

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

4th European GSE / IBM Technical University for z/VSE, z/VM and Linux on System z
25-27 October 2010, Hotel The Westin Grand München, Germany



IBM collaborates with the Linux community

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

Linux Kernel & Subsystem Development

Kernel Base
Security
Systems Mgmt
Virtualization
Filesystems,
and more...

Expanding the Open Source Ecosystem

Apache
Eclipse
Mozilla Firefox
OpenOffice.org,
and more...

Promoting Open Standards & Community Collaboration

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

Foster and Protect the Ecosystem

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



Facts on Linux

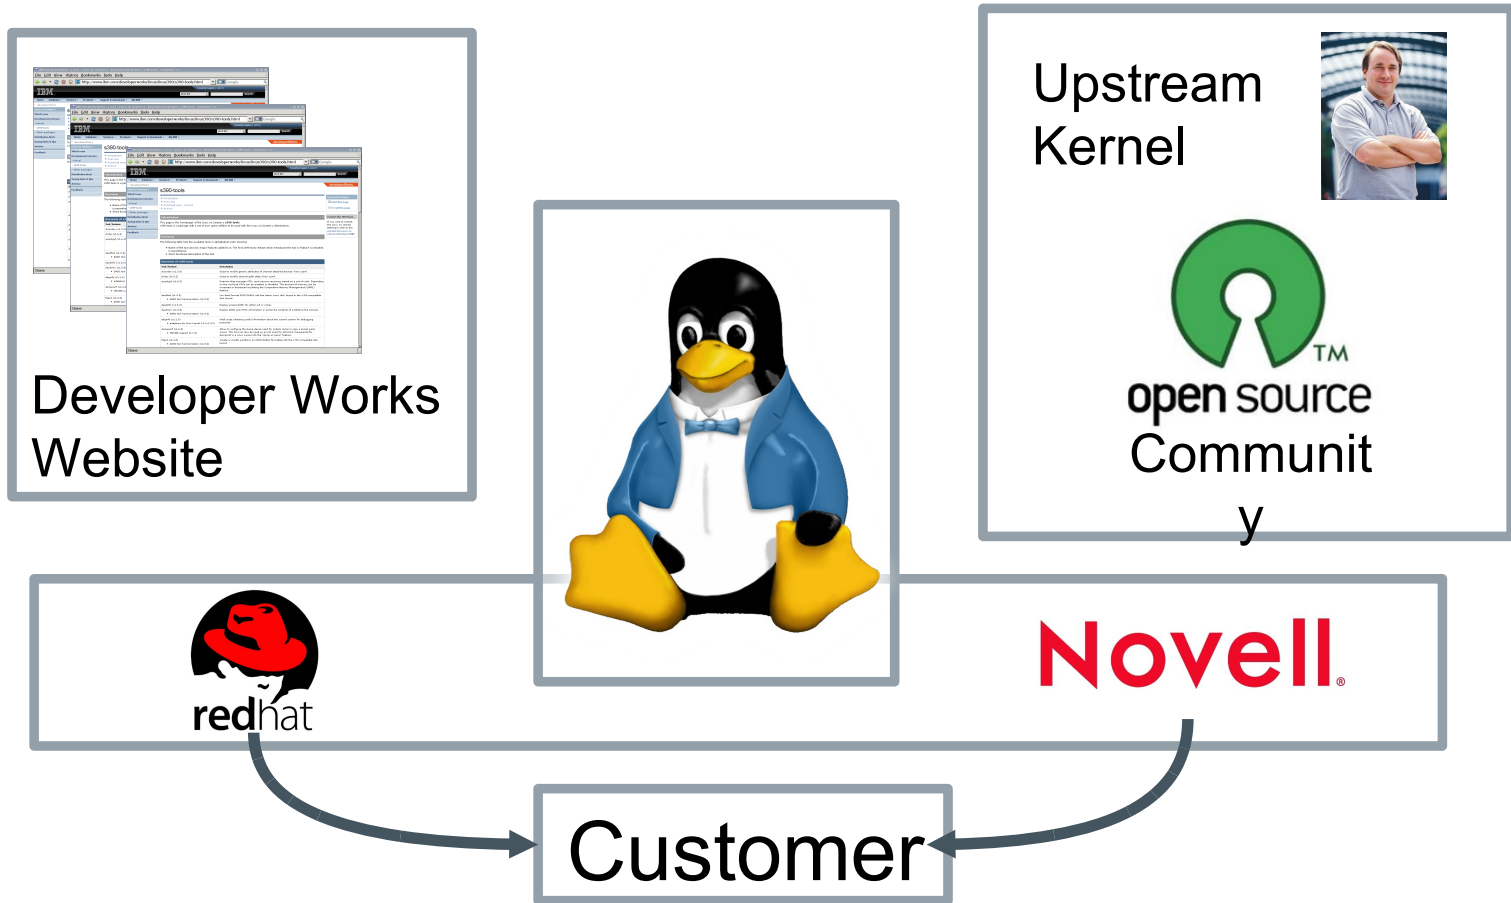
- Last year, **75%** of the Linux code was developed by **programmers working for corporations.**
- **\$7.37 billion:** projected cost to produce the 283 million lines of code which are contained in Linux Distribution **in a commercial environment.**
- IDC forecasts show that **Linux server revenue will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

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



IBM Linux on System z Development

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




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



Linux kernel development: System z contributions

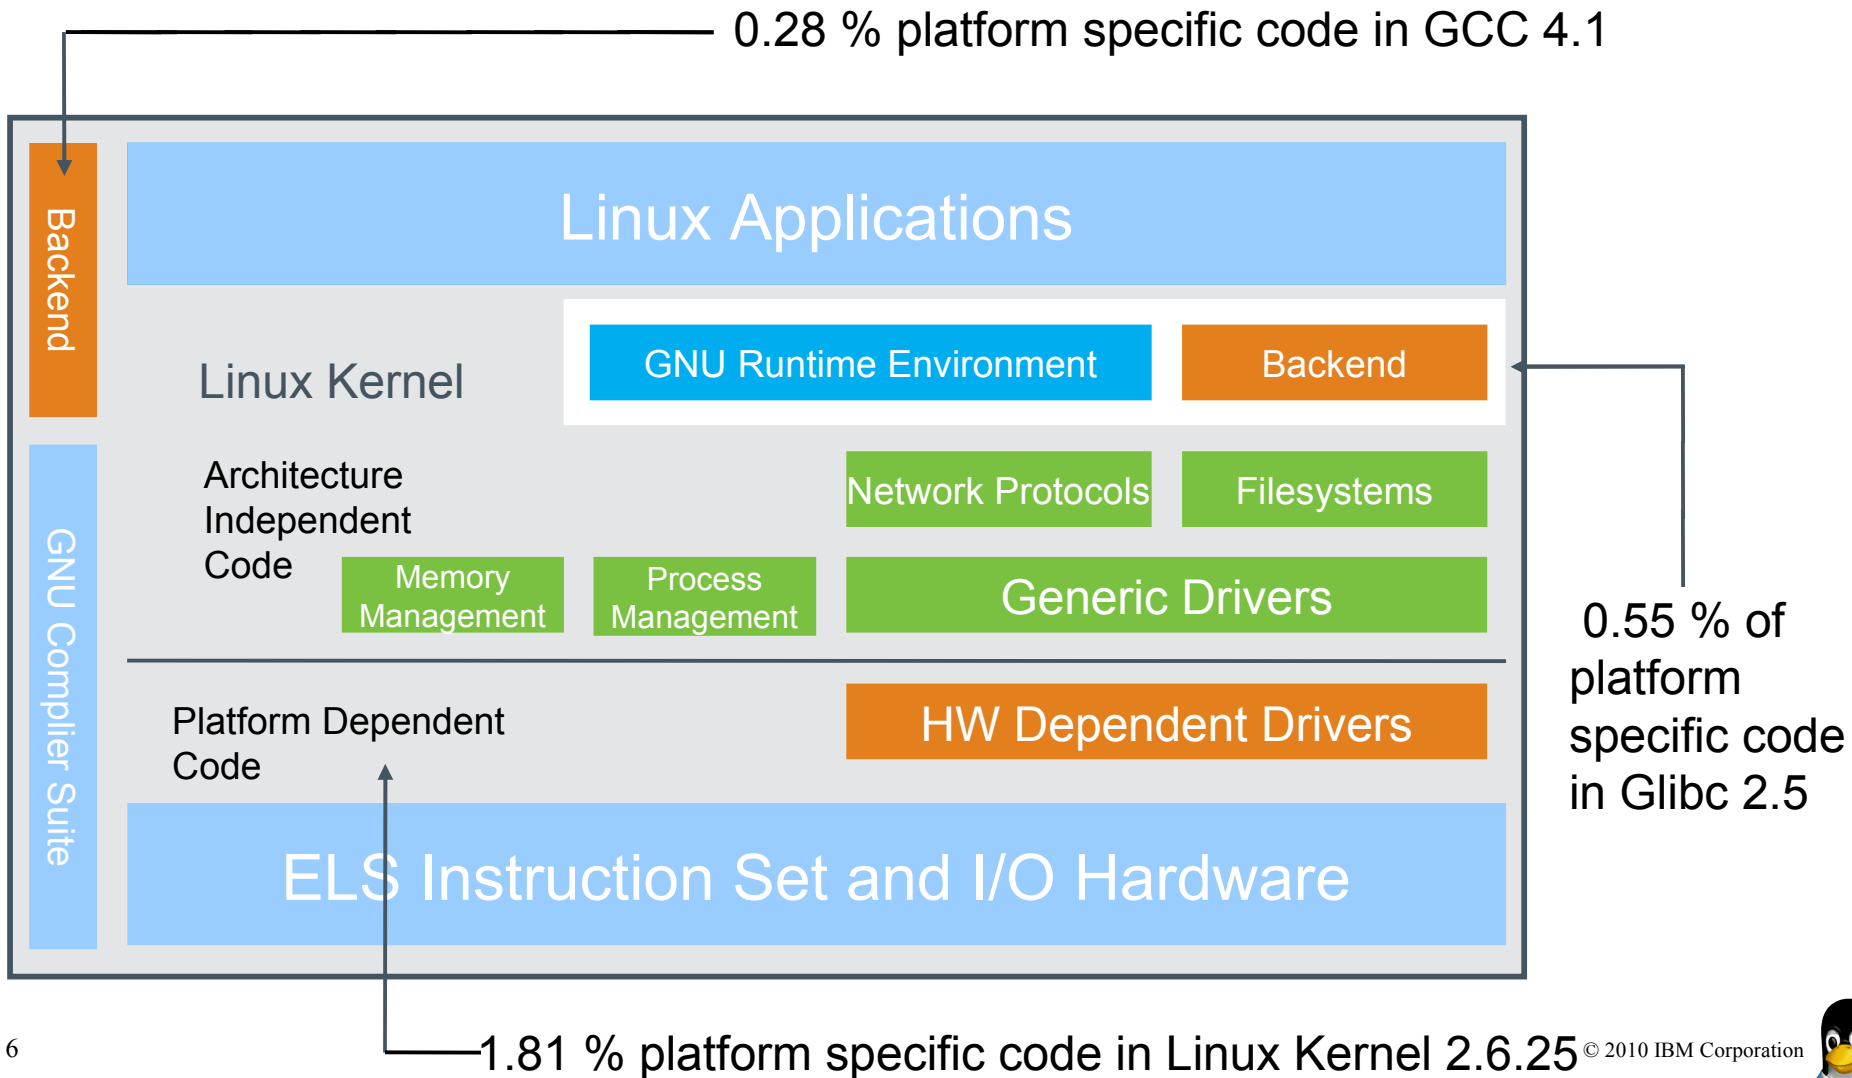
Changesets per release



Average: 6,422 lines added, 3,285 lines removed, and 1,687 lines changed every day for the past 4 1/2 years.

Structure of Linux on System z

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



Kernel news – Common code

Linux version 2.6.32 (2009-08-03)

- Per-backing-device based writeback (pdflush replaced by flush <major>)
- Kernel Samepage Merging (memory Deduplication)
- CFQ I/O scheduler low latency mode
- S+core architecture support

Linux version 2.6.33 (2010-02-24)

- DRDB
- Swappable KSM pages
- Compcache: memory compressed swapping

Linux version 2.6.34 (2010-05-16)

- Ceph distributed network file system
- LogFS flash memory file system
- Asynchronous suspend / resume
- 'Perf' performance analysis improvements, cross architecture support

Linux version 2.6.35 (2010-08-01)

- BKL removal work (incomplete)
- Intergration of AppArmor
- New out-of-memory killer
- Concurrency managed workqueues

Linux version 2.6.36 (2010-10-21)

- CIFS local caching



Linux on System z Development Focus

Integration

Application Serving

- z/OS & z/VSE integration

Data Hub

- Database Consolidation

Virtualization

Virtualization & Virtualization Management

- Ease of Use
- Serviceability
- Hosting capacity

Security

Security

- Certifications
- Data security & privacy

RAS

Continuous Availability & Data Replication

- RAS
- Differentiation for mission critical workloads

Base Tasks

Customer Requirements

- Address customer observed deficiencies

Competitiveness

- Close competitive gaps
- Differentiation / innovation that matters

Hardware Support

- Exploitation of new System z HW
- Storage exploitation

Linux

- Maintainership & code currency



IBM Supported Linux Distributions for System z

| Distribution | zEnterprise 196 | System z10 | System z9 | zSeries |
|---|-----------------|------------|-----------|---------|
| RHEL 5 | ✓ | ✓ | ✓ | ✓ |
| RHEL 4 (1) | — | ✓ | ✓ | ✓ |
| RHEL 3 (1) | — | — | * | ✓ |
| SLES 11 | ✓ | ✓ | ✓ | ✗ |
| SLES 10 | ✓ | ✓ | ✓ | ✓ |
| SLES 9 (1) | — | ✓ | ✓ | ✓ |
| (1) Also available as 31-bit distributions. | | | | |

✓ 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.

✗ Indicates that the distribution is not supported by IBM.

— Indicates that the distribution has not been tested by IBM.

* Supported on customer request (RPQ).



Future Linux on System z Technology

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



The IBM zEnterprise System -- A new dimension in computing ...is supported by Linux on System z



- The IBM zEnterprise 196 (z196) was announced 2010-07-22
- The following new features of the z196 are supported in the upstream kernel:
 - Third subchannel set
 - Up to 32 HiperSockets
 - Support for the new OSA CHPID types OSX and OSM is accepted upstream for integration into kernel 2.6.35 -- appropriate s390-tools updates will be published shortly.

Get all the important Hardware details at:

http://www.ibm.com/common/ssi/rep_ca/0/897/ENUS110-170/ENUS110-170.PDF



Upcoming zEnterprise Linux features

- **System z196 toolchain support (binutils + gcc)**
Add new instructions to the gnu assembler and disassembler. Teach the gnu compiler about the new pipeline structure
- **System z196 enhanced node affinity (> kernel 2.6.36)**
Create appropriate scheduling groups for the 2-level CPU topology of System z196. The extended CPU topology information is provided by a new STSI sub-command.
- **Performance indicator bytes (> kernel 2.6.36)**
The static power savings mode of System z196 allows to lower the energy consumption of the machine. This reduces the performance and is indicated via STSI.
- **Access-Exception-Fetch/Store-Indication (> kernel 2.6.36)**
Provide indication if an access exception has been due to a fetch or a store operation. This is used to avoid protection faults if the first access to a mapped page has been a write.



Virtualization

- **Extra kernel parameter for SCSI IPL (kernel 2.6.32)**
Modify the SCSI loader to append extra parameters specified with the z/VM VMPARM option to the kernel command line.
- **Deliver z/VM CP special messages as uevent (kernel 2.6.34)**
Allows to forward SMSG messages starting with “APP” to user space. udev rules can be used to trigger application specific actions
- **Automatic detection of read-only devices (2.6.34)**
Improve usability by automatically detection of read-only dasd devices with diagnose 210
- **CMSFS user space file system support (s390-tools 1.9.0 for the read-only cmsfs support)**
Implement a FUSE file system that allows to read from and write to CMSFS minidisks. Writing is difficult, the record based CMSFS does not fit well into the byte stream oriented Linux VFS



hyptop - Display hypervisor performance data

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

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

The following things are required to run hyptop:

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



hyptop – Displaying hypervisor performance data

Displaying performance data for the z/VM hypervisor

```
10:11:56 CPU-I: 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 | 500 |
| VTAM | 1 | 0.00 | 0:01 | 62:21:46 | 0.01 | 0.03 | 100 |
| T6360023 | 2 | 0.00 | 0:04 | 6:21:20 | 0.46 | 0.75 | 100 |
| PERFSVM | 1 | 0.00 | 2:12 | 7:18:04 | 0.05 | 0.06 | 0 |
| AUTOVM | 1 | 0.00 | 0:03 | 62:21:46 | 0.00 | 0.03 | 100 |
| FTPSEVE | 1 | 0.00 | 0:00 | 62:21:47 | 0.01 | 0.03 | 100 |
| TCPIP | 1 | 0.00 | 0:01 | 62:21:47 | 0.01 | 0.12 | 3000 |
| DATAMOVE | 1 | 0.00 | 0:06 | 62:21:47 | 0.00 | 0.03 | 100 |
| VMSERVU | 1 | 0.00 | 0:00 | 62:21:47 | 0.00 | 0.03 | 1500 |
| OPERSYMP | 1 | 0.00 | 0:00 | 62:21:47 | 0.00 | 0.03 | 100 |



hyptop – Displaying hypervisor performance data

Displaying performance data for a single LPAR

```
10:16:59 H05LP30 CPU-I: IFL(18) CP(3) UN(2)                                     ?=help
```

| cpu_id (#) | type (str) | cpu (%) | mgm (%) | visual (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 | |



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

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

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

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

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



Networking

- **OSA QDIO Data Connection Isolation (kernel 2.6.33)**
Isolate data traffic from Linux on System z guests sharing an OSA card
Communication between guests needs to go over via external entity
- **HiperSockets Network Traffic Analyser (kernel 2.6.34)**
Trace HiperSockets network traffic for problem isolation and resolution.
Supported for layer 2 and layer 3
- **Offload outbound checksumming (kernel 2.6.35)**
Move calculation of checksum for non-TSO packets from the driver to the OSA network card
- **Toleration of optimized latency mode (kernel 2.6.35)**
OSA devices in optimized latency mode can only serve a small number of stacks / users. Print a helpful error message if the user limit is reached.
Linux does not exploit the optimized latency mode



Networking (cont)

- **NAPI support for QDIO and QETH (> kernel 2.6.35)**
Convert QETH to the NAPI interface, the “new” Linux networking API
NAPI allows for transparent GRO (generic receive offload)
- **QETH debugging per single card (> kernel 2.6.35)**
Split some of the global QETH debug areas into separate per-device areas
Simplifies debugging for complex multi-homed configurations
- **Configuration tool for System z network devices (s390-tools 1.8.4)**
Provide a shell script to ease configuration of System z network devices



znetconf network device configuration tool

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

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

- Configure device 0.0.f503

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

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

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

- Remove network device 0.0.f503

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



New Linux on System z Storage Features

- **FCP adjustable queue depth (kernel 2.6.31)**
Customizable queue depth for SCSI commands in zfc
- **Resume reordered devices (kernel 2.6.34)**
Allow resume of a guest with different subchannels for individual devices. Allow suspend of a system with devices in the disconnected state
- **Unit check handling (kernel 2.6.35)**
Improve handling of unit checks for internal I/O started by the common- I/O layer
After a unit check certain setup steps need to be repeated, e.g. for PAV
- **Store I/O status and initiate logging (SIOSL) (> kernel 2.6.35)**
Enhance debug capability for FCP attached devices
Enables operating system to detect unusual conditions on a device of channel path



New Linux on System z Storage Features (cont)

- **Automatic LUN scanning (> kernel 2.6.35)**
Scan and attach accessible LUNs automatically
Available only for a NPIV FCP attachment
- **Automatic menu support in zipl (> s390-tools 1.9.0)**
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 1.9.0)**
The automatic re-IPL function only works with a physical device
Enhance the zipl support for device-mapper devices to provide the name of the physical device if the zipl target is located on a logical device



Usability / RAS

- **Dump on panic – prevent reipl loop (s390-tools 1.8.4)**
Delay arming of automatic reipl after dump
Avoids dump loops where the restarted system crashes immediately
- **Add support for makedumpfile tool (kernel 2.6.34, s390-tools 1.9.0)**
Convert Linux dumps to the ELF file format
Use the makedumpfile tool to remove user data from the dump
Multi-volume tape dump will be removed
- **Breaking event address for user space (kernel 2.6.35)**
Store the breaking-event-address for user space programs
Valuable aid in the analysis of wild branches
- **Precise process accounting (> kernel 2.6.36)**
Extend the taskstats interface to provide better process accounting values
Quality goal is a resolution of 10ths of microseconds



System z toolchain

- **64 bit register in 31 bit compat mode**

Make use of 64 bit registers in 31 bit application running in z/Architecture mode.
Allows to use instruction operating on 64 bits, e.g. 64 bit multiplication
Needs kernel support for asynchronous signals

- **Oprofile hardware customer mode sampling**

Provide CPU measurement data to applications for performance tuning
Based on hardware counters and samples built into the CPU
Use oprofile to communicate the information to user space programs

- **Valgrind System z support**

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



Valgrind System z support

- **valgrind --tool=memcheck [--leak-check=full] [--track-origins] <program>**

Detects if your program accesses memory it shouldn't

Detects dangerous uses of uninitialized values on a per-bit basis

Detects leaked memory, double frees and mismatched frees

- **valgrind --tool=cachegrind**


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

Identifies the number of cache misses

- **valgrind --tool=massif**

Profile heap usage, takes regular snapshots of program's heap

Produces a graph showing heap usage over time



System z kernel features - Misc

- **Kernel image compression (kernel 2.6.34)**

The kernel image size can be reduced by using one of three compression algorithms: gzip, bzip2 or lzma.

- **KULI (2009-06-24)**

kuli is experimental userspace sample to demonstrate that KVM can be used to run virtual machines on Linux on System z.

This experimental proof of concept is unsupported and should not be used for any production purposes.

- **Oprofile**

Starting with version 0.9.4, oprofile supports sampling of Java byte code applications for Linux on System z.

- **Eclipse 3.3**

Starting with Eclipse 3.3, Linux on System z is officially supported.



Current Linux on System z Technology

Features & Functionality contained in the Novell
& Red Hat Distributions



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 3 (GA 09/2009)
- 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 8 (GA 05/2009)
- Red Hat Enterprise Linux AS 5 (GA 03/2007)
 - Kernel 2.6.18, GCC 4.1.0, Update 5 (GA 03/2010)
- Others
 - Debian, Slackware, Fedora
 - Support may be available by some third party

Red Hat Enterprise Linux 6 Beta 2 has been published as a public download as of July 21, 2010 at <http://www.redhat.com/rhel/beta/> using Kernel 2.6.32 GCC 4.4.3



Integration

- **AF_IUCV SOCK_SEQPACKET support**

Introduce AF_IUCV sockets of type SOCK_SEQPACKET that map read/write operations to a single IUCV operation.

The socket data is not fragmented.

The intention is to help application developers who write applications using the native IUCV interface, e.g. Linux to z/VSE.



- **HiperSockets Layer3 support for Ipv6**

Providing Layer3 IPv6 communication, for communication to z/OS



- **Linux to add Call Home data if running in LPAR**

Also referred to as Control Program Identification (CPI) or SCLP_CPI

Allows the user to set information about the LPAR which will be displayed on the HMC/SE



Virtualization

- **TTY terminal server over IUCV**

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


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

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



- **Dynamic memory attach/detach**

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




- **Extra kernel parameter via VMPARM**

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

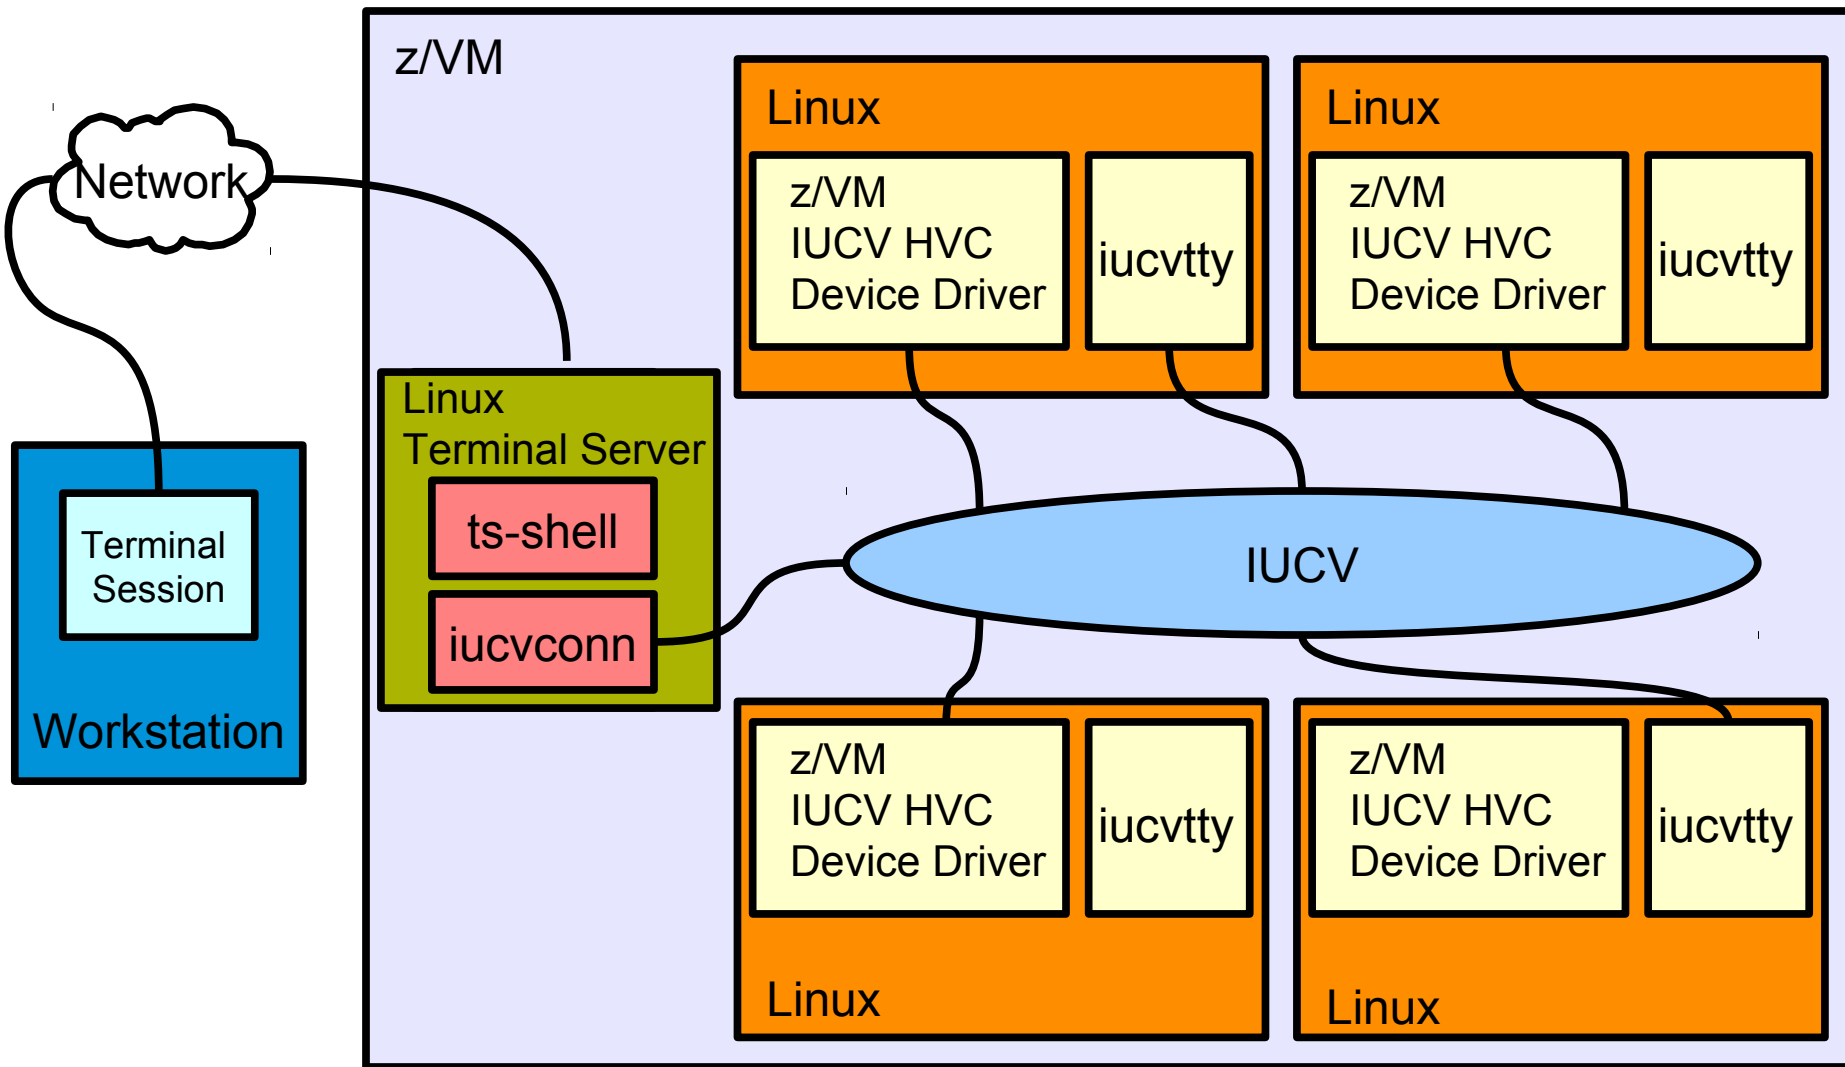


- **Provide CMS script for initial IPL**

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





IUCV terminal environment



Virtualization (cont.)

- **Exploitation of DCSSs above 2G**
Solves restriction to use DCSS above or greater than 2GB.
- **Provide service levels of HW & Hypervisor in Linux**
Improves serviceability by providing uCode and z/VM levels via /proc interface

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



Security

- **Long Random Numbers Generation**
Provide access to the random number generator feature on the Crypto card (high volume random number generation, compared to a CPU based solution)
- **Crypto Express3 cards enablement**
Support for Crypto Express3 Accelerator (CEX3A) and Crypto Express3 Coprocessor (CEX3C)
- **Crypto device driver use of thin interrupts**
Provides better performance and lower CPU consumption.



RAS

- **Shutdown Actions Interface**

The shutdown actions interface allows the specification of a certain shutdown action (stop, ipl, reipl, dump, vmcmd) for each shutdown trigger (halt, power off, reboot, panic)

Possible use cases are e.g. to specify that a vmdump should be automatically triggered in case of a kernel panic or the z/VM logoff command should be executed on halt.



- **Automatic IPL after dump**

The new shutdown action `dump_reipl` introduces a system configurations which allows to create a dump in case of a Linux panic, followed by a re-ipl of the system, once the dump was successfully created.

Allows to configure system to re-ipl after a dump is taken.



RAS

- **Suspend / resume support (kernel 2.6.31)**
Add the ability to stop a running Linux system and resume operations later on. The image is stored on the swap device and does not use any system resource while suspended.
Only suspend to disk is implemented, suspend to RAM is not supported.
- **Add Call Home data on halt and panic if running in LPAR (kernel 2.6.32)**
Report system failures (kernel panic) via the service element to the IBM service organization. Improves service for customers with a corresponding service contract. (by default this features is deactivated)
- **Large image dump on DASD**
Solves restriction to dump only 48GB of memory to DASD. Now up to 32 ECKD DASDs can be used in a multiple volume configuration



Suspend / resume support

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

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

- Suspend operation is started with a simple echo:

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

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

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



Storage

- **FICON DS8000 Large Volume (EAV) Support**
Large Volume Support is a feature that allows to use ECKD devices with more than 65520 cylinders (>50GB).
This features is available with DS8000 R4.0 Allows to exploit
- **DS8000 Disk Encryption Support**
Shows the encryption status of the DS8000 Storage.
- **EMC Symmetrix DASD Format Record 0**
Allows to initialized unformatted disks on EMC storage arrays
- **FCP LUN discovery tool**
New LUN discovery tool: Isluns (e.g. used by yast)
- **FCP performance data collection & reports:**
Fibre Channel Protocol (FCP) performance data can now be measured.



Red Hat Enterprise Linux 5 Update 5

- **GA since 03/30/2010**
 - Kernel GA: 2.6.18-194
- **New Features:**
 - ***FICON DS8000 Large Volume (EAV) Support:*** Allows to exploit DS8000 Storage feature to use DASD volumes >50GB.
 - ***AF_IUCV SOCK_SEQPACKET support:*** Enhances existing AF_IUCV to allow customer to develop using SOCK_SEQPACKET.
 - ***Provide CMS script for initial IPL:*** Avoids having to create an script to start a new installation under z/VM.
 - ***Installer re-IPL support:*** Solves past restriction and allows the installer to direct reboot in the installed system right after installation
- **Bugfixes**



SuSE Linux Enterprise Server 11 SP1

- **GA since June 2, 2010**
 - Kernel GA:kernel 2.6.32
- **New Features:**
 - **ALL Linux on System z upstream kernel features since 2.6.27**
 - **Suspend / Resume support:** stop a running instance and later continue operations. A suspended Linux instance does not require memory or processor cycles. gives you better performance, resource utilization, and power savings
 - **Automatic IPL after dump:** extension to the shutdown action interface which combines the actions dump and re-ipl, helps increase availability and minimize downtime, as well as keepmanagement and service costs low
 - **DS8000 support - Large volume support architecture:** use one large volume, instead of multiple small volumes, for your large amount of data. You no longer need to combine and manage various small disks anymore.
 - **Support of HPF:** increases performance for database serving
 - **Next generation crypto HW device driver exploitation:** new System z crypto hardware features and performance improvements are exploited by SUSE Linux Enterprise Server for System z. Hardware-driven crypto acceleration functions help reduce operations and maintenance costs.



SuSE Linux Enterprise Server 11 SP1 (cont)

New Features:

- **AF_IUCV SOCK_SEQPACKET support**: improves close collaboration between SUSE Linux Enterprise Server for System z and z/VM in the networking area. This provides better performance for intra machine / VM communication.
- **TTY terminal server over IUCV**: provides central access to the Linux console for the different guests of a z/VM. Fullscreen applications like vi are usable on the console.
- **System z kernel message documentation**: Cleanup messages in System z related code, script to generate a man page for every kernel message
- **FCP adjustable queue depth**: Customizable queue depth for SCSI commands in zfcpl. In the past was at constant 32 queue entries. Improves performance

- **Bugfixes**

- **More Information:**

http://www.ibm.com/developerworks/linux/linux390/documentation_novell_suse.html



More Information

The screenshot shows the IBM developerWorks website. The main heading is "Documentation for Development stream". Below it, there are tabs for "Development stream", "Novell SUSE", and "Red Hat". The "Development stream" tab is highlighted with a red circle. A blue arrow points from the top right towards this tab. On the left side, there is a navigation menu with categories like "What's new", "Development stream", "Distribution hints", "Documentation", "Tuning hints & tips", and "Archive".

Linux on System z

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

March, 2010

Linux on System z

Using the Dump Tools

Development stream (Kernel 26.33)

Linux on System z

Kernel Messages

Development stream (Kernel 26.33)

Linux on System z

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

Development stream (Kernel 26.33)

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

Device Drivers, Features, and Commands

Development stream (Kernel 26.33)

New: Distribution specific Documentation



Questions?



Hans-Joachim Picht
Linux on System z Initiatives

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

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



Your Linux on System z Requirements?

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



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

Send your input to hans@de.ibm.com



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

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



Trademarks & Disclaimer

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

The following are trademarks or registered trademarks of other companies

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

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

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions. This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

