


What's New in Linux on System z



Martin Schwidefsky
IBM Lab Böblingen, Germany
August 8 2012



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:

IBM, the IBM logo, BladeCenter, Calibrated Vecteded Cooling, ClusterProven, Cool Blue, POWER, PowerExecutive, Predictive Failure Analysis, ServerProven, System p, System Storage, System x, System z, WebSphere, DB2 and Tivoli are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see <http://www.ibm.com/legal/copytrade.shtml>.

The following are trademarks or registered trademarks of other companies: Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both. 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. UNIX is a registered trademark of The Open Group in the United States and other countries or both. Linux is a trademark of Linus Torvalds in the United States, other countries, or both. Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc. InfiniBand is a trademark of the InfiniBand Trade Association.

Other company, product, or service names may be trademarks or service marks of others.

NOTES: Linux penguin image courtesy of Larry Ewing (lewing@isc.tamu.edu) and The GIMP

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment. IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided "AS IS" without warranty of any kind. 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 are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography. 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

How Linux on System z is developed

How does the “community” work.

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

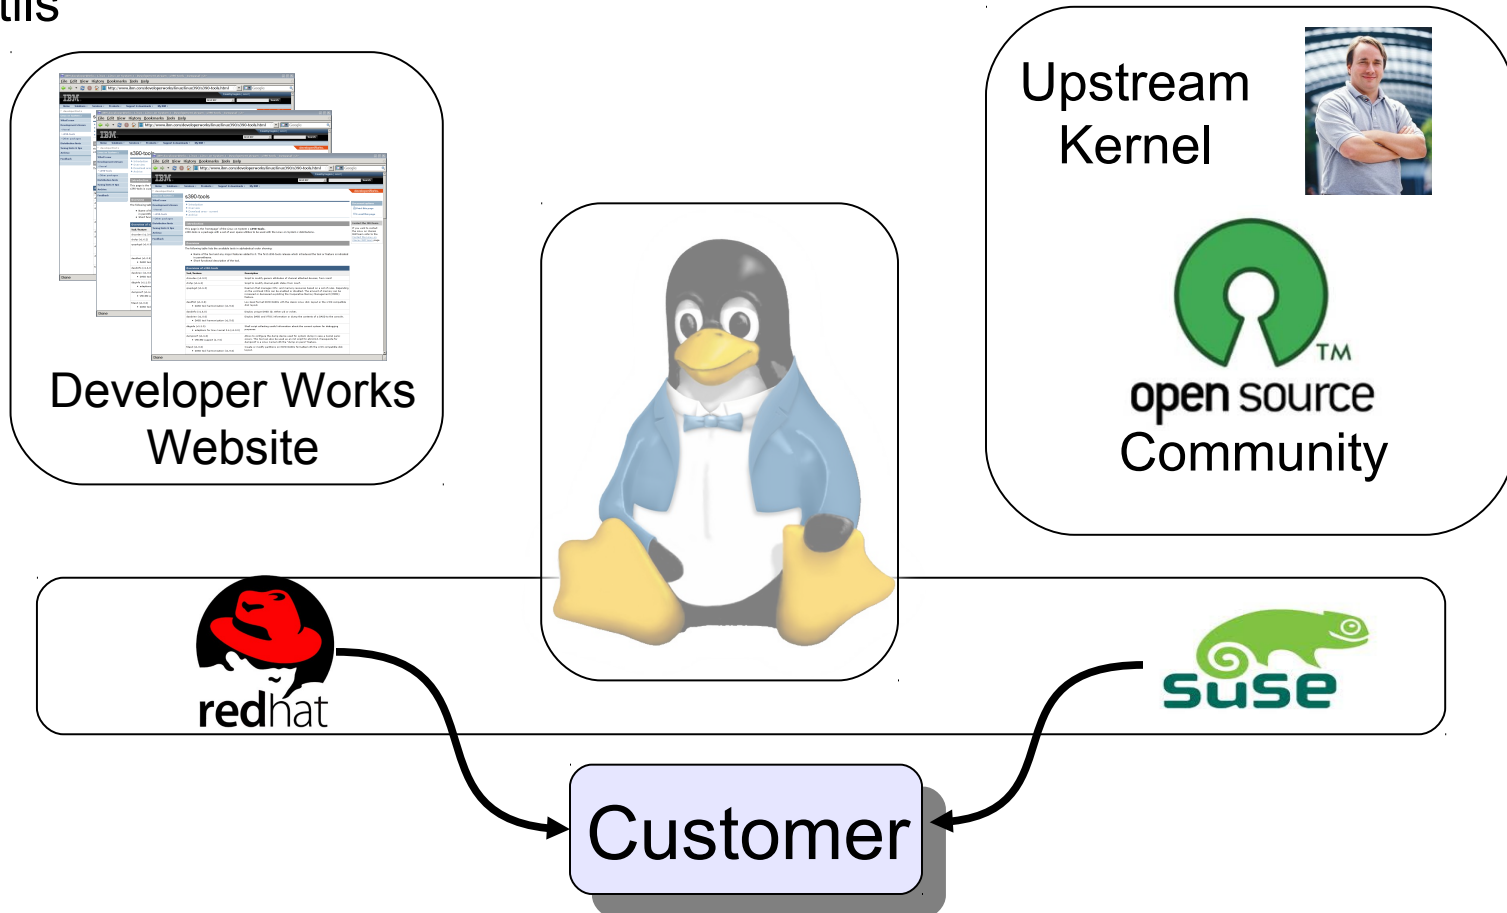
The Linux Foundation
Linux Standards Base
Common Criteria certification,
and more...

Foster and Protect the Ecosystem

Software Freedom Law Center
Free Software Foundation (FSF),
and more...

The IBM Linux development process

IBM Linux on System z development contributes in the following areas:
Kernel, s390-tools, open source tools (e.g. eclipse, ooprofile), gcc, glibc, binutils




Facts on Linux

- Linux kernel 1.0.0 was released with 176,250 lines of code
How many lines of code has the kernel version 3.2 ?
14,998,737 lines of code
- How many of the world's top 500 supercomputers run Linux (Jan 2012)?
457 / 91.4%
- What percentage of web servers run Linux (Jan 2012) ?
63.6% run Unix, of those 51.6% run Linux (46.5% unknown) = 32.8%
- What percentage of desktop clients run Linux (Jan 2012) ?
1.6%
- What is the largest Linux architecture in number of devices ?
ARM, > 100 million activated android devices
- **Linux is Linux**, but ...features, properties and quality differ dependent on your platform and your use case

Source: <http://kernel.org>
<http://top500.org/stats>
<http://w3techs.com>
<http://www.w3counter.com>
<http://googleblog.blogspot.com/2011/05/android-momentum-mobile-and-more-at.html>


Linux kernel development: rate of change

Average for the last 7 1/2 years (without renames): 79 days per release, 5910 lines added, 2621 lines removed and 1251 lines modified **per day**



Linux kernel development: System z contributions

Changesets per 2.6.x / 3.x kernel release



Linux on System z distributions (Kernel 2.6 based)

- **SUSE Linux Enterprise Server 9 (GA 08/2004)**
 - Kernel 2.6.5, GCC 3.3.3, Service Pack 4 (GA 12/2007), end of regular life cycle
- **SUSE Linux Enterprise Server 10 (GA 07/2006)**
 - Kernel 2.6.16, GCC 4.1.0, Service Pack 4 (GA 05/2011)
- **SUSE Linux Enterprise Server 11 (GA 03/2009)**
 - Kernel 2.6.27, GCC 4.3.3, Service Pack 1 (GA 06/2010), Kernel 2.6.32
 - Kernel 3.0.13, GCC 4.3.4, Service Pack 2 (GA 02/2012)
- **Red Hat Enterprise Linux AS 4 (GA 02/2005)**
 - Kernel 2.6.9, GCC 3.4.3, Update 9 (GA 02/2011), end of regular life cycle
- **Red Hat Enterprise Linux AS 5 (GA 03/2007)**
 - Kernel 2.6.18, GCC 4.1.0, Update 8 (GA 02/2012)
- **Red Hat Enterprise Linux AS 6 (GA 11/2010)**
 - Kernel 2.6.32, GCC 4.4.0 Update 2 (GA 12/2011)
- **Others**
 - Debian, Slackware,
 - Support may be available by some third party

Supported Linux Distributions

Distribution	zEnterprise – z114 and z196	System z10	System z9	zSeries
RHEL 6	✓	✓	✓	X
RHEL 5	✓	✓	✓	✓
RHEL 4 (*)	✓ (1)	✓	✓	✓
SLES 11	✓	✓	✓	X
SLES 10	✓	✓	✓	✓
SLES 9 (*)	✓ (2)	✓	✓	✓



Two options for zSeries machines



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.

(1) 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.

(2) SLES 9 SP4 + 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.



Also available as 31-bit distribution.

Current Linux on System z Technology

Features & Functionality contained in the SUSE
& Red Hat Distributions

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



System z kernel features – Core

- **QDIO outbound scan algorithm (kernel 2.6.38)**



- Improve scheduling of QDIO tasklets, OSA / HiperSockets / zfcps need different thresholds

- **Enabling spinning mutex (kernel 2.6.38)**



- Make use of the common code for adaptive mutexes.
- Add a new architecture primitive `arch_mutex_cpu_relax` to exploit sigp sense running to avoid the mutex lock retries if the hypervisor has not scheduled the cpu holding the mutex.

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
```

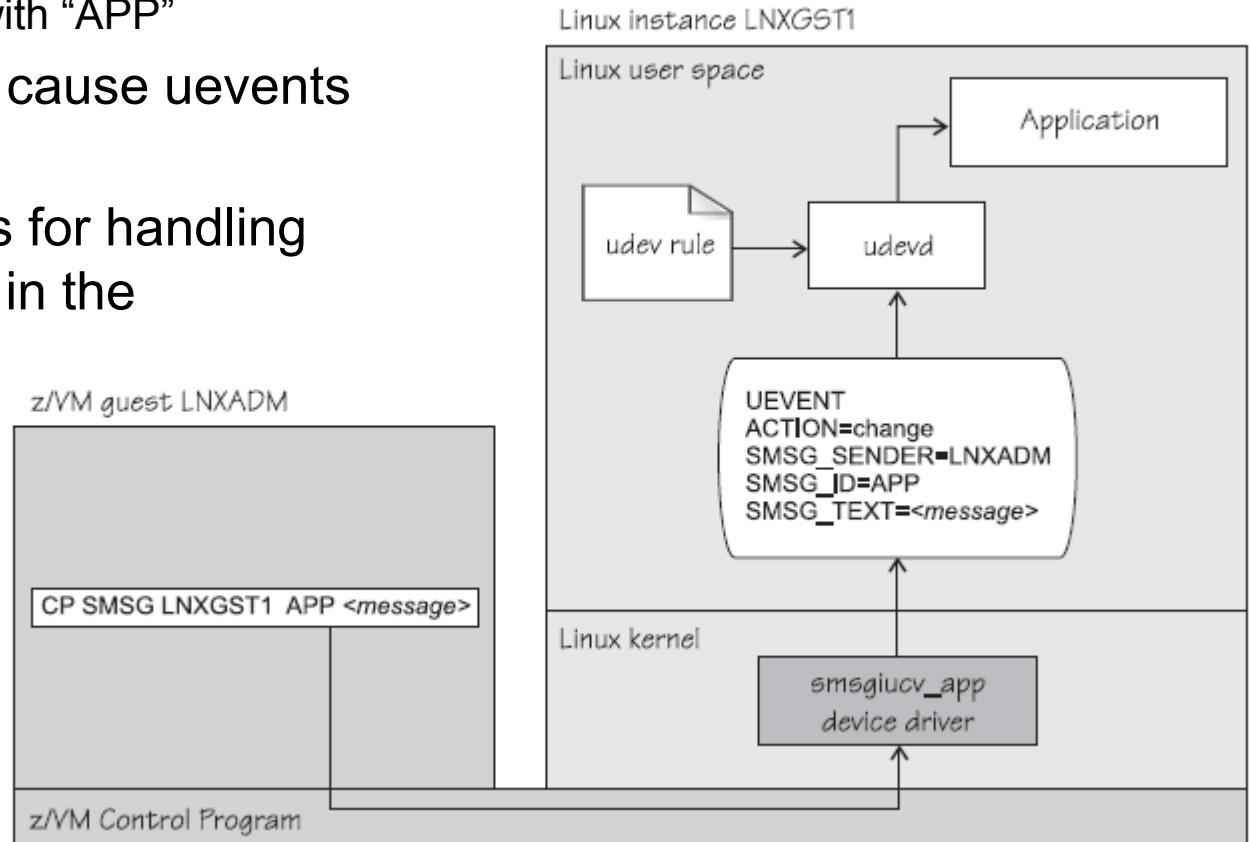
- Write support is work in progress, almost completed
 - use “vi” to edit PROFILE.EXEC anyone ?
- Use fusermount to unmount the file system again

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

Deliver z/VM CP special messages as uevent



- 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 – FICON

■ **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

■ **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

■ **Tunable default grace period for missing interrupts in DASD (kernel 2.6.36)**



- Provide a user interface to specify the timeout for missing interrupts for standard I/O operations on DASD

■ **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 system

System z kernel features – FICON

■ **Multi-track extension for HPF (kernel 2.6.38)**

- Allows to read from and write to multiple tracks with a single CCW



■ **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 to create a boot menu for all eligible non-menu sections in zipl.conf



■ **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



System z kernel features – FCP

▪ **Store I/O and initiate logging (SIOSL) (kernel 2.6.36)**



6.1



11.2

- Enhance debug capability for FCP attached devices
- Enables operating system to detect unusual conditions on a FCP channel

▪ **Add NPIV information to symbolic port name (kernel 2.6.39)**



11.2

- Add the device bus-ID and the network node to the symbolic port name if the NPIV mode is active.

▪ **SAN utilities (kernel 2.6.36, lib-zfcp-hbaapi 2.1)**



6.1



11.2

- Two new utilities have been added: zfcp_ping and zfcp_show
- They are useful to discover a storage area network

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)

System z kernel features – Networking

▪ QETH debugging per single card (kernel 2.6.36)

- Split some of the global QETH debug areas into separate per-device areas
- Simplifies debugging for complex multi-homed configurations



▪ Support for assisted VLAN null tagging (kernel 2.6.37)

- Close a gap between OSA and Linux to process null tagged frames correctly
- z/OS may sent null-tagged frames to Linux



▪ New default qeth configuration values (kernel 2.6.39)

- Receive checksum offload, generic receive offload & number of inbound buffers



▪ IPv6 support for the qetharp tool (kernel 2.6.38)

- Extend the qetharp tool to provide IPv6 information in case of a layer 3 setup.
- This is required for communication with z/OS via HiperSockets using IPv6.



System z kernel features – Networking

- **Add OSA concurrent hardware trap (kernel 3.0)**

- To ease problem determination the qeth driver requests a hardware trace when the device driver or the hardware detect an error
- Allows to correlate between OSA and Linux traces.



- **Configuration tool for System z network devices (s390-tools 1.8.4)**

- Provide a shell script to ease configuration of System z network devices

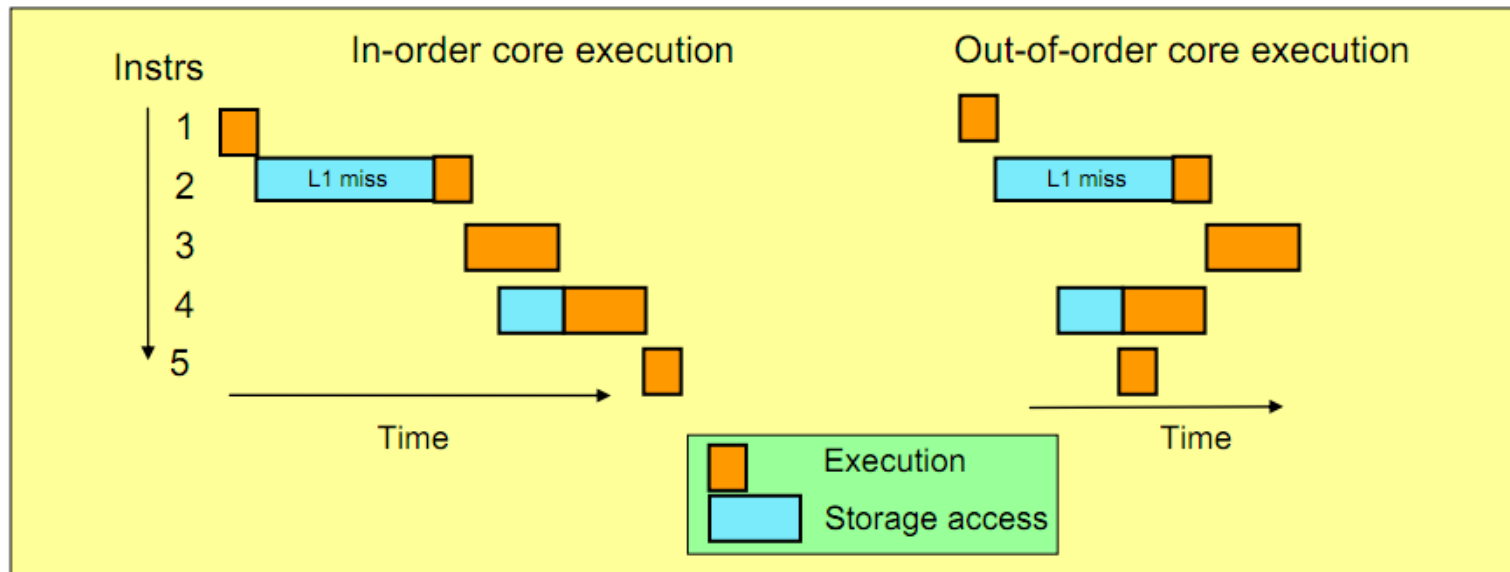


System z toolchain

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-order pipeline of z196
- Re-compiled code/apps get further performance gains through 110+ new instructions



System z kernel features – Crypto

■ 4096 bit RSA fast path (kernel 2.6.38)

- Make use of 4096 bit RSA acceleration available with Crypto Express 3 GA2 cards.



■ CP ACF exploitation of System z196 (kernel 3.0)

- Add support for new HW crypto modes:
cipher feedback mode (CFB), output feedback mode (OFB),
counter mode (CTR), Galois counter mode (GCM),
XEX based Tweaked Code Book with Cipher Text Stealing (XTS),
cipher based message authentication mode (CMAC),
and counter with cipher block chaining message authentication (CCM)



■ New libica APIs for supported crypto modes (libica 2.1.1)

- 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



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
- <http://lnxhc.sourceforge.net/>

Future Linux on System z Technology

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

System z kernel features – core kernel

- **Add support for physical memory > 4TB (kernel 3.3)**
 - Increase the maximum supported memory size from 4TB to 64TB.
- **Two stage dumper / kdump support (kernel 3.2, s390-tools-1.17.0)**
 - Use a Linux kernel to create a system dump
 - Use a preloaded crashkernel to run in case of a system failure
 - Can be triggered either as panic action or by the stand-alone dumper, integrated into the shutdown actions framework
 - Pro
 - Enhanced dump support that is able to reduce dump size, shared disk space, dump to network, dump to a file-system etc.
 - The makedumpfile tool can be used to filter the memory of the crashed system
 - Con
 - kdump is not as reliable as the stand-alone dump tools
 - kdump cannot dump a z/VM named saved system (NSS)
 - For systems running in LPAR kdump consumes memory

System z kernel features – two stage dumper / kdump support

- Add a `crashkernel=` to the kernel command line

```
crashkernel=<size>@<offset>
```

- Boot your system and check the reservation



```
#cat /proc/iomem
00000000-3fffffff : System RAM
00000000-005f1143 : Kernel code
005f1144-00966497 : Kernel data
00b66000-014c4e9f : Kernel bss
40000000-47fffffff : Crash kernel
48000000-7fffffff : System RAM
```

- Load the kdump kernel with `kexec`

```
# kexec -p kdump.image -initrd kdump.initrd
--command-line="dasd=1234 root=/dev/ram0"
```

- Manually trigger for kdump under z/VM

```
#cp system restart
```



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 kernel features – Storage FCP

- **FICON Express8S hardware data router support for FCP (kernel 3.2)**
 - Hardware data router support requires an adapted qdio request format.
 - Improves performance by reducing the path length for data.
- **FCP support for DIF/DIX (kernel 3.2)**
 - End to end data checking (aka data integrity extension) is no longer experimental.
 - Can be used with either direct I/O with a file system that fully supports end-to-end data consistency checking. Currently XFS only.
- **SCSI device management tool (> s390-tools 1.14.0)**
 - Implement a tool analog chccwdev which allows to enable/disable a SCSI LUN addressed by HBA/target port/LUN.

System z kernel features – Networking

- **Add support for AF_IUCV HiperSockets transport (kernel 3.2)**
 - Use HiperSockets with completion queues as transport channel for AF_IUCV sockets
- **Allow multiple paths with netiucv between z/VM guests (kernel 3.3)**
 - Speed up netiucv by using parallel IUCV paths.

System z toolchain

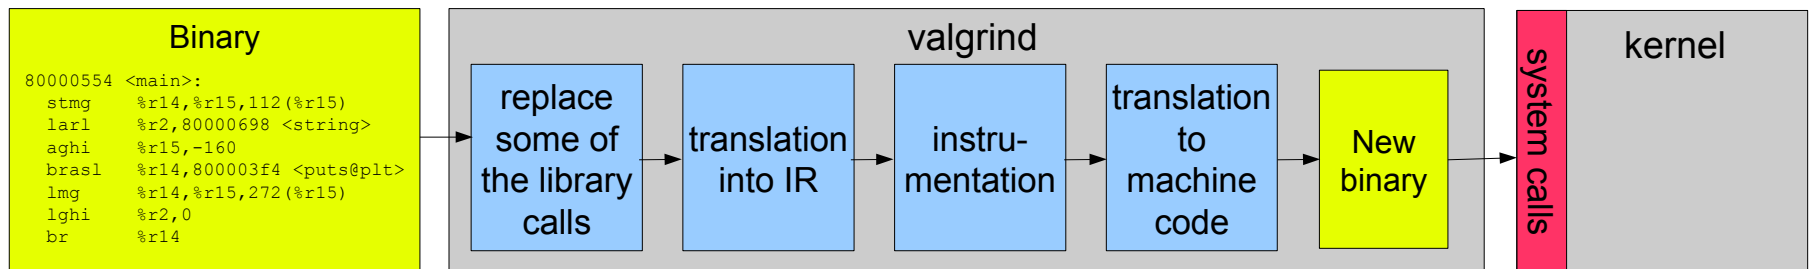
- **64 bit register in 31 bit compat mode**
 - Make use of 64 bit registers in 31 bit application running in z/Architecture mode.
 - Allows to use instruction operating on 64 bits, e.g. 64 bit multiplication
 - Needs kernel support for asynchronous signals
- **ATLAS support (libatlas 3.9.52)**
 - Add 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

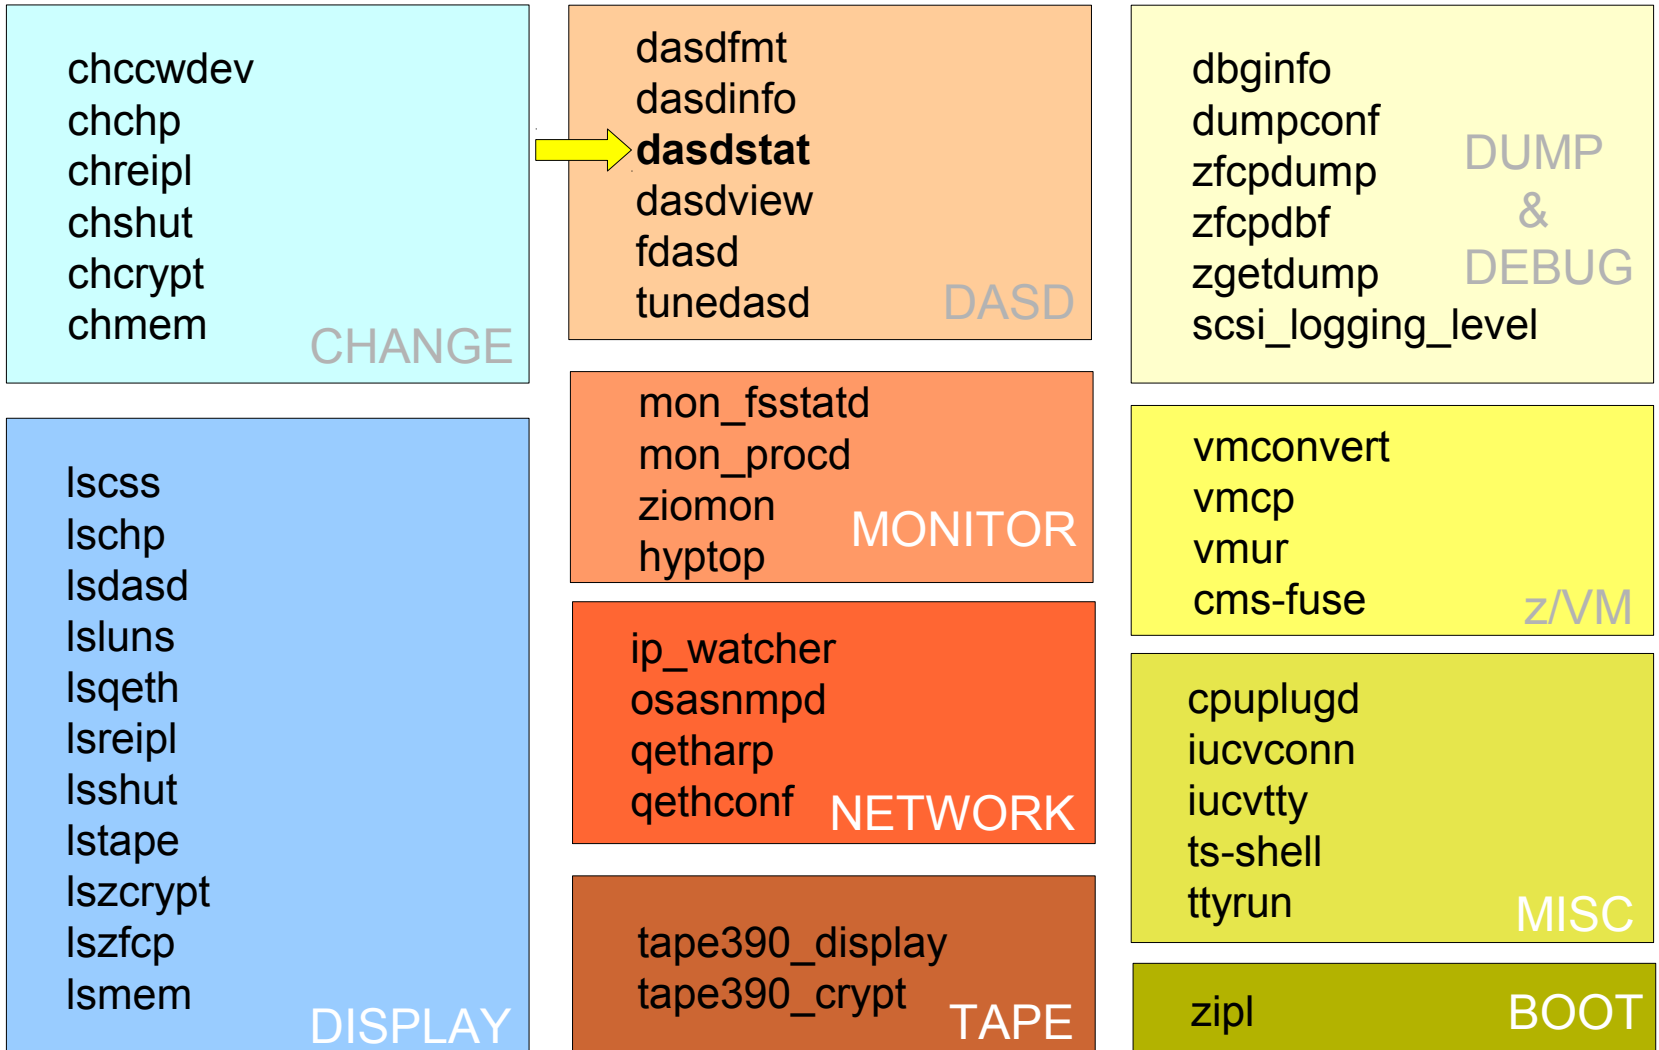
- `valgrind --tool=memcheck [--leak-check=full] [--track-origins] <program>`
 - Detects if your program accesses memory it shouldn't
 - Detects dangerous uses of uninitialized values on a per-bit basis
 - Detects leaked memory, double frees and mismatched frees
- `valgrind --tool=cachegrind`
 - Profile 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



Kernel news – Common code

▪ **Linux version 3.0 (2011-07-21)**

- New kernel version numbering scheme
- Cleancache (was transcendent memory) support for ext4, btrfs and XFS
- Preemptible mmu_gather for reduced latency
- Enhancements for the memory cgroup controller

▪ **Linux version 3.1 (2011-10-24)**

- New architecture: OpenRISC
- Dynamic writeback throttling
- Slab allocator speedups
- VFS scalability improvements
- New iSCSI implementation
- Software RAID: Bad block management

Kernel news – Common code

■ **Linux version 3.2 (2012-01-04)**

- btrfs improvements: faster scrubbing, automatic backup of tree roots
- ext4: support for bigger block sizes up to 1MB
- Process bandwidth controller
- I/O less dirty throttling
- New architecture: Hexagon

■ **Linux version 3.3 (2012-03-18)**

- Open vSwitch – software multilayer network switch
- Teaming – better bonding of network interfaces
- Network transmit queue limits – bufferbloat fighting
- New architecture: TI C6X

Kernel news – Common code

■ **Linux version 3.4 (2012-05-20)**

- Btrfs improvements: better error handling, repair and recovery tools
- New X32 ABI: 64-bit mode with 32-bit pointers
- perf tool improvements
- YAMA security module
- QNX6 filesystem

■ **Linux version 3.5 (2012-07-21)**

- Filesystem improvements: ext4 metadata checksums, btrfs latency reduction
- Uprobes: userspace probes
- Bufferbloat improvements: CoDel queue management
- TCP: connection repair, early retransmit
- SCSI over FireWire and USB

More information: Developer Works

The screenshot shows the IBM Developer Works website interface. At the top, there's a navigation bar with the IBM logo, a search bar, and a 'Country/region [select]' dropdown. Below the navigation bar, there's a breadcrumb trail: 'Home Solutions Services Products Support & downloads My IBM'. The main content area is titled 'Documentation for Development stream'. On the left, there's a sidebar with navigation links: 'Linux on System z', 'What's new', 'Development stream', 'Distribution hints', 'Documentation', 'Tuning hints & tips', 'Archive', and 'Feedback'. The main content area has a sub-header 'Development stream' with links to 'Novell SUSE' and 'Red Hat'. Below this, there's an 'Introduction' section with a paragraph of text and a link to 'Development stream'. There are three main sections: 'Base documentation', 'How to documents', and 'Reference documentation'. Each section contains a list of documents with titles, PDF links, and dates.

Document Title	File Size	Date
Device Drivers, Features, and Commands (kernel 2.6.33) - SC33-8411-05	PDF, 4.4MB	March 2010
Using the Dump Tools (kernel 2.6.33) - SC33-8412-04	PDF, 0.6MB	March 2010
How to Improve Performance with PAV - SC33-8414-00	PDF, 0.1MB	May 2008
How to use FC-attached SCSI devices with Linux on System z (kernel 2.6.33) - SC33-8413-04	PDF, 1.0MB	March 2010
How to use Execute-in-Place Technology with Linux on z/VM - SC34-2594-01	PDF, 0.5MB	March 2010
Download a tarball with sample scripts.		
How to Set up a Terminal Server Environment - SC34-2596-00	PDF, 0.3MB	June 2009
Kernel Messages (Kernel 2.6.33)	PDF, 0.4MB	March 2010
libica Programmer's Reference - SC34-2602-00	PDF, 0.3MB	June 2009

Linux on System z

How to use Execute-in-Place Technology with Linux on z/VM

March, 2010

Linux on System z

How to use FC-attached SCSI devices with Linux on System z

June 2009

Linux on System z

How to Set up a Terminal Server Environment on z/VM

June 2009

Linux Kernel 2.6 - Development stream

Linux on System z

Using the Dump Tools

Development stream (Kernel 2.6.33)

Linux on System z

Kernel Messages

Development stream (Kernel 2.6.33)

Linux on System z

Device Drivers, Features, and Commands

Development stream (Kernel 2.6.33)

<http://www-03.ibm.com/systems/z/os/linux/resources/index.html>

New Redbooks




z/VM and Linux on IBM System z The Virtualization Cookbook for Red Hat Enterprise Linux 6.0

Hands-on instructions for installing z/VM and Linux on the mainframe

Updated information for z/VM V6.1 and Red Hat Enterprise Linux 6.0

New, more versatile file system layout

4



Brad Hinson
Michael MacIsaac

Redbooks

ibm.com/redbooks




z/VM and Linux on IBM System z The Virtualization Cookbook for SLES 11 SP1

Hands-on instructions for installing z/VM and Linux on the mainframe

Updated information for z/VM 6.1 and Linux SLES 11 SP1

A new, more versatile file system layout



Michael MacIsaac
Marian Gasparovic

Redbooks

ibm.com/redbooks

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

Questions?



Martin Schwidefsky

*Linux on System z
Development*



*Schönaicher Strasse 220
71032 Böblingen, Germany*

*Phone +49 (0)7031-16-2247
schwidefsky@de.ibm.com*