


Linux on System z update



Wilhelm Mild
IT Architect
IBM Lab, Boeblingen, Germany
wilhelm.mild@de.ibm.com


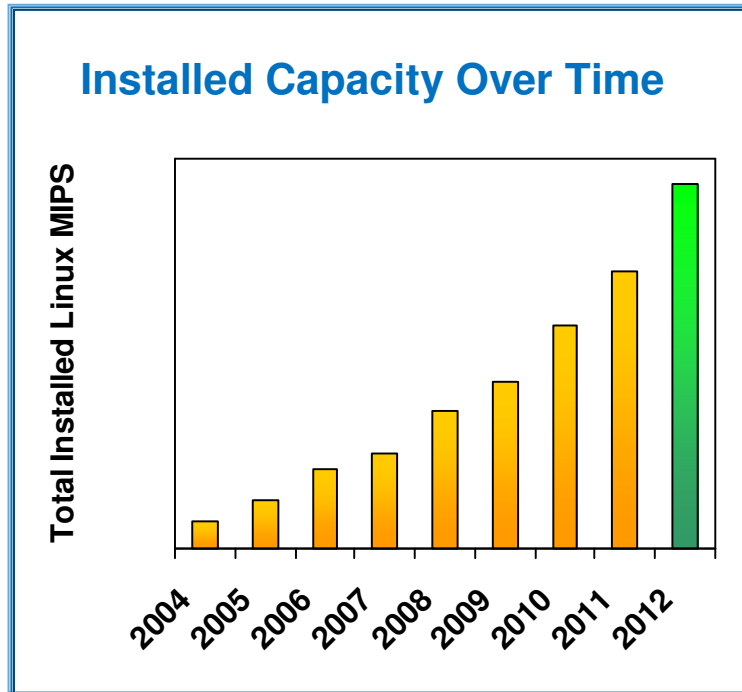
Trademarks

- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

<http://www.ibm.com/legal/copytrade.shtml>

Linux on System z - a growing business !

Linux on IBM System z at Year-End-2012



- 22,8% of Total installed System z MIPS run Linux
- Installed IFL MIPS increased 32% from 4Q11 to 4Q12

- 36% of System z Clients have IFL's installed
- 84 of the Top 100 System z Clients are running Linux on System z*

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

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 3.7 ?
16,191,784 lines of code (+1,193,047 since 3.2)
- How many of the world's top 500 supercomputers run Linux (Nov 2012)?
453 / 90.6%
- What percentage of web servers run Linux (Nov 2012) ?
64.9% run Unix, of those 50.7% run Linux (47.5% unknown) = 32.9%
- What percentage of desktop clients run Linux (Dec 2012) ?
1.75%
- What is the largest Linux architecture in number of devices ?
ARM, > 100 million activated android devices (Spinnaker – UK Univ.)
- Linux is Linux, but ...features, properties and quality differ dependent on your platform and 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>

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,

Open Cloud Consortium

and more...

Promoting Open Standards & Community Collaboration

The Linux Foundation

Linux Standards Base

Common Criteria certification,

Open Virtualization Format

and more...

Foster and Protect the Ecosystem

Software Freedom Law Center

Free Software Foundation (FSF),

and more...

IBM focus: Linux on all IBM Systems


Linux on System z

- Applications close to the data
- Applications need "Mainframe" criteria
- Vertical and horizontal scaling (hundreds of virtual servers at the same time)
- Great workload-management
- excellent virtualization



Linux on System x

- Low-priced entry
- Simultaneous operation of Windows and Linux via virtualization
- High reliability and simple manageability by Enterprise X-Architektur
- Very good price / performance ratio via Clustering
- Big performance density by BladeCenter



Linux everywhere Linux on System i

Linux on System p

- Applications demanding high performance
- Consolidation of multiple servers / applications via LPARs
- Native Linux or AIX




Linux on special Hardware

- Applications demanding highest performance
- special purpose



Linux on System i

- If an integrated solution is required (hardware / software / network)
- Simple administration and operation
- If midrange system is required
- Consolidation of applications via LPARs (up to 31)
- High degree of security



Software Product Compatibility Report

IBM Middleware for Linux Information

Software product compatibility reports
Reports about compatible software combinations

The support shown in the reports generated by this tool may require a particular maintenance for the products.

Operating system reports

- a) Operating systems for a specific product**
Usage: Find out all the operating systems on which a product you have in mind runs.
[View sample report](#)
- b) Products that use a specific operating system**
Usage: Find out all the products that are supported by an operating system you have in mind.
[View sample report](#)
- c) Matrix between specific product(s) and desired operating systems**
Usage: Find out if a set of software products runs on operating systems you have in mind.
[View sample report](#)

Prerequisite reports

- Prerequisites of a specific product**
Usage: Find out all the software that a specific product needs or uses per selected capabilities.
[View sample report](#)
- Products that use a specific prerequisite**
Usage: Find out all the software that use or support a specific product.
[View sample report](#)
- Matrix between specific product(s) and desired prerequisites**
Usage: Find out if a set of prerequisites will work with a set of specific products.
[View sample report](#)

Virtualization environment reports

- Virtualization environments supporting a product**
Usage: Find out all the virtualization environments that support a product you have in mind.
[View sample report](#)
- Products supported by a virtualization environment**
Usage: Find out all the products that are supported by a virtualization environment you have in mind.
[View sample report](#)
- Matrix between specific product(s) and desired virtualization environment**
Usage: Find out if a set of products is supported by virtualization environments you have in mind.
[View sample report](#)

Miscellaneous reports

- d) End of service report for products**
Usage: Find out the end of service dates for the products you have in mind.
[View sample report](#)

- Software product compatibility reports provides information about compatible software combinations based on Clearinghouse

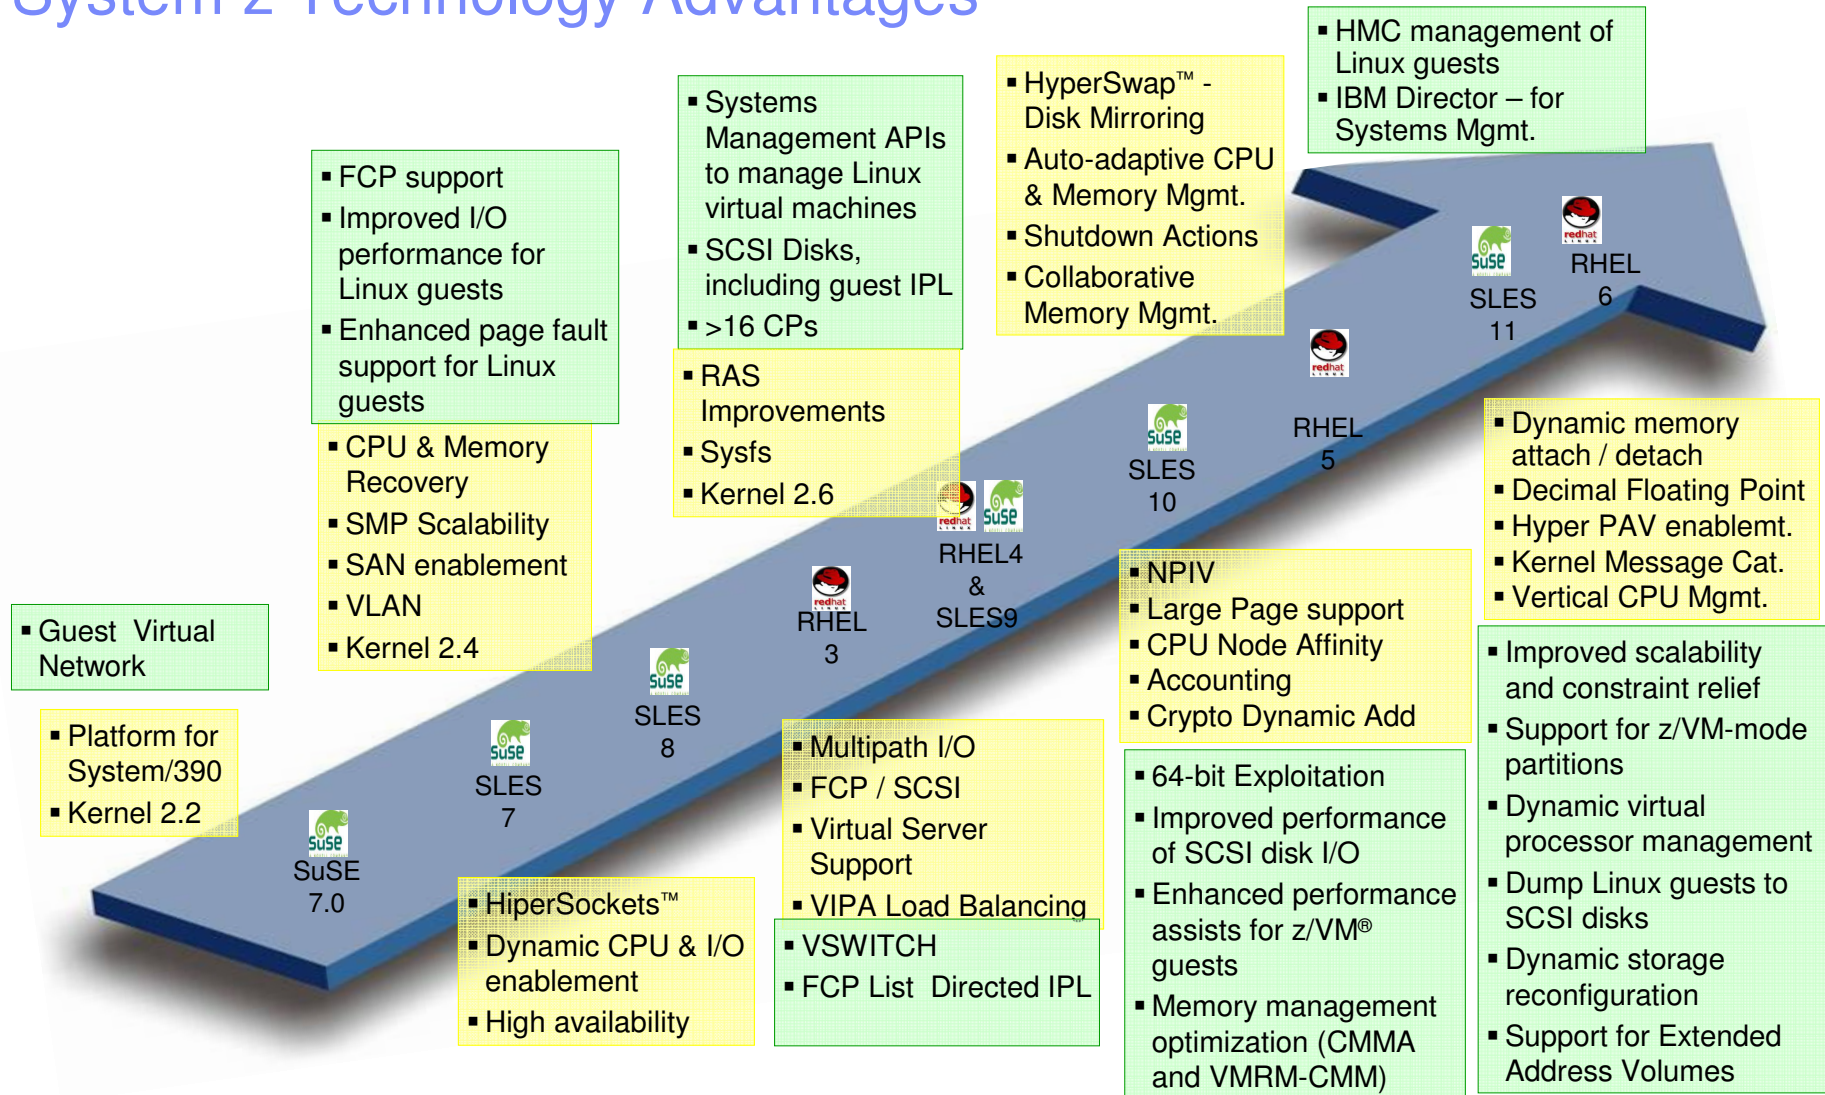
Clarity:

<http://publib.boulder.ibm.com/infocenter/prodguid/v1r0/clarity/index.jsp>

- a) Provide Operating Systems for a specific Product
- b) Provide Products that use a specific Operating System
- c) Provide a Matrix between specific product(s) and desired prerequisites (like Operating Systems)
- d) Provide Product End of Service dates

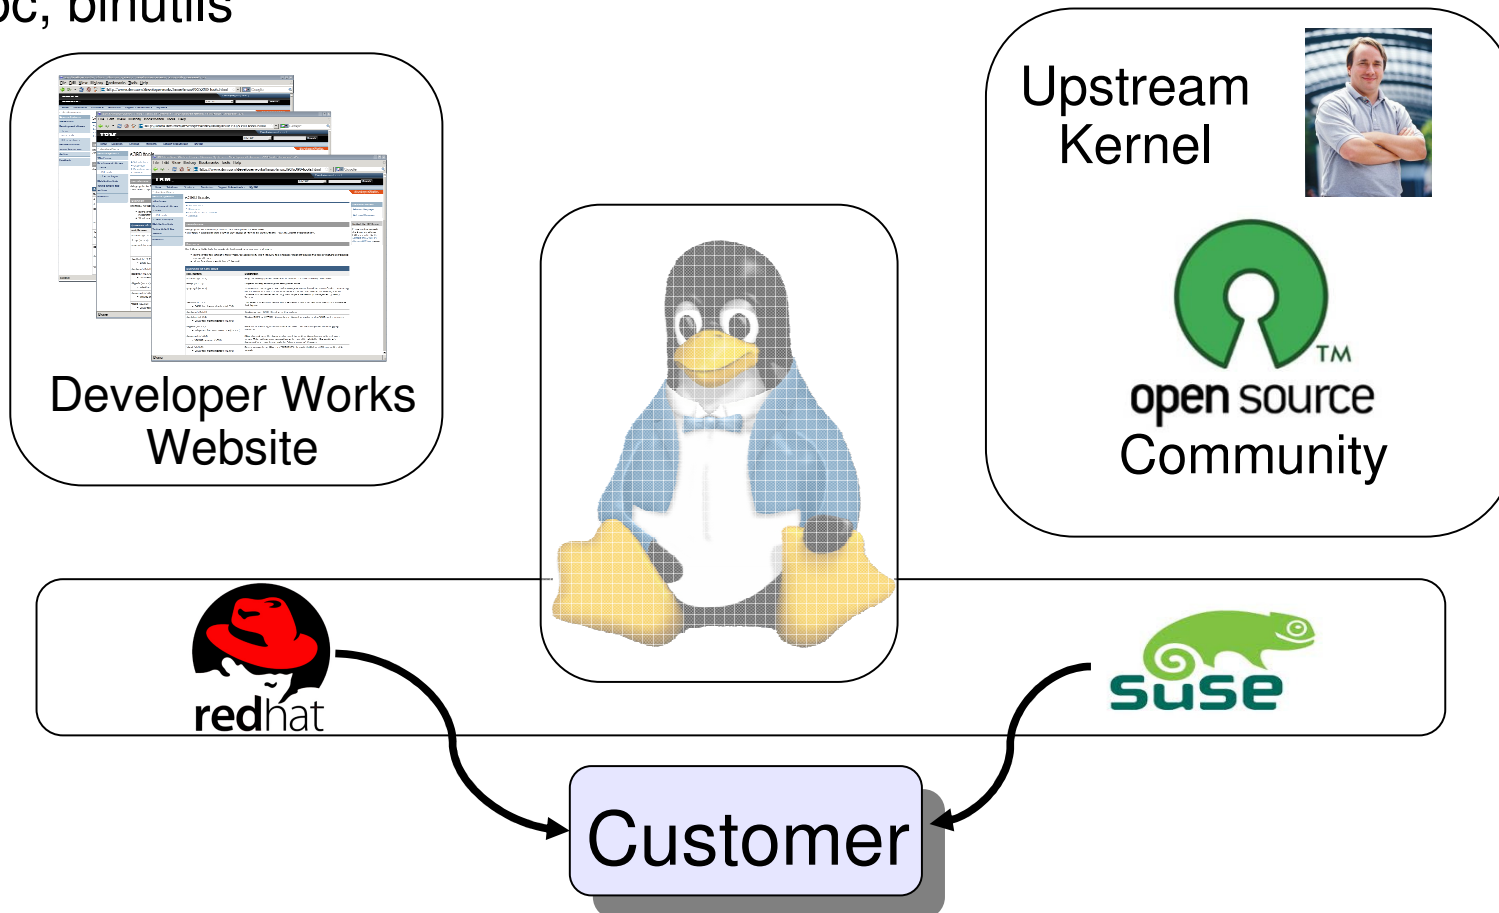


Unique Linux Extensions to Leverage Powerful IBM System z Technology Advantages

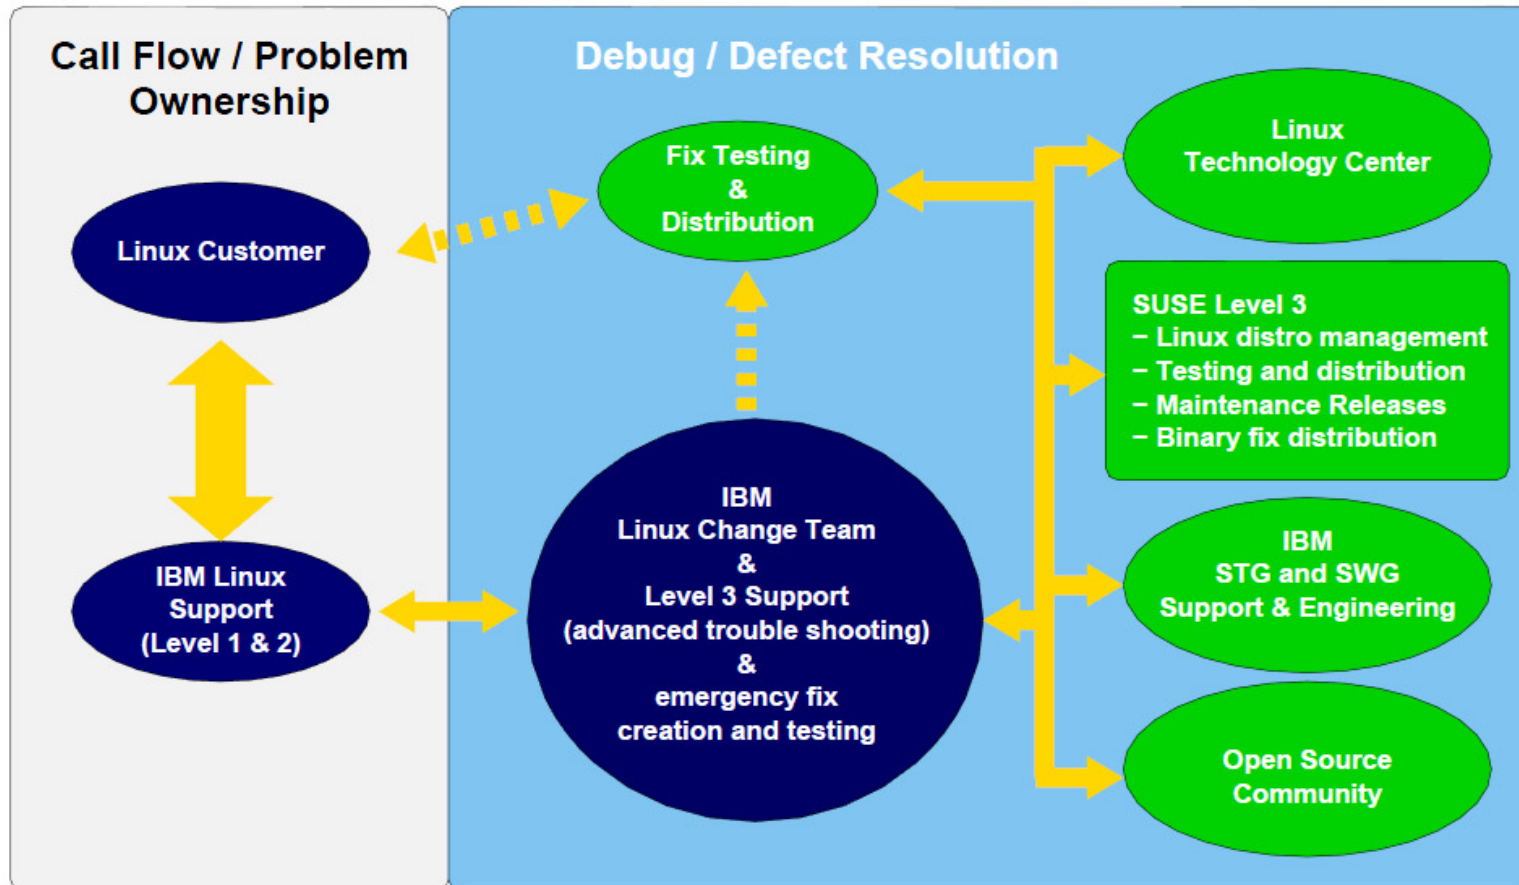


The IBM Linux on z development process

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



Linux Support structure



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 9 (GA 01/2013)

Red Hat Enterprise Linux AS 6 (GA 11/2010)

Kernel 2.6.32, GCC 4.4.0, Update 3 (GA 05/2012)

Others

Debian, Slackware,

Support may be available by some third party

Supported Linux Distributions

Distribution	zEnterprise EC12	zEnterprise - z114 and z196	System z10	System z9	zSeries
RHEL 6	✓ ⁽¹⁾	✓	✓	✓	X
RHEL 5	✓ ⁽²⁾	✓	✓	✓	✓
RHEL 4 (*)	X	✓ ⁽⁵⁾	✓	✓	✓
SLES 11	✓ ⁽³⁾	✓	✓	✓	X
SLES 10	✓ ⁽⁴⁾	✓	✓	✓	✓
SLES 9 (*)	X	✓ ⁽⁶⁾	✓	✓	✓

Two options for zSeries machines

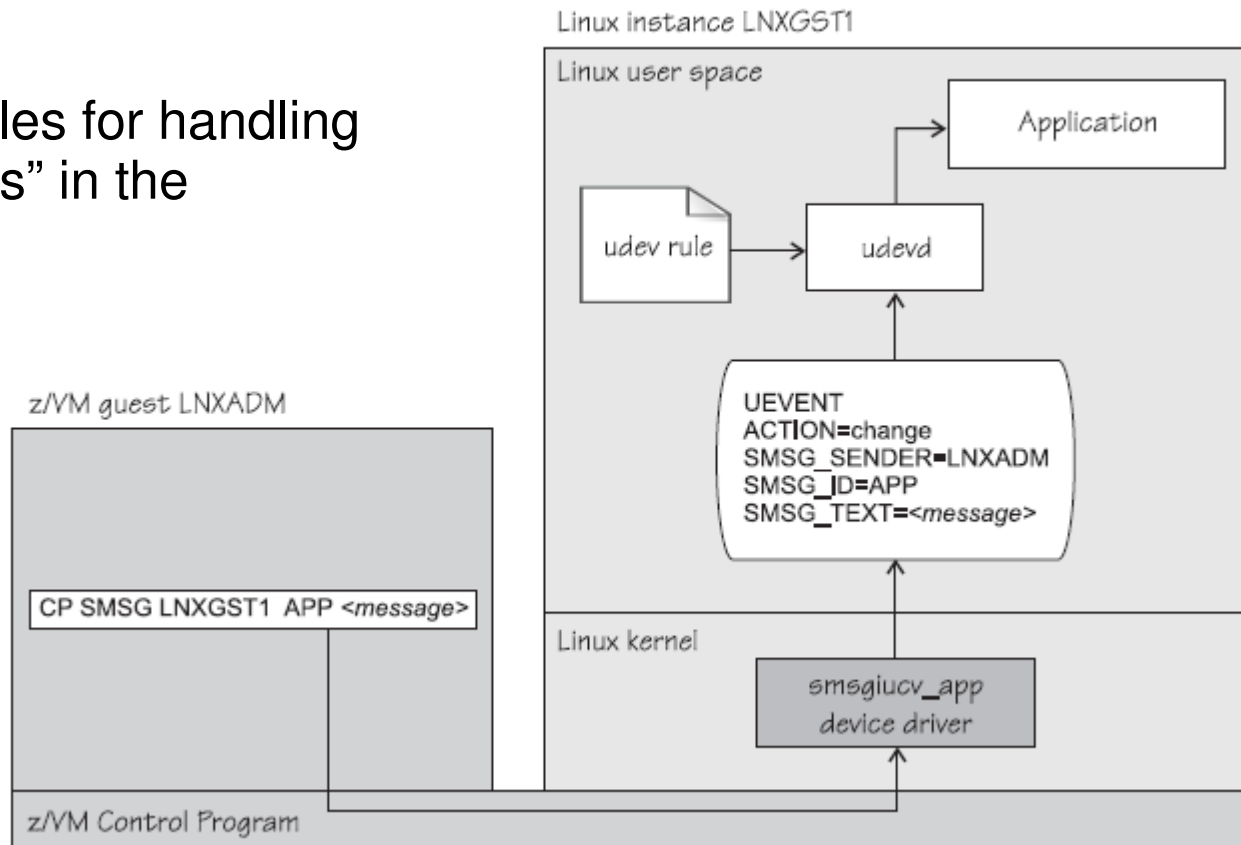
- Indicates that the distribution (version) has been tested by IBM on the hardware platform, will run on the system, and is an IBM supported environment. Updates or service packs applied to the distribution are also supported.
 - (1) Recommended level: RHEL 6.3
 - (2) Recommended level: RHEL 5.8
 - (3) Recommended level: SLES 11 SP2
 - (4) Recommended level: SLES 10 SP4 with latest maintenance updates
 - (5) RHEL 4.8 only. Some functions have changed or are not available with the z196, e.g. the Dual-port OSA cards support to name one of several. Please check with your service provider regarding the end of service.
 - (6) SLES 9 SP4 with latest maintenance updates only. Some functions have changed or are not available with the z196, e.g. the Dual-port OSA cards support to name one of several. Please check with your service provider regarding the end of service.
- X Indicates that the distribution is not supported by IBM on this server.
- (*) The distribution is out of service, extended support is required.

Current Linux on System z Technology

Features & Functionalities included in
SuSE and Red Hat
Distributions

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



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 completed (integrated in SLES11.2 only)
 - use "vi" to edit PROFILE.EXEC anyone ?
- Use fusermount to unmount the file system again



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


System z kernel features – Core


- **z196 enhanced node affinity support (kernel 2.6.37)**

 6.1
  11.2



 - Allows the Linux scheduler to optimize its decisions based on the z196 topology
- **Performance indicator bytes (kernel 2.6.37)**

 6.1
  11.2


 - Display capacity adjustment indicator introduced with z196 via /proc/sysinfo
- **QDIO outbound scan algorithm (kernel 2.6.38)**

 11.2

 - Improve scheduling of QDIO tasklets, OSA / HiperSockets / zfcip need different thresholds
- **Enabling spinning mutex (kernel 2.6.38)**

 6.3
  11.2

 - 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.
- **Detailed IRQ statistics (kernel 2.6.38)**

 11.2

 - A detailed, per-cpu list of interrupts in /proc/interrupts
 - Useful to debug interrupt imbalances

Linux on System z features - usability RAS

▪ **hyptop (kernel 2.6.35, s390-tools 1.12.0)**



- 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

▪ **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

▪ **Module read-only protection (kernel 2.6.39)**



- The text and the read-only data section of modules are write protected in the kernel page table, this avoids unwanted modification

System z kernel features – FICON

- **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

- **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.

- **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

System z kernel features – FCP

Store I/O and start logging (SIOSL) (kernel 2.6.36)



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)



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)



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

- 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

- **Improve memory ballooning with cpuplugd (s390-tools 1.15.0)**



- A number of improvements to the memory balloon part to cpuplugd to make it more useful.

- **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

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 toolchain

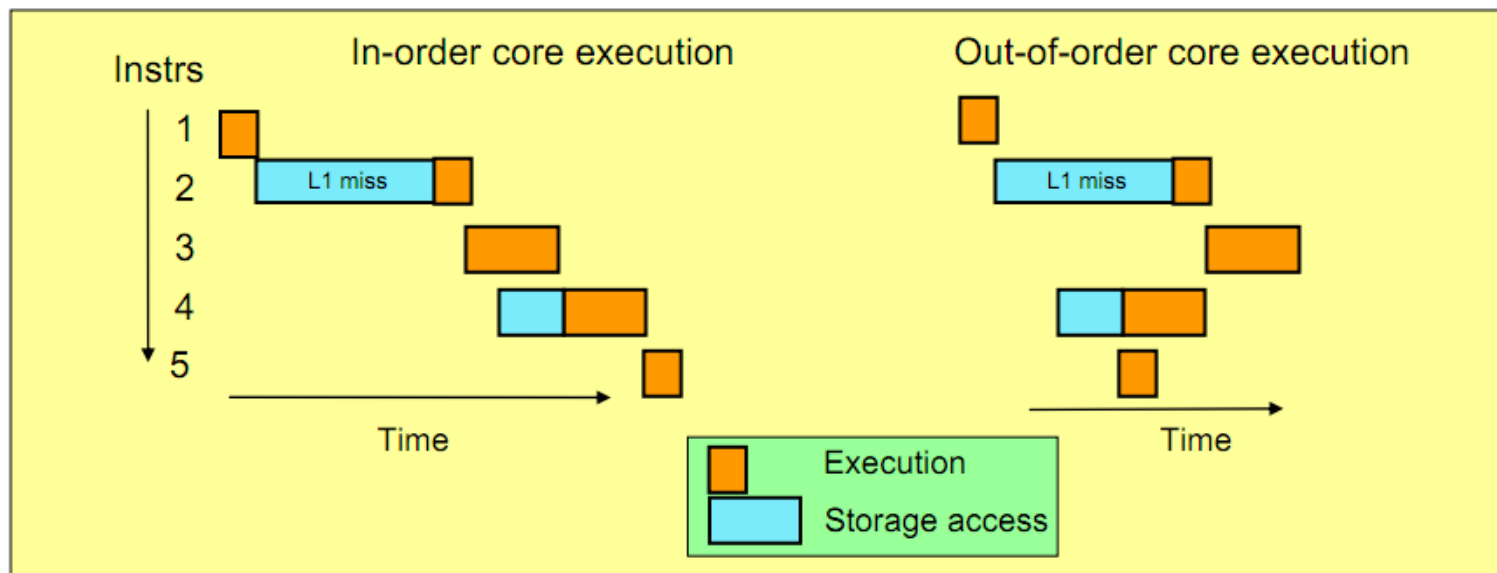
zEnterprise 196 exploitation (gcc 4.6)



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

Function enhancements already developed and integrated into the upstream Linux Kernel
- but **not** yet available in any Enterprise Linux Distribution

Linux on System z features - zEC12 support

- **Transactional execution (kernel 3.7)**
 - Also known as hardware transactional memory
 - CPU features that allows to execute a group of instructions atomically
 - Optimistic execution, if a transaction conflicts a rollback to a saved state is done
- **Storage class memory – Flash Express (kernel 3.7)**
 - Internal Flash Solid State Disk (SSD)
 - Accessed via Extended Asynchronous Data Mover (EADM) sub-channels
 - Up to four pairs of cards with a maximum capacity of 6.4TB
- **Support for Crypto Express 4S cards (kernel 3.7)**
 - New generation of crypto adapters plug-able into the I/O drawer
 - New type 10 which uses a bit field to indicate capabilities of the crypto card
- **Toolchain support (gcc 4.8)**
 - Add new instructions to the compiler
 - Add description of the new pipeline to generate optimal code

System z kernel features - core

- **Add support for physical memory > 4TB (kernel 3.3)**
 - Increase the maximum supported memory size from 4TB to 64TB.
 - Memory sizes large than 4TB require a 4-level page table
 - Makes memory accesses by the kernel slightly slower, the kernel will automatically use a 3-level page table for memory sizes \leq 4TB
- **Add support for System z hardware counters (kernel 3.4)**
 - Enables hardware counters for Linux running in an LPAR
 - Available counters are: basic counter set, problem-state counter set, crypto-activity counter set, extended counter set with System z10
 - System zEC12 counter set with kernel version 3.7
- **Allow to compare dump system with boot system**
 - With z/VM 6.2 Single-System-Image it is possible to move active Linux instances between different z/VM instances
 - To aid debugging a log of past live-guest-relocations is made available in both the live system and in the dump of a system

System z kernel features - core

▪ **Fuzzy Live Dump (kernel 3.5)**

- Add the capability to generate a dump of a live system.
- Not all data structures will be consistent but the dump may still be useful.

▪ **BPF JIT compiler for System z (kernel 3.7)**

- The Berkeley Packet Filter is an interface and a language definition that allows to pass a filter to the kernel to select network packets to send on a socket
- The BPF JIT compiler in the kernel translates the interpreted BPF code to System z code.
- A secondary use of the BFP language is system call filtering.

▪ **Expose CPU cache topology in sysfs (kernel 3.7)**

- Add an interface to expose the CPU cache topology to user space.
- System z only provides information about CPU caches which are private to a CPU, information about shared caches is not exposed.

▪ **Add page table dumper (kernel 3.7)**

- Add a sysfs interface to read the current layout of the kernel address space.
- Useful information for the kernel developer.

System z kernel features – Core kernel / crypto

Transparent huge page support (kernel 3.7)

Make the common code transparent huge page support available for Linux on System z.

With THP 1MB pages will be used to back normal anonymous memory mappings.

Any application will benefit from using huge pages.

Not as effective as using the large pages directly, no memory savings for page tables due to huge page splitting.

AP adapter resiliency (kernel 3.7)

Improve RAS capabilities of the AP bus and the zcrypt devices.

External AP bus configuration changes are now handled accordingly.

System z kernel features – Storage FICON / FCP

DASD sanity check to detect path connection errors (kernel 3.3)

- Add a check in the DASD driver to make sure that each available channel path leads to the same storage server

FICON Express8S hardware data router support for FCP (kernel 3.2)

- Hardware data router support with an adapted qdio request format to improve 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 and can be used with either direct I/O with a file system that fully supports end-to-end data consistency checking. Currently XFS only.

No automatic port rescan on events (kernel 3.7)

- The rescan of a zfcf port following a fabric change event can cause high fabric traffic, especially when many Linux images share an FCP channel over multiple subchannels with NPIV enabled. This can lead to errors due to timeouts.
- Ports are still scanned when the adapter is set online and on manual user triggered writes to the port_rescan sysfs attribute.

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)

- Adds support for multiple point-to-point netiucv interfaces between guests
- Speed up netiucv by using parallel IUCV paths.

Add query OSA address table support (kernel 3.4)

- The new qethqoat command queries the OSA address table and displays physical and logical device information.
- Viewing the OSA address table is a low level diagnostic requirement.

Linux on System z features – tools

Extend lscpu tool and add new chcpu tool (util-linux 2.21)

- Improve the lscpu tool to display CPU topology and CPU state
- Add the new chcpu tool to change CPU state, rescan CPUs and change the CPU dispatching mode (horizontal vs. vertical polarization)

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 toolchain

64 bit register in 31 bit compat mode (gcc 4.6, glibc 2.11)

- 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

Exploitation of zEC12 instructions (gcc 4.8)

- Use option `-march=zEC12` to utilize the new instructions added with zEC12
- Use `-mtune=zEC12` to schedule the instruction appropriate for the pipeline of zEC12

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 application development tools

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`

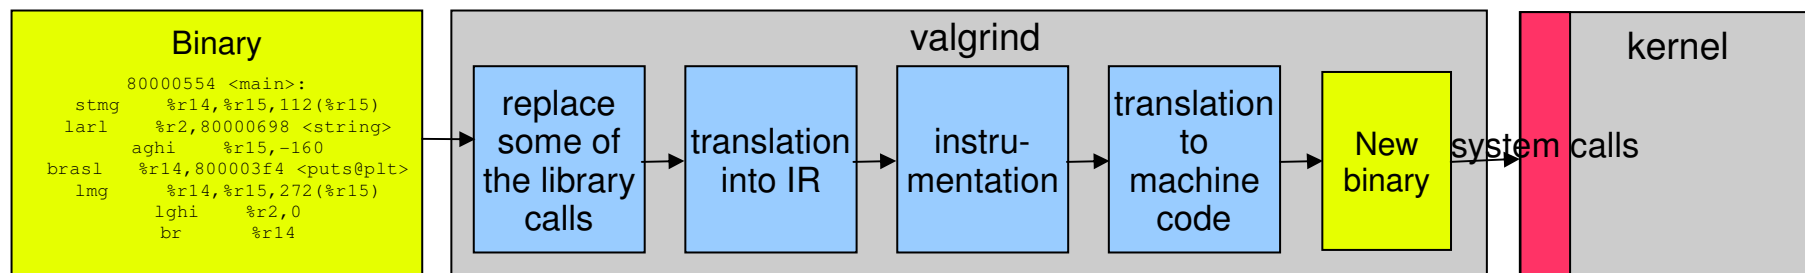
Profile cache usage, simulates instruction and data cache of the cpu

Identifies the number of cache misses

`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

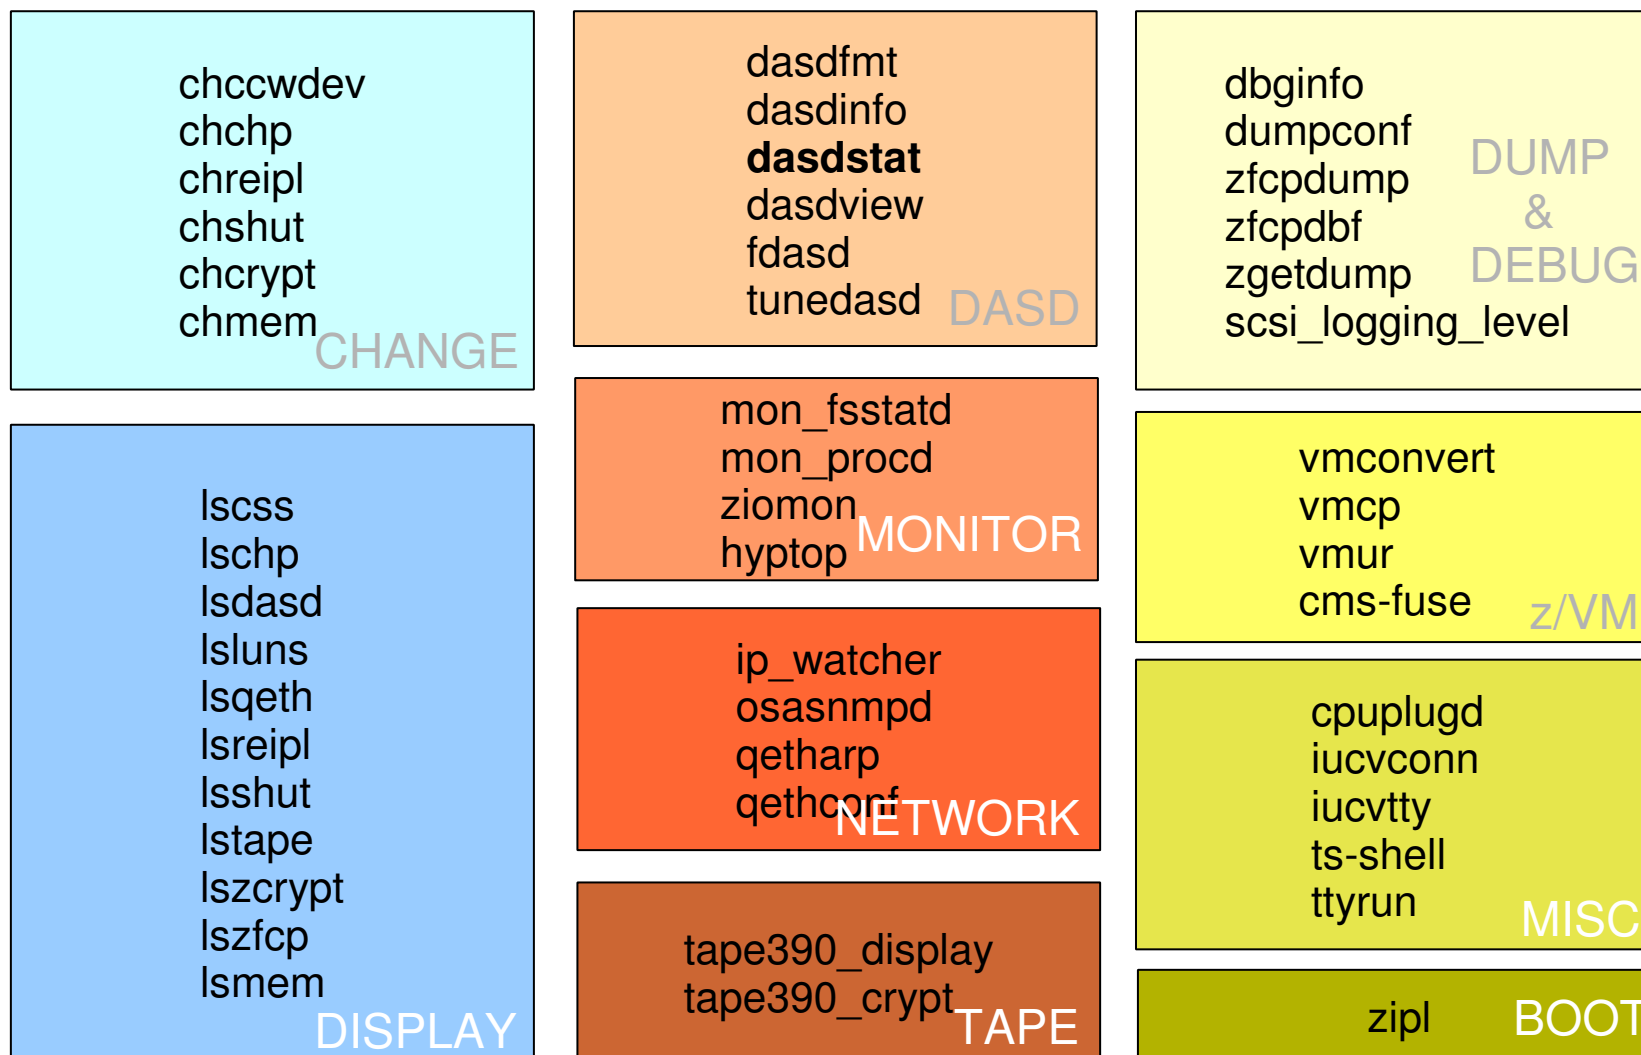
SuSE Linux Enterprise Server 11

Website:

<http://www.ibm.com/developerworks/linux/linux390/s390-tools.html>

Feedback: linux390@de.ibm.com

s390-tools package: the content



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	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	100
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
VMSEVRU	1	0.00	0:00	62:21:47	0.00	0.03	1500
OPERSVMP	1	0.00	0:00	62:21:47	0.00	0.03	100



system z Information

IBM Linux Information Center:

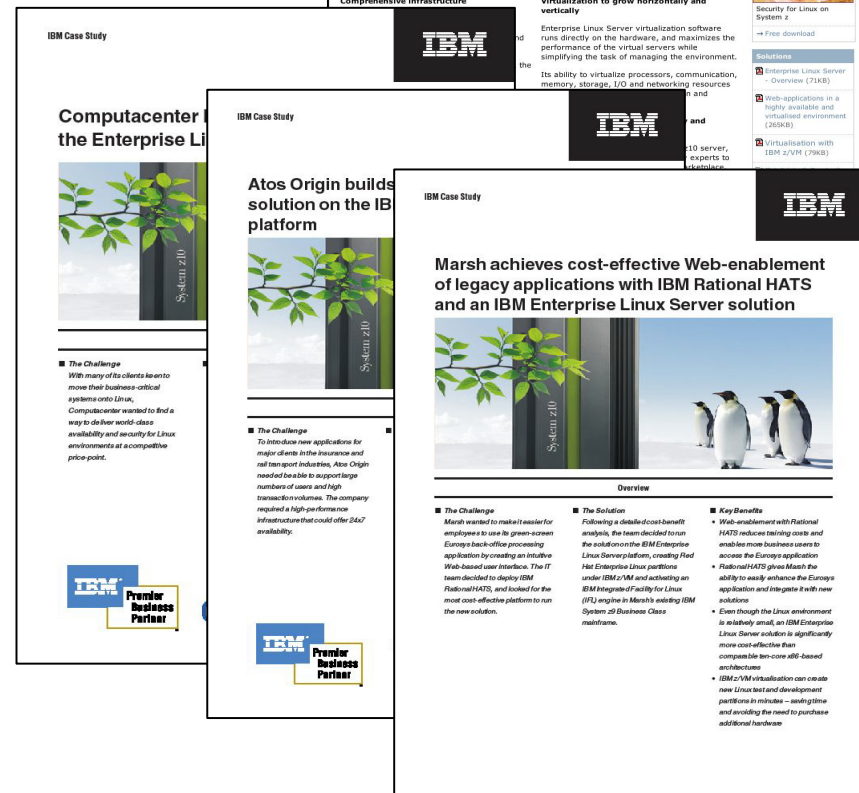

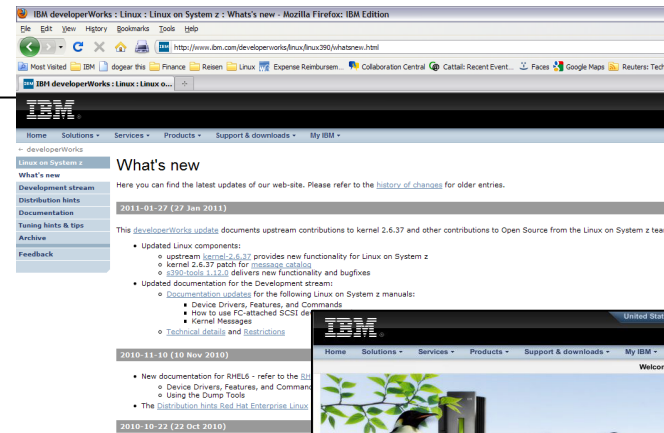
http://publib.boulder.ibm.com/infocenter/lnxinfo/v3r0m0/index.jsp?topic=%2Fliag%2F0lgr00_2012.htm

IBM developerWorks – What’s new

<http://www.ibm.com/developerworks/linux/linux390/whatsnew.html>

RedBooks

<http://www.redbooks.ibm.com/portals/linux>



IBM Continues to Support the Open Source and Linux Partner Ecosystem



- Over 6,000 IBM Business Partners
- Leaders for Linux Program
 - Custom co-marketing campaigns
- PartnerWorld Industry Networks
- Solution Builders Express
- Value Networks
 - Creating repeatable solutions
- Chiphopper
 - Porting ISV Applications

- **developerWorks Linux zone**
 - Tutorials, Redbooks, best practices, etc.




You are welcome to participate on ISV Products and Solutions !

- IBM PartnerWorld -> Global Solutions Directory

<http://www-304.ibm.com/partnerworld/gsd/homepage.do>

<http://www-304.ibm.com/partnerworld/gsd/search.do>



- IBM PartnerWorld -> Business Partner Application Showcase

<http://www-304.ibm.com/partnerworld/gsd/showcase.do>

- (W3) STG ISV Central

<http://w3.ibm.com/ibm/resource/isv.html>

- Red Hat ISV Partner Program

<http://www.redhat.com/partners/isv/>

<https://www.redhat.com/wapps/partnerlocator/web/home.html>



- Novell SUSE Linux Enterprise Software Catalog


<http://www.novell.com/partner/isv/isvcatalog>



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



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



Questions?



Wilhelm Mild
IBM IT Architect



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

Office: +49 (0)7031-16-3796
mildw@de.ibm.com

Acknowledgement

My best thanks belong to
Martin Schwidefsky
for the input and support to this
presentation

