

Update That SUSE® Linux Enterprise Server VM

Don Vosburg
Systems Engineer
DVosburg@suse.com




Agenda

- SUSE® Linux Enterprise Server for System z update
- Upgrading and updating your SUSE VM

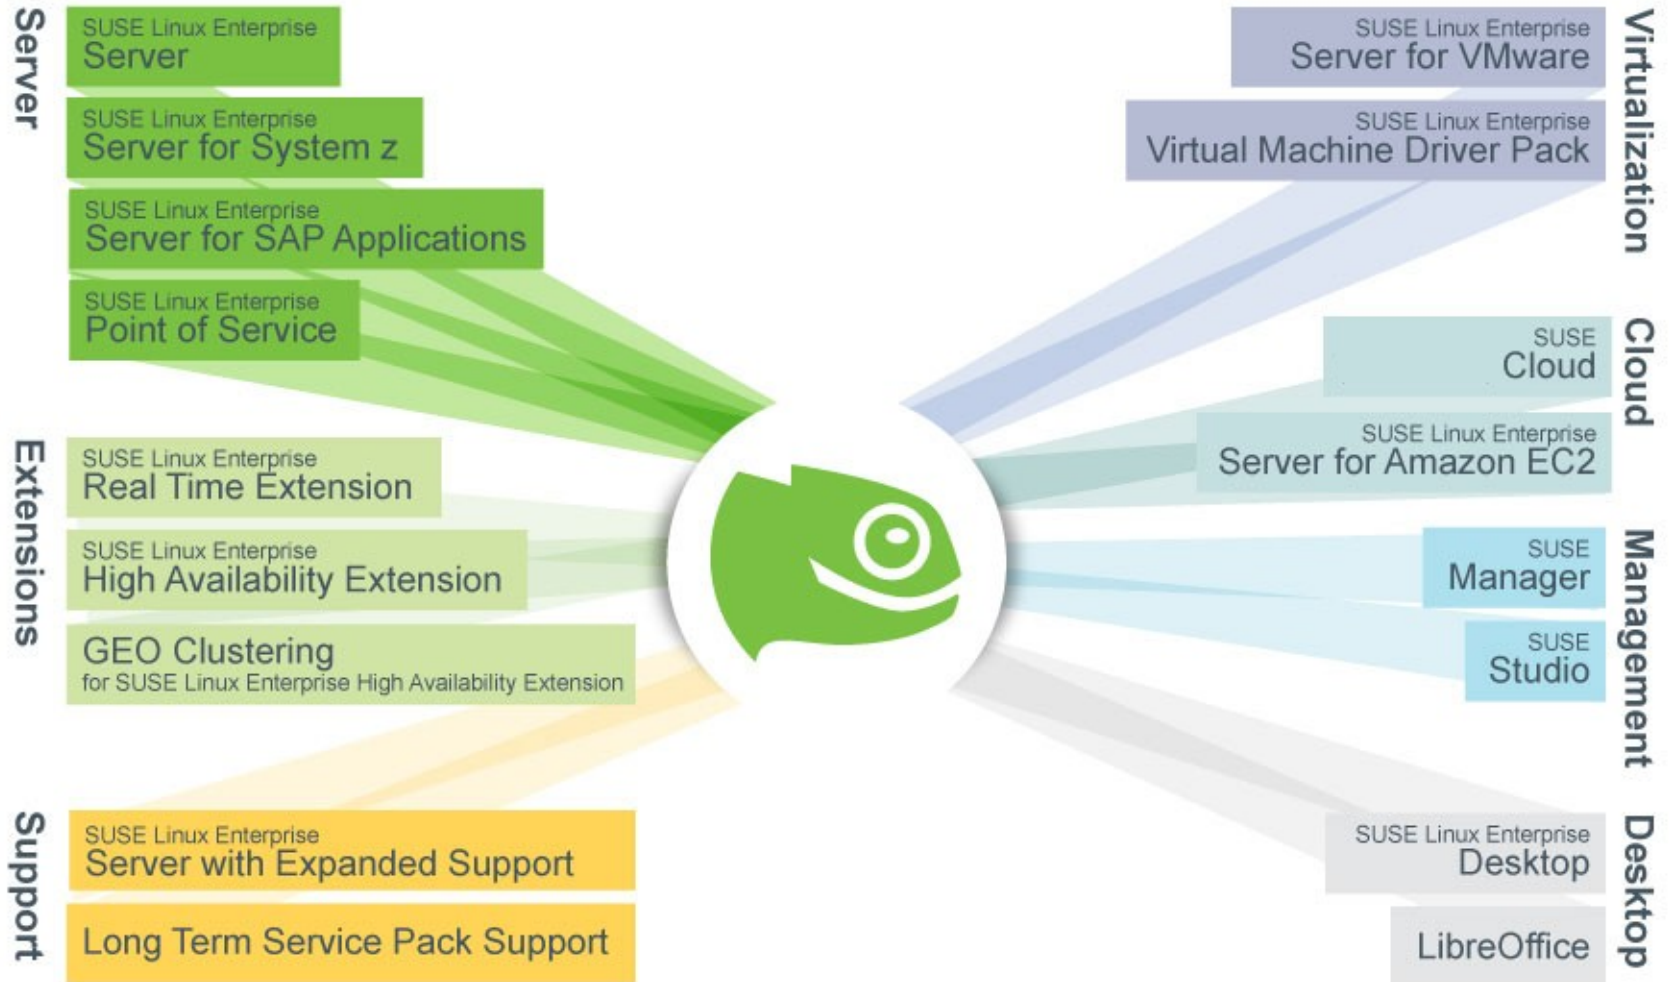


SUSE® Linux Enterprise Server

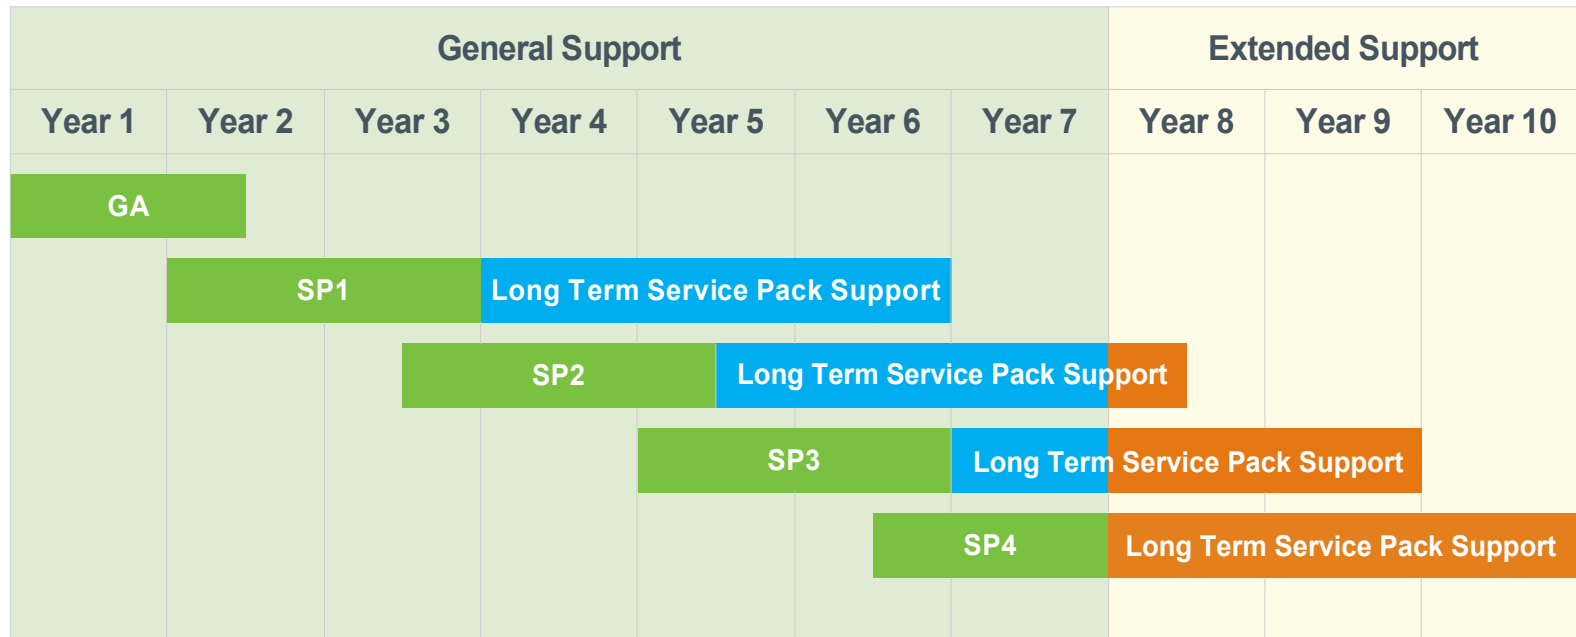
- SUSE Linux Enterprise Server 10/2000
- SUSE Linux Enterprise Server 7 08/2001
- SUSE Linux Enterprise Server 8 10/2002
- SUSE Linux Enterprise Server 9 08/2004
- SUSE Linux Enterprise Server 10 07/2006
- SUSE Linux Enterprise Server 11 03/2009
- SUSE Linux Enterprise Server 12 ~2014



Comprehensive Portfolio




Standard 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



Current SUSE® Linux Enterprise Streams



- Dependable release timing
- Predictability for planning rollouts and migrations
 - Service Pack releases, development and product schedules announced to customers and partners
- Major releases every 4-5 years.

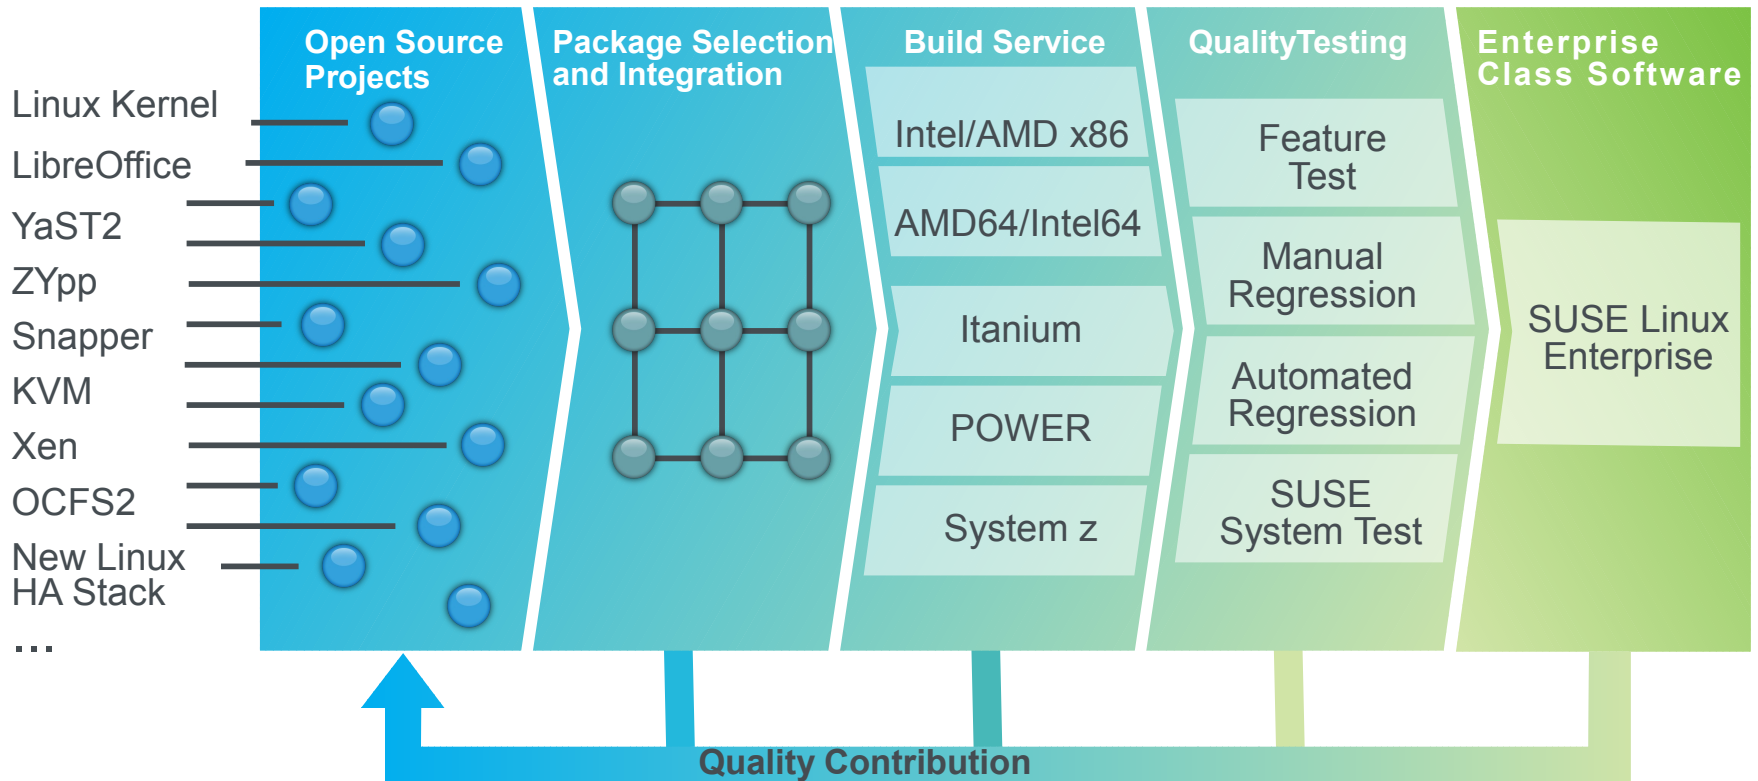


SUSE® Linux Enterprise

The SUSE® Build Service* Advantage

Development
Contribution

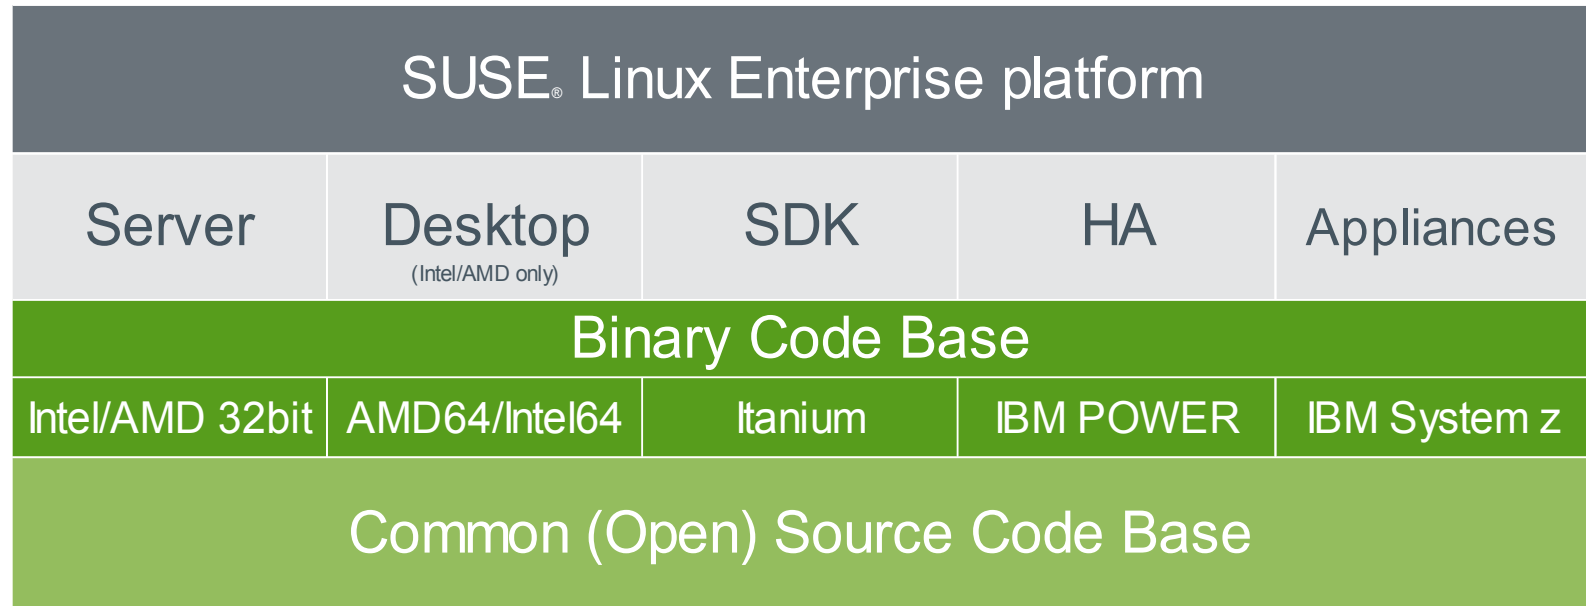
Infrastructure
Contribution



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

- Reduces production problems
- Consolidates IT skills across disparate systems
- Delivers critical updates in hours – not days or weeks

Common Code Base

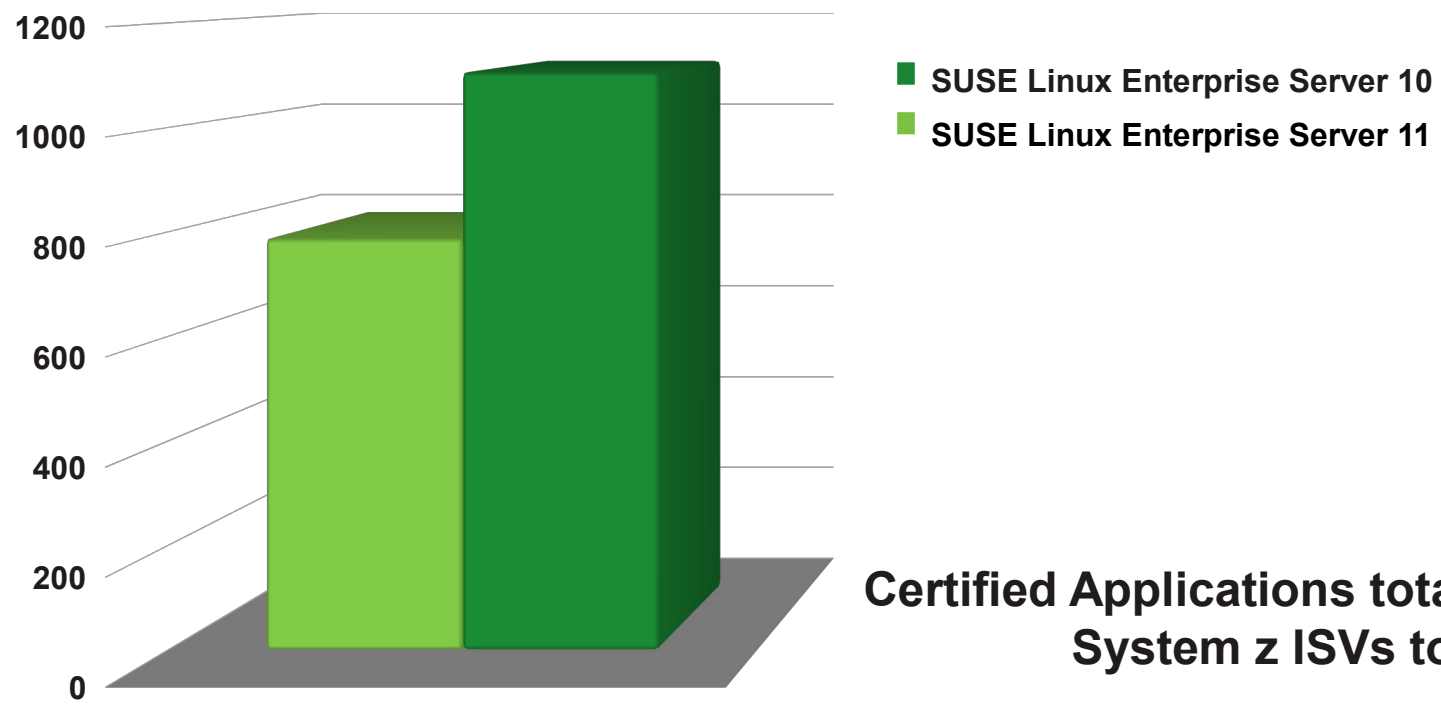


- Foundation for SUSE Linux Enterprise products
- Fully supported core system
- Choose the right Architecture for your workload



Certified ISV Applications For

SUSE® Linux Enterprise Server for System z



Certified Applications total: 1894
System z ISVs total: 516



SUSE® Linux Enterprise Server 11 SP1

- 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, z/OS and z/VSE 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 uniform resources support robust DR recovery setups
 - SUSE Linux Enterprise High Availability Extension is included



SUSE® Linux Enterprise Server 11 SP2

- 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
 - SUSE Linux Enterprise Server supported for both hardware architectures
- Improved tools and z specific support
 - Disk storage & crypto enhancements
 - Linux RAS support, s390-tools update



SUSE® Linux Enterprise 11 SP3

- Consolidation release
- **Hardware** Support
 - AMD64/Intel64
 - IBM System z and POWER
- **Unix to SUSE Linux Enterprise** migration
 - Improve existing infrastructure and solutions
- **Standards** Compliance
 - IPv6
- Solid infrastructure for SUSE Cloud and Integrated Systems offerings
- Select important customer-specific issues with business relevance


First customer shipment: June 2013 (Plan)

SUSE® Linux Enterprise Server

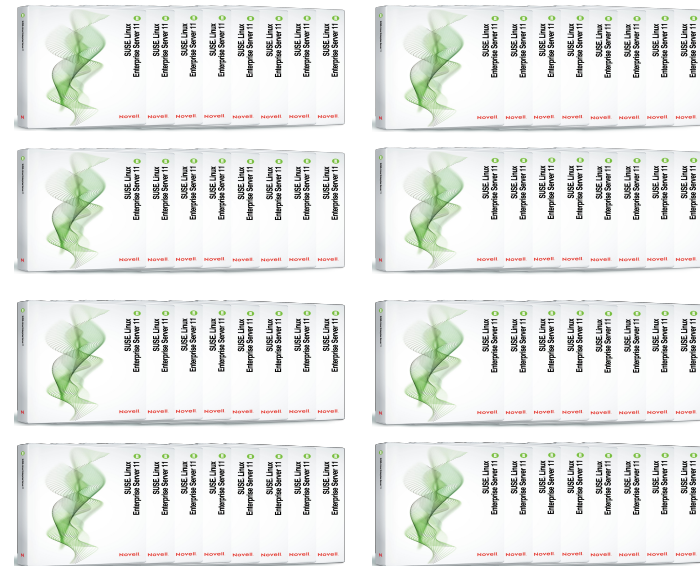
for IBM zEnterprise and zBX

If you own or buy 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 Summary

for SUSE Linux Enterprise Server for System z


Attractive subscriptions and services are available: talk to us!

- **Multi year, multiple IFL subscriptions offer monetary value – scale & save**
 - For larger volumes, you get up to 45% of discount per subscription
 - <http://www.suse.com/products/systemz/how-to-buy/>
- **SUSE Linux Enterprise High Availability Extension included – prevent outages**
 - Additionally, save more than \$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>



Upgrading your SUSE portfolio: Software Lifecycle Management

Current SUSE® Linux Enterprise Streams



- Dependable release timing
- Predictability for planning rollouts and migrations
 - Service Pack releases, development and product schedules announced to customers and partners
- Major releases every 4-5 years.





zEnterprise 196



z BladeCenter Extension

Upgrade options for SUSE® Linux

- Major version upgrades – for SLES9 and SLES10
- Service pack migration – within a major release
- “Do nothing” option – consider the success of Washington, DC

Consequences of “Do Nothing”

- Past the SUSE support timetable – and isn't that what you are paying for?
- Long Term Support is cost prohibitive
- Upgrading does not get easier over time
- Feature frozen, increasingly less functional and potentially less secure
- Vulnerability for other solutions to look “shinier”
- Outdated portfolio becomes harder for your organization to manage and support

Before you start upgrading...

- Start with low-hanging fruit
- Check any third-party applications' compatibility matrix
- Get buy-in from your team
- Test and validate
- Inform SUSE support in advance for readiness
- Do it again
- Weigh the options for more difficult loads - “functional migration”
- Plan wider deployment – understand the tools

Major Version Upgrades

- Can be done by booting from the newer ISO
- No “jumping” versions
- Consider deprecated packages or feature changes
- Pros:
 - In-place makes use of the same resources
 - Continuity – migration of configurations
 - Supported
- Cons:
 - Disk partitioning/filesystems inherited
 - “Ugly stuff” may still be there
 - May be difficult to return if trouble happens



SLES9 to SLES10 Considerations

- Validate at the most current set of patches/updates
- If you fully update, SLES9 will take you to SP4
 - There was only one update channel
- Review your usage of ReiserFS
- Boot to SLES10 SP4 media to get to a supported platform
- Once there, you migrate from SLES10 to SLES11

SLES10 to SLES 11 Considerations

- Need to be at SLES10 SP3 minimum
- Again – review ReiserFS usage
- Two methods to upgrade:
 - Boot from the SLES 11SP2 media and upgrade (For the first few systems or special cases)
 - Automated upgrade script – Makes use of AutoYaST (For the rest of your portfolio)

Service Pack Migrations

- Repository changes
- On SLES11, will also change kernel majors
 - 11 SP1 kernel version 2.6.32
 - 11 SP2, SP3 kernel version 3.0
- Upgrade options
 - Boot from media/ISO (any version)
 - YaST2 wagon
 - zypper dup (SLES11 SP1+) and repository changes
 - SUSE Manager

**and now:
Show time!**



Corporate Headquarters
Maxfeldstrasse 5
90409 Nuremberg
Germany

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

Join us on:
www.opensuse.org

Unpublished Work of SUSE. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE. Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

General Disclaimer

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. SUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for SUSE products remains at the sole discretion of SUSE. Further, SUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All SUSE marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.



Appendix

Unique Tools Included

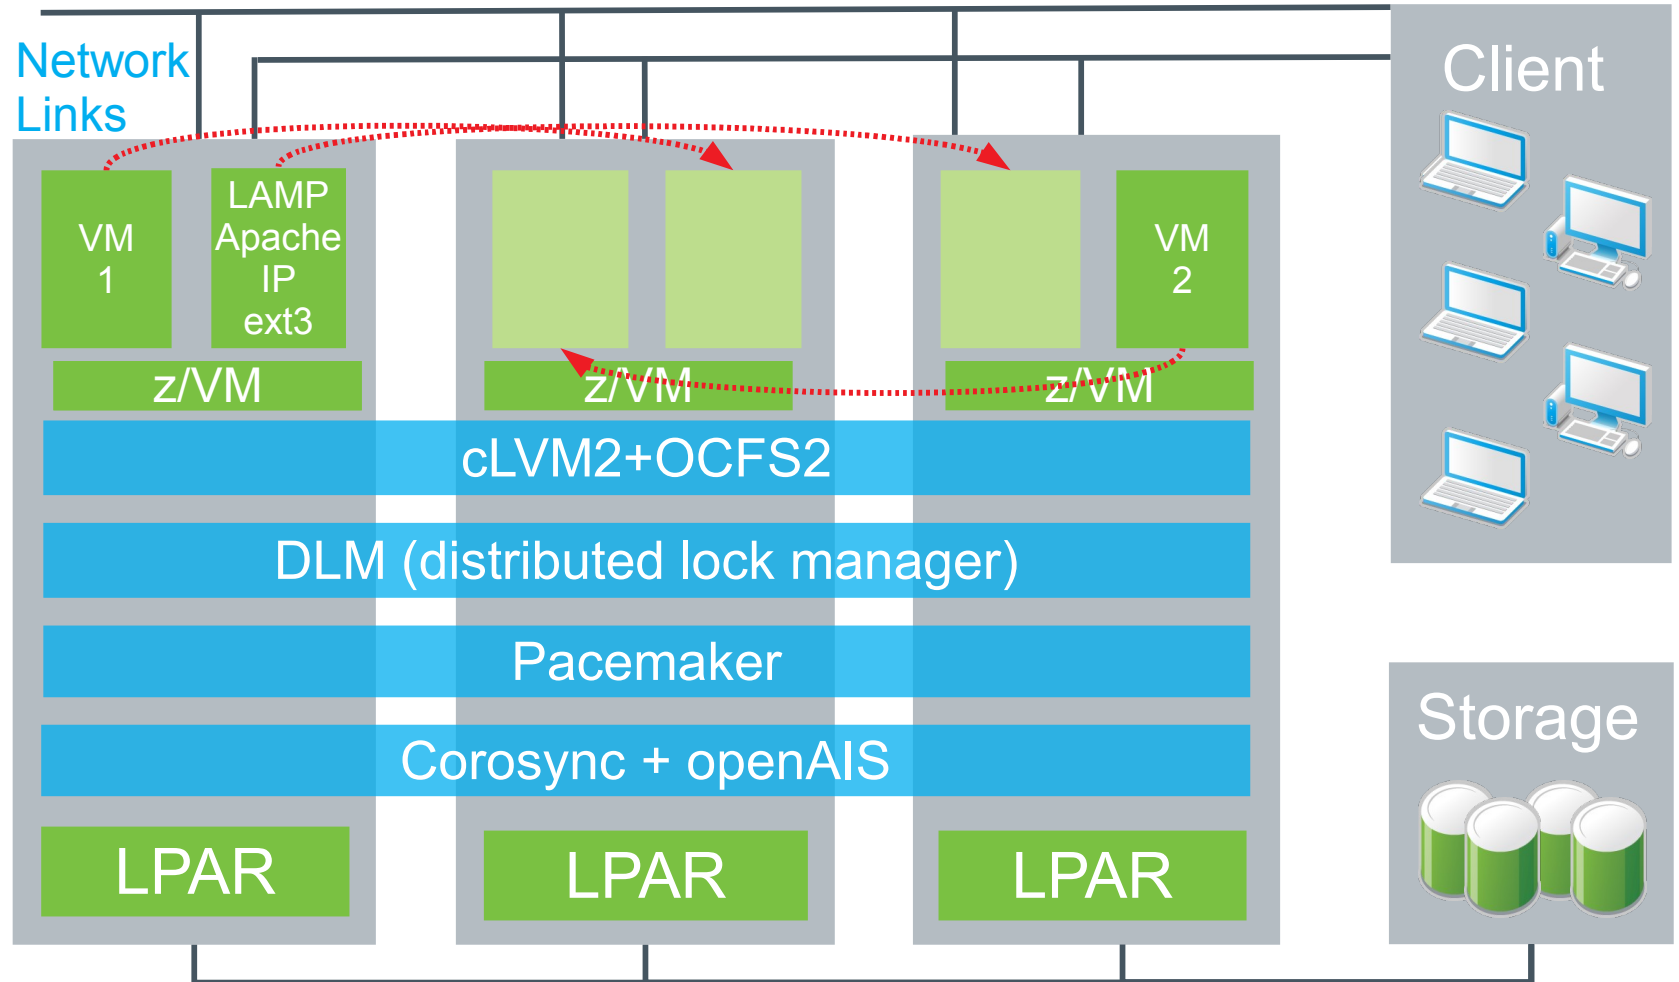
- High Availability Extension for SUSE® Linux Enterprise Server
 - Cluster Framework, ClusterFS, DRBD, tools, GEO-cluster
- AppArmor Security Framework
 - Application confinement
- Yast and Integrated Systems Management
 - Install, deploy, and configure every aspect of the server
- Subscription Management Tool
 - Subscription and patch management, proxy/mirroring/staging
- Starter System for System z
 - A pre-built installation server, deployable with z/VM tools

Overview

- **Service availability 24/7**
 - Policy-driven clustering
- **Shared and scaled data access**
 - Cluster file system
 - Clustered Samba
- **Geo clustering**
 - Cluster across unlimited distance
- **Virtualization agnostic**
 - Platform independent setup
- **Disaster tolerance**
 - Data replication via IP
 - Node recovery
- **Scale network services**
 - IP load balancing
- **User friendly tools**
 - Graphical user interface
 - Unified command line interface
- **Free resource agents**

Cluster Example

SUSE® Linux Enterprise High Availability Extension



Enhance and Secure Your Applications

Examples

- SUSE® Linux Enterprise High Availability Extension: make your applications high availability ready
 - Resource agents examples
 - `/usr/lib/ocf/resource.d/heartbeat/*` → examples of resource agents
 - <http://www.opencf.org>
- AppArmor: secure your applications
 - Easy to use GUI tools with statics analysis and learning-based profile development
 - Create custom policy in hours, not days

AppArmor: usr.sbin.vsftpd

/etc/apparmor/profiles/extras/

```
#include <tunables/global>

/usr/sbin/vsftpd {
  #include <abstractions/base>
  #include <abstractions/nameservice>
  #include <abstractions/authentication>


  /dev/urandom                r,
  /etc/fstab                   r,
  /etc/hosts.allow             r,
  /etc/hosts.deny              r,
  /etc/mtab                    r,
  /etc/shells>                 r,
  /etc/vsftpd.*                r,
  /etc/vsftpd/*                r,
  /usr/sbin/vsftpd>            rmix,
  /var/log/vsftpd.log          w,
  /var/log/xferlog             w,
  # anon chroots
  /                             r,
  /pub                          r,
  /pub/**                       r,
  @{HOMEDIRS}                   r,
  @{HOME}/**                    rwl,
}
```

Subscription Management Tool

Overview

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.



Starter System for System z

- A pre-built installation server that can be installed on your z/VM system using CMS tools
- Eliminates the need for coordinating access to a separate Linux or UNIX system elsewhere on your network
- Minimizes the impact of network-based installation on your internal and external networks
- Most current release is based on SUSE® Linux Enterprise Server 10 SP4 (Jul 2012)
- Future refresh based on SUSE Linux Enterprise Server 11 SP2 (H2 2012)



SUSE® Linux Enterprise Server 11 SP2 Specific Features

Forward Looking Development Model

Benefits



New Features and
Functionality
Faster

Improved RAS,
Scalability and
Security



More Hardware
Choice (especially
newer)

Application
Compatibility



Smart Innovation, Enterprise Quality
and Investment Protection

Kernel 3.0

Select benefits

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

SUSE® Linux Enterprise 11 SP2

- Hardware enablement and RAS improvements
- Equivalent or exceeding proprietary UNIX capabilities
 - btrfs: **file system** with “Copy on Write”, checksums, snapshotting, reduce cost of storage management by providing an integration of logical volume management and filesystem, checksums on data and metadata ensure data integrity
 - LXC: **container** support based on control groups
- Snapshot / rollback for package and configuration updates
 - YaST2 + ZYPP + btrfs
- SUSE Linux Enterprise High Availability Extension: Geo-cluster, automated and pre-configuration
- Unattended upgrade from SUSE Linux Enterprise 10 to SUSE Linux Enterprise 11



Why btrfs?


Why Another File System?

- Solve Storage Challenges
 - Scalability
 - Data integrity
 - Dynamic resources (expand and shrink)
 - Storage management
 - Server, cloud – desktop, mobile
- Match and Exceed other Operating Systems


Technology Overview

Subvolume (1)

Normal File system



With Subvolumes



Snapshots in SUSE Linux Enterprise 11 SP2

YaST2 Management

The image shows two overlapping windows from the YaST2 management tool. The background window is titled "Snapshots" and displays a table of system snapshots. The foreground window is titled "Selected Snapshot Overview" and provides a detailed view of the selected snapshot (ID 10-11), including a file tree and a diff of the modified file `printers.conf`.

ID	Type	Start Date	End Date	Description
1	Single	Wed 17 Aug 2011 04:30:01 PM CEST		timeline
2 - 3	Pre & Post	Wed 17 Aug 2011 04:31:54 PM CEST	Wed 17 Aug 2011 04:32:46 PM CEST	yast lan
4 - 5	Pre & Post	Wed 17 Aug 2011 04:32:48 PM CEST	Wed 17 Aug 2011 04:32:59 PM CEST	yast lan
6 - 7	Pre & Post	Wed 17 Aug 2011 04:36:10 PM CEST	Wed 17 Aug 2011 04:36:11 PM CEST	zypp(zypper)
8 - 9	Pre & Post	Wed 17 Aug 2011 04:36:16 PM CEST	Wed 17 Aug 2011 04:36:19 PM CEST	zypp(zypper)
10 - 11	Pre & Post	Wed 17 Aug 2011 04:36:26 PM CEST	Wed 17 Aug 2011 04:37:21 PM CEST	yast printer
12	Single	Wed 17 Aug 2011 05:20:01 PM CEST		timeline
13	Single	Wed 17 Aug		
14	Single	Wed 17 Aug		
15	Single	Wed 17 Aug		
16	Single	Wed 17 Aug		
17	Single	Wed 17 Aug		
18	Single	Wed 17 Aug		
19	Single	Thu 18 Aug 2		
20	Single	Thu 18 Aug 2		

Selected Snapshot Overview

10 - 11

- etc
 - cups
 - printers.conf
 - printers.conf.O
 - var

Show the difference between first and second snapshot
 Show the difference between first snapshot and current system
 Show the difference between second snapshot and current system

File content was modified.

```
--- /snapshots/10/snapshot/etc/cups/printers.conf 2011-08-17 16:20:38.325347599 +0200
+++ /snapshots/11/snapshot/etc/cups/printers.conf 2011-08-17 16:36:54.936184604 +0200
@@ -1,12 +1,12 @@
# Printer configuration file for CUPS v1.3.9
-# Written by cupsd on 2011-08-17 16:20
+# Written by cupsd on 2011-08-17 16:36

Info HP LaserJet 4050 Series Postscript (recommended)
```

Restore From First Restore From Second

Help Cancel Restore Selected

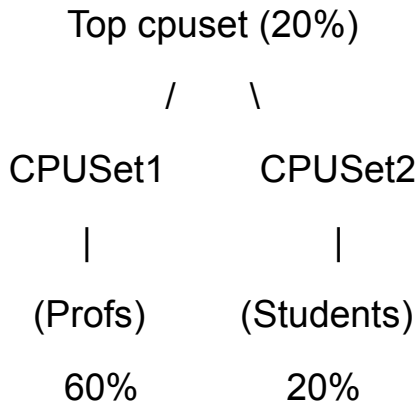
What Are Control Groups?

- Control Groups provide a mechanism for aggregating and partitioning sets of tasks, and all their future children, into hierarchical groups with specialized behavior.
 - cgroup is another name for **Control Groups**
 - **Partition tasks** (processes) into a one or many groups of **tree hierarchies**
 - **Associate** a set of tasks in a group to a set of subsystem parameters
 - **Subsystems** provide the parameters that can be assigned
 - Tasks are **affected** by the assigning parameters

Capabilities of a cgroup

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:

CPU



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>



Memory Subsystem

Resource Control

- 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 disabled at boot time with kernel parameter "cgroup_disable=memory"

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



Device Subsystem

Isolation

- A system administrator can provide a list of devices that can be accessed by processes under cgroup
 - Allow/Deny Rule
 - Allow/Deny : READ/WRITE/MKNOD
- Limits access to device or file system on a device to only tasks in specified cgroup

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



IBM System z - 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).

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Hardware exploitation z196 processor topology and cache hierachy, increase cache hit ratio and therefore overall performance	<ul style="list-style-type: none">• Increase of application workload density per system• Increased performance

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



Crypto CPACF 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
<ul style="list-style-type: none">• Cryptographically secured connections• devices maximum request size is adjusted based on a test	<ul style="list-style-type: none">• 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



Fill entropy with hwrng 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
<ul style="list-style-type: none">• Use /dev/random as a source of random numbers generated by hardware at a high rate• Avoids stalling of processes querying for randomness	<ul style="list-style-type: none">• 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



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
<ul style="list-style-type: none">• Hardware exploitation of the z196 and zBX for hybrid computing	<ul style="list-style-type: none">• Fit-to-purpose workload placement support

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
<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• 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



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
<ul style="list-style-type: none">• Guided setup and configuration of a suitable dump device	<ul style="list-style-type: none">• Improved serviceability

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
<ul style="list-style-type: none">• More flexible I/O configuration for Linux running in LPARs	<ul style="list-style-type: none">• Non disruptive change of IO configuration

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
<ul style="list-style-type: none">• Provides a new cio layer function using an interface to get the maximum usable data size for zHPF requests on a given subchannel	<ul style="list-style-type: none">• Maximize I/O performance with FICON, zHPF and DS8000 storage servers

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
<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• 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



Provide Linux Process Data into z/VM Monitor Stream

Fate 302787 / [LTC 200920]

<http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.x.html> -> z/VM

Device Drivers, Features, and Commands as available with SUSE Linux Enterprise Server 11 p.179ff

- **z/VM infrastructure exploitation:** monitor the performance of SLES or other VM guests with performance monitoring tools on z/VM or on Linux. Use own, IBM Performance Toolkit for VM, or third party tools. VMs being monitored require agents that write monitor data.

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Health and performance check of applications in multiple VMs• Adapt to our own requirements	<ul style="list-style-type: none">• Keep SLAs, resource utilization control

SLES	10	11
GA	-	yes
SP1	-	yes
SP2	yes	n/a
SP3	yes	n/a



Kernel Message Catalog Autogeneration

Fate 306999 / [LTC 201727]

<http://www.ibm.com/developerworks/linux/linux390/kernel-2.6.29.html#message-catalog>

Kernel Messages (lk32km00.pdf)

- **Additional documentation and improved problem resolution:** the kernel components are instrumented in a way that a message catalog as well as related man pages are automatically generated and can be retrieved by the administrator.

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Documentation of kernel related messages• Faster error isolation and issue resolution	<ul style="list-style-type: none">• Keep SLAs by faster problem resolution

SLES	10	11
GA	-	yes
SP1	-	yes
SP2	-	n/a
SP3	-	n/a

 = SLES unique feature

More Features ...

- About ~ 80 more new z specific features included in SUSE® Linux Enterprise Server 11SP2
- Contact Marcus Kraft <mkraft@suse.com> to request a full and detailed list

Tools / SDK

zPDT

IBM System z Personal Development Tool

https://www.ibm.com/partnerworld/page/pw_com_zpdt



- zPDT is a software-based application tool
 - Low cost IBM System z platform for ISV application development, testing, demo
 - A virtual System z architecture environment that allows select mainframe operating systems, middleware and software to run unaltered on x86 processor-compatible platforms.
 - Portable System z platform for training & education of applications and operating system environments
- Supports openSUSE 10.3+, SUSE® Linux Enterprise Server 11SP1 x86_64, and others
 - SUSE's evaluation versions for x86_64 and s390x available at <http://www.suse.com/products/server/eval.html>



Tools

Dynamic Analysis Tools

- valgrind
 - Memcheck
 - Cachegrind
 - Massif
 - Helgrind
 - DRD
 - None
 - Exp-ptrcheck
 - Callgrind
- <http://valgrind.org>



Tools

memcheck

- Invalid accesses
 - Non allocated memory
 - Recently freed memory
 - Access beyond boundaries (eg. Arrays)
 - Invalid stack areas
- Use of uninitialized data
- Memory leaks
- False memory release
 - freeing memory twice
 - (not: heap overrun)
- Parameters: `--leak-check=full, --track-origins=yes`

Tools

drd/helgrind

- Discover potential races of threads
 - Read/write without locking
 - Lock contention
- Wrong use of pthread library
 - Invalid pointers (eg. Mutex instead of condition variable)
 - Unlock without previous lock
 - Unlock of a lock done by another thread
 - Unlock of a lock which is still in use
 - ...
- DRD / helgrind do overlap in the function
 - Each tool covers areas, which the other cannot detect
- Parameters: `--leak-check=full, --track-origins=yes`



Tools

callgrind

- Call statistic profiling
 - Callgraphs
 - Dependencies of functions
 - Cache accesses
 - Branch prediction
- Generates callgrind.out.<pid> data
- GUIs help to analyze the data
 - Assembly data `–dump-instr=yes`
 - Branches `–trace-jump=yes`
- Includes simple branch prediction simulation

Tools

callgrind profiling GUI (e.g. <http://kcachegrind.sourceforge.net>)

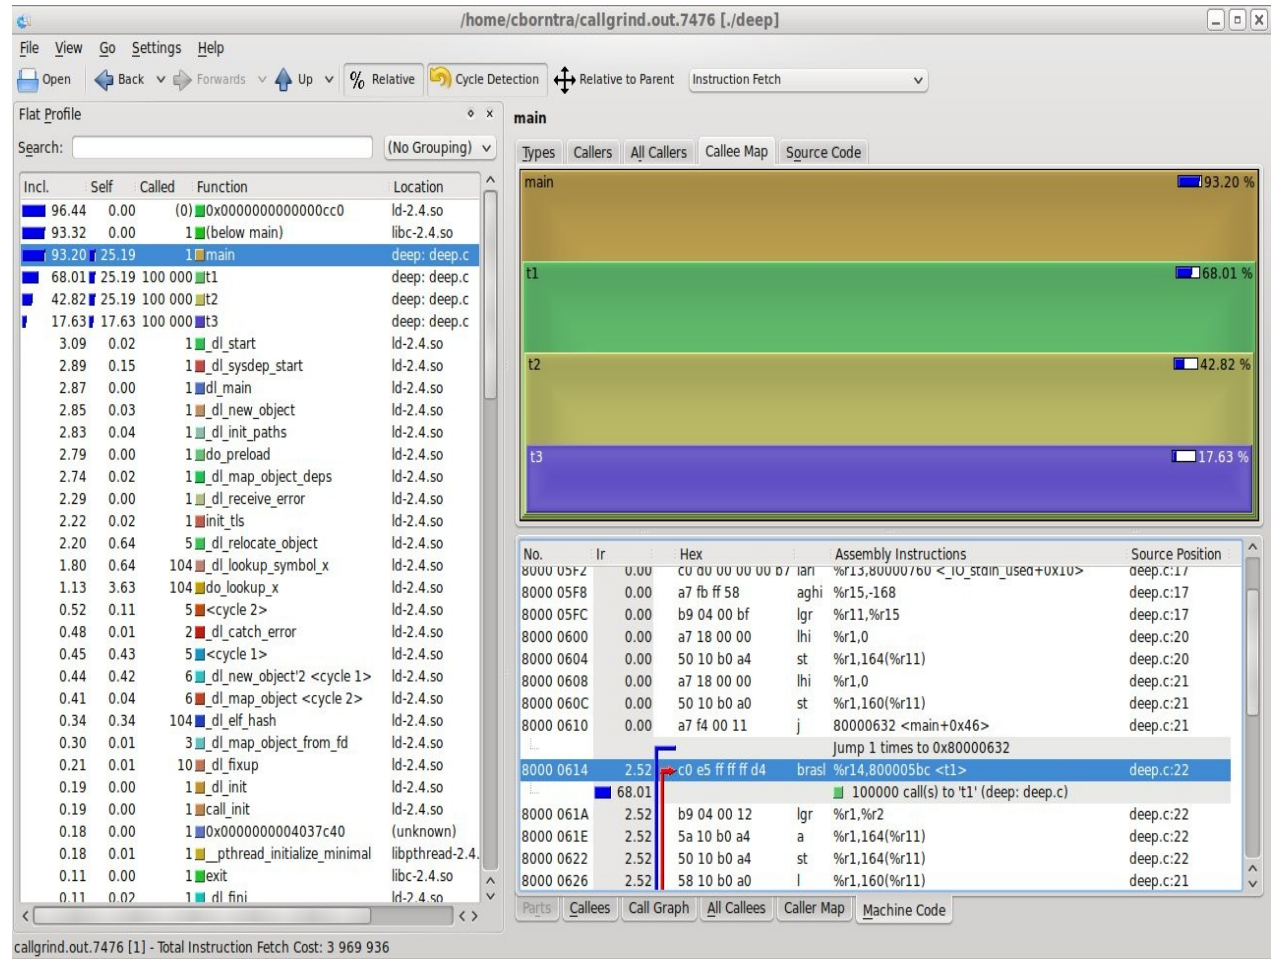
```
int t3(void)
{
    return 1;
}

int t2(void)
{
    return t3();
}

int t1(void)
{
    return t2();
}

int main()
{
    int i, sum;

    sum = 0;
    for (i=0; i<100000;
i++)
        sum+=t1();
    return sum;
}
```



valgrind --tool=callgrind --dump-instr=yes --trace-jump=yes ./deep

Tools

cachegrind

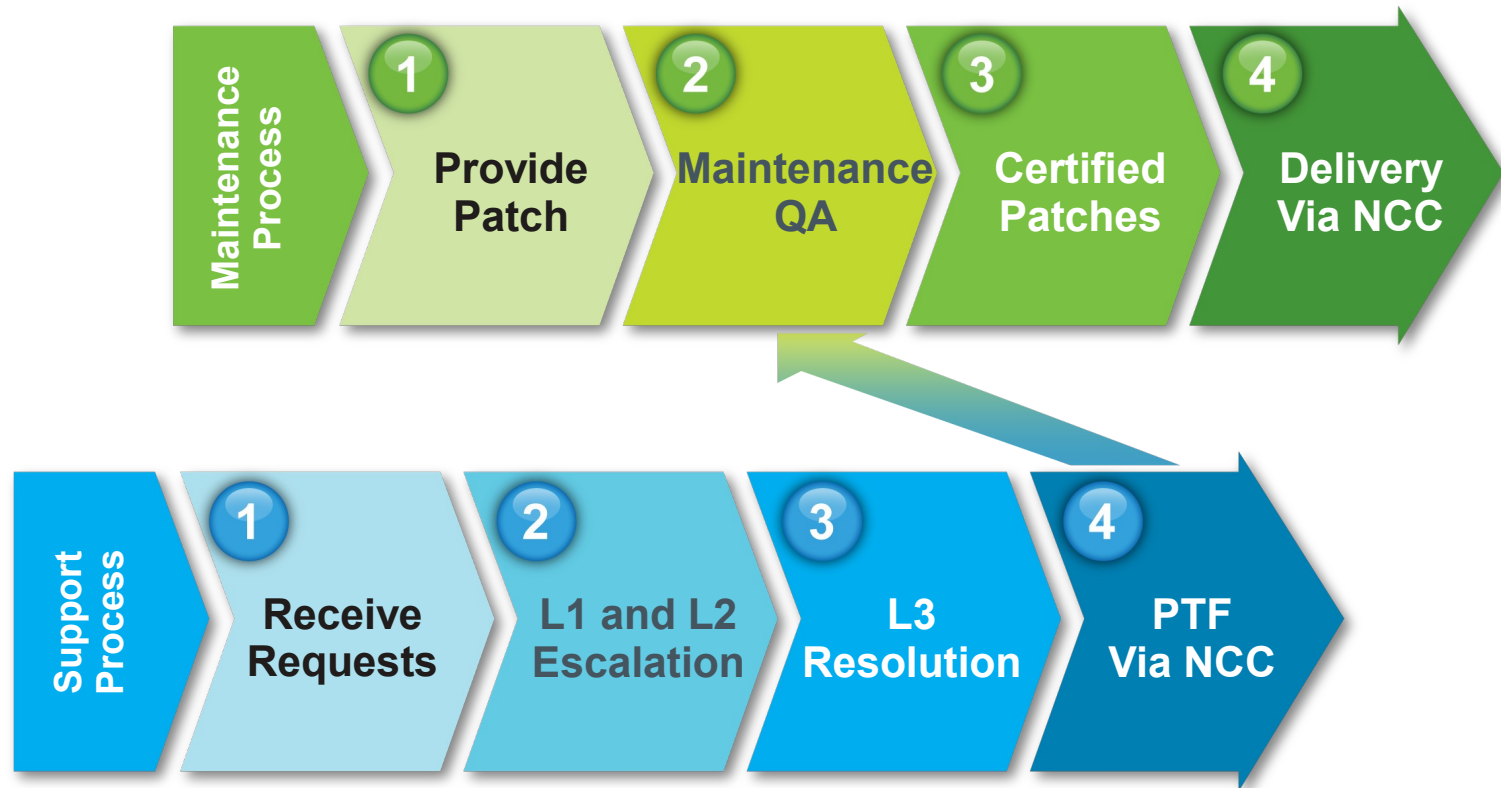
- Analysis of cache behaviour of applications
 - z10 cache sizes used as default, changeable (z9, z196)
 - Two cache levels (1st and last level) for instructions & data
 - Writes cachegrind.out.<pid> files

```
r1745045:~ # valgrind --tool=cachegrind ls
==21487== Cachegrind, a cache and branch-prediction profiler
==21487== Copyright (C) 2002-2010, and GNU GPL'd, by Nicholas Nethercote et al.
==21487== Using Valgrind-3.6.1 and LibVEX; rerun with -h for copyright info
==21487== Command: ls
==21487==
--21487-- Warning: Cannot auto-detect cache config on s390x, using one or more defaults
bin inst-sys repos testtools
==21487==
==21487== I refs:          656,270
==21487== I1 misses:         792
==21487== LLi misses:        656
==21487== I1 miss rate:    0.12%
==21487== LLi miss rate:   0.09%
==21487==
==21487== D refs:          453,124 (361,066 rd + 92,058 wr)
==21487== D1 misses:       1,869 ( 1,589 rd + 280 wr)
==21487== LLd misses:     1,313 ( 1,061 rd + 252 wr)
==21487== D1 miss rate:    0.4% ( 0.4% + 0.3% )
==21487== LLd miss rate: 0.2% ( 0.2% + 0.2% )
==21487==
==21487== LL refs:          2,661 ( 2,381 rd + 280 wr)
==21487== LL misses:       1,969 ( 1,717 rd + 252 wr)
==21487== LL miss rate:   0.1% ( 0.1% + 0.2% )
```



Additional Information

Maintenance and Support Process



Documentation and Release Notes

- Product Pages
 - <http://www.suse.com/products/server/>
 - <http://www.suse.com/products/highavailability/>
 - <http://www.suse.com/products/realtime/>
- Unix to Linux Migration
 - <http://www.novell.com/linux/unixtolinux/>
- Documentation
 - <http://www.novell.com/documentation/suse.html>
- Product Life-cycle
 - <http://support.novell.com/lifecycle/linux.html>
- Release Notes
 - <http://www.novell.com/linux/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



Documentation and Release Notes

- Product Pages
 - <http://www.suse.com/products/server/>
 - <http://www.suse.com/products/highavailability/>
- Release Notes
 - <http://www.novell.com/linux/releasenotes/>
- Product Life-cycle
 - <http://support.novell.com/lifecycle/linux.html>
- ISV Partner Program
 - suse_isv@suse.com
- Unix to Linux Migration
 - <http://www.novell.com/linux/unixtolinux/>
- Documentation
 - <http://www.novell.com/documentation/suse.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>



How to Find Certified Applications

Visit our catalog:
www.suse.com/partner/isv/isvcatalog

SUSE Linux Enterprise Software Catalog

Get the most out of SUSE® Linux Enterprise. **Search or browse our Independent Software Vendor catalog for applications certified on SUSE Linux Enterprise.**

Browse or Search

Browse the catalog by	Search for an ISV
+ Application	<input type="text"/>
+ ISV	Search for an application
+ Industry	<input type="text"/>
+ Platform	+ Advanced Search
+ Solution Area	<input type="button" value="Search"/>
+ Workload	

Request a call
Talk to our SUSE Linux Enterprise certification team to learn more

Get Certified
Find out how to become certified

Get help
Support for your existing deployment

This list includes all Novell partners that provide applications supported on current SUSE Linux Enterprise Platforms.

Browse by ISV

[ALL](#) [123](#) [A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) [J](#) [K](#) [L](#) [M](#) [N](#) [O](#) [P](#) [Q](#) [R](#) [S](#) [T](#) [U](#) [V](#) [W](#) [X](#) [Y](#) [Z](#) ▲ Top

- ▣ 3BVIEW LTD
- ▣ ABSOFT
- ▣ ACCESSVIA INC

Security

System Hardening	YaST2 Security Center
Application confinement	AppArmor
System Confinement	SE Linux (Stack Support)
Intrusion Detection (Filesystem)	AIDE
Fine-grained access rights	Filesystem POSIX capabilities
Encryption capabilities	Three ways: “Full Disk” – Volume – File System (eCryptFS)
Certifications	CGL 4.0 IPv6 (refresh)
Measure and monitor system integrity during (re)boot	Trusted Platform Modules (TPM) – Trusted Computing

SUSE® Linux Enterprise

Phases of Flexible Lifecycle

Feature	General Support for most recent Service Pack		General Support for former Service Pack with LTSS option	Extended Support with LTSS option Year 8-10
	Year 1-5	Year 6-7		
Support for existing stacks and workloads	yes	yes	yes	yes
Support for new deployments	yes	yes (1)	yes (1)	no
Enhancement Requests	yes	yes (1)	no	no
Hardware Enablement/Optimization	yes	no	no	no
Driver Updates via PLDP	yes	yes(1)	yes (1)	no
Backport of fixes from more recent Service Pack	yes	n/a	yes (2)	n/a
Critical Security Updates	yes	yes	yes	yes
Defect Resolution	yes	yes(2)	yes (2)	yes (2)
Technical Support	yes	yes	yes	yes
Access to Patches and Fixes	yes	yes	yes	yes
Access to Documentation and Knowledgebase	yes	yes	yes	yes
(1) Limited. Based on partner and customer requests.				
(2) Limited. Severity Level 1 and 2 defects only.				
LTSS option is available for an extra price per codestream, per architecture, per year				

SUSE® Linux Enterprise 11 SP2

Kernel Capabilities

SLE 11 SP 1 (2.6.32)	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-/ kernel space	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



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
<ul style="list-style-type: none">• Hardware exploitation of the z10 and z196 instruction set for user land applications (ISV and self compiled applications)	<ul style="list-style-type: none">• z10 and z196 optimized code and efficient execution 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



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
<ul style="list-style-type: none">• New sophisticated handling of mutexes and scheduler decisions to improve performance also for z/VM based workloads	<ul style="list-style-type: none">• 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



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
<ul style="list-style-type: none">• Determine behaviour via parameter and depending on workload	<ul style="list-style-type: none">• 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 now changed to retrieve CPC and LPAR information dynamically. With the new function, GDPS always resets exactly the LPAR in which the OS is running.

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• avoid possible failures from manual or forgotten changes• avoid resetting a LPAR due to incorrect configuration definitions GDPS	<ul style="list-style-type: none">• additional protection against outages

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
<ul style="list-style-type: none">• Can be used to synchronize log gathering between the operating system and the channel firmware.	<ul style="list-style-type: none">• Concurrent data collection for problem resolution, minimizing customer operation impact

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



Access to raw ECKD data from Linux (DASD)

Fate 311973 / [LTC 66951]

- **Interoperability:** 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
<ul style="list-style-type: none">• Includes Linux ECKD disks used with LVM, Linux ECKD disks that are used directly, and z/OS ECKD disks	<ul style="list-style-type: none">• 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



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
<ul style="list-style-type: none">• 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.	<ul style="list-style-type: none">• 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



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
<ul style="list-style-type: none">• 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.	<ul style="list-style-type: none">• Optimized or no downtime when resuming a Linux instance

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
<ul style="list-style-type: none">• 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.	<ul style="list-style-type: none">• 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



Access to raw ECKD data from Linux (DASD)

Fate 311973 / [LTC 66951]

- **Interoperability:** 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
<ul style="list-style-type: none">• Includes Linux ECKD disks used with LVM, Linux ECKD disks that are used directly, and z/OS ECKD disks	<ul style="list-style-type: none">• 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
<ul style="list-style-type: none">• Can be used to synchronize log gathering between the operating system and the channel firmware.	<ul style="list-style-type: none">• Concurrent data collection for problem resolution, minimizing customer operation impact

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
<ul style="list-style-type: none">• Hardware exploitation of the z196 instruction set for user land applications (ISV and self compiled applications), recompile programs with <code>--march=z196</code> and/or <code>-mtune=z196</code>	<ul style="list-style-type: none">• z196 optimized code and efficient execution 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



DS8000 Disk Encryption

Fate 307004 / [LTC 201740]

<http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.1.html> -> dasdview
Device Drivers, Features, and Commands as available with SUSE Linux Enterprise Server 11 p.350

- **Hardware support:** enhances s390-tools to be able to display if the disk storage has its disk encrypted or not.
- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Retrieve info on encryption status of device	<ul style="list-style-type: none">• Secure data storage

SLES	10	11
GA	-	yes
SP1	-	yes
SP2	-	yes
SP3	yes	n/a

