



IBM Linux and Technology Center

# What's New in Linux on System z

Martin Schwidefsky  
IBM Lab Böblingen, Germany  
August 10, 2011 – Session 9473



# 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](http://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.

# 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

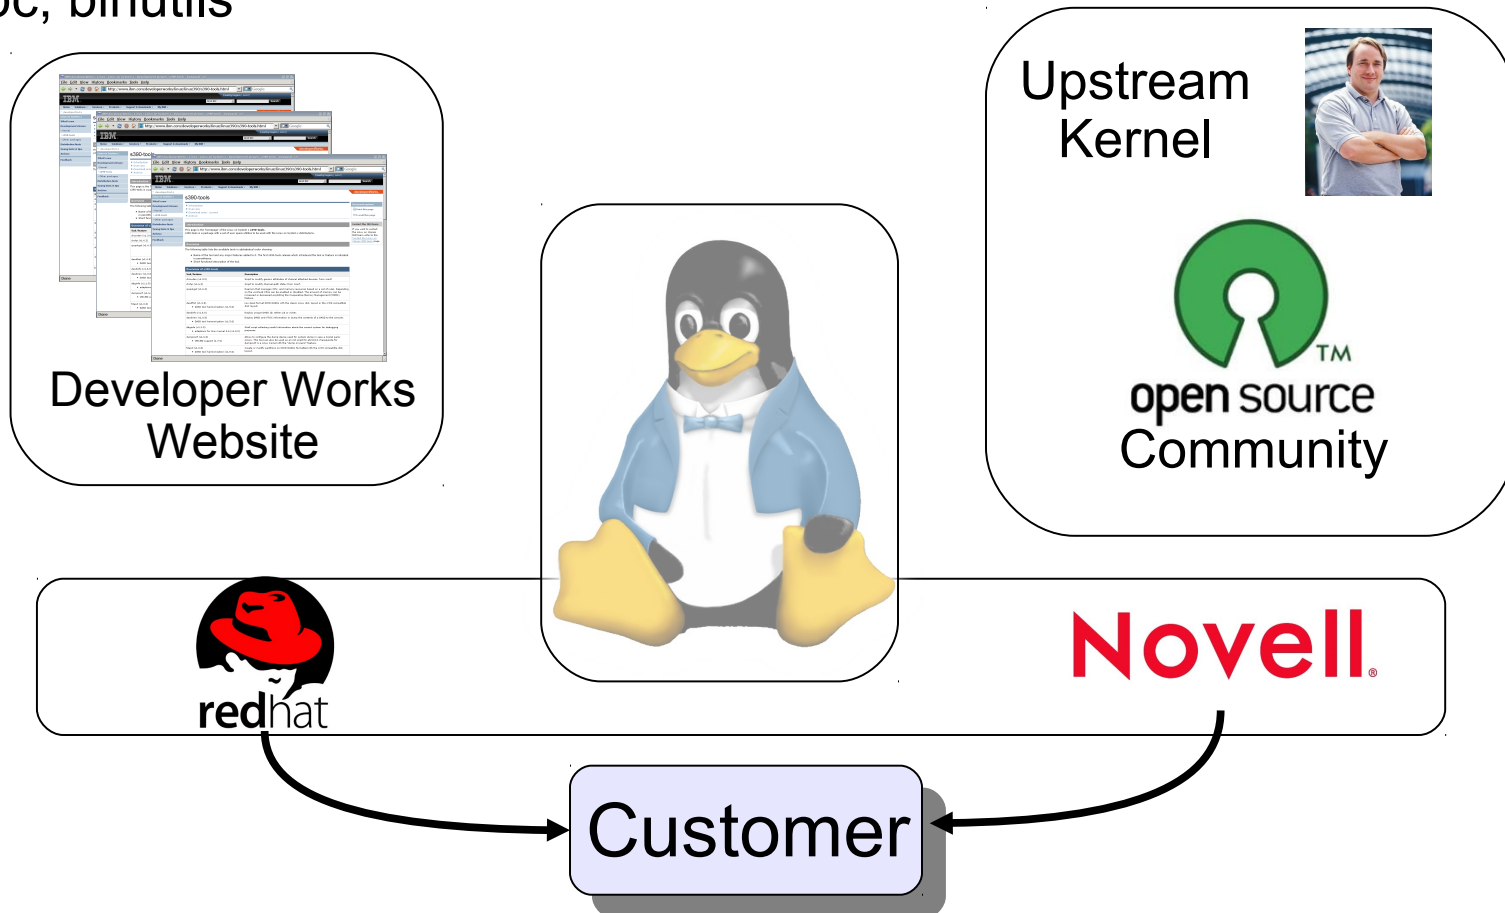
The Linux Foundation  
Linux Standards Base  
Common Criteria certification,  
and more...

## 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

- 2009, **75%** of the Linux code was developed by **programmers working for corporations**.
- **\$1.14 billion**: estimated total value of the **Linux kernel 2.6.30**
- **\$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: <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](http://public.dhe.ibm.com/software/au/downloads/IBM_zLinux_DAG_FINAL.pdf)  
<http://www.h-online.com/open/news/item/Study-Linux-kernel-R-D-worth-over-1-billion-euros-940817.html>

# 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](http://www.linuxfoundation.org/docs/lf_linux_kernel_development_2010.pdf)

# Linux kernel development: System z contributions

- Changesets per 2.6.x/3.0 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)
- 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
- Red Hat Enterprise Linux AS 4 (GA 02/2005)
  - Kernel 2.6.9, GCC 3.4.3, Update 9 (GA 02/2011)
- Red Hat Enterprise Linux AS 5 (GA 03/2007)
  - Kernel 2.6.18, GCC 4.1.0, Update 6 (GA 01/2011)
- Red Hat Enterprise Linux AS 6 (GA 11/2010)
  - Kernel 2.6.32, GCC 4.4.0, Update 1 (GA 05/2011)
- Others
  - Debian, Slackware,
  - Support may be available by some third party



# Supported Linux Distributions

Distribution	zEnterprise 196	System z10	System z9	zSeries
RHEL 6	✓	✓	✓	X
RHEL 5	✓	✓	✓	✓
RHEL 4 (*)	✓ (1)	✓	✓	✓
SLES 11	✓	✓	✓	X
SLES 10	✓	✓	✓	✓
SLES 9 (*)	✓ (2)	✓	✓	✓

✓ 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.

X Indicates that the distribution is not supported by IBM.

(\*) Also available as 31-bit 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)**

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

## Kernel news – Common code

### **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 (2011-03-14)**

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

## Kernel news – Common code




### **Linux version 2.6.39 (2011-05-18)**

- Ext4 SMP scalability
- IPset network resource groups
- Transcendent memory
- Unicore32 architecture

### **Linux version 3.0 (2011-07-21)**

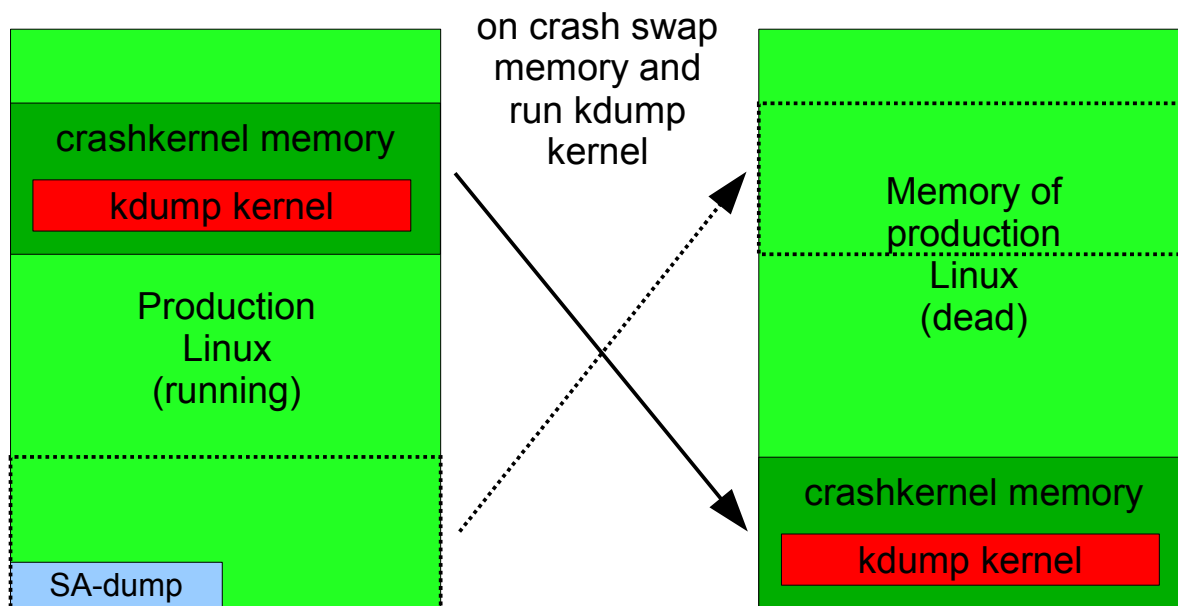
- New kernel version numbering scheme
- Cleancache (was transcendent memory) support for ext4, btrfs and XFS
- Preemptible mmu\_gather for reduced latency
- Enhancements for the memory cgroup controller

## System z kernel features – Core

- Breaking event address for user space programs (kernel 2.6.35)  6.1
  - 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)  6.1
  - Allows the Linux scheduler to optimize its decisions based on the z196 topology.
- Performance indicator bytes (kernel 2.6.37)  6.1
  - 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 / zfcps need different thresholds.
- Two stage dumper / kdump support (> kernel 3.0)
  - Enhanced dump support that is able to reduced dump size, share disk space, dump to network, etc.
  - Integrated into the System z stand-alone dump tools and shutdown actions framework

## Two stage dumper / kdump support

- 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
- Use the makedumpfile tool to filter the memory of the crashed system



## System z kernel features – z/VM

- CMSFS user space file system support (s390-tools 1.9.0/1.12.0)
  - 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 stream oriented Linux VFS
  - Attention: you can inadvertently damage files and lose data when directly operating on files within the cmsfs file system. To avoid damaging files, copy the file to a location outside the cmsfs-fuse file system, edit the file, and then copy it back to its original location.
- 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.



## 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
```

- Use fusermount to unmount the file system again

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



# System z kernel features – Storage 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 missings 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 Linux



## System z kernel features – Storage FICON

- 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
- 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
  - Storage array needs to support the read-track and write-full-track commands.

## System z kernel features – Storage FICON

- 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

## System z kernel features – Storage FCP

- 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 an FCP channel
- SAN utilities (trace, ping, ..) (kernel 2.6.36, lib-zfcp-hbaapi 2.1)
  - Two new utilities have been added: zfcp\_ping and zfcp\_show
  - They are useful to discover a storage area network
- Automatic LUN scanning (kernel 2.6.37)
  - Scan and attach accessible LUNs automatically
  - Available only for a NPIV FCP attachment
- Add NPIV information to symbolic port name (kernel 2.6.39)
  - Add the device bus-ID and the network node to the symbolic port name if the NPIV mode is active.
- 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.



## SAN Utilities: zfcplib

- Query Fibre Channel nameserver about ports available for my system:

```
# zfcplib -n
Local Port List:
    0x500507630313c562 / 0x656000 [N_Port] proto = SCSI-FCP  FICON
    0x50050764012241e4 / 0x656100 [N_Port] proto = SCSI-FCP
    0x5005076401221b97 / 0x656400 [N_Port] proto = SCSI-FCP
```

- Query SAN topology, requires FC management server access:

```
# zfcplib
Interconnect Element Name      0x100000051e4f7c00
Interconnect Element Domain ID 005
Interconnect Element Type      Switch
Interconnect Element Ports     256
  ICE Port 000  Online
    Attached Port [WWPN/ID] 0x50050763030b0562 / 0x650000 [N_Port]
  ICE Port 001  Online
    Attached Port [WWPN/ID] 0x50050764012241e5 / 0x650100 [N_Port]
  ICE Port 002  Online
    Attached Port [WWPN/ID] 0x5005076303008562 / 0x650200 [N_Port]
  ICE Port 003  Offline
  ...
```

## SAN Utilities: zfc\_ping

- Check if remote port responds (requires FC management service access):





```
# zfc_ping 0x5005076303104562
Sending PNG from BUS_ID=0.0.3c00 speed=8 GBit/s
    echo received from WWPN (0x5005076303104562) tok=0 time=1.905 ms
    echo received from WWPN (0x5005076303104562) tok=1 time=2.447 ms
    echo received from WWPN (0x5005076303104562) tok=2 time=2.394 ms

----- ping statistics -----
min/avg/max = 1.905/2.249/2.447 ms
-----
```

- zfc\_show and zfc\_ping are part of the zfc-hbaapi 2.1 package:

<http://www.ibm.com/developerworks/linux/linux390/zfc-hbaapi-2.1.html>

## System z kernel features – Networking

- Offload outbound checksumming (kernel 2.6.35)  6.1
  - 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
- OSX (OSM) CHPIDs for hybrid data network (kernel 2.6.35)  10.4<sup>N</sup>  
11.1  5.6  
6.1
  - The OSA cards for the zBX Blade Center Extension will have a new CHPID type
  - Allows communication between zBX and Linux on System z
- NAPI support for QDIO and QETH (kernel 2.6.36)  6.1
  - Convert QETH to the NAPI interface, the “new” Linux networking API
  - NAPI allows for transparent GRO (generic receive offload)
- New default qeth configuration values (kernel 2.6.39)
  - Receive checksum offload, generic receive offload & number of inbound buffers

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





## znetconf network device configuration tool

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

```
# 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

```
znetconf -a 0.0.f503
```

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

```
znetconf -a 0.0.f503 -o layer2=1 -o portname=myport
```

- Remove network device 0.0.f503

```
znetconf -r 0.0.f503
```

## System z kernel features – Usability / RAS / Security

- CHPID reconfiguration handling (kernel 2.6.37)
  - Update data structures after channel-path related information change
  - Inform device drivers about relevant changes
- 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

## System z kernel features – Usability / RAS / Security

- Get CPC name (kernel 2.6.39)
  - Useful to identify a particular hardware system in a cluster
  - The CPC name and HMC network name are provided
- CP ACF exploitation of System z196 (kernel 2.6.39)
  - 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)
- Removal of data execution protection
  - “no execute” support relies on the secondary space mode for data separation
  - With System z10 the new instructions LRL, LGRL and LGFRL for pc-relative data access have been added
  - These new instructions access the memory operand in the same address space from where the instructions has been fetched.

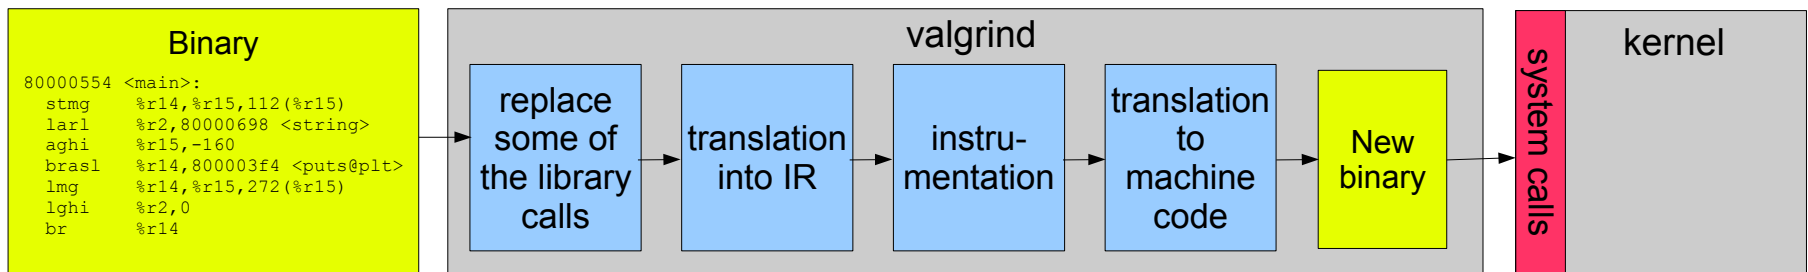
# System z toolchain

- 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
- 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 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
- 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



## 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



## 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](mailto:linux390@de.ibm.com)

# s390-tools package: the content

chccwdev  
 chchp  
 chreipl  
 chshut  
 chcrypt  
 chmem

CHANGE

dasdfmt  
 dasdinfo  
 dasdview  
 fdasd  
 tunedasd

DASD

dbginfo  
 dumpconf  
 zfcpdump  
 zfcpdbf  
 zgetdump  
 scsi\_logging\_level

DUMP  
 &  
 DEBUG

lscss  
 lschp  
 lsdasd  
 lsluns  
 lsqeth  
 lsreipl  
 lsshut  
 lstape  
 lszcrypt  
 lszfcp  
 lsmem

DISPLAY

mon\_fsstatd  
 mon\_procd  
 ziomon  
**hyptop**

MONITOR

vmconvert  
 vmcp  
 vmur  
 cms-fuse

z/VM

ip\_watcher  
 osasmpd  
 qetharp  
 qethconf

NETWORK

cpuplugd  
 iucvconn  
 iucvty  
 ts-shell  
 ttyrun

MISC

tape390\_display  
 tape390\_crypt

TAPE

zipl

BOOT

## s390-tools package

- Version 1.10.0 (2010-09-16)
  - chchp: Use /proc/cio\_settle
  - znetconf: Add support for new CHPIDs OSX and OSM
- Version 1.11.0 (2010-10-22)
  - cmsfs-fuse: Add write support
  - zipl: Add support for automatic menus
- Version 1.12.0 (2011-01-27)
  - **hyptop: Provides real-time view of System z hypervisor environment**
  - chreipl: Various enhancements
  - cio\_ignore: Add query option
  - cmsfs-fuse: Configurable code page conversion
  - tunedasd: Add option to query reservation status of a device
  - zgetdump: Add kdump support for –info option
  - zfcpdump/zipl: Disable automatic activation of LUNs



## s390-tools package

- Version 1.13.0 (2011-05-19)
  - qetharp: Support IPv6 for query ARP cache for HiperSockets
  - zfcpdbf: Adjust to 2.6.38 zfc driver changes
- Version 1.14.0 (2011-06-30)
  - fdasd: Implement new partition types

## hyptop: Display hypervisor utilization data

- The hyptop command is a top-like tool that displays a dynamic real-time view of the hypervisor environment
  - It works with both the z/VM and the LPAR hypervisor
  - Depending on the available data it can display information about CPU and memory
  - running LPARs or z/VM guest operating systems
- The following is 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 guests your instance requires additional privileges
    - For z/VM: The user ID requires privilege class B
    - For LPAR: The global performance data control box in the LPAR activation profile needs to be selected

# hyptop: Display hypervisor utilization data

- Example of z/VM utilization data

```
10:11:56 CPU-T: UN(16)                                     ?=help
```

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

# hyptop: Display hypervisor utilization data

- Example of single LPAR utilization data

```

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

```

# More information

The screenshot shows the IBM developerWorks website interface. At the top, there is a navigation bar with the IBM logo, a search bar, and menu items like Home, Solutions, Services, Products, Support & downloads, and My IBM. Below the navigation bar, the page title is "Documentation for Development stream". A left sidebar contains a table of contents with categories like "What's new", "Development stream", "Distribution hints", "Documentation", "Tuning hints & tips", "Archive", and "Feedback". The main content area is titled "Introduction" and contains a paragraph of text. Below this, there are two tables: "Base documentation" and "How to documents".

Document Title	File Size	Date
Device Drivers, Features, and Commands (kernel 2.6.33) - SC33-8411-05 (PDF, 4.4MB)	4.4MB	March 2010
Using the Dump Tools (kernel 2.6.33) - SC33-8412-04 (PDF, 0.6MB)	0.6MB	March 2010

Document Title	File Size	Date
How to Improve Performance with PAV - SC33-8414-00 (PDF, 0.1MB)	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)	1.0MB	March 2010
How to use Execute-in-Place Technology with Linux on z/VM - SC34-2594-01 (PDF, 0.5MB)	0.5MB	March 2010
Download a tarball with sample scripts.		
How to Set up a Terminal Server Environment - SC34-2596-00 (PDF, 0.3MB)	0.3MB	June 2009

Document Title	File Size	Date
Kernel Messages (Kernel 2.6.33) (PDF, 0.4MB)	0.4MB	March 2010
Ibica Programmer's Reference - SC34-2602-00 (PDF, 0.3MB)	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

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)*

# 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](http://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](http://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](mailto:schwidefsky@de.ibm.com)*