

What's new with SUSE[®] Linux Enterprise Server 11 SP1

for IBM System z

Mark Post

Technical Support Engineer IV
mpost@suse.com



Agenda

- Introduction
- New maintenance model
- System z specific
 - Hardware support
 - Performance improvements
 - Usability improvements
 - Software additions/upgrades/removals
 - Miscellaneous changes

Agenda (2)

- Non-System z specific
 - Performance improvements
 - Usability improvements
 - Software additions/upgrades/removals
 - Miscellaneous changes
- Questions (I'll also take questions during the talk, unless we start to run short of time.)
- Backup/general information slides

Introduction

- As the name indicates, SLES11 SP1 is the first service pack of the SLES11 version.
- If you remember the SLES11 announcement, there was a notice of an “Architecture Level Set to only support z9 or later” hardware for System z.
 - The ALS was *not* done, even though the support terms remained as "z9 and later."
 - The ALS may still be done, but perhaps not until SLES12.
- Note that I will not be talking about any features that were introduced in SLES11. This is only about changes introduced in SLES11 SP1.
- Includes lots and lots of bug fixes created for SLES11.

SUSE Product Lines and Products

Server Product Line

- SUSE Linux Enterprise Server
- SUSE Linux Enterprise Server for System z
- SUSE Linux Enterprise High Availability Extension
- SUSE Linux Enterprise Real Time Extension
- SUSE Linux Enterprise Mono Extension
- SUSE Linux Enterprise Point of Service
- SUSE Linux Enterprise Server for SAP Applications

SUSE Manager

Desktop Product Line

- SUSE Linux Enterprise Desktop
- LibreOffice

Virtualization Product Line

- SUSE Linux Enterprise Server

Appliance Product Line

- SUSE Studio Online
- SUSE Appliance Toolkit

Specialized Support Offers

- SUSE Linux Enterprise Server Subscription with Expanded Support (for Red Hat migrations)
- SUSE Linux Enterprise Server Priority Support for SAP

SUSE Product Lines and Products

Server Product Line

- SUSE Linux Enterprise Server
- SUSE Linux Enterprise Server for System z
- SUSE Linux Enterprise High Availability Extension
- SUSE Linux Enterprise Real Time Extension
- SUSE Linux Enterprise Mono Extension
- SUSE Linux Enterprise Point of Service
- SUSE Linux Enterprise Server for SAP Applications

SUSE Manager

Desktop Product Line

- SUSE Linux Enterprise Desktop
- LibreOffice

Virtualization Product Line

- SUSE Linux Enterprise Server

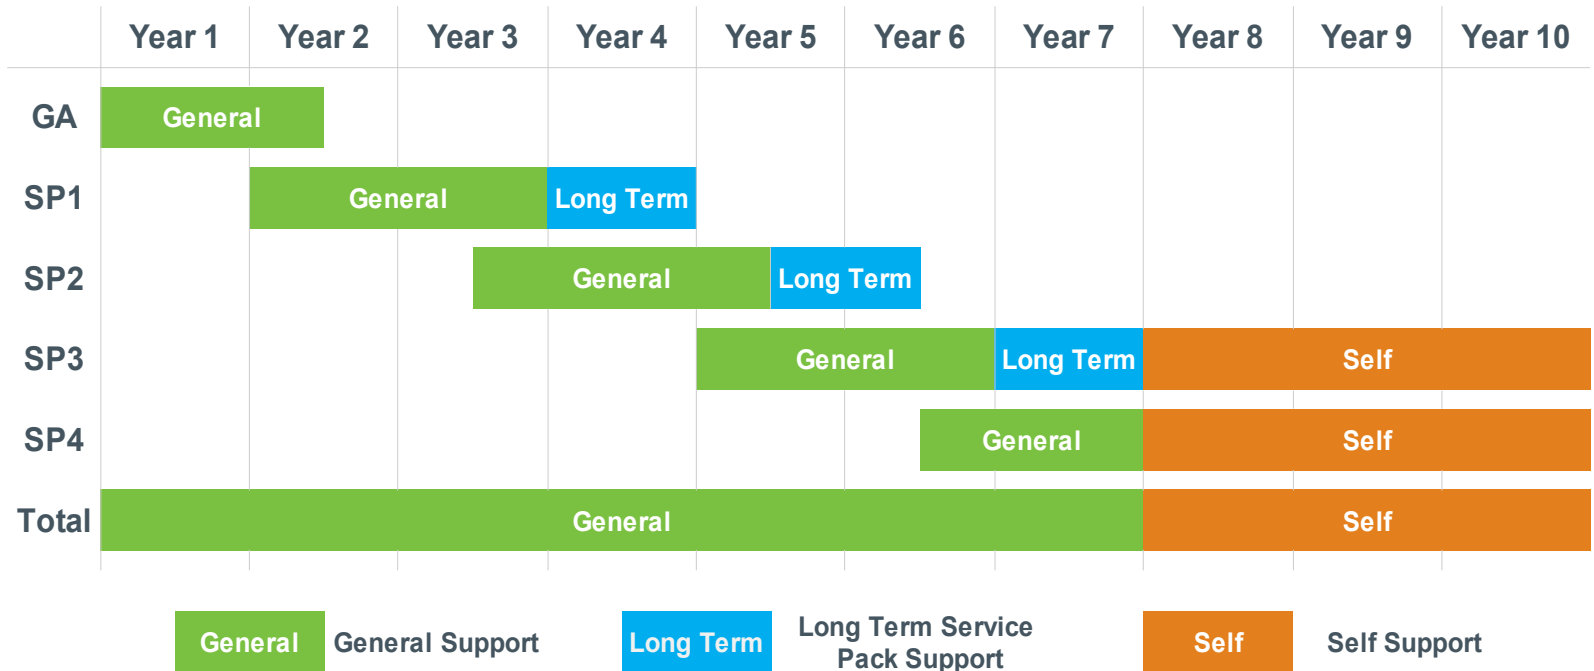
Appliance Product Line

- SUSE Studio Online
- SUSE Appliance Toolkit

Specialized Support Offers

- SUSE Linux Enterprise Server Subscription with Expanded Support (for Red Hat migrations)
- SUSE Linux Enterprise Server Priority Support for SAP

Standard Product Life-cycle



- 10-year life-cycle (7 years general support, 3 years self support)
- Service packs every 16-18 months, major releases every 3-4 years
- Six month upgrade window
- Long-term service pack support option available

Current Product Life-cycle Overview

	2009	2010	2011	2012	2013	Notes
SLE 9	Green			Orange		CR: SLE 9 SP4 (2007) EOGS: Q3 2011 EOSS: Q3 2014
SLE10	Green (SP3)	Green (SP4)		Green	Green (end of 2013) / Orange (start of 2014)	CR: SLE 10 SP3 (2009) EOGS: Q3 2013 EOSS: Q3 2016
SLE 11	Green (GA)	Green (SP1)	Green (SP2)	Green	Green (SP3)	CR: SLE 11 GA (2009) EOGS: Q2 2016 EOSS: Q2 2019
SLE 12					Green (GA)	GA: 2H 2013

Note: Milestones in future years are planned dates, and subject to change
Copyright 2011 SUSE



New Maintenance Model

Benefits to Customers


- Improve Customer Satisfaction by
 - Making **Service Packs** more lightweight
 - Easier to test and deploy ("consume")
 - Allowing for **staying with older versions**
 - And remaining supported for the full system
- Answer market needs **in between Service Packs** by
 - Allowing **more Updates** in the GU repository
 - Delivering selective **Enhancements**
(Mitigates the longer distance between SPs)

Implementation

- One primary repository per major version
 - Starting with the SLES11 SP1 updates repository
 - You may hear references to “GU” (General Update repository) here.
Keep in mind: GU repo = existing SLES11 SP1 updates repo
- Service Pack repositories contain dependencies that are strictly bound together “patch sets” each with its own repository
- Updates are released when they are ready
- Quality for the enterprise – test entire service pack

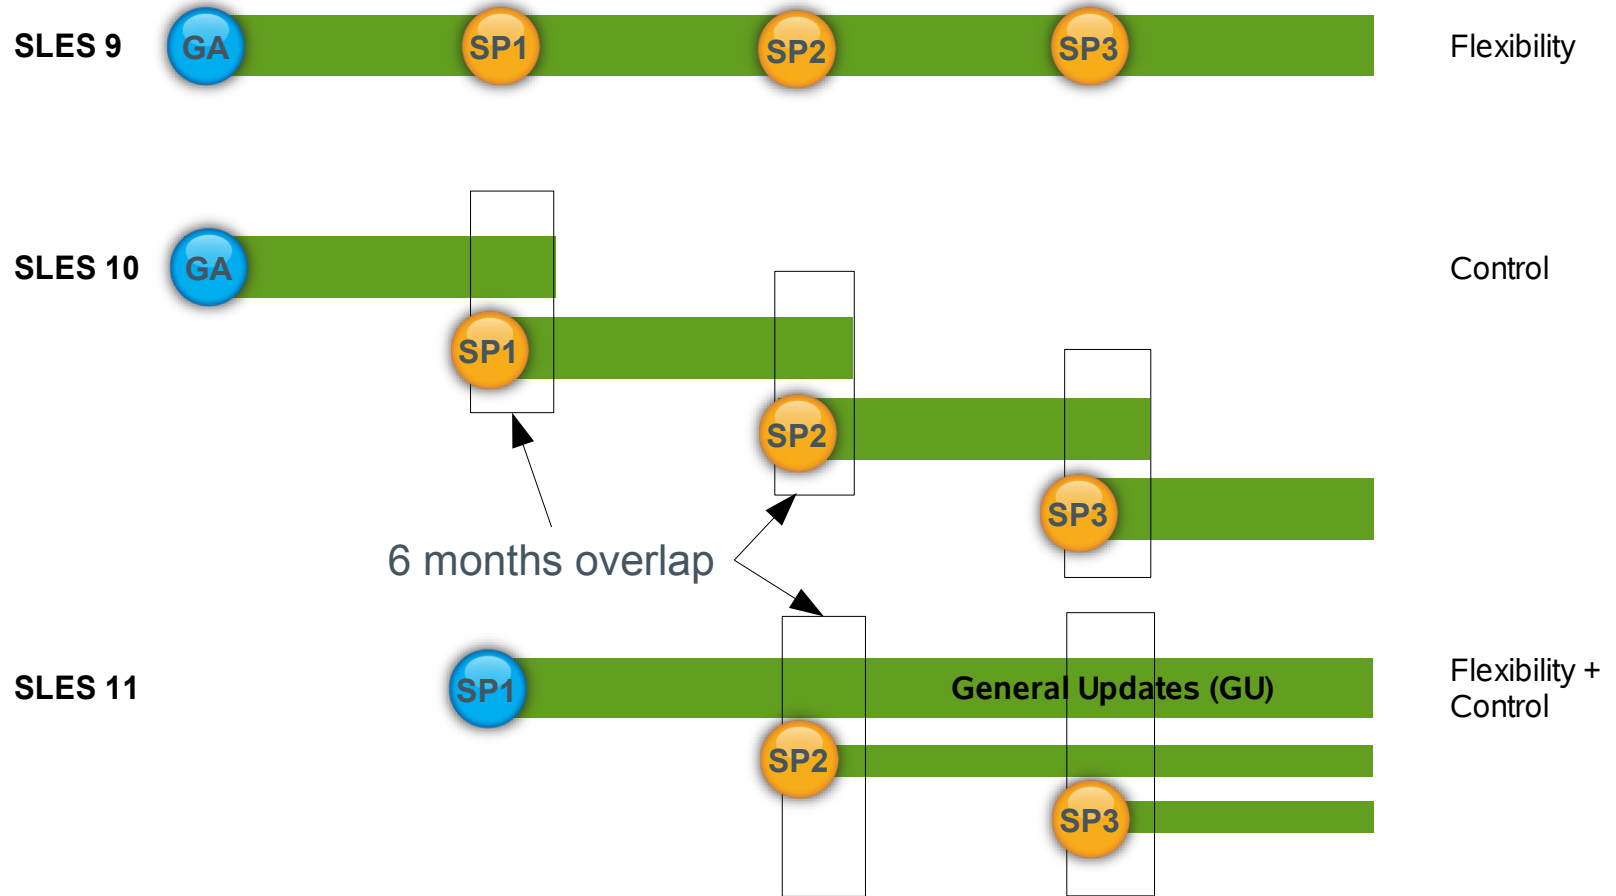
New Maintenance Model

SLES 11



- Service Pack
 - Dependencies – ABI/API and other
 - Major impact changes with high configuration QA requirement
 - Examples: intrusive kernel changes; invasive hardware updates
- GU + Previous Service Pack(s)
 - Service pack independent updates
 - Examples: Firefox
 - No delay on features and updates

Maintenance Delivery Evolution



Where do SP Patches End Up?

- Final decision by SUSE Release Managers and Maintenance Team depending on technical feasibility
- General Rules:
 - Independent new package or fully compatible update, bugfixes or improvements with minor impact
→ SLE 11 SP1 repository = GU Repository
 - New package or fully compatible update, bugfixes or improvements with major impact or high QA need
→ SLE 11 SP1 repository together with Service Pack
 - Change which might affect compatibility or where (new or existing) package depends on other packages in an SP
→ SLE 11 SP(n) repository together with Service Pack

Pick and Choose

- Customers can choose to remain on an arbitrary version of a package in the GU repository
- NTS will support them, i.e.
 - Accept a Service Request on those versions
 - Analyze with the customer if his issue is already fixed
 - Trigger fixing of an obvious issue
 - Try to minimize the need of customers to upgrade
- Fixes to the package will only be delivered on top of the latest version
- If the latest version is in a Service Pack repository, customer either must migrate to that Service Pack or buy LTSS (Long Term Service Pack Support)

System z specific

Hardware support

Hardware Support

- DS8000 large volume support
 - Allows use of volumes with more than 64K cylinders.
 - Uses the CCCHR (vs. CCHHR) addressing scheme (Fate #304026, 307003, IBM # 201597, 201738)
- High Performance FICON
 - DASD driver will use new transport mode control words (TCW) instead of CCWs.
 - Requires z10 or later hardware
 - Only for LPAR installations, not under z/VM (Fate # 307005, IBM # 201742)

Hardware Support (2)

- Format Record 0 on ECKD devices
 - Corrects a problem with formatting non-IBM storage devices (Fate # 307011, 307012, IBM # 201834, 201835)
- OSA QDIO Data Connection Isolation
 - Enables isolating QDIO data connections on a shared OSA-card
 - supports multi-tier security zones (Fate # 307015, IBM # 201899)
- Next generation crypto hardware exploitation
 - That is, crypto cards for the z196 and z114 (CEX2, CEX3)
 - Needs libica 2.0.1 or later (Fate # 307112, IBM # 201995)

Performance improvements

Performance Improvements

- Decimal Floating Point library support: libdfp
 - Includes required glibc infrastructure for I/O
 - Provides everything that is essential to support DFP (Fate # 304044, 307062, IBM # 201180, 201933)
- Crypto Device Driver Thin Interrupts
 - Eliminates polling for work (Fate # 306998, IBM # 201174)
- Kernel vdso support
 - Speeds up system calls to `gettimeofday`, `clock_getres` and `clock_gettime`
 - If you've ever done an `strace` on a command/program, you'll know why this is very helpful. (Fate # 307013, IBM # 201857)

Performance Improvements (2)

- Iconv character conversion routines speedup
 - Uses the new translate hardware instruction to speedup character conversion routines in glibc
 - Very important if you're doing a lot of EBCDIC ↔ ASCII conversions or running in a non-English or UTF8 locale. (Fate # 304074, IBM # 201185)
- AF_IUCV SOCK_SEQPACKET support
 - Speeds up communications between guests or with z/VM
 - Adds AF_IUCV datagram stream-oriented sockets in addition to the existing AF_IUCV byte stream-oriented sockets (Fate # 307014, IBM # 201885)

Performance Improvements (3)

- Suspend / Resume support
 - Allows you to stop a running Linux instance and later continue operations.
 - A suspended Linux instance does not require memory or processor cycles.
 - Gives you better performance, and resource utilization.
 - Read Chapter 37, Suspending and resuming Linux in the “Device Drivers, Features, and Commands on SUSE Linux Enterprise Server 11 SP1” , SC34-2595-01 manual.
 - > **There is a litany of restrictions and cautions there, as well as “how to” information. (Fate # 307017, IBM # 201901)**

Performance Improvements (4)

- FCP performance data reports
 - The item comprises facilities (included in s390-tools) which present the data stored by 'FCP - Performance data collection' in a human readable way.
 - Data presented comprises:
 - > **Performance relevant data about stack components as Linux devices, SCSI LUNs and HBA and storage controller information**
 - > **Per stack component: current values of relevant measurements as throughput, utilization and other applicable measurements**
 - > **Statistical aggregations (min, max average, histogram) of data associated with I/O requests such as size, latency per component and in total.**

Usability improvements

Usability Improvements

- Automatic IPL after dump
 - Similar in concept to z/VM's re-IPL after a CP abend
 - Increase availability, minimize downtime
 - Nice to have in combination with dump on panic (Fate # 306973, 307009, IBM # 201169, 201757)
- TTY terminal server over IUCV
 - Some Linux admins just can't get past the 3215 console interface when the network interface isn't working. This solves that problem
 - Requires a z/VM Linux guest to act as the terminal server.
 - Changes need to be made on each Linux client for access. (Fate # 307001, 307002, IBM # 201734, 201735)

Usability Improvements (2)

- Automatic scan of network devices
 - Introduces znetconf (included in s390-tools) to automatically configure network devices (load proper module, group and bring them online) (Fate # 307016, IBM # 201900)
- Add vmconvert option to vmur tool
 - Allows you to convert on-the-fly a z/VM dump to Linux dump format while reading from the spool and writing to disk. (Fate # 307010, IBM # 201758)
- Provide hardware and hypervisor service levels in Linux
 - Could be helpful in problem diagnosis
 - Introduces a new pseudo-file, /proc/service_levels (Fate # 307008, IBM # 201753)

Usability Improvements (3)

- Report status of DS8000 disk encryption
 - Allows you, without having to ask the storage administrator, if the physical disks where your data is being stored are encrypted or not. (Fate # 307004, IBM 201740)
- FCP/SCSI error recovery hardening
 - Improves error recovery cooperation between SCSI-mid-layer and zFCP by allowing the SCSI error recovery to wait for completion of the zFCP error recovery. (Fate # 307006, IBM # 201743)
- Even more kernel message man pages
 - Now installed by default (Fate # 306999, 307018, IBM # 201727, 201969)

Usability Improvements (4)

- A read-only root file system is now a supported configuration.
 - I haven't had the time to figure out just exactly what this entails.
 - Supposedly based on the IBM Redpaper at <http://www.redbooks.ibm.com/abstracts/redp4322.html>
 - > **I haven't really read it or tried to implement what's documented there.**
 - > **I do know that some customers have tried, and had many problems.**
 - I'm trying to get a project going with a customer willing to experiment to determine just how that might work.
 - The intent is to do this across all architectures, not just System z. This may or may not limit what we can do on any particular architecture. (Fate # 307235, IBM # 201983)

Software additions/upgrades/removals

Software Add/Upgrade/Removal

- libica 2.0.2 (Fate # 307007, IBM # 201744)
- mt_st 0.9b (Fate # 307019, IBM # 201972)

Miscellaneous changes

FCP Code Cleanup Stage 2

- Stage 2 of the code cleanup covers functional aspects of zFCP driver
- state machine for adapter, port and unit
- layering violation issues regarding request processing
- modify exchange-port, -config data to not use recovery semaphore
- kernel thread API adaption
- remove forced port recovery when port recovery is already running
- lock consolidation (config lock vs. other locks)
- Fate # 304038, IBM # 201546

Non-System z specific

Performance improvements

Performance Improvements

- Tunable sched_mc_powersavings=n framework
 - The power saving loadbalancer CONFIG_SCHED_MC attempts to run the workload in the system on a minimum number of CPU “packages” and tries to keep rest of the CPU packages idle for longer duration.
 - > **The existing implementation is very conservative and does not work effectively across different workloads, hence this change.**
 - IBM is considering enhancing cpuplugd with this functionality. Final results on effectiveness on System z are not yet available. (Fate # 307097, IBM # 201924)

Performance Improvements (2)

- Timer migration from idle CPUs
 - Timers from IRQ context (like the ones in device drivers) are “stuck” on the CPU where they were initialized.
 - This framework will help move all 'movable timers' using a sysctl interface.

Usability improvements

Usability Improvements

- Allow key acceptance in zypper when in non-interactive modes
 - The current version prompts the user to agree to importing a new key into the keyring when a new repository is added.
 - This eliminate those prompts during automated installs. (Fate # 307087, IBM # 201907)

Software additions/upgrades/removals

Software Add/Upgrade/Removal

- The kernel was updated to 2.6.32
 - Done for a number of good reasons, but hopefully a once in a lifetime event.
 - Extensive testing was done to ensure kernel application binary interface (kABI) didn't change.
 - > **This means ISV software shouldn't have to be re-certified.**
- SystemTap 0.9.7 (Fate # 307061, IBM # 201932)
- Java 6 SR6 (Fate # 307083, IBM # 201892)
- CIM SBLIM Java Client 2.0.9.2 (Fate # 307105, IBM # 201984)

Software Add/Upgrade/Removal (2)

- CIM SBLIM packages (Fate # 307089, IBM # 201909)
 - sblim-cmpi-base 1.5.9 (Bug fixes)
 - sblim-cmpi-devel 2.0.0 (CMPI 2.0 compliance)
 - sblim-cmpi-fsvol 1.5.0 (Licence change to EPL)
 - sblim-cmpi-network 1.4.0 (License change to EPL)
 - sblim-cmpi-nfsv3 1.1.0 (License change to EPL)
 - sblim-cmpi-nfsv4 1.1.0 (License change to EPL)
 - sblim-cmpi-params 1.3.0 (License change to EPL)
 - sblim-cmpi-sysfs 1.2.0 (License change to EPL)
 - sblim-cmpi-syslog 0.8.0 (License change to EPL)
 - sblim-gather 2.1.9 (KVM support via libvirt)
 - sblim-testsuite 1.2.5 (bug fixes)
 - sblim-wbemcli 1.6.0 (bug fixes)

Software Add/Upgrade/Removal (4)

- libcggroup 0.34
 - Numerous enhancements - configuration updates and classification updates
 - Numerous bug fixes
 - Newer API (iteration)
 - Improved tools and man pages (cgclear, cgcreate, etc) (Fate # 307064, IBM # 201982)

Post Service Pack 1 Enhancements

- Common Code Base:
 - Firefox
 - **SUSE Manager integration**
 - Ruby on Rails 2.3.8
 - PHP 5.2.14
 - **Unattended Upgrade** from SLE 10 to SLE 11
- Server
 - **OFED 1.5.2**
 - PureFTPd (Netware compatibility)
- Desktop
 - LibreOffice

Miscellaneous changes

Miscellaneous Changes

- CIM sfcb rebuilt with `–enable-slp`
 - Ensures IBM Director can work with SLES (Fate # 307088, IBM # 201908)
- CIM SBLIM-HBA
 - A set of Storage Management providers in SBLIM helps system management products. (Fate # 307091, IBM # 201917)
- fallocate support added to glibc
 - Having a direct call to fallocate improves cross-platform portability. (Fate # 307179, IBM # 201881)
- IPv6 support in NFS
 - Wasn't quite ready in time for SLES11. (Fate # 307240, IBM # 201026)

Questions

Backup Material

SUSE Linux Enterprise Kernels

- SUSE Linux Enterprise Server 9 (GA 08/2004)
Kernel 2.6.5, GCC 3.3.3, Service Pack 4 (GA 12/2007)
- SUSE Linux Enterprise Server 10 (GA 07/2006)
Kernel 2.6.16, GCC 4.1.0, Service Pack 4 (GA 03/2011)
- SUSE Linux Enterprise Server 11 (GA 03/2009)
Kernel 2.6.27, GCC 4.3.3
Service Pack 1 Kernel 2.6.32, GCC 4.3.4 (GA 06/2010)

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				

File Systems

Feature	Ext 3	reiserfs	XFS	OCFS 2	btrfs
Data/Metadata Journaling	•/•	○/•	○/•	○/•	N/A [2]
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	•	•	•	•	• [3]
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	Technology Preview

Provide service levels of HW and Hypervisor in Linux

Fate 307008 / [LTC 201753]

<http://www.ibm.com/>-> search developerworks

- **Identify hardware and MLC levels:** admin can query machine name, MCL level, zVM machine data and version information from a Linux guest

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Query machine a software information• Faster error isolation and issue resolution	<ul style="list-style-type: none">• Keep SLAs by faster problem resolution

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

TTY terminal server over IUCV

Fate 302001 / [LTC 201734]

<http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.1.html> -> iucvterm
Device Drivers, Features, and Commands as available with SUSE Linux Enterprise Server 11 p.219ff

- **Terminal access without IP:** Implement client & server for a TTY terminal server under z/VM using IUCV as communication vehicle.

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Optional full terminal to access a guest even if TCP/IP stack is down• Alternative to 3215/3270 block-mode based console access	<ul style="list-style-type: none">• Ease of use• Lower learning curve for administrators

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

Support Decimal Floating Point Library

Fate 307062 / [LTC 201933]

<http://www.ibm.com/developerworks/linux/linux390/toolchain.html> -> decimal floating

- **Decimal floating point hardware exploitation:** Most financial transactions require decimal computation. Supporting the emerging standard for Decimal Floating Point.

- **Customer benefit**

technical	business
<ul style="list-style-type: none">• Faster business related math operations due to hardware acceleration	<ul style="list-style-type: none">• Offloads processor from software emulation for the benefit of other workloads

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

Category: toolchain



= SLES unique feature



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.



What's new with SUSE® Linux Enterprise Server 11 SP1

for IBM System z

Mark Post

Technical Support Engineer IV
mpost@suse.com













Let's first review Novell's standard platform lifecycle policy.

Customers have told us that they want longer product lifecycles so that they can build, test and deploy stable, more durable solution stacks, especially for their mission critical workloads. For this reason, the typical lifecycle for a Novell Linux platform product spans 10 years.

- We provide general support for our Linux platform products for seven full years after their initial release.
- For 3 years after that, we continue to make available all our support content – things like technical informations documents (TIDs), briefs, and guides – in our customer facing support knowledge base

During the seven year general support period for a particular platform product, Novell will typically release 4 services packs, on average about 1 every 16-18 months.

- Typically contains all of the patches, fixes and updates issued since the previous release of the product
- Also includes product enhancements and drivers to support new hardware.

Novell strongly encourages its customers to upgrade to the most current service pack available, so that they are using what we know to be the most reliable, secure and best performing version for their enterprise.

- Novell provides an overlapping 6 month support period – a support window when the latest service pack, as well as its prior release, are both simultaneously supported.
- After this six month window, customers calling Novell about a particular support issue will typically be asked to apply all of the patches, fixes and updates required to bring them up to the current release, in order to efficiently and effectively diagnose and resolve their support issue.

Because some customers require more time than 6 months to upgrade to the latest service pack, Novell offers optional long term service pack support options.

- For an additional fee, Novell will extend the overlapping support window by an additional year, allowing customers a full 18 months to upgrade.

With that overview, let's review the present lifecycle status for SUSE Linux Enterprise (SLE).

- As a reminder - these milestones are our current planned dates, and subject to change.
- This is our best information at this time, and we provide you with these estimates so that you can make better informed plans and roll-out decisions.

Because Novell typically releases a major version of SLE every 3-4 years, at any given point in time, there may be 2 or 3 generally available and supported versions of the platform in the market.

- In fact today, customers with a valid subscription can choose to deploy SLE 11, 10 or 9, or any combination of the three,
- And they are entitled to receive support for all of the versions they deploy. (You should relate that to the information regarding the service packs the page before!! Supported version if updated to the latest service pack)

SLE 9 was first released in Aug 2004

- General support for SLE 9 is currently scheduled to end 7 years after its GA date, in July 2011.
- The current release of this platform is Service Pack 4 (SP4). Another service pack for version 9 is not planned. {note: many customers are using SLES 9 SP4, released in Dec 2007}
- Support information for this platform is scheduled to remain in our knowledge base until July 2014.

SLE 10 was first released in July 2006

- General support for version 10 will be offered until July 2013.
- SUSE Linux Enterprise 10 Service Pack 3 (SLE 10 SP3) was recently released in Oct 2009, and Service Pack 4 (SP4) is planned to be released in the first half of 2011. {note: many customers are using SLES 10 SP2, released in May 2008}

SLE 11 is our most current enterprise Linux platform, and was launched in Mar 2009.


- It is our most versatile, interoperable platform for mission critical computing, and represents the cumulative experience and combined expertise of our in-house engineers and extended community of open source developers.
- Development is presently underway on SUSE Linux Enterprise 11 Service Pack 1 (SLE 11 SP1), and is scheduled to be released later this year.
- SLE 11 is the platform where we will be concentrating the majority of our development efforts. Customers currently using version 9 or 10 should consider upgrading to SLE 11 SP1, to get maximum reliability, security, performance and innovation.




































[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] 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.”




Add/remove: This increases the uptime and availability of the Linux guest and the provided applications.

Disk subsystem: Performance analysts would like to have the same view into SCSI over Fibre Channel performance. This feature allows performance analysts to more quickly determine if anything in the SCSI subsystem is causing a performance problem or not.

Many of these features are available only when running SLES in an LPAR.


The existing OSA layer-2 support is utilized to enable HiperSockets layer-2. This includes IPv6 support for HiperSocket layer-2. This support is available only on z10, running Linux on system z in an LPAR or as a z/VM guest.



IUCV allows establishing point to point communication channels, either between 2 virtual machines or between a virtual machine and hypervisor services.

In IUCV terminology, the session between 2 end points is called a PATH. It is identified at each end by a Path ID which is only relevant to the virtual machine that owns the session end. A path is always a connected channel - meaning there are no connectionless paths.

This feature allows for data exchange between VM s without the need to protocol overhead, e.g. like TCPIP. The user is kind of allowed to specify his own protocol.



IUCV allows establishing point to point communication channels, either between 2 virtual machines or between a virtual machine and hypervisor services.

In IUCV terminology, the session between 2 end points is called a PATH. It is identified at each end by a Path ID which is only relevant to the virtual machine that owns the session end. A path is always a connected channel - meaning there are no connectionless paths.

This feature allows for data exchange between VM s without the need to protocol overhead, e.g. like TCPIP. The user is kind of allowed to specify his own protocol.



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.

