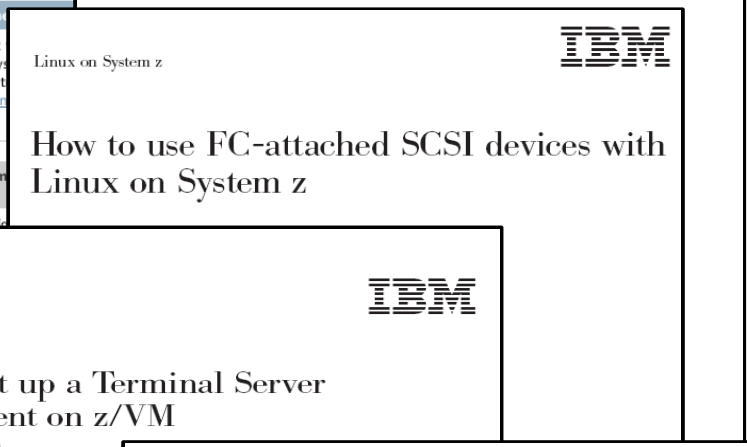

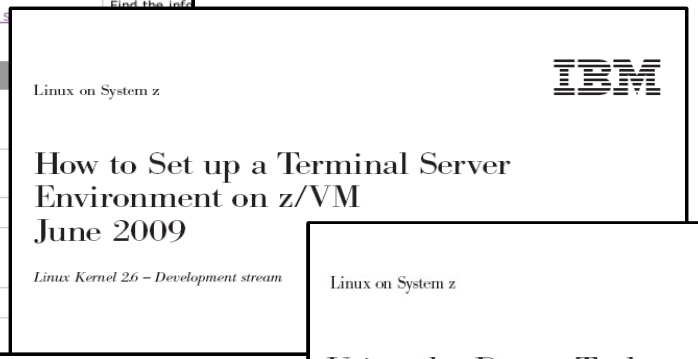

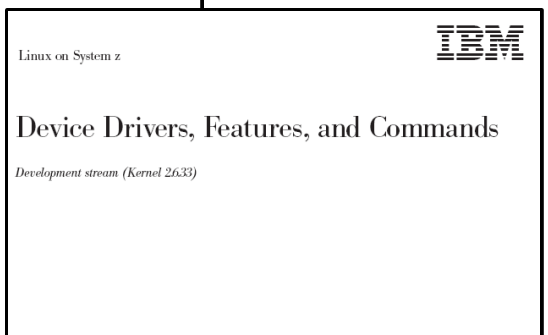

2 new Redbooks have just been released!

Visit http://www.redbooks.ibm.com

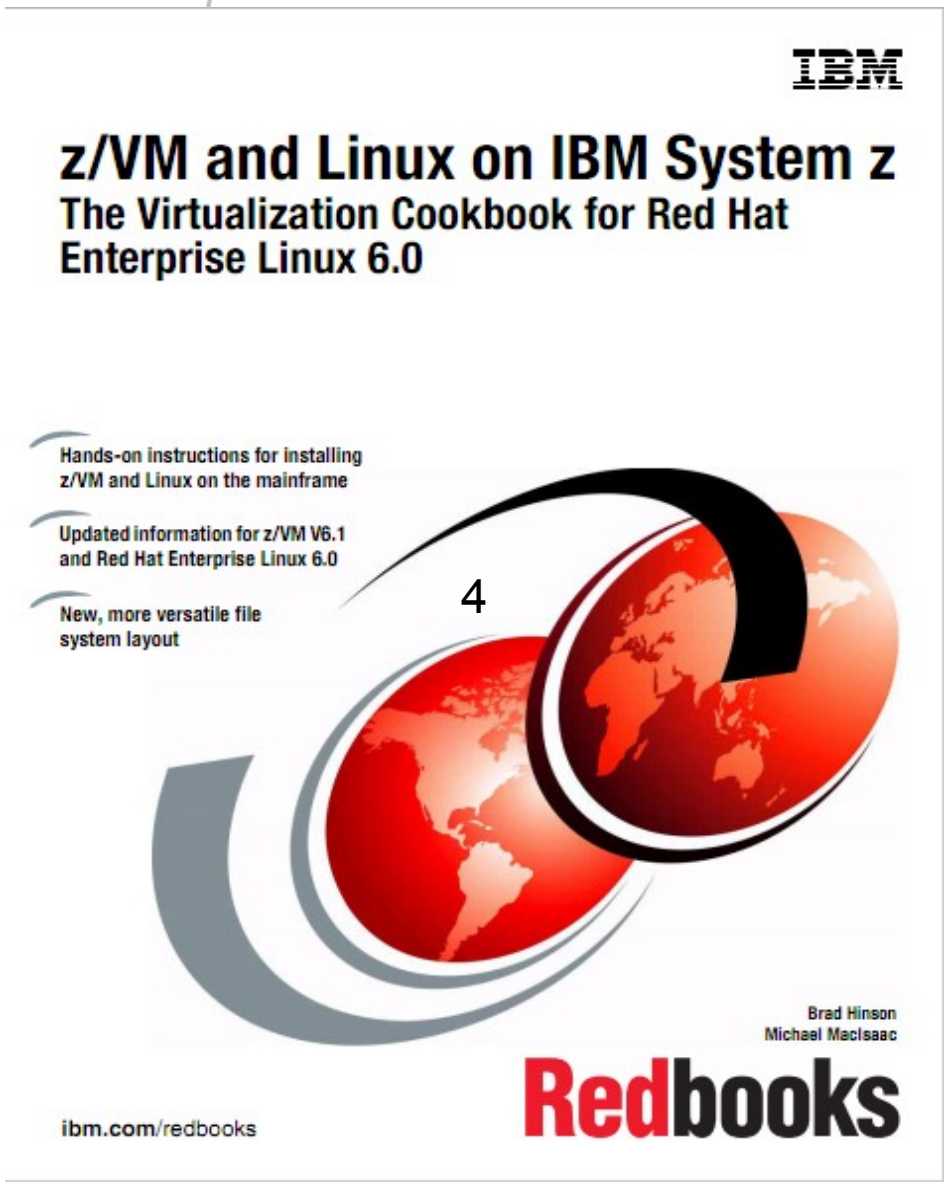

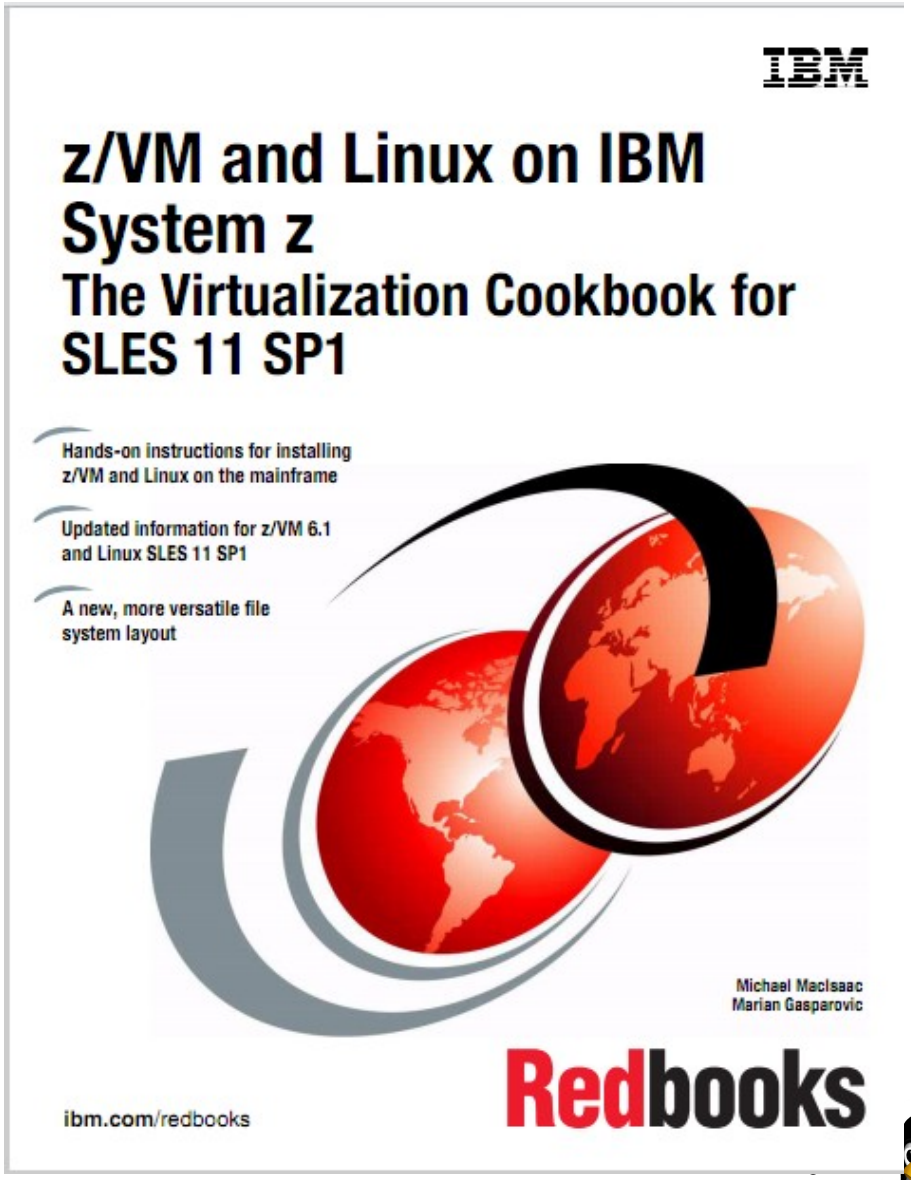

Your Linux on System z Requirements?

Are you missing a certain feature, functionality or tool? **We'd love to hear from you!**

We will evaluate each request and (hopefully) develop the additional functionality you need.

Send your input to hans@de.ibm.com

Questions?

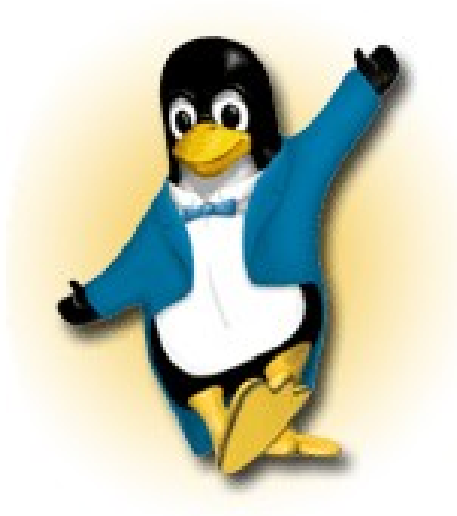

Hans-Joachim Picht

Linux on System z Initiatives

IBM Deutschland Research & Development GmbH
Schönaicher Strasse 220
71032 Böblingen, Germany

Mobile +49 (0)175 - 1629201
hans@de.ibm.com

How to explain the benefits of running Linux on System z in 2:39?

http://www.youtube.com/watch?v=0i7kBnhN3Lg

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.