

SUSE® Linux Enterprise SP2 Overview

Ralf Dannert

Systems Engineer

rdannert@suse.com

OSL Technologietage
2012-09-12

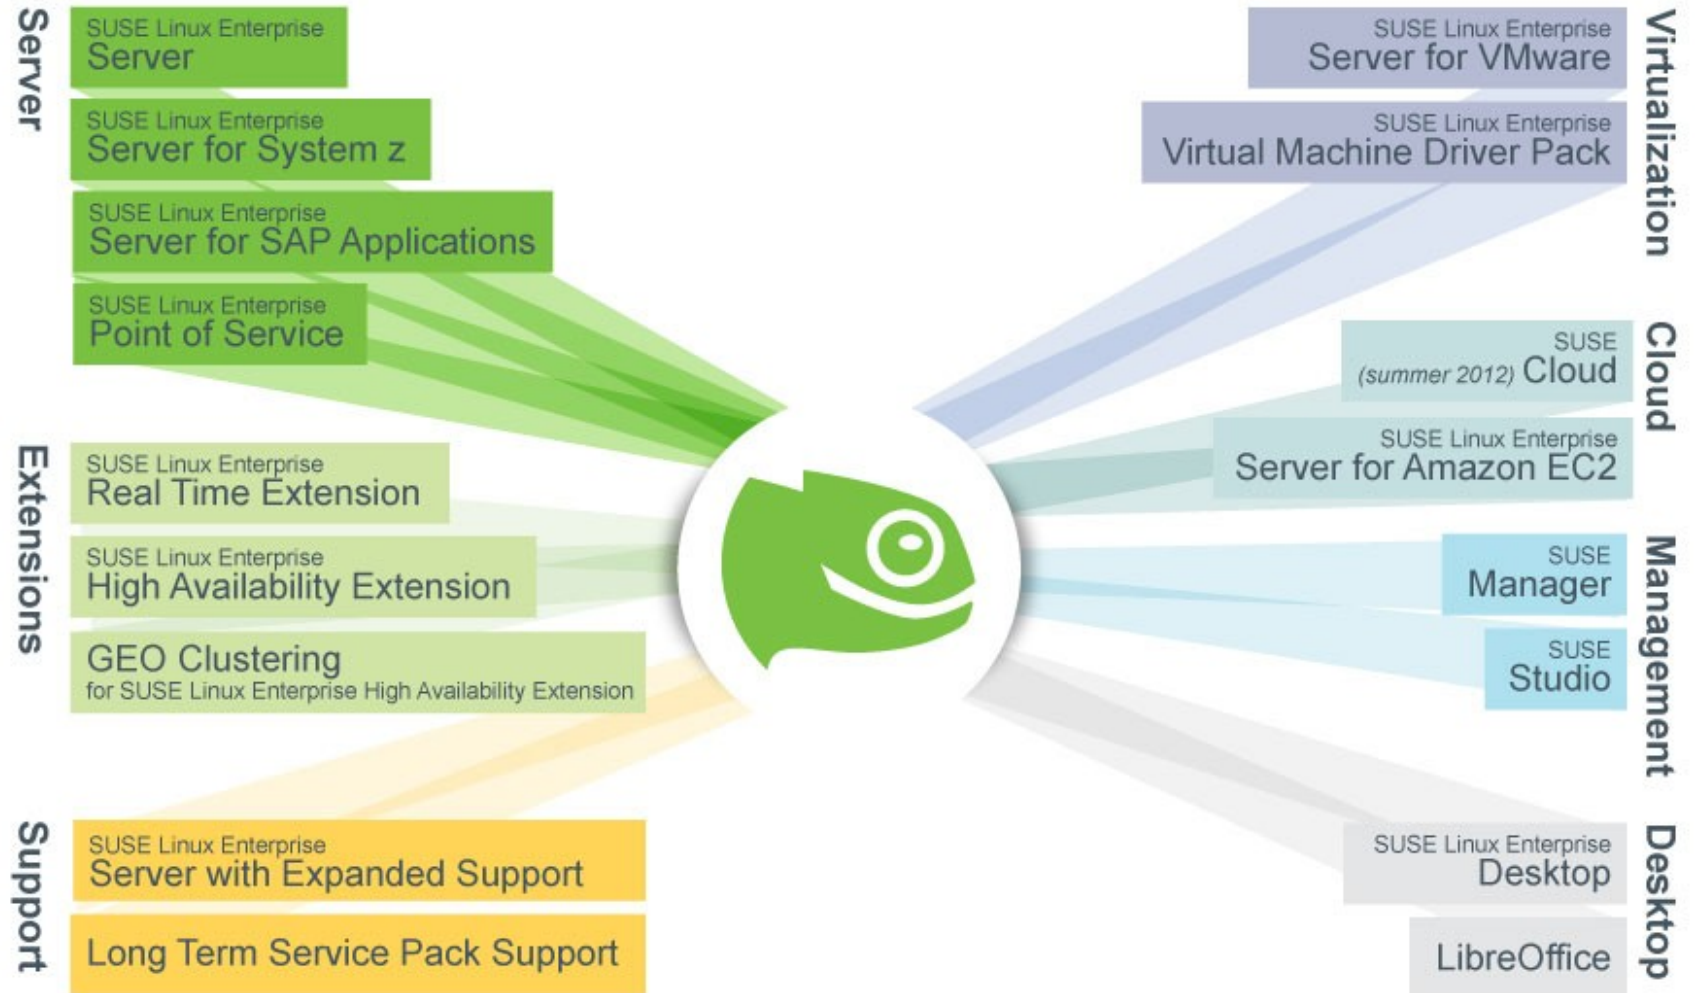


Agenda

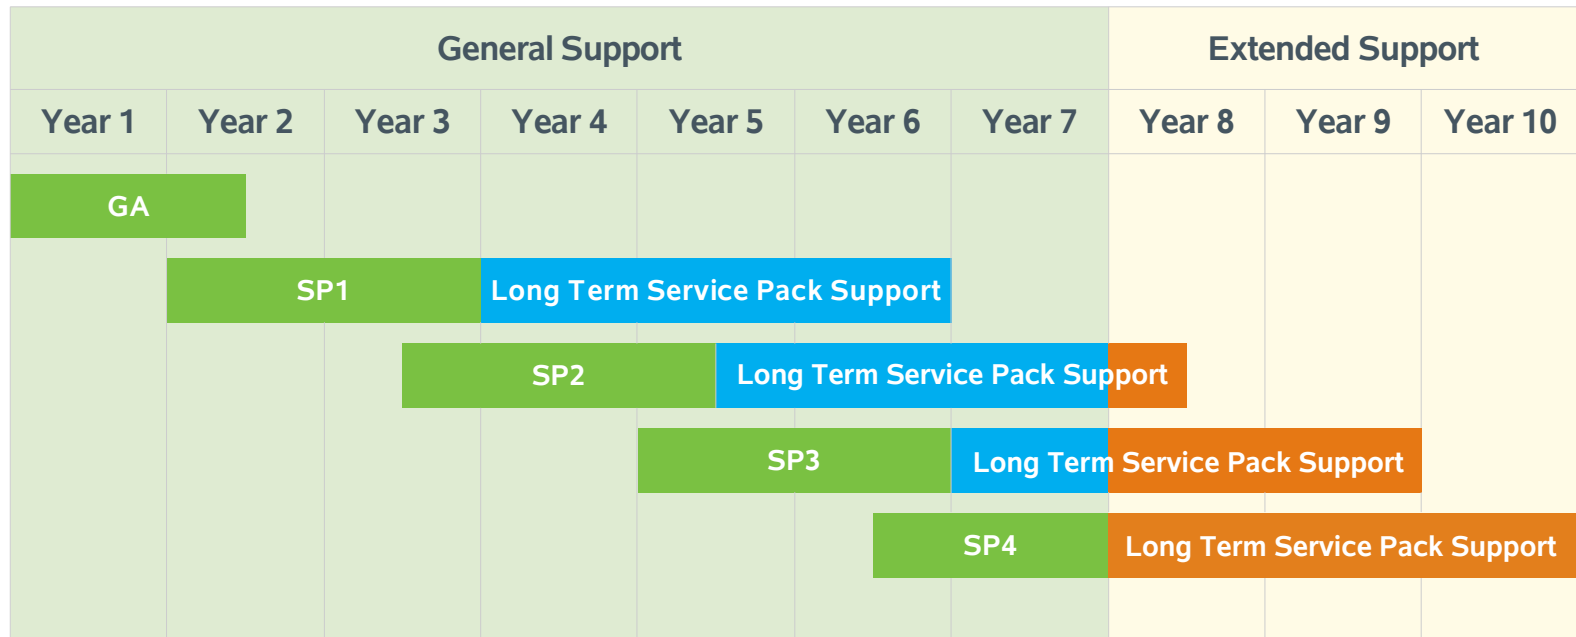
- Product Family
- SLE 11 SP2
 - Forward Looking Development Model
 - Kernel 3.0 (selected benefits)
 - Hardware enablement
 - Virtualization
 - Storage and Networking
 - Performance, Scalability, Reliability
 - Desktop
 - interoperability
 - High Availability
 - Security
 - Systems Management
 - Match and exceed Solaris

Product Overview

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 Platform Lifecycle

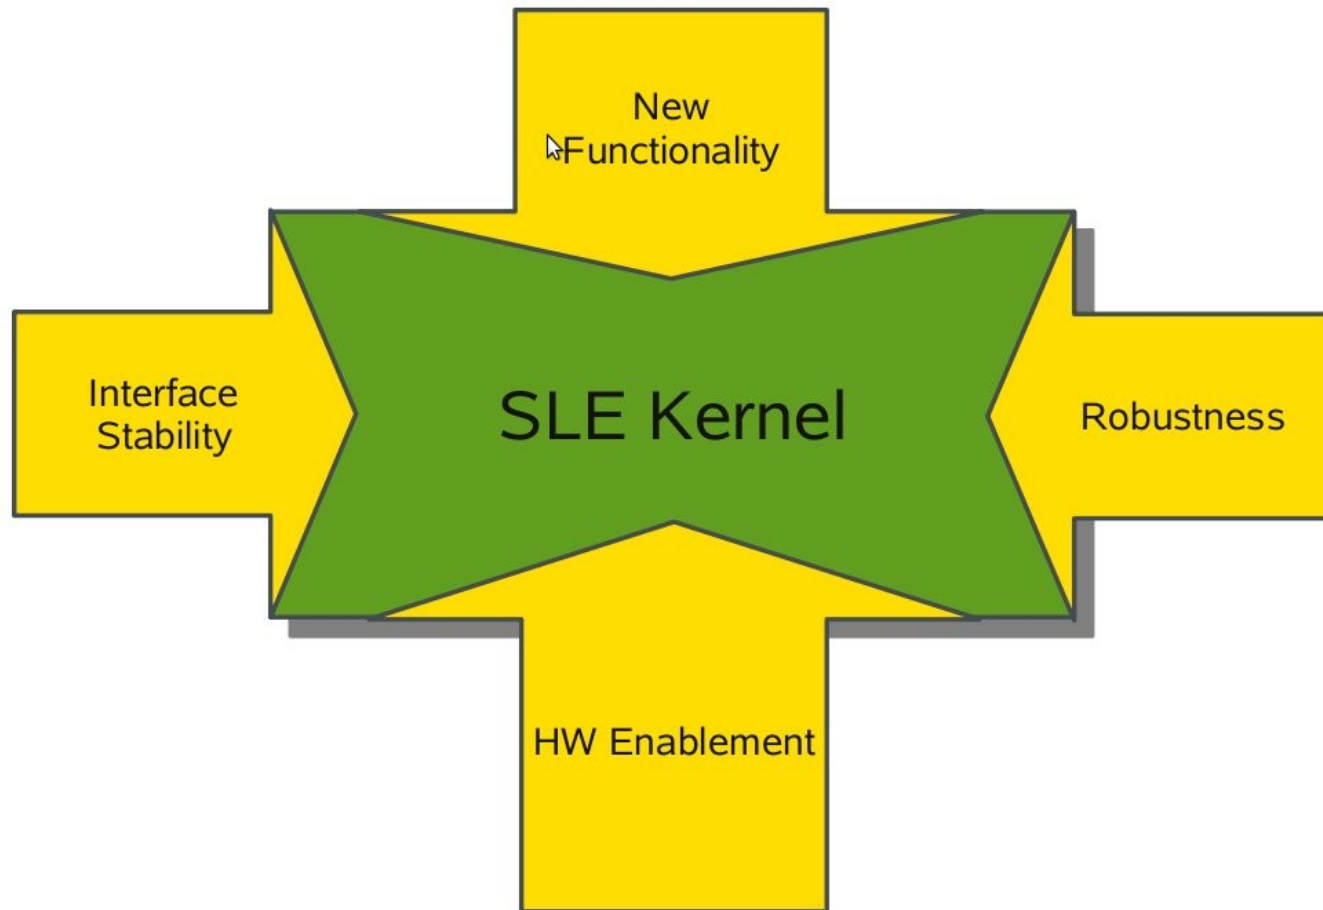


- SUSE announces service pack releases and development and product schedules to customers and partners
- Dependable release timing
- Predictability for planning rollouts and migrations
- Major releases every 4-5 years.

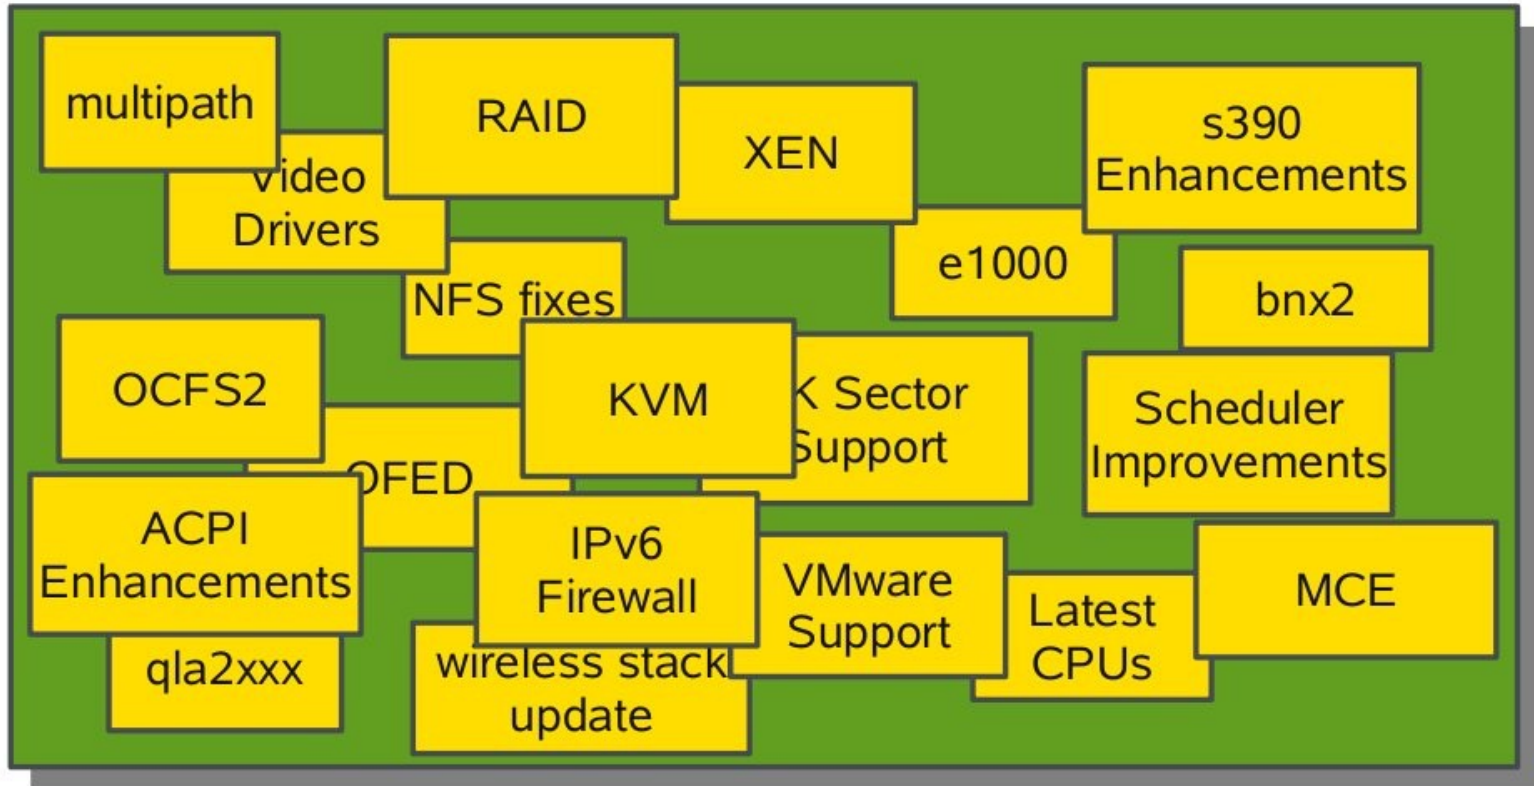


SUSE Linux Enterprise 11 SP2

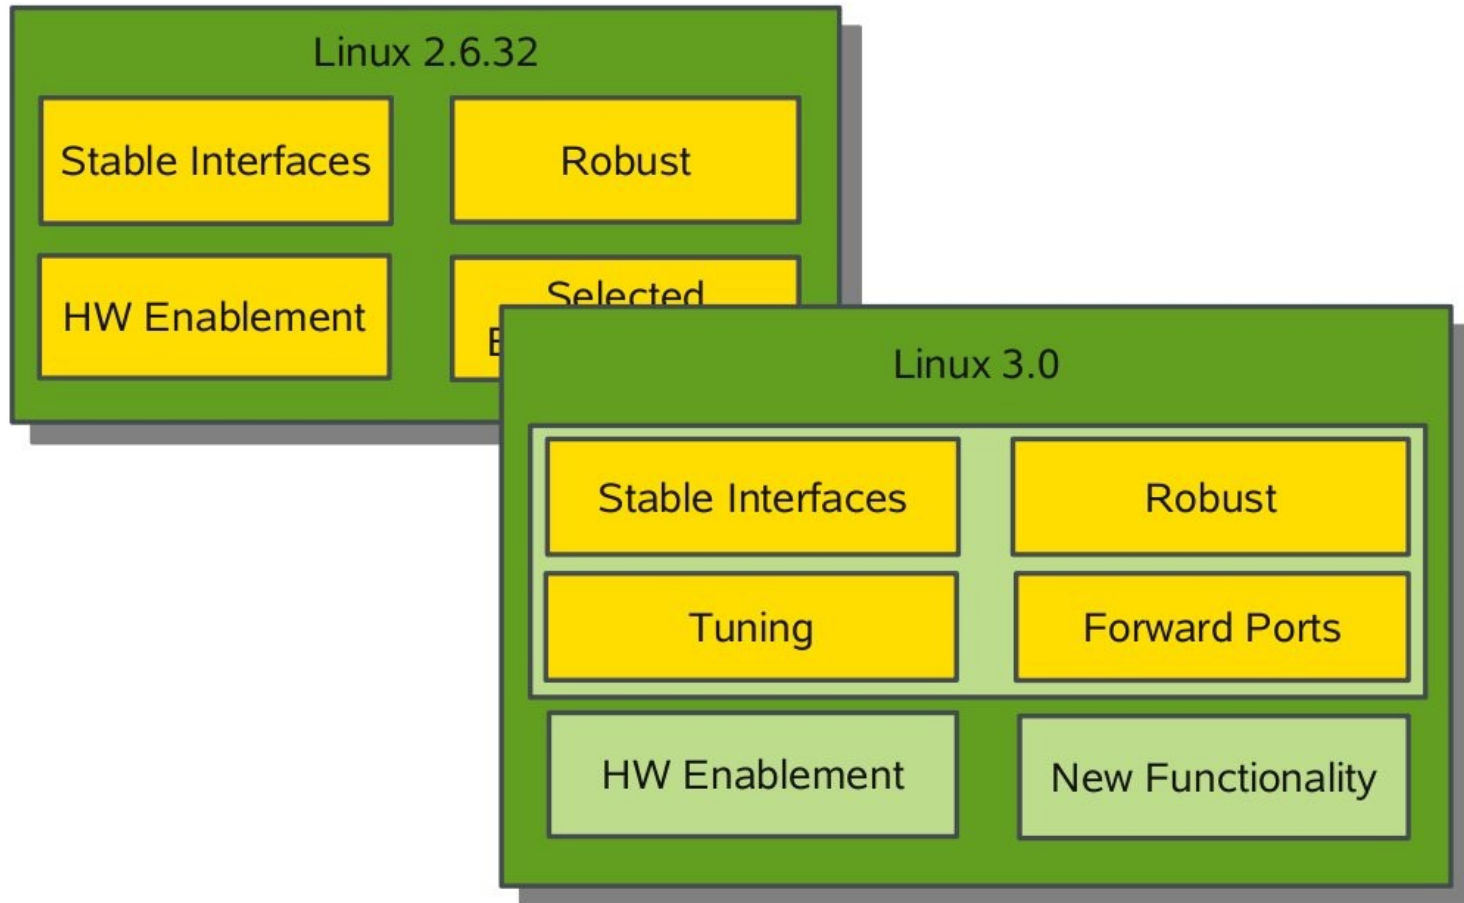
Four Forces of a Service Pack Kernel



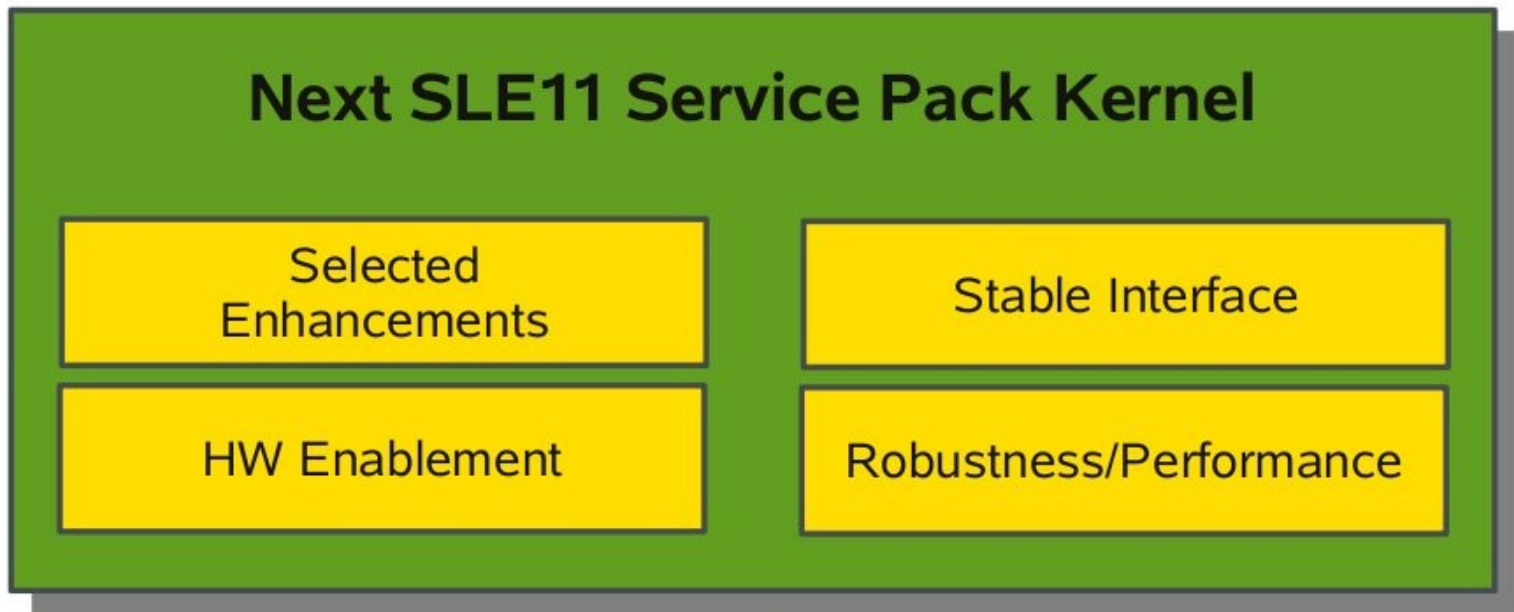
Kernel Components Changed in a Service Pack



What does this change mean?



How do we make this all work?



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

Example 1: Kernel / Process Scheduler

- Kernel needs to know about HW architectures and their internals
 - Optimize process-CPU-scheduling
 - Optimize use of L2/L3 caches
 - Examples: AMD “Bulldozer”, Intel “Sandybridge”
 - Customer benefits
 - Fully exploit most recent hardware development
 - Higher performance / less power consumption
- SUSE Linux Enterprise 11 SP2 automatically gets this with Kernel 3.0

Example 2: Interoperability / samba

- SUSE Linux Enterprise 11 SP1: Samba 3.4
- SUSE Linux Enterprise 11 SP2: Samba 3.6

Benefits of Samba 3.6:

- Support for SMB2 protocol
- Changed ID Mapping algorithms
- Changed security defaults

Changes introduce risk of breaking existing installations

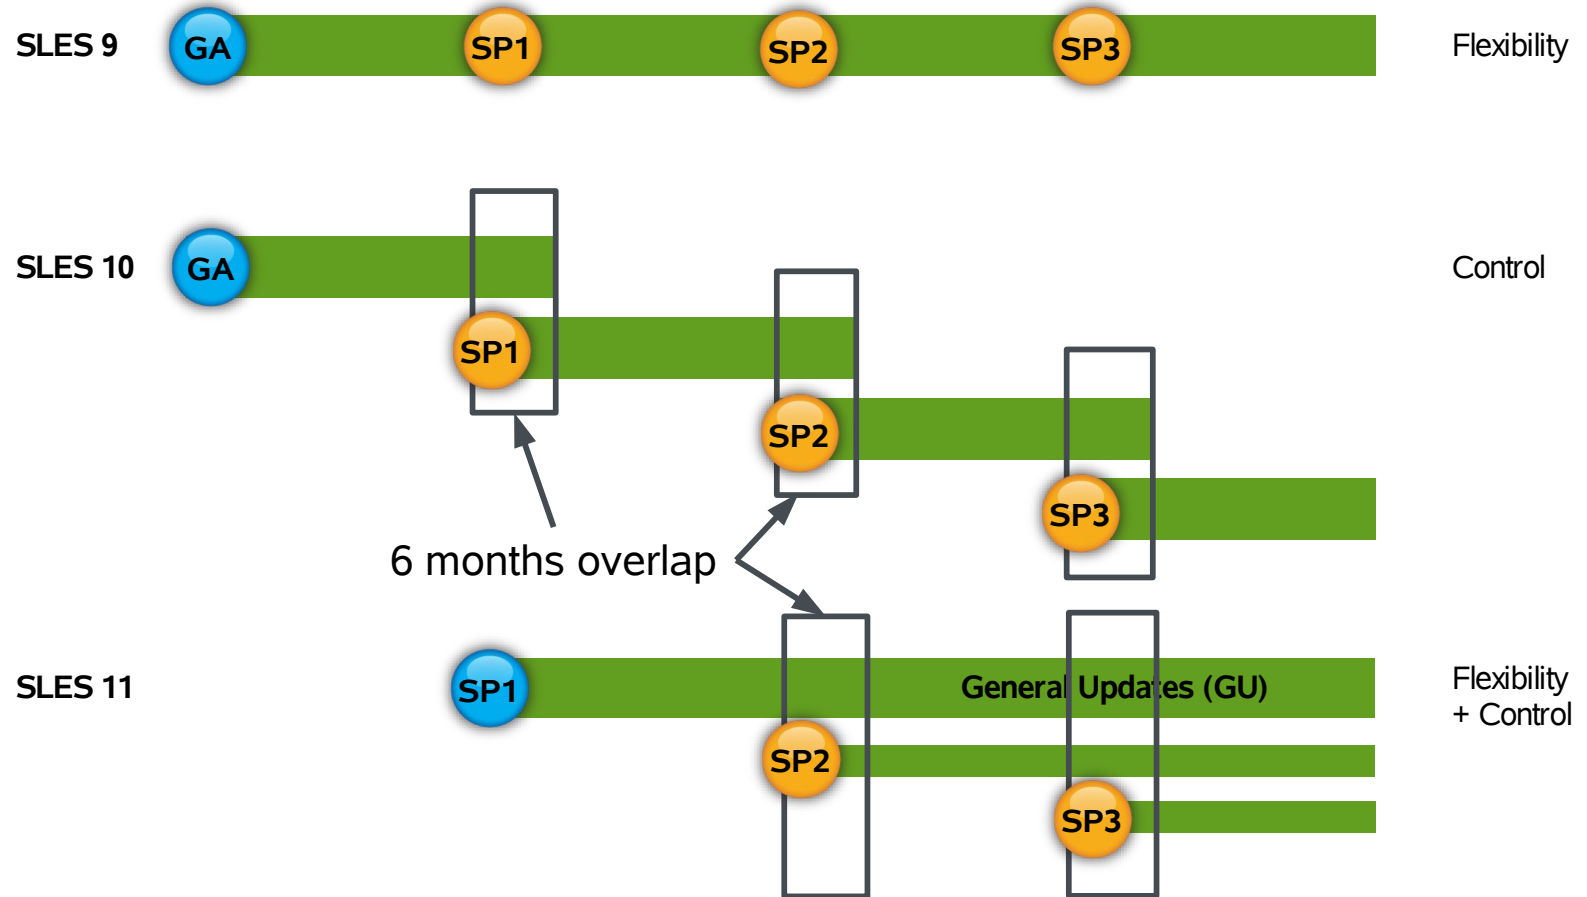
→ Samba 3.6 only available in SLE 11 SP2

Example 3: Certifications / openssl

- OpenSSL is a base library of every Linux system
 - SUSE Linux Enterprise 11 SP1: OpenSSL-0.9.8i
 - Latest upstream version: OpenSSL-1.0.0
 - Issue: OpenSSL-1.0.0 highly incompatible to 0.9.8
 - Requirements
 - Implement infrastructure needed for upcoming FIPS 140-2 certification
 - Do not change API or ABI
 - Approach
 - Start with OpenSSL-0.9.8j
 - Adopt ABI to become fully backward compatible
- Maintenance Update to SUSE Linux Enterprise 11 SP1 (and SP2 respectively): OpenSSL-0.9.8j



SUSE Linux Enterprise 11 Maintenance Delivery Evolution



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 1-5	Year 6-7		
Technical Support	yes	yes	yes	yes
Access to Patches and Fixes	yes	yes	yes	yes
Access to Documentation and Knowledgebase	yes	yes	yes	yes
Support for existing stacks and workloads	yes	yes	yes	yes
Support for new deployments	yes	limited (1)	limited (1)	no
Enhancement Requests	yes	limited (1)	no	no
Hardware Enablement/Optimization	yes	no	no	no
Driver Updates via PLDP	yes	limited (1)	limited (1)	no
Backport of fixes from more recent Service Pack	yes	n/a	limited (2)	n/a
Critical Security Updates	yes	yes	yes	yes
Defect Resolution	yes	limited (2)	limited (2)	limited (2)
(1) Based on partner and customer requests.				
(2) Severity Level 1 and 2 defects only.				

SUSE Maintenance Model Benefits

- **Deliver new functionalities in between Service Packs**
 - Starting with SLES 11 SP1, new features can be delivered anytime, making the SLES products competitive
- **Usage of older version of packages**
 - Customers can pick and use any compatible package available in the SP 1 Updates repository
 - The update of a package in SP1 update repositories might trigger the installation of other packages within this repository or in the SP2 repository.
- **Lightweight Service Packs**
 - Customers can choose between a minimal and a full migration
 - Minimal migration = install all the packages available in SP2 Updates repository
 - Full migration = install all the packages available in SP2 and SP1 Updates repository



Partner Linux Driver Process (PLDP)

- What?
 - Allow partners to **package and provide kernel device drivers** for SUSE Linux Enterprise – independent of SUSE releases
 - Fully supported due to mutual support agreements
- Why?
 - Quick hardware enablement
- How?
 - Build infrastructure (Open Build Service)
 - Developer support
 - <http://drivers.suse.com/>

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

Kernel 3.0 (selected benefits)

- Most recent hardware enablement
 - including USB 3.0
- Removal of BKL (Big Kernel Lock)
- Transparent Huge Pages (THP)
- Control Groups enhancements
 - I/O throttling support for process groups
 - memory cgroup controller optimization
- Upstream Integration of AppArmor
- Network:
 - Transparent per-CPU network load balancing on multiqueue devices (up to 20% performance increase)
 - initcwnd/initrwnd
 - More powerful firewalls based on faster packet filtering
- Support for fanotify, allowing clean antivirus solutions



Hardware enablement

- Intel Sandybridge and "Romley"
- AMD Bulldozer ...
- AVX support (GCC 4.6 in SLE SDK 11 SP2)
- IBM System z and POWER exploitation for new hardware and systems
- General RAS enhancements in AMD64/Intel64 architecture, e.g.:
 - CPU offlining
 - Memory Offlining

Virtualization

- **KVM 0.15**
 - I/O improvements, storage and network device hotplugging
 - **Microsoft Windows support**
- **Xen 4.1**
 - Latency improvements, flexible partitioning
 - Better fault handling, improved scalability and performance
- **Perfect Guest** strategy, tuned to run great as a guest on all major hypervisors, including:
 - Microsoft Hyper-V
 - VMware ESX
 - Citrix XenServer
 - SUSE Linux Enterprise with Xen and KVM
 - LPAR and z/VM for IBM System z

Virtualization

- Virtual machine driver pack: paravirtualized drivers
 - Increase performance, improve reliability and stability
- SUSE Appliance Program and **SUSE Studio™** for easy building, testing, distributing and managing of stacks
- **Unique cost efficiency** – unlimited number of virtual machines per physical server with one subscription
- **Container technology** / soft partitioning
 - Kernel resource management (control groups, **LXC**)

Virtualization Future

- Virtualization host
 - Improve interoperability and establish virtual machine migration between Xen and KVM for smooth migration
 - KVM inherits from the Linux kernel
 - Power saving capabilities and
 - Scalability improvements
- Container technology:
 - More detailed control and tuning for single processes and process groups

KVM - tuning

- high performance networking: switch to a bridged setup
- Storage:
 - raw partition instead of raw and qcow2
 - Avoid IDE, use if=virtio
 - Avoid btrfs on host for image files(low IO performance)
- Networking:
 - avoid caching mode(reduce data copies and bus traffic)
- Snapshots
- Derived images qcow2
- KSM: `/usr/src/linux/Documentation/vm/ksm.txt`
 - `/sys/kernel/mm/ksm/`

Storage and Networking

- Maintain existing capabilities (e.g. XFS)
- FCoE integration
- iSCSI improvements
- IPv6 readiness and certification
 - IPv6 ready logo 2 / USGv6 v.10
 - Achieved for SLES 11 SP1, refresh planned for SP2
 - incl. IPv6 IPSECv 1.2 , IKE v1.1 & v2.0
- pNFS client support (delivery in 2012)
- Update to OFED 1.5.2, including
 - Better hardware support
 - More features in the Infiniband stack
 - Improved support for 10G networks

Performance, Scalability, Reliability

- Keep and extend high level of scalability in terms of
 - Number of CPUs (4096)
 - Filesystem and Storage: XFS (8 EiB), BTRFS (16 EiB)
 - Memory (16 TiB and beyond)
- Improve support for hardware based RAS capabilities
- Scheduler and Memory management optimizations
- Faster and more powerful control groups, which allow for resource isolation between process groups.
- Linux Container support for soft partitioning of large physical systems (LXC)

Interoperability

- Remote Desktop Protocol (RDP, Windows environments)
 - Support xrdp server on AMD64/Intel64 (SLES)
 - Introduce FreeRDP client: better Performance and Compatibility (SLED only)
- CIFS (also known as SMB)
 - Samba 3.6: SMB2 protocol support
 - Improve CIFS kernel module
- IPv6
 - Improve DHCP server and client (ISC dhcp)
 - IPv6 support in NFS
 - IPv6 capable Squid proxy (most other network daemons are IPv6 ready since SLE 10 times)
 - Ensure IPv6 capabilities with UEFI network boot

High Availability

- GEO Cluster
- Management
 - Guided wizard / Templates
 - HA boot strap
 - Cluster Test Drive
 - Granular Cluster Access, Map customer organization to access roles
 - Enhanced scalability for the web based cluster management
- Supportability
 - History Explorer
- Cloud Integration
 - Automatic discovery for CPU & memory requirements of VMs
 - Load limiting for concurrent live migrations



SUSE® Linux Enterprise High Availability Extension 11 SP1

Key Capabilities Today

- **Service Availability 24/7**
 - Policy driven clustering
 - OpenAIS messaging and membership layer
 - Pacemaker cluster resource manager
- **Sharing and Scaling Data-access by Multiple Nodes**
 - **OCFS2 Clusterfile system**
 - Clustered logical volume manager
- **Disaster Tolerance**
 - Continuous data replication via IP
 - Improved Distributed replicated block device (**DRBD**) capabilities
- **Scale Network Services**
 - IP load-balancing
- **User-friendly Tools**
 - Graphical user interface / Unified command line interface



SUSE[®] Linux Enterprise High Availability Extension 11 SP1

Key Capabilities Today

- **Storage Quorum Coverage**
 - Enabling the use of a storage device as a quorum instance to match traditional Unix setups and to prevent split brain scenarios
- **Integrated Samba Clustering**
 - Integration of Samba with OCFS2 for higher throughput and scale out of SMB access
- **Metro-Area Clusters**
 - Supporting clustering between different data center locations
- **DR Framework**
 - A tool set (ReaR) for node recovery



Security

- Answer customer demand for security certifications: SUSE Linux Enterprise Server is in evaluation for a Common Criteria Certification in Evaluation Assurance Level 4 with augmentation according to the BSI OSPP (CC/OSPP EAL 4+)
- Enhanced Hardening Guide

Security and Certifications Tomorrow

- **Common Criteria Certification**
in Evaluation Assurance Level 4 with augmentation according to the BSI OSPP (CC/OSPP EAL 4+)
- **FIPS 140-2 Certification**
 - “libopenssl” module
 - tentative/future modules based on customer demand
- **Trusted Execution Technology (TXT)**
 - Enhances Trusted Computing with processor-based separation functions on a page-level in memory
 - Keep systems in consistent and proven (“measured”) state
 - in virtual environments and cloud

Systems Management

- Unattended upgrade from SUSE Linux Enterprise 10 to SUSE Linux Enterprise 11
 - This is a key differentiator: other Enterprise Linux distributions do not support this, not even manually
- SUSE developed open source "snapper" tool to manage btrfs snapshots – including:
- YaST2 and zypper integration with snapper:
 - Rollback changes to the system, which have been unwanted (administrator error) or did show unwanted results or side effects
- Allow to work with several kernels installed in parallel implemented in ZYpp

Match and exceed Solaris (1)

- Tracing
 - LTTng
 - to be delivered after SLES 11 SP2 in an extra Extension
- LXC - Linux Container support
 - So called “OS Level virtualization”
 - Use Cases
 - Hosting business
 - Datacenter use
 - Outsourcing business
 - Built upon Linux Kernel Control Groups and specific settings
 - Build “System Containers” for SLES 11 SP2 on SLES 11 SP2
 - Support for “Application Containers” planned for a later release
- SWAP over NFS (since SLES 11 SP1)
 - SLES is the only Enterprise Linux offering this functionality

Match and exceed Solaris (2) – btrfs

- btrfs
 - Support for Copy on Write
 - Integrated Volume Management
 - Powerful Snapshot capabilities
 - Scalability (16 EiB) including effective shrink
 - Supports offline in-place migration from ext2+
 - Other Capabilities:
 - Compression
 - Data integrity (checksums)
 - SSD optimization
 - Btrfs builds the foundation for the Ceph distributed filesystem and its RADOS object store layer → Cloud
- snapper
 - Single file rollback (“undochange”)
 - User interface: cmdline and YaST2 integration
 - Snapshots for YaST2 and zypper activities
 - Automated snapshots (time based) and snapshot cleanup



SUSE Linux Enterprise Server for SAP Applications

Key Capabilities

- Installation Wizard for OS and SAP installation
 - Simplified installation of Operating System and SAP solution with minimized manual interaction
- “Vanilla” Operating System installation
 - For custom deployments of SAP solutions
- Configuration and package selection based on SAP requirements
- Supported Page Cache Feature
 - Allows for better performance of memory intensive applications with rarely accessed memory pages
- Extended Service Pack Overlap Support
 - Extends service pack upgrade window to 18 months

Appendix

Snapshots in SUSE® Linux Enterprise 11 SP2

snapper – example

2011-08-18 12:33 (0) ~

ios root (0) # snapper list

Type	#	Pre #	Date	Cleanup	Description
single	0				current
single	1		Wed 17 Aug 2011 04:30:01 PM CEST	timeline	timeline
pre	2		Wed 17 Aug 2011 04:31:54 PM CEST	number	yast lan
post	3	2	Wed 17 Aug 2011 04:32:46 PM CEST	number	
pre	4		Wed 17 Aug 2011 04:32:48 PM CEST	number	yast lan
post	5	4	Wed 17 Aug 2011 04:32:59 PM CEST	number	
pre	6		Wed 17 Aug 2011 04:36:10 PM CEST	number	zypp(zypper)
post	7	6	Wed 17 Aug 2011 04:36:11 PM CEST	number	
pre	8		Wed 17 Aug 2011 04:36:16 PM CEST	number	zypp(zypper)
post	9	8	Wed 17 Aug 2011 04:36:19 PM CEST	number	
pre	10		Wed 17 Aug 2011 04:36:26 PM CEST	number	yast printer
post	11	10	Wed 17 Aug 2011 04:37:21 PM CEST	number	
single	12		Wed 17 Aug 2011 05:30:01 PM CEST	timeline	timeline
single	13		Wed 17 Aug 2011 06:30:01 PM CEST	timeline	timeline
single	14		Wed 17 Aug 2011 07:30:01 PM CEST	timeline	timeline
single	15		Wed 17 Aug 2011 08:30:01 PM CEST	timeline	timeline
single	16		Wed 17 Aug 2011 09:30:01 PM CEST	timeline	timeline
single	17		Wed 17 Aug 2011 10:30:01 PM CEST	timeline	timeline
single	18		Wed 17 Aug 2011 11:30:01 PM CEST	timeline	timeline
single	19		Thu 18 Aug 2011 11:30:02 AM CEST	timeline	timeline
single	20		Thu 18 Aug 2011 12:30:01 PM CEST	timeline	timeline

Snapshots in SUSE Linux Enterprise 11 SP2

Snapshotting “/” – Partitions/Subvolumes

The screenshot displays the YaST2 Expert Partitioner interface. The top window shows the 'System View' on the left and a table of 'Available Storage on ios' on the right. The table lists various partitions and subvolumes with their sizes, file systems, and mount points.

Device	Size	F	Enc	Type	FS Type	Label	Mount Point	Mount By	Used By
/dev/sda	1.82 TB			WDC-WD2002FYPS-0					
/dev/sda1	1.82 TB			Extended					
/dev/sda5	155.00 MB			Linux native	Ext3	boot	/boot	Label	
/dev/sda6	2.01 GB			Linux swap	Swap	swap	swap	Label	
/dev/sda7	20.00 GB			Linux native	Btrfs		/	UUID	BTRFS ef6d384d-b9e4-444c-a5fa-c2380dd8b93d
/dev/sda8	1.80 TB			Linux native	Btrfs		/testing	UUID	BTRFS aece9681-8263-4fb3-ab00-8224a8f5b6e9
tmpfs	1.80 GB			TMPFS	TmpFS		/dev/shm	Kernel	
tmpfs	1.80 GB			TMPFS	TmpFS		/dev	Kernel	

The bottom-left window is titled 'Edit Btrfs ef6d384d-b9e4-444c-a5fa-c2380dd8b93d' and shows 'Formatting Options' and 'Mounting Options'. The 'Mounting Options' section has 'Mount partition' selected, with the 'Mount Point' set to '/'. A 'Subvolume Handling' button is visible at the bottom of this window.

The bottom-right window is titled 'Subvolume Handling' and shows a list of 'Existing Subvolumes' in a tree view:

- @
- @/tmp
- @/opt
- @/srv
- @/var/spool
- @/var/log
- @/var/run
- @/var/tmp
- @/home
- @/.snapshots
- @/.snapshots/2/snapshot
- @/.snapshots/3/snapshot

Below the list is a 'New Subvolume' input field and buttons for 'Add new', 'Remove', 'OK', 'Cancel', and 'Help'.

Your Thoughts?

Thank you.





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.

