SUSE® Linux Enterprise Server

IBM System z

Kai Dupke
Senior Product Manager
kdupke@suse.com

Date: 2012-08-02
Distribution: PDF any

SUSE and the Attachmate Group

- SUSE, headquartered in Nürnberg / Germany, is an independently operating business unit of the Attachmate Group, Inc.
- The Attachmate Group is a privately held 1 billion+ $ revenue software company with four brands:

SUSE® Leadership

MAINFRAME LINUX

OVER 80%

of all Linux running on mainframe computers is SUSE Linux Enterprise Server

SAP ON LINUX

OVER 70%

of all SAP running on Linux runs on SUSE Linux Enterprise Server

LINUX IN CHINA

SUSE Linux Enterprise Server is the most widely used commercial enterprise Linux distribution in China — more than Red Hat

LINUX IN AEROSPACE AND DEFENSE

Nearly 80% of the US Fortune 500 aerospace and defense companies use SUSE Linux Enterprise Server

MOST CERTIFIED APPLICATIONS

Over 8500 applications are certified and supported on SUSE Linux Enterprise Server, more than any other Linux distribution

OVER 8500

LINUX IN RETAIL

NEARLY **70%**

of the US Fortune 100 general merchandisers, specialty retailers, and food and drug stores use SUSE Linux Enterprise Server

LINUX IN AUTOMOTIVE

SUSE Linux Enterprise Server is used by nearly all of the world's major automobile manufacturers

LINUX IN HPC

Half of the world's largest supercomputer clusters use SUSE Linux Enterprise Server

BEST LINUX SUPPORT

SUSE offers better Linux support than Red Hat or Oracle

LINUX IN GLOBAL FORTUNE 100

Over two-thirds of the global Fortune 100 use SUSE Linux Enterprise Server

MOST CERTIFIED HARDWARE

Over 13,500 hardware systems are certified and supported on SUSE Linux Enterprise, more than any other Linux distribution

13,500

SUSE

SUSE® Strategy

Cloud Infrastructure

SUSE®

Enterprise Computing

Integrated Systems

SUSE Linux Enterprise

Product Lifecycle

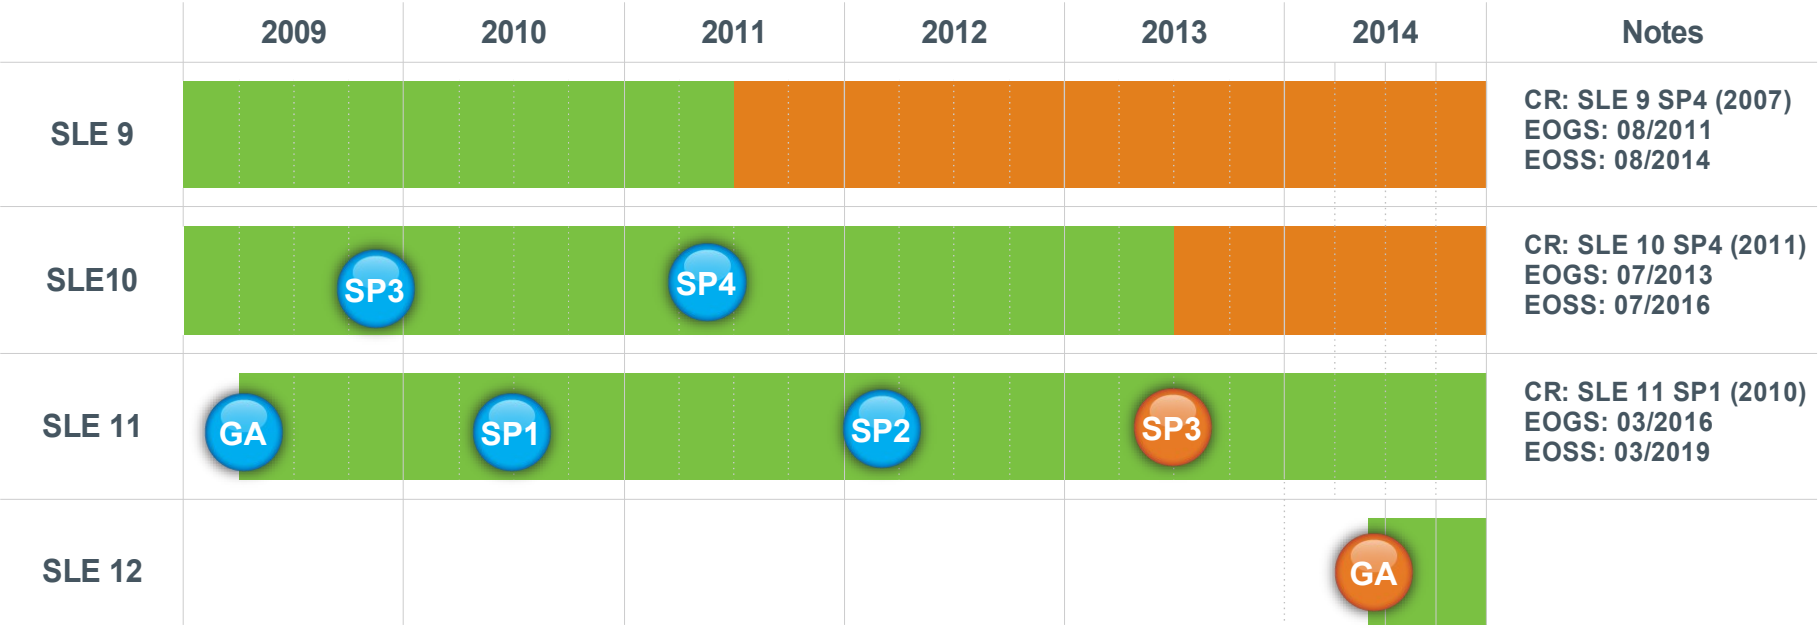

- SUSE announces service pack releases and development and product schedules to customers and partners
- Dependable release timing
- Predictability for planning rollouts and migrations
- Major releases every 4-5 years.

The SUSE® Build Service* Advantage

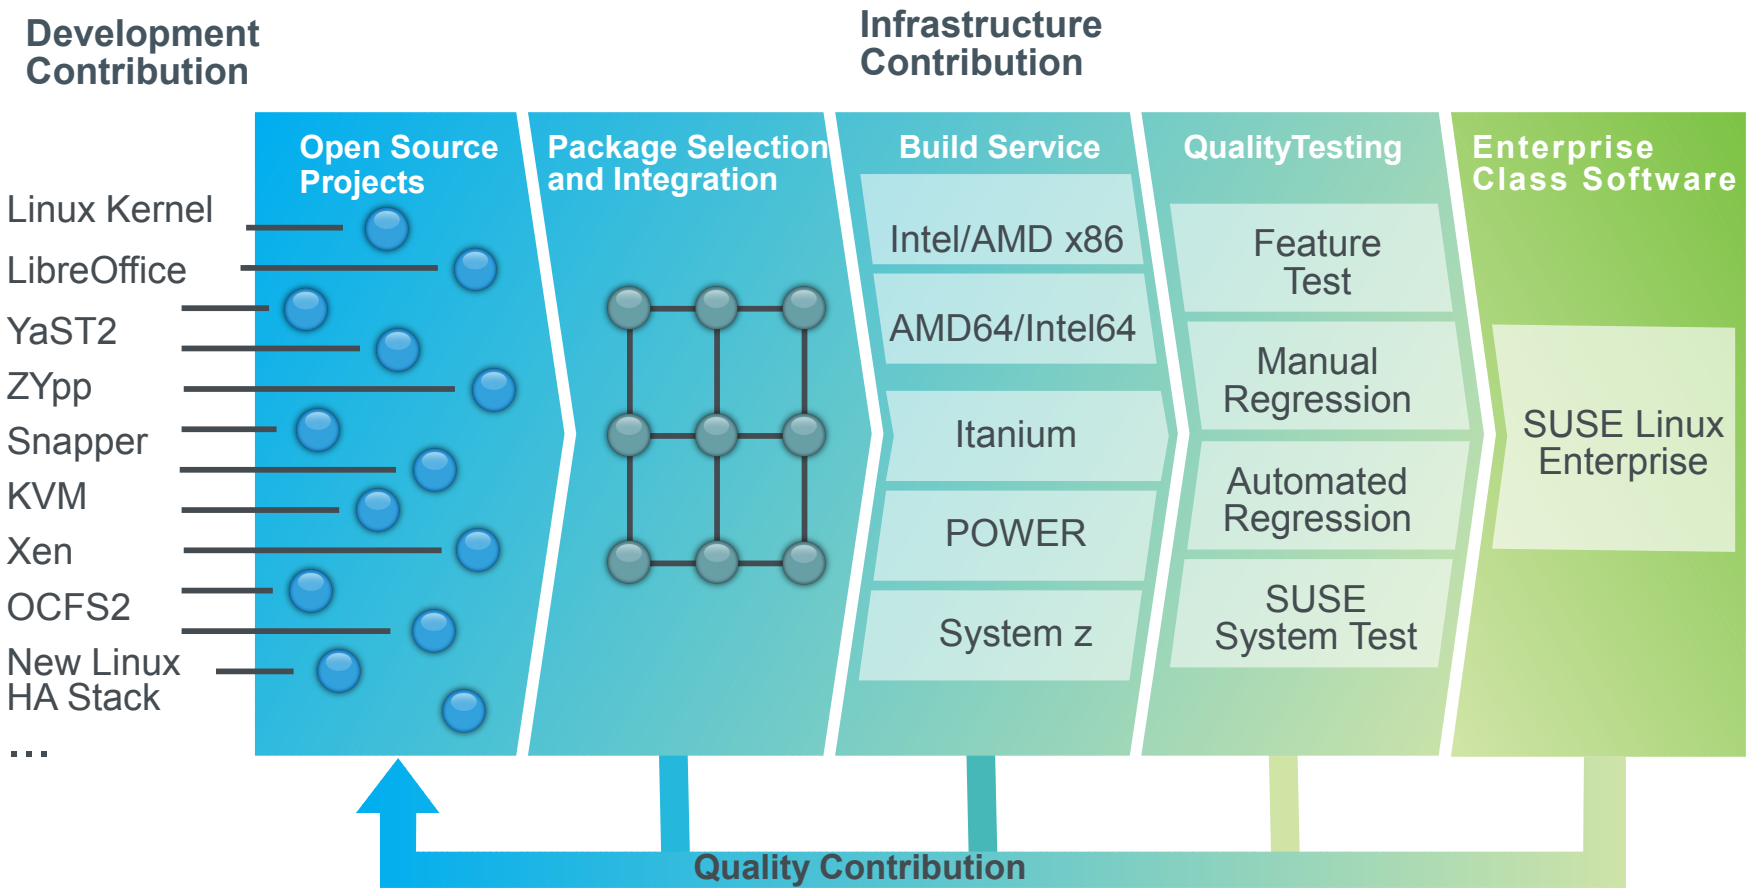

* SUSE Build Service is the internal entity of the Open® Build Service

Platform Lifecycle

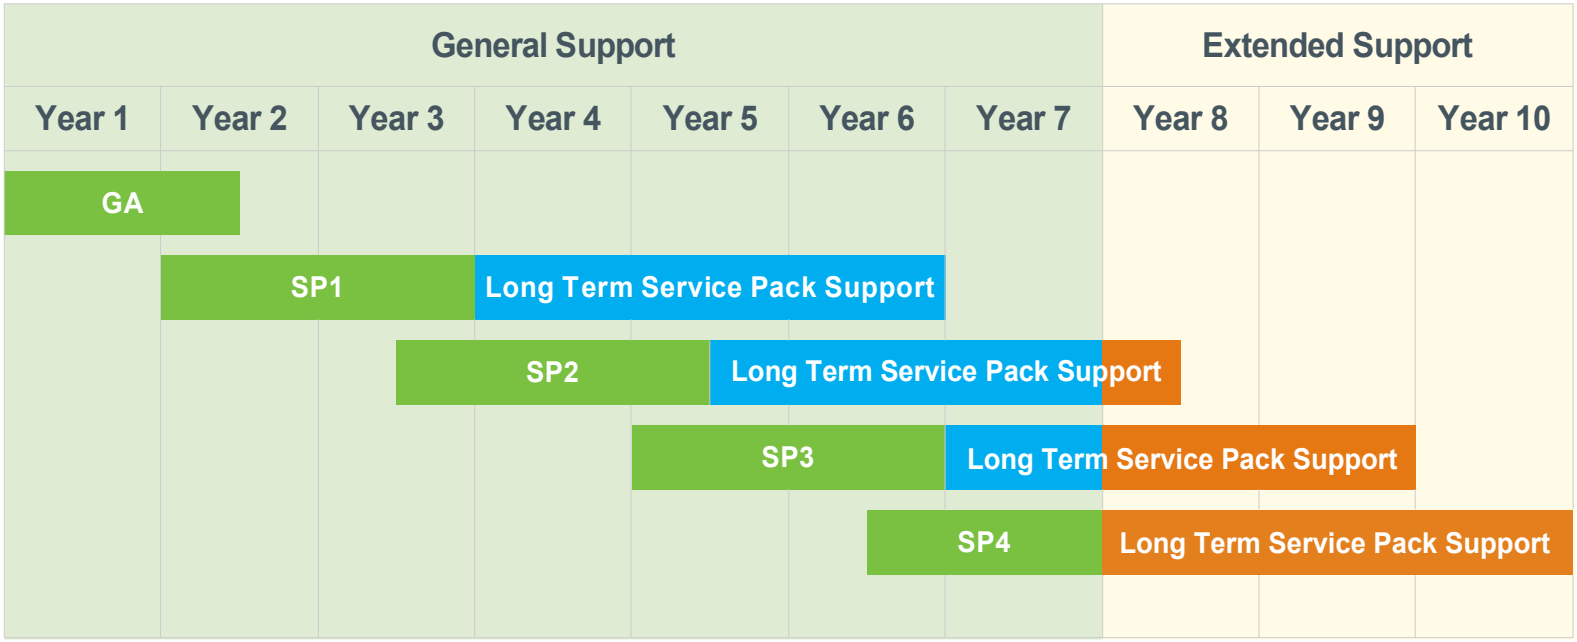

- 10-year lifecycle (7 years general support, 3 years extended support)
- Service Packs are released every ~18 months
  - 5 years lifetime with
  - ~2 years general support per Service Pack
  - 6 month upgrade window after release of the next Service Pack
- Long Term Service Pack Support (LTSS)
  - Extend upgrade window or extend major release lifecycle

Long Term Service Pack Support (LTSS)

Use Cases

I want to run my software stack unchanged for a very long time

- Updating OS does not improve my business process
- Updates can be very expensive to deploy
- Any change may impose additional risk

I need more time to move to the next Service Pack

- Approval process from stake holders
- QA processes
- Very large and/or distributed environment

SUSE Linux Enterprise 11 SP2

SUSE Linux Enterprise 11 SP2

- Hardware enablement and RAS
- Solaris compete
  - btrfs: file system with *Copy on Write*, checksums, snapshotting
  - LXC: container support based on control groups
  - LTTng (Linux Trace Toolkit) capabilities
- Snapshot / rollback for package and configuration updates
  - YaST2 + ZYPP + btrfs
- SUSE Maintenance Model
  - Increased flexibility while retaining full control
- SUSE Linux High Availability
  - Geo-cluster, easy and automated installation

Shipping since February 2012

SUSE® Linux Enterprise 11 SP2

Kernel 3.0

- Most recent HW enablement
- Removal of BLK (Big Kernel Lock)
- Control Groups enhancements
  - I/O throttling support for process groups
  - memory cgroup controller
- Integration of AppArmor
- More powerful firewalls with faster packet filtering
- Transparent Huge Pages (THP)

Snapshots in SUSE® Linux Enterprise 11 SP2

YaST2 Management

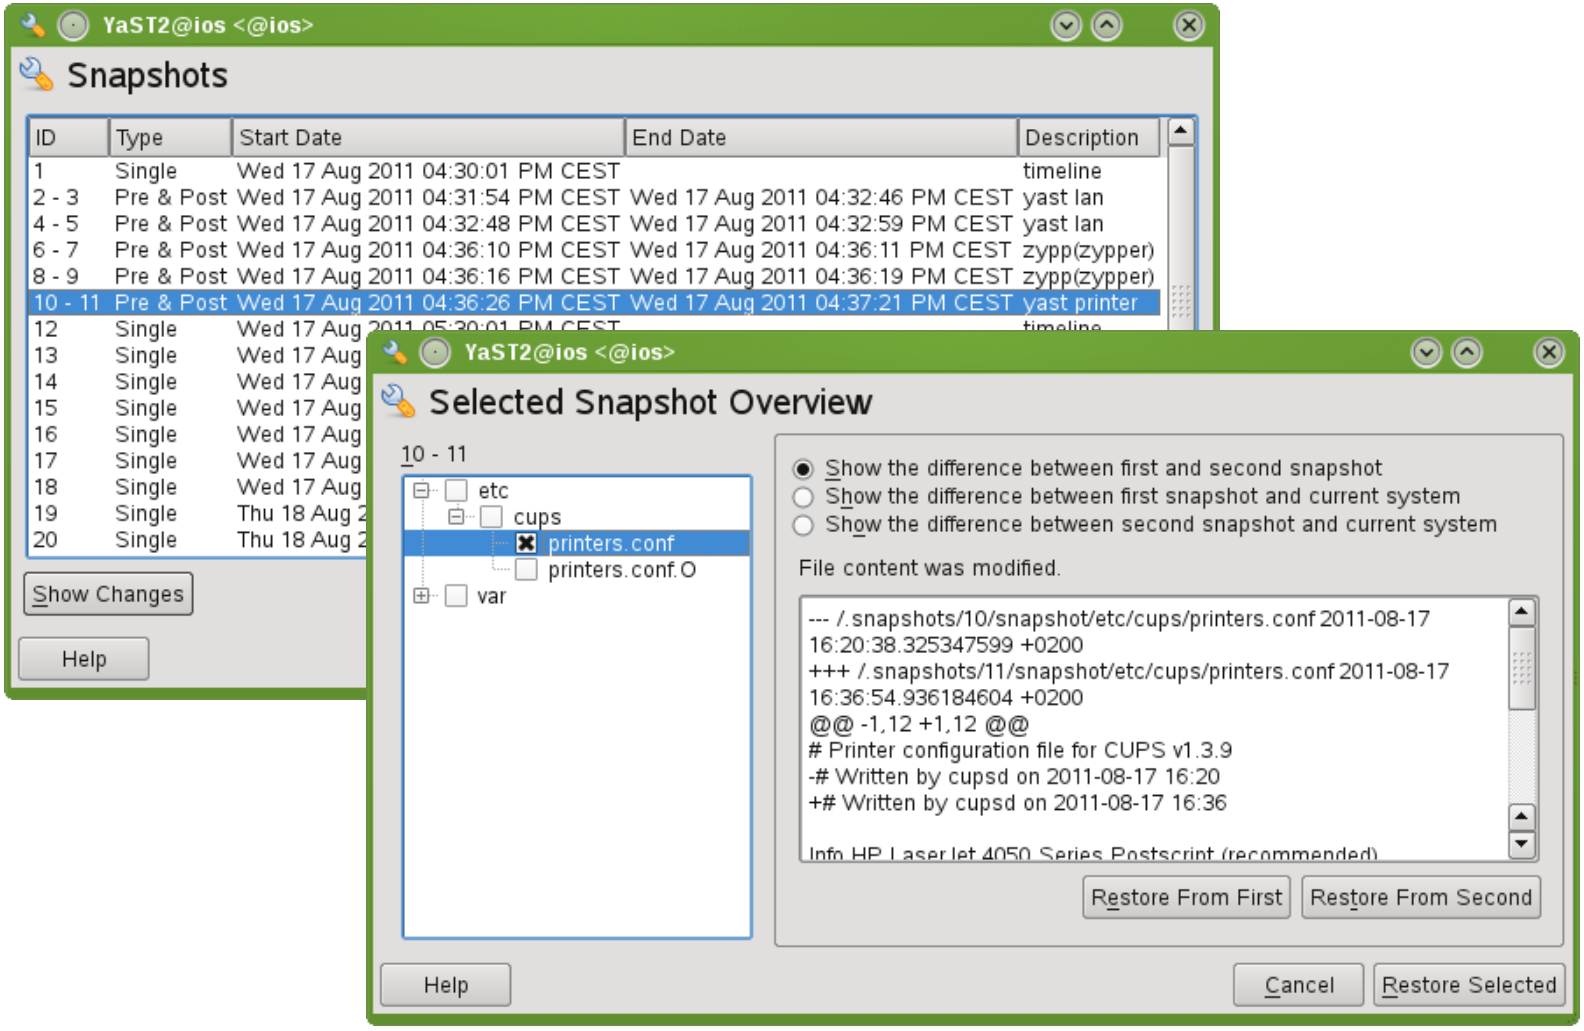

SUSE® Linux Enterprise 11

Maintenance Model Benefits

1. Make service packs more lightweight
   - easier to test and deploy
2. Allow for staying with older versions
   - with support for the full system
3. Answer market needs in between Service Packs by
   - Selective enhancements
   - Allowing more updates in the General Update repository

cgroups - Resource Control

Consider a large university server with various users - students, professors, system tasks etc. The resource planning for this server could be along the following lines:

CPUs

Top cpuset (20%)

/ \

CPUSet1 CPUSet2

| |

(Profs) (Students)

60% 20%

Memory

Professors = 50%

Students = 30%

System = 20%

Disk I/O

Professors = 50%

Students = 30%

System = 20%

Network I/O

WWW browsing = 20%

/ \

Prof (15%) Students (5%)

Network File System (60%)

Others (20%)

Source: /usr/src/linux/Documentation/cgroups/cgroups.txt

cgroups - Memory Subsystem

- For limiting memory usage of user space processes.
- Limit LRU (Least Recently Used) pages
  - Anonymous and file cache
- No limits for kernel memory
  - Maybe in another subsystem if needed
  - Note: cgroups need ~2% of (resident) memory
    - can be disable at boot time with kernel paramenter "cgroup_disable=memory"

Source: http://jp.linuxfoundation.org/jp_uploads/seminar20081119/CgroupMemcgMaster.pdf

Key Facts

SUSE® Linux Enterprise

Application Choice

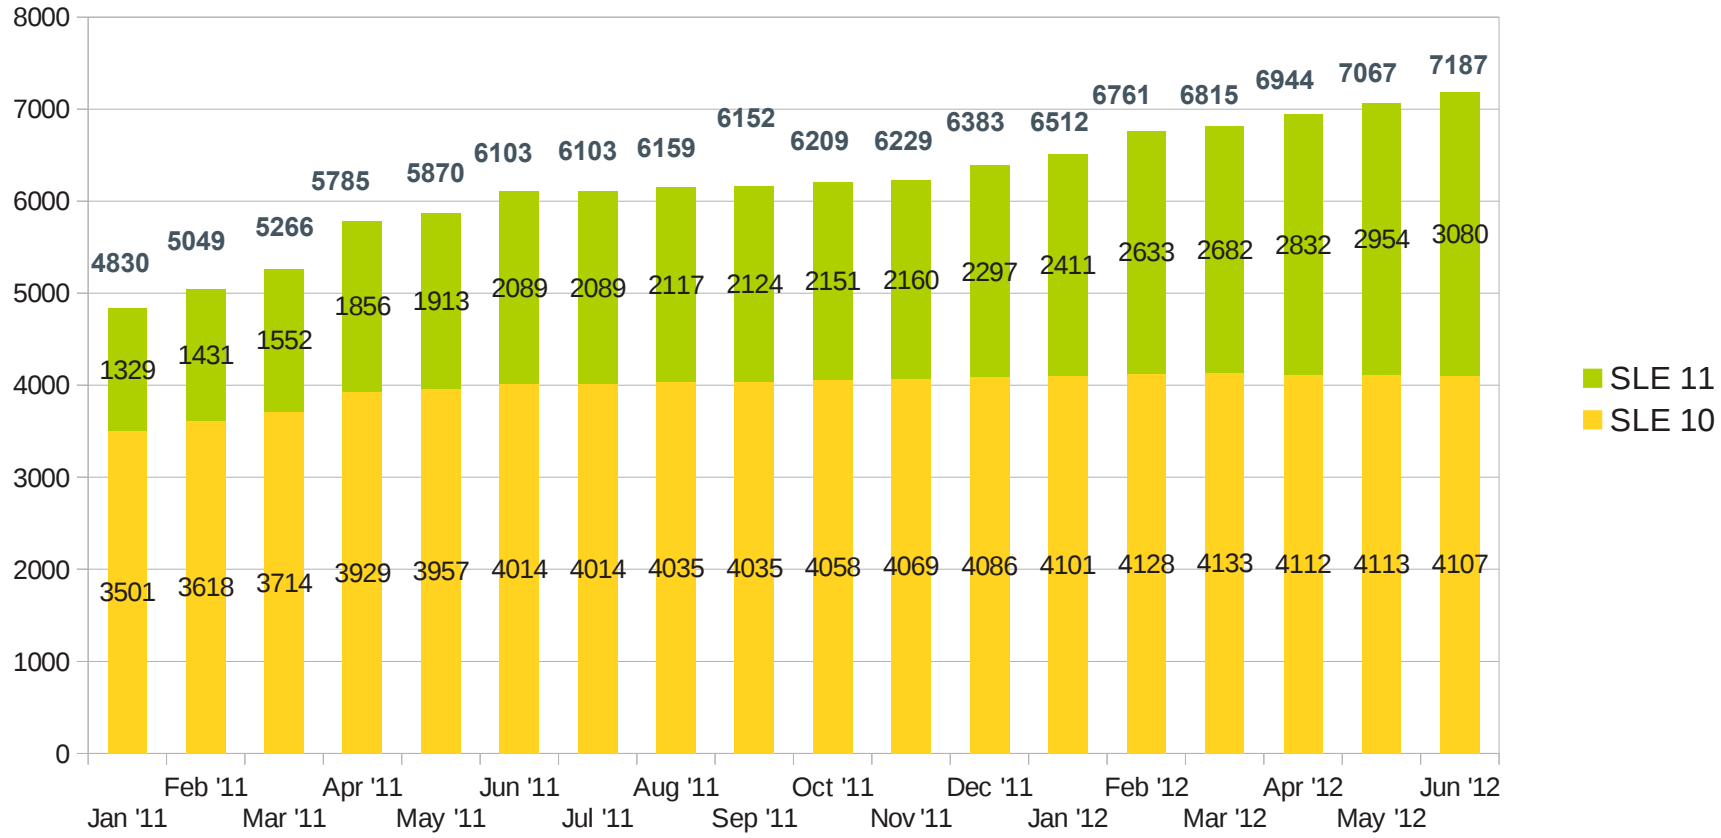

SUSE® Linux Enterprise Server

Technical Facts


Highly scalable


Available cross architecture


Windows and Unix interoperability and integration


High availability included


Systems management included

SUSE Linux Enterprise Server for System z


IBM
IBM
zEnterprise 196
z BladeCenter Extension

SUSE Linux Enterprise Server 11 for System z

- Full Dynamic Resource Handling
  - Two levels of virtualizations available: LPAR and z/VM
    - Choose the level of isolation mandated by compliance
    - Flexible resource allocation and reallocation without downtime
  - CPU, memory, I/O hotplug
    - Provide the resource where they are needed in LPAR and z/VM guest
- Abundant memory, IO bandwidth and transaction capability
  - Hipersocket support connects Linux and z/OS applications and data
  - I/O fan out and transaction workload capacity is unmatched
- RAS
  - I/O device and other performance statistics
  - Dump generation, handling and inspection tools
  - Centralized and uniform resources support DR recovery setups
  - SUSE Linux Enterprise High Availability Extension included
  - System z specific kernel messages with documentation

SUSE Linux Enterprise Server 11 SP2 for System z

- z196 / z114 + zBX = IBM zEnterprise exploitation
  - CPU topology and instruction set exploitation of z196 (SDK)
  - New CHPID support connecting both environments
- Choose the right environment for the right workload
  - ISVs application support might mandate the platform
  - SLES supported for both hardware architectures
- Improved tools and z specific support
  - Disk storage & crypto enhancements
  - Linux RAS support, s390-tools update

SUSE Linux Enterprise Server 11 SP2 for System z

IBM zEnterprise System

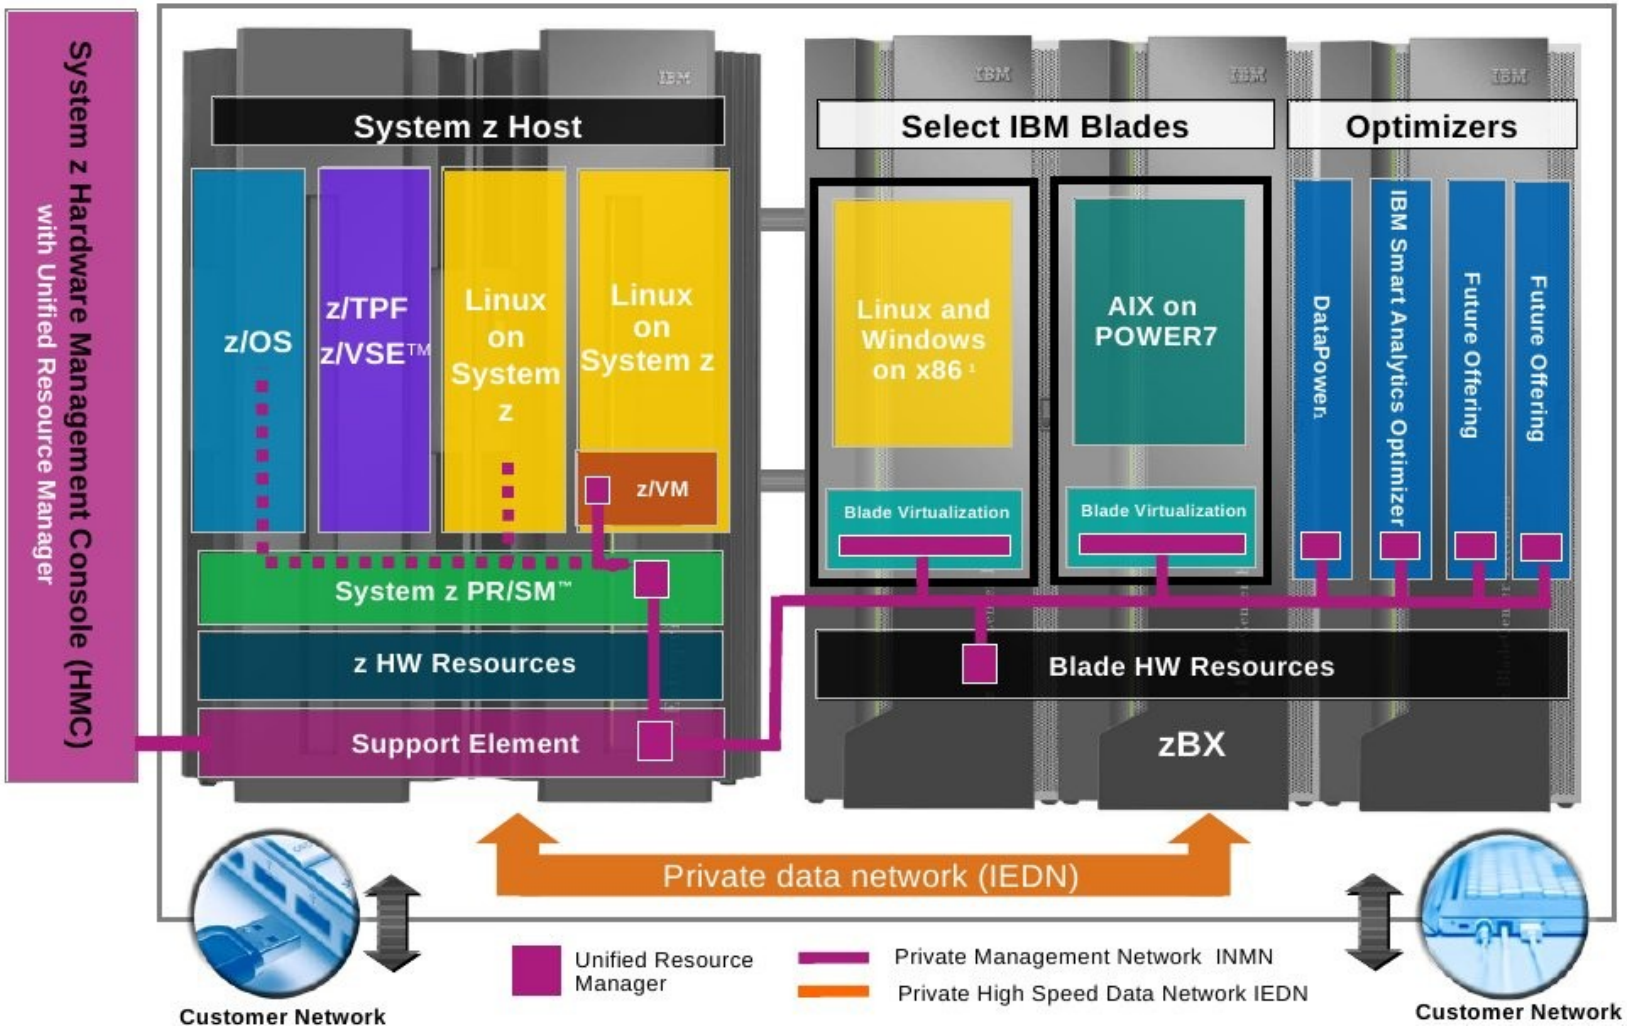

Benefits

Unique Tools Included

- Starter System for System z
  - A pre-built installation server, deployable with z/VM tools
- ***free*** High Availability Extension
  - Cluster Framework, Cluster FS, DRBD, GEO-cluster
- AppArmor Security Framework
  - Application confinement
- YaST2 systems management
  - Install, deploy, and configure every aspect of the server
- Subscription Management Tool
  - Subscription and patch management, proxy/mirroring/staging

SUSE® Linux Enterprise Server for System z

Benefits

Starter system for System z

- Pre-built installation server that can be installed on z/VM system using CMS tools
- Eliminates the need for a separate installation system elsewhere
- Minimizes the network load for network-installations
- Available for SLES 10 SP4 today,
  for SLES 11 SP2 soon

Free for every System z customer

SUSE® Linux Enterprise Server for System z

Benefits

SUSE Linux Enterprise High Availability

- Most modern and complete open source solution for implementing high available Linux clusters
- A suite of robust open source technologies that is:
  - Affordable
  - Integrated
  - Virtualization agnostic
- Used with SUSE Linux Enterprise Server, it helps to:
  - Maintain business continuity
  - Protect data integrity
  - Reduce unplanned downtime for your mission-critical Linux workloads

Part of every System z subscription

SUSE® Linux Enterprise Server for System z

Benefits

AppArmor

- Easy security on application level
- GUI tools with statics analysis and learning-based profile development
- Create custom policy in hours, not days
- Example:

```
#include <tunables/global>
/usr/sbin/vsftpd {
#include <abstractions/base>
<...>
/dev/urandom                r,
/usr/sbin/vsftpd     rmix,
/                           r,
/pub                        r,
/pub/**                     r,
@{HOMEDIRS}                 r,
@{HOME}/**                  rwl,
}
```

SUSE® Linux Enterprise 11

Systems Management Today

- YaST – unique, highly integrated local management tool
  - Ease of use, effective learning curve; reduces training efforts
  - Automation via AutoYaST datacenter mass deployments
  - WebYaST delivers web-based administration
- Fastest Open Source update stack (ZYPP)
  - Reduce management time, effort and cost
  - Improve reliability and availability by reducing downtimes
- Unattended migration from SUSE Linux Enterprise 10 to 11
- Snapshot-Rollback for package updates with
  - Btrfs – zypper – snapper
  - Rollback unwanted system changes

Subscription Management Tool

SMT is a proxy and auditing tool that mirrors the Customer Center and tightly integrates with it. It allows you to accurately register and manage an entire SUSE Linux Enterprise deployment, guaranteeing the subscription compliance and secure IT process flow organizations require.

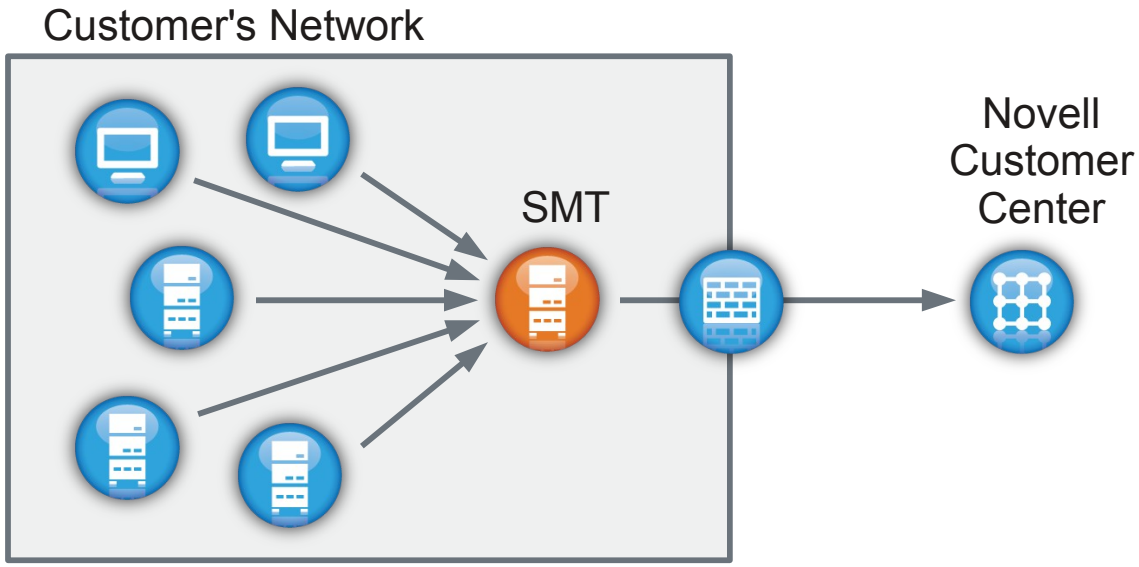

SUSE Linux Enterprise Server for IBM System z

Available today

- IBM zEnterprise System exploitation
- Enhanced tools and z support
- Choose the right environment for the right workload

Meet us at booth 324

Thank you.

Appendix

Offers

SUSE Linux Enterprise Server
for IBM zEnterprise and zBX

With a SUSE Linux Enterprise Server for System z subscription you get full subscription entitlements for your x86_64 zBX blades.

If you own this...

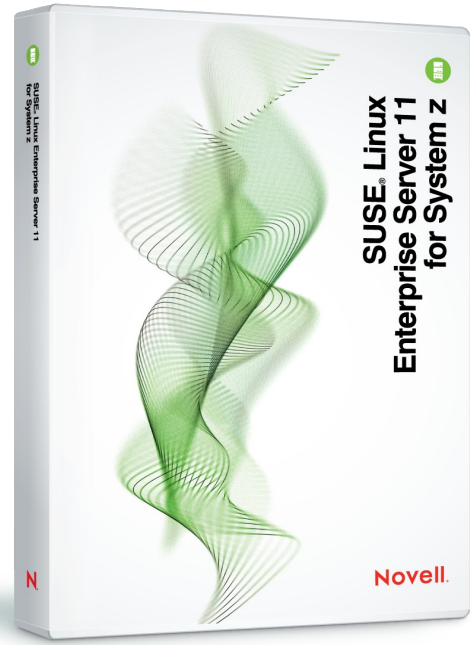

… you get these !

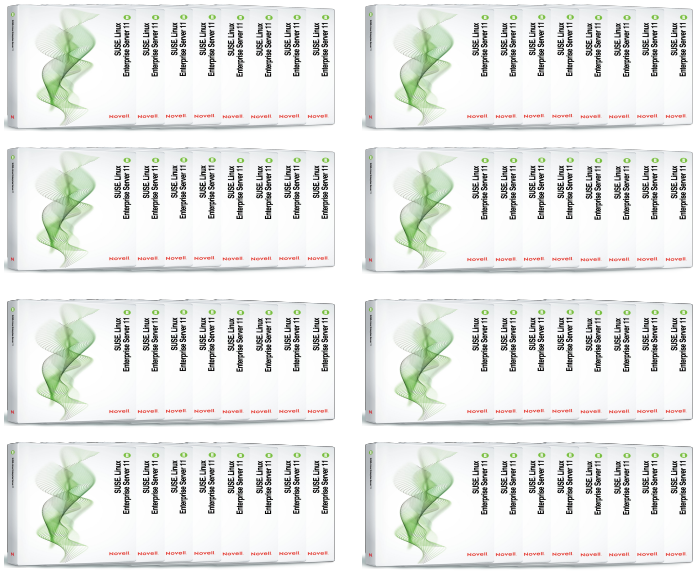

SUSE offerings

for SUSE Linux Enterprise Server for System z

- Attractive subscriptions & services are available: talk to us !
  - **Multi year, multiple IFL subscriptions offer monetary value – scale & save**
    - For larger volumes, you get a scaled price
    - *http://www.suse.com/products/systemz/how-to-buy/*
  - **SUSE Linux Enterprise High Availability Extension included – prevent outages**
    - Save $700+ per IFL compared to distributed systems
    - *http://www.suse.com/products/systemz/features/linux_server_clustering.html*
  - **zBX offering – 1 IFL, multiple zBX subscriptions – boost consolidation**
    - Save up to $39k running your ensemble with SUSE Linux Enterprise Server
    - *http://www.suse.com/promo/zbx.html*
  - **Long Term Support Service available – preserve environments**
    - *http://www.suse.com/support/programs/long-term-service-pack-support.html*

SUSE Linux Enterprise Server 11 SP2 s390x specific features

z196 enhanced node affinity support

Fate 311860 / [LTC 66807]

http://www.ibm.com/developerworks/linux/linux390/ -> affinity

Device Drivers, Features, and Commands (Kernel 2.6.38) – Chapter 27, p.292

- **CPU node affinity support for z196:** allowing the Linux kernel scheduler to optimize its decisions based on the z196 processor, cache and book topology (in LPAR, enable with 'topology=on' kernel parameter).

- **Customer benefit**

technical	business
• Hardware exploitation z196 processor topology and cache hierachy, increase cache hit ratio and therefore overall performance	• Increase of application workload density per system • Increased performance

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	yes
SP4	-	n/a

OSX (OSM) chpids for hybrid data (management) network

Fate 311898 / [LTC 66966]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.35.html → OSX

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.10.0.html → znetconf

- **z196 and zBX exploitation:** enhancement in the network device configuration tool znetconf (s390-tools) by updating internal tables to handle OSX and OSM CHPIDs.

- **Customer benefit**

technical	business
• Hardware exploitation of the z196 and zBX for hybrid computing	• Fit-to-purpose workload placement support

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Performance Indicator Bytes

Fate 311854 / [LTC 66802]

- **Serviceability:** two new fields in /proc/sysinfo (Capacity-Change Reason (CCR) and Capacity-Adjustment Indication (CAI) can help to ease problem determination eg. If machine performs slower than expected.

- **Customer benefit**

technical	business
• Faster problem determination	• Reduce time spend on issues

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Spinning mutex performance enhancement

Fate 312075 / [LTC 70029]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.38.html → spinning mutex

- **Performance:** The status of a thread owning a locked mutex is examined and waiting threads are not scheduled unless the first is scheduled on a virtual *and* physical processor.

- **Customer benefit**

technical	business
• New sophiticated handling of mutexes and scheduler decisions to improve performance also for z/VM based workloads	• Performances benefits for workloads making usage of parallel processing in an SMP environment of virtual CPUs

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

z196 exploitation via alternate GCC 4.6

Fate 311859 / [LTC 66797]

http://www.ibm.com/developerworks/linux/linux390/ -> toolchain (pending)

http://gcc.gnu.org/gcc-4.6/changes.html -> z196

- **Performance improvement for applications:** exploitation of new z196 processor instructions and optimized alignment of code (out-of-order pipeline architecture, conditional load/store instructions, new 3 register operand instructions, new atomic instructions, etc)

- **Customer benefit**

technical	business
• Hardware exploitation of the z196 instruction set for user land applications (ISV and self compliled applications), recompile programs with --march=z196 and/or –mtune=z196	• z196 optimized code and efficient excecution use less time and cycles for same workload • Increase of application workload density per system

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	yes
SP4	-	n/a

Exploitation of z10 prefetching instructions in GCC

Fate 311845 / [LTC 66745]

http://www.ibm.com/developerworks/linux/linux390/ -> gcc 4.5.1 (upstream)
http://gcc.gnu.org/gcc-4.5/

- **Toolchain based performance improvement for applications:** z10 introduced pre-fetching instructions to enhance memory access like copying memory, zeroing out memory and exploiting predictable loops by help of the compiler.

- **Customer benefit**

technical	business
• Hardware exploitation of the z10 and z196 instruction set for user land applications (ISV and self compliled applications)	• z10 and z196 optimized code and efficient excecution use less time and cycles for same workload • Increase of application workload density per system

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	yes
SP4	-	n/a

cio: handle channel path description changes

Fate 311913 / 311911 / [LTC 69631]

http://www.ibm.com/developerworks/linux/linux390/kernel-3.0.html -> dynamic IODF

- **Dynamic resource allocation:** the common I/O layer handles dynamic IODF changes that result in changed capabilities of channel paths. Applies for LPAR installations only, since the required channel subsystem notifications are not supported on current z/VM versions.

- **Customer benefit**

technical	business
• More flexible I/O configuration for Linux running in LPARs	• Non disruptive change of IO configuration

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

cio resume handling for reordered devices

Fate 311876 / [LTC 66907]

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.368

- **Usability:** Improves cio resume handling to cope with devices that were attached on different subchannels prior to the suspend operation.

- **Customer benefit**

technical	business
• If the subchannel changes for disk device, the configuration is changed to reflect the new subchannel. This change is accomplished without de-registration. Device name and device configuration are preserved.	• Optimized or no downtime when resuming a Linux instance

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

User space handle to wait for cio pending work

Fate 311763 / 311850 / [LTC 66753]

- **Reliability:** With this feature, user space processes can delay I/O operations until all pending requests against the common I/O layer have been completed.

- **Customer benefit**

technical	business
• Enables user space processes to wait for completion of any pending actions affecting device availability..	• Increased reliability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

cmsfs support for kernel 2.6

Fate 311847 / 311858 / [LTC 60032] / [LTC 66799]

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.441

- **s390-tools:** read and write configuration files stored on CMS disks directly from Linux. CMS disk can be mounted so the files on the disk can be accessed by common Linux tools

- **Customer benefit**

technical	business
• cmsfs-fuse tool translates the record-based EDF file system on the CMS disk to UNIX semantic • Text files can be automatically converted from EBCDIC to ASCII	• Access data (config, files, dumps,...) from the z/VM CMS filesystem, from Linux during operation (no shutdown required)

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Enhanced Yast to support SCSI tape devices

Fate 311718 / [LTC 70284]

- **Usability:** Support for SCSI tapes in yast for interactive handling. Udev rules are adjusted to make changes persistent in case of reboot

- **Customer benefit**

technical	business
• Interactive handling	• Usability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

FICON DS8k support – solid state disk flag

Fate 311756 / [LTC 60095]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.2.html → solid

- **Solid State Drive support:** transparent to the DASD device driver, no change is needed to use solid state disks. A new flag in the device characteristics will show the administrator if a device is a solid state disk.

- **Customer benefit**

technical	business
• Storage servers can be queried if they provide solid state disks • Device characteristics are already exported per ioctl and can be read as binary data with the dasdview tool.	• Workloads can be placed on storage which support best their I/O characteristics • Acceleration of random I/O, cost effective placement of sequential I/O

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

FICON Dynamic PAV toleration

Fate 311760 / [LTC 66751]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.35.html → dynamic PAV

- **Dynamic PAV:** the DASD device driver tolerates dynamic Parallel Access Volume (PAV) changes for base PAV. PAV changes in the hardware configuration are detected and the mapping of base and alias devices in Linux is adjusted accordingly

- **Customer benefit**

technical	business
• If the mapping of an alias to a base device is changed another device the DASD driver will tolerate this change • Change in the base/alias mapping is automatically discovered by the DASD device driver	• Improve the flexibility and availability of SLES for System z, by allowing to tolerate changes in the PAV infrastructure without need to restart the system

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

FICON Multi-Track extensions for High Performance

Fate 311870 / [LTC 66846]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.38.html → multitrack

- **Hardware exploitation:** exploit DS8000 storage systems support for multi-track High Performance FICON requests (read or write data to more than one track).

- **Customer benefit**

technical	business
• Provides a new cio layer function using an interface to get the maximum usable data size for zHPF requests on a given subchannel	• Maximize I/O performance with FICON, zHPF and DS8000 storage servers

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

FICON IPL & device discovery hardening

Fate 311848 / 311759 / [LTC 66748]

- **Description:** Improves the DASD error recovery procedures used in the early phases of IPL and DASD device initialization with additional error recovery procedures.

- **Customer benefit**

technical	business
• More robust and reliable recovery in early phases of IPL	• Reduced downtime in case of an error

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

FICON API & Tool to query DASD reservation status

Fate 311905 / 311910 / [LTC 67827]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.37.html → reservation

- **Description:** allows the DASD device driver to determine the reservation status of a given DASD in relation to the current Linux instance

- **Customer benefit**

technical	business
• More detailed information on reservation status • Use tunedasd tool in s390-tools 1.12.0 (or later) to displays the DASD reservation status.	• Optimized handling

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

FICON Improved handling of stolen DASD reservation

Fate 311906 / [LTC 67830]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.38.html → stolen

- **Description:** Allows to specify a policy for the DASD device driver behavior in case of a lost device reservation. The policy can be specified via a new DASD sysfs attribute reservation_policy.

- **Customer benefit**

technical	business
• More options to handle device assignments for ISVs	• N/A

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Tunable grace period for missing interrupts (DASD driver)

Fate 311880 / [LTC 66911]

- **Reliability:** The default value for this timeout was 300 seconds for standard ECKD and FBA I/O operations and 50 seconds for DIAG I/O operations. This feature provides a sysfs interface to allow users to change the timeout value.

- **Customer benefit**

technical	business
• Allow HA solutions to react earlier on stalled I/O connectivity.	• Shorter reaction time on issue, thus faster resolution – reduced pot. downltime

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Access to raw ECKD data from Linux (DASD)

Fate 311973 / [LTC 66951]

- **Interoperabilty:** allows to access ECKD disks in raw mode. Linux dd command can copy the disk level content of an ECKD disk to a Linux file, and vice versa. Works independent of the operating system or file system that is on the ECKD disk.

- **Customer benefit**

technical	business
• Includes Linux ECKD disks used with LVM, Linux ECKD disks that are used directly, and z/OS ECKD disks	• Use case for Linux by eliminating the need for data transfers from z/OS to Linux via network

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Store I/O Operation Status and initiate logging (SIOSL)

Fate 311917 / [LTC 66847]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.36.html

- **Description:** interface for the store-I/O-operation-status-and-initiate-logging (SIOSL) CHSC command and its exploitation by the FCP device driver

- **Customer benefit**

technical	business
• Can be used to synchronize log gathering between the operating system and the channel firmware.	• Concurrent data collection for problem resolution, minimizing customer operation impact

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Automatic detection of read only DASDs

Fate 311878 / [LTC 66909]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.34.html

- **RAS:** prevents unintentional write requests and subsequent I/O errors, by detecting if a z/VM attached device is read-only using the z/VM DIAG 210 interface and setting the respective Linux block device to read-only as well.

- **Customer benefit**

technical	business
• Extends the DASD device driver to automatically detect read-only devices and set the Linux read-only flag respectively	• Reduce chance for unintentional data loss

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Unit Check handling

Fate 311874 / [LTC 66908]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.35.html

- **Description:** improves handling of unit checks reported during CIO-internal operations. Control units such as the DS8000 storage server are using Unit Checks as a means to inform Linux of events which may affect the operational state of the devices provided.

- **Customer benefit**

technical	business
• Better analysis possible by the system on how to react to external I/O storage events.	• Improved RAS capabilities of System z HW

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

DASD Tools - implement new partition types

Fate 311921 / [LTC 69643]

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.467

- **s390-tools:** This feature introduces new partition types like RAID and LVM. Today fdasd only supports "Linux native" and "swap"

- **Customer benefit**

technical	business
• Enhanced usability	• N/A

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Automated LUN scanning for NPIV only

Fate 311953 / [LTC 75664]

- **Description:** Linux SCSI midlayer to scan and automatically attach SCSI devices that are available for the NPIV WWPN. The manual configuration of LUNs in zfcp is now only required for non-NPIV FCP subchannels.

- **Customer benefit**

technical	business
• Automatted access to SCSI devices • NOTE: use with case in an environment weer different OS share one FCP channel.	• important for server consolidation

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

End-To-End data consistency checking [zfcp]

Fate 311915 / [LTC 69627]

- **RAS:** introduces the zfcp-specific part in the Linux on System z I/O stack for E2E data consistency checking. (T10DIF)

- **Customer benefit**

technical	business
• Improved service and control of data flow between adapter and storage device by introducing the zFCP specific part of the enhanced SCSI standard for E2E data consistency checking	• Increased RAS: End (storage) to end (application) data protection

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Support for 4096 bit RSA FastPath (zcrypt)

Fate 311901 / [LTC 66955]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.38.html → RSA

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.282

- **Description:** extends the support for current hardware acceleration of RSA encryption and decryption from 2048-bit keys to the new maximum of 4096-bit keys in zcrypt Linux device driver

- **Customer benefit**

technical	business
• Crypto Express3 hardware exploitation and stronger encryption	• More secure network connections due to possible use of stronger encryption

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Crypto CP ACF exploitation

Fate 311914 / 311924 311091 / [LTC 69628 etc]

http://opencryptoki.git.sourceforge.net/ -> libica

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.38.html → zcrypt

- **Cryptography:** hardware based acceleration of complex cryptographic algorithms, support for 4096 bit RSA FastPath (support zEnterprise Crypto Express3 card RSA mod expo operations with 4096-bit RSA keys in ME (Modulus Exponent) and CRT (Chinese Remainder Theorem) format)
- **Customer benefit**

technical	business
• Cryptographically secured connections • devices maximum request size is adjusted based on a test	• Improve security for data transfers over the network • reduce the cost of SSL acceleration replacing expensive and MIPS intensive mathematical calculations in software

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Switchable default compression on/off in OpenSSL

Fate 312076 / [LTC 70054]

- **Choice:** OpenSSL usually compresses data before encryption with a performance impact (throughput decrease, CPU load increase) on platforms with cryptographic hardware. The setting is made tunable via a /etc/ssl/openssl.cnf parameter.

- **Customer benefit**

technical	business
• Determine behaviour via parameter and depending on workload	• Improves throughput on encrypted connections if compression is turned of and reduces CPU usage at the same time

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Get CPC name (xDR)

Fate 311920 / [LTC 69632]

- **RAS/DR:** enables for dynamic changes in the GDPS environment definition to avoid possible failures from manual or forgotten changes. GDPS now changed to retrieve CPC and LPAR information dynamically.

- **Customer benefit**

technical	business
• Avoid resetting a LPAR due to incorrect configuration definitions • With the new function, GDPS always resets exactly the LPAR in which the OS is running.	• Additional protection to avoid handling errros in HA/DR scenario

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

snIPL support for z/VM 6

Fate 311909 / 311904 / [LTC 69380]

http://www.ibm.com/developerworks/linux/linux390/snipl-0220.html

- **Description:** This feature offers socket-based (AF_INET) remote system management of z/VM 6 guests with snipl and stonith if SMAPI support is available.

- **Customer benefit**

technical	business
• snIPL can now access the z/VM systems management API through a SMAPI request server. .	• Contiued support for latest z/VM versions preserve or extends existing APIs

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

snIPL enhanced to trigger SCSI dump on remote container

Fate 311758 / [LTC 63604]

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.383

- **Description:** This feature enhances snIPL to take a remote SCSI dump using the snIPL interface.

- **Customer benefit**

technical	business
• Remote IPL and dump to a zfcp scsi device	• More flexible IPL and dump targets • Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

zipl to automatically calculate boot device ramdisk address (s390-tools)

Fate 311877 / [LTC 69650]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.4.html → ramdisk

- **Description:** relax the need for a default address for the initial ramdisk on the boot device. The address is now calculated dependent on the locations of the other components.

- **Customer benefit**

technical	business
• If the user provides an initrd_addr then this one is used, if not then - instead of a fixed value (0x800000) - a suitable calculated value is used.	• N/A

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

zipl automatic menu support (s390-tools)

Fate 311908 / [LTC 69625]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.11.0.html → automatic

- **Description:** adds support for automatic menu generation to zipl package. That is to provide a user interface that keeps the default behavior of zipl where possible and minimizes changes for the user.

- **Customer benefit**

technical	business
• Consitent behaviour of zipl	• Efficient administration

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

reipl tool chreipl enhancements

Fate 311861 / [LTC 66811]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.12.0.html → chreipl

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.430

- **Description:** add support to re-IPL from device-mapper devices, including mirror devices and multipath devices, add support to re-IPL from named saved systems (NSS), add support to specify additional kernel parameters for the next re-IPL, add "auto target" support

- **Customer benefit**

technical	business
• More reipl configuration options	• Allow more choice for business needs and policies

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Intuitive dump device configuration

Fate 304024 / [LTC 201624]

- **Description:** provide a yast dialog to prepare I/O devices for dump, during the installation and post-installation.

- **Customer benefit**

technical	business
• Guided setup and configuration of a suitable dump device	• Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Dump on panic - Prevent reipl loop

Fate 311757 / [LTC 63603]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.4.html → re-IPL loop

- **Description:** configuration parameter DELAY_MINUTES is introduced in the dumpconf configuration file. Using this keyword the activation of dumpconf can be delayed in order to prevent potential re-IPL loops.

- **Customer benefit**

technical	business
• Tunable dump behaviour in case of a dump loop	• Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Efficient dump file analysis with makedumpfile

Fate 311641 / [LTC 65511]

http://www.ibm.com/developerworks/linux/linux390/toolchain.html#binutils2.20.51.0.7

- **Description:** makedumpfile port to Linux on System z and infrastructure to convert the Linux on System z dumps to ELF core format.

- **Customer benefit**

technical	business
• Compressed dumps • Convert s390 format dumps into ELF format dumps • Suitable dumps even if memory footprint increases over time	• Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Crash Utility support to read compressed/filtered dumpfile
Fate 311654 / [LTC 67487]

- **Description:** read and analyse compressed and filtered dumpfile generated by makedumpfile tool for s390. Related and dependent on makedumpfile port.

- **Customer benefit**

technical	business
• Reduced size of dumps without loss of essential information	• Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

OSA concurrent software/hardware trap

Fate 312954 / [LTC 73705]

- **Description:** enable collective problem analysis through consolidated dumps of software and hardware.

- **Customer benefit**

technical	business
• Command line interface can be used to generate qeth/qdio trace data as well as trigger the internal dump of an OSA device.	• Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Extend and Improve zFCP trace utilities

Fate 311918 / 311919 / 311928 / [LTC 69664/69665/69663]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.13.0.html → zfcpdbf

How to use FC-attached SCSI devices with Linux on System z → lk32ts03.pdf

- **Description:** new tracing concept of the zfcp driver. It moves the dependency of the data analysis from the driver to an external utility.

- **Customer benefit**

technical	business
• Extended lgging information • Improved data interpretation utilities • Simplyfied maintainablility of logging data	• Improved performance analytics

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Breaking-event-address for user space programs

Fate 311851 / [LTC 69650]

- **Description:** A breaking-event-address is the last instruction to cause a break in the sequential execution of a CPU. The recorded addresses are valuable debugging data and can be read with the ptrace interface.

- **Customer benefit**

technical	business
• Records breaking-event-addresses for user space processes using the PER-3 facility introduced with z10 • Any breaking-event in the range from 0 to 8MB will not be recorded	• Useful for application development

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Remove support for multi-volume tape dumps (s390-tools)

Fate 311895 / [LTC 66989]

- **Description:** The multi-volume tape dump support will be removed from zipl and zgetdump. The reason for this decision is that current tape cartridges have hundreds of gigabyte capacity and therefore the multi-volume support is not needed any more

- **Customer benefit**

technical	business
• s.a.	• N/A

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Virtualtop (aka hyptop)

Fate 311766 & 311764 / [LTC 66756 / 66775]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.12.0.html ->hyptop

- **LPAR and z/VM hypervisor monitoring:** provides a dynamic real-time view of a System z hypervisor environment. Depending on the available data it shows eg. CPU and memory consumption of active LPARs or z/VM guests.

- **Customer benefit**

technical	business
• Provides a curses based user interface similar to the popular Linux "top" command. • Using kernel infrastructure via debugfs	• Real time workload monitoring using Linux tools • Improve problem resolution time

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Optimal qeth default settings

Fate 311917 / [LTC 69650]

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.108

- **Performance:** optimal default settings for several qeth parameters to avoid that customers work with semi-optimal network setup.

- **Customer benefit**

technical	business
• Improves performance by providing optimal network parameter for qeth • Improve RAS by avoiding manual system tuning except for special setups.	• Optimized network related performance

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Improved QDIO performance statistics

Fate 311853 / [LTC 66790]

http://www.kernel.org/pub/linux/kernel/v2.6/ChangeLog-2.6.33 -> qdio performance

- **Device statistics:** Improves serviceability in situations where the customer needs a detailed performance analysis of network (qeth) or SCSI (zfcp) usage.

- **Customer benefit**

technical	business
• Converts global statistics to per-device statistics and adds new counter for the input queue fill condition	• More detailed information on device I/O characteristics can facilitate faster problem resolution times → improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Offload outbound checksumming to OSA (qeth driver)

Fate 311871 / [LTC 66956]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.35.html → Checksumming Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.125

- **Description:** introduces OSA adapter hardware support for the checksum calculations which TCP and UDP use to ensure data integrity.

- **Customer benefit**

technical	business
• Offloading this calculation will reduce the processor load compared to software implementation (automatic fallback) • reduces TCP/IP path-length	• Saved processor cycles can be reused for other workloads • More efficient use of the available hardware

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Deliver z/VM CP special messages with udev events

Fate 311855 / [LTC 66956]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.34.html → smsgiucv_app

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.249

- **Description:** a new kernel device driver for receiving z/VM CP special messages (SMSG) and delivering these messages to user space as udev events (uevents).

- **Customer benefit**

technical	business
• Deliver z/VM CP messages to a specific guest or application (filtering messages starting with "APP")	• Enhances the systems management capabilities for virtualized environments

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Improve memory ballooning with cpuplugd (s390-tools)

Fate 312069 / [LTC 69630]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.15.0.html → cpuplugd

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.446

- **Description:** improved defaults to deal with limited memory resources in virtualized environments

- **Customer benefit**

technical	business
• Rule based scheme to control the size of the CMM1 memory balloon	• More efficient use of memory resources

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Tool to safely start getty through init

Fate 311894 / [LTC 66991]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.11.0.html → ttyrun

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.309

- **Description:** prevent failing re-spawns during initial startup. The tool ttyrun should be integrated per default for all possible terminals supported in a Linux on System z environment in inittab.

- **Customer benefit**

technical	business
• More robust startup	• Improved system reliabilty

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Disable cpuplugd by default

Fate 312295 / [LTC 71776]

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.446

- **Description:** the latest Linux scheduler is optimized to achieve simular results scheduling processes without the cost intensive operation of CPU plug and unplug. If the use case is not fully exploited, it is advisable to disable the cpuplugd by default.

- **Customer benefit**

technical	business
• Simplify the setup of a Linux image	• N/A

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

IPv6 support for qetharp tool

Fate 311912 / [LTC 69639]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.13.0.html → qetharp

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.523

- **Description:** adds IPv6 support to the qetharp tool

- **Customer benefit**

technical	business
• Inspection and modification of the ARP cache of OSA cards or HiperSockets (real and virtual) operated in layer 3 mode	• IPv6 support

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Support assisted VLAN null tagging support

Fate 311899 / [LTC 66983]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.37.html → VLAN Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.139

- **Description:** the qeth device driver for OSA-Express (QDIO) and HiperSockets can now handle tagged frames with VLAN ID 0.

- **Customer benefit**

technical	business
• Enables qeth to receive priority-tagged frames on a base interface without configuring an additional VLAN interface.	• IBM z/OS and Linux on System z

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

NAPI support for qeth and qdio

Fate 311890 / [LTC 66967]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.37.html → NAPI

- **Description:** the qdio interface is extended to allow direct processing of inbound data in qeth.

- **Customer benefit**

technical	business
• Using NAPI, the device driver can disable interrupts to reduce CPU load under high network traffic	• Increased throughput and less CPU consumption allow for more efficient use

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Enhancement in the config tool for z network devices

Fate 312067 / [LTC 66958]

http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.4.html → qethconf

Device Drivers, Features, Commands on SUSE Linux Enterprise Server 11 SP2, p.526

- **Description:** enhances the qethconf tool by providing improved information messages

- **Customer benefit**

technical	business
• More detailed informational messages in unsuccessful actions	• Improved serviceability

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Optimized Latency Mode (OLM) toleration (qeth)

Fate 311892 / [LTC 66982]

http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.35.html → OLM

- **Description:** enhances the qeth driver with a meaningful message in case an OSA-connection fails due to an active OLM-connection on the shared OSA-adapter

- **Customer benefit**

technical	business
• OLM may be activated by z/OS on an OSA Express3 adapter, which reduces the number of allowed concurrent connections, if adapter is used in shared mode	• Serviceability when sharing the system with Linux and z/OS

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Fill entropy with hwrandom for z10

Fate 310591 / [LTC -]

- **Description:** z10 processor and successors have a random number generator built in, that can be accessed at /dev/hwrng if active. However, with z90crypt device driver and crypto express cards /dev/random delivers hardware generated random numbers at high rate.

- **Customer benefit**

technical	business
• Use /dev/random as a source of random numbers generated by hardware at a high rate • Avoids stalling of processes querying for randomness	• Better scalability for workloads with lots of processes requiring randomness to execute or proceed

SLES	10	11
GA	-	-
SP1	-	-
SP2+3	-	**yes**
SP4	-	n/a

Specs

SUSE® Linux Enterprise 11 SP2

Kernel Capabilities

SLE 11 SP 2 (3.x)	x86	ia64	x86_64	s390x	ppc64
CPU bits	32	64	64	64	64
max. # logical CPUs	32	up to 4096	up to 4096	64	up to 1024
max. RAM (theoretical/practical)	64/ 16 GiB	1 PiB/ 8+ TiB	64 TiB/ 16TiB	4 TiB/ 256 GiB	1 PiB/ 512 GiB
max. user-/ kernelspace	3/1 GiB	2 EiB/φ	128 TiB/ 128 TiB	φ/φ	2 TiB/ 2 EiB
max. swap space	up to 31 * 64 GB				
max. #processes	1048576				
max. #threads per process	tested with more than 120000; maximum limit depends on memory and other parameters				
max. size per block device	up to 16 TiB	and up to 8 EiB on all 64-bit architectures			

Supported on certified hardware only

SUSE® Linux Enterprise 11 SP2

Filesystems

Feature	Ext 3	reiserfs	XFS	OCFS 2	btrfs
Data/Metadata Journaling	•/•	○/•	○/•	○/•	N/A [3]
Journal internal/external	•/•	•/•	•/•	•/○	N/A
Offline extend/shrink	•/•	•/•	○/○	•/○	•/•
Online extend/shrink	•/○	•/○	•/○	•/○	•/•
Inode-Allocation-Map	table	u. B*-tree	B+-tree	table	B-tree
Sparse Files	•	•	•	•	•
Tail Packing	○	•	○	○	•
Defrag	○	○	•	○	•
ExtAttr / ACLs	•/•	•/•	•/•	•/•	•/•
Quotas	•	•	•	•	○
Dump/Restore	•	○	•	○	○
Blocksize default	4KiB				
max. Filesystemsize [1]	16 TiB	16 TiB	8 EiB	4 PiB	16 EiB
max. Filesize [1]	2 TiB	1 EiB	8 EiB	4 PiB	16 EiB
Support Status	SLES	SLES	SLES	SLE HA	SLES

SUSE® Linux Enterprise was the first enterprise Linux distribution to support journaling filesystems and logical volume managers back in 2000. Today, we have customers running XFS and ReiserFS with more than 8TiB in one filesystem, and the SUSE Linux Enterprise engineering team is using our 3 major Linux journaling filesystems for all their servers. We are excited to add the OCFS2 cluster filesystem to the range of supported filesystems in SUSE Linux Enterprise. For large-scale filesystems, for example for file serving (e.g., with with Samba, NFS, etc.), we recommend using XFS. (In this table "+" means "available/supported"; "-" is "unsupported")

[1] The maximum file size above can be larger than the filesystem's actual size due to usage of sparse blocks. It should also be noted that unless a filesystem comes with large file support (LFS), the maximum file size on a 32-bit system is 2 GB (2^{31} bytes). Currently all of our standard filesystems (including ext3 and ReiserFS) have LFS, which gives a maximum file size of 2^{63} bytes in theory. The numbers given in the above tables assume that the filesystems are using 4 KiB block size. When using different block sizes, the results are different, but 4 KiB reflects the most common standard.

[2] 1024 Bytes = 1 KiB; 1024 KiB = 1 MiB; 1024 MiB = 1 GiB; 1024 GiB = 1 TiB; 1024 TiB = 1 PiB; 1024 PiB = 1 EiB (see also http://physics.nist.gov/cuu/Units/binary.html)

[3] Btrfs is a copy-on-write logging-style file system, so rather than needing to journal changes before writing them in-place, it writes them in a new location, and then links it in. Until the last write, the new changes are not “committed.”

[4] Btrfs quotas will operate differently than traditional quotas. The quotas will be per-subvolume rather than operating on the entire filesystem at the user/group level. They can be made functionally equivalent by creating a subvolume per- user or group.

Resources

SUSE® Linux Enterprise

Documentation and Release Notes

- Product Pages
  - http://www.suse.com/products/server/
  - http://www.suse.com/products/sles-for-sap/
  - http://www.suse.com/products/highavailability/
  - http://www.suse.com/products/realtime/
- Unix to Linux Migration
  - http://www.suse.com/solutions/enterprise-linux-servers/unixtolinux.html
- Documentation
  - http://www.suse.com/documentation/
- Release Notes
  - http://www.suse.com/releasenotes/

Resources

- Product website
  www.suse.com/products/systemz
- Customer References
  www.suse.com/success → extended search for SUSE Linux Enterprise Server for System z
- Download SUSE Linux Enterprise Server for System z
  www.suse.com/products/server/eval.html
- Promotion Website
  www.novell.com/products/systemz/els.html
- Partner Website
  www.suse.com/mainframe
- Starter System for System z
  www.suse.com/partner/ibm/mainframe/startersystem.html

Resources

- SUSE Linux Enterprise Server and IBM zEnterprise
  http://www.novell.com/docrep/2010/11/suse_linux_enterprise_server_and_ibm_zenterprise_system.pdf
- zBX entitlement for SUSE Linux Enterprise Server offering
  http://www.suse.com/promo/zbx.html
- SUSE Linux Enterprise Server for System z
  http://www.suse.com/products/systemz/
- IBM zEnterprise Success Story: Sparda-Datenverarbeitung eG
  http://www.novell.com/success/sparda.html
- Chalk Talk: Server consolidation on IBM System z
  http://www.novell.com/media/content/chalktalk-server-consolidation-on-system-z.html
- SUSE Manager
  http://www.suse.com/products/suse-manager
- SUSE Studio
  http://www.susestudio.com

Corporate Headquarters
Maxfeldstrasse 5
90409 Nuremberg
Germany

+49 911 740 53 0 (Worldwide)
www.suse.com

Join us on:
www.opensuse.org