




Release Notes for SUSE Linux Enterprise Server 11 Service Pack 4 (SP4)



Release Notes for SUSE Linux Enterprise Server 11 Service Pack 4 (SP4)

ABSTRACT

This document provides guidance and an overview to high level general features and updates for SUSE Linux Enterprise Server 11 Service Pack 4 (SP4). Besides architecture or product-specific information, it also describes the capabilities and limitations of SLES 11 SP4. General documentation may be found at: <http://www.suse.com/documentation/sles11/> .

Publication date: 2015-06-19 , Version: Version 11.4.16 (2015-06-18)

SUSE Linux Products GmbH
Maxfeldstr. 5
90409 Nürnberg
GERMANY




<http://www.suse.com/documentation> 

Contents

1	How to Obtain Source Code	1
2	SUSE Linux Enterprise Server	2
3	Important Upgrade Information	3
4	Support Statement for SUSE Linux Enterprise Server	5
5	Installation	12
6	Features and Versions	17
7	Driver Updates	29
8	Other Updates	34
9	Software Development Kit	39
10	Update-Related Notes	41
11	Deprecated Functionality	52
12	Infrastructure, Package and Architecture Specific Information	55
13	Resolved Issues	87
14	Technical Information	88

15	Documentation and Other Information	102
16	Miscellaneous	103
17	Legal Notices	104

1 How to Obtain Source Code

This SUSE product includes materials licensed to SUSE under the GNU General Public License (GPL). The GPL requires SUSE to provide the source code that corresponds to the GPL-licensed material. The source code is available for download at <http://www.suse.com/download-linux/source-code.html> . Also, for up to three years after distribution of the SUSE product, upon request, SUSE will mail a copy of the source code. Requests should be sent by e-mail to mailto:sle_source_request@suse.com  or as otherwise instructed at <http://www.suse.com/download-linux/source-code.html> . SUSE may charge a reasonable fee to recover distribution costs.

2 SUSE Linux Enterprise Server

SUSE Linux Enterprise Server is a highly reliable, scalable, and secure server operating system, built to power mission-critical workloads in both physical and virtual environments. It is an affordable, interoperable, and manageable open source foundation. With it, enterprises can cost-effectively deliver core business services, enable secure networks, and simplify the management of their heterogeneous IT infrastructure, maximizing efficiency and value.

The only enterprise Linux recommended by Microsoft and SAP, SUSE Linux Enterprise Server is optimized to deliver high-performance mission-critical services, as well as edge of network, and web infrastructure workloads.

Designed for interoperability, SUSE Linux Enterprise Server integrates into classical Unix as well as Windows environments, supports open standard CIM interfaces for systems management, and has been certified for IPv6 compatibility,

This modular, general purpose operating system runs on five processor architectures and is available with optional extensions that provide advanced capabilities for tasks such as real time computing and high availability clustering.

SUSE Linux Enterprise Server is optimized to run as a high performing guest on leading hypervisors and supports an unlimited number of virtual machines per physical system with a single subscription, making it the perfect guest operating system for virtual computing.

SUSE Linux Enterprise Server is backed by award-winning support from SUSE, an established technology leader with a proven history of delivering enterprise-quality support services.

With the release of SUSE Linux Enterprise Server 11 Service Pack 4 the former SUSE Linux Enterprise Server 11 Service Pack 3 enters the 6 month migration window, during which time SUSE will continue to provide security updates and full support and maintenance. At the end of the six-month parallel support period, on 2015-mm-dd, general support for SUSE Linux Enterprise Server 11 Service Pack 3 will be discontinued. Long Term Service Pack Support (LTSS) for SUSE Linux Enterprise Server 11 Service Pack 2 is available as a separate add-on.

3 Important Upgrade Information

3.1 What's New in SUSE Linux Enterprise Server 11 SP4

- New CPU enablement, such as Intel® Xeon® processor E7-8800/4800 v3 product family, IBM z13™ (z13), and IBM POWER8 BE.
- Public Cloud module and Security module are now available for SP4. These modules are independent repository channels and are included in subscription without additional cost:

Public Cloud Module

The Public Cloud Module is a collection of tools that enables you to create and manage cloud images from the command line on SUSE Linux Enterprise Server. When building your own images with KIWI or SUSE Studio, initialization code specific to the target cloud is included in that image. The tools and initialization code in this module will be updated whenever a new version is ready, always giving you the freshest.

Security Module


The Security Module was introduced a while ago and it allows customers and partners to build TLS 1.2 compliant infrastructures beyond the https protocol.

3.2 Upgrade Information


For users upgrading from a previous SUSE Linux Enterprise Server release it is recommended to review:

- *Chapter 4, Support Statement for SUSE Linux Enterprise Server*
- *Chapter 10, Update-Related Notes*
- *Chapter 14, Technical Information*

Installation Quick Start and Deployment Guides can be found in the docu language directories on the media. Documentation (if installed) is available below the /usr/share/doc/ directory of an installed system.

These Release Notes are identical across all architectures, and the most recent version is always available online at <http://www.suse.com/releasenotes/> . Some entries are listed twice, if they are important and belong to more than one section.

4 Support Statement for SUSE Linux Enterprise Server

To receive support, customers need an appropriate subscription with SUSE; for more information, see <http://www.suse.com/products/server/services-and-support/> .

4.1 General Support Statement

The following definitions apply:

L1

Problem determination, which means technical support designed to provide compatibility information, usage support, on-going maintenance, information gathering and basic troubleshooting using available documentation.

L2

Problem isolation, which means technical support designed to analyze data, duplicate customer problems, isolate problem area and provide resolution for problems not resolved by Level 1 or alternatively prepare for Level 3.

L3

Problem resolution, which means technical support designed to resolve problems by engaging engineering to resolve product defects which have been identified by Level 2 Support.

For contracted customers and partners, SUSE Linux Enterprise Server 11 will be delivered with L3 support for all packages, except the following:

- technology previews
- sound, graphics, fonts and artwork
- packages that require an additional customer contract
- packages provided as part of the Software Development Kit (SDK)

SUSE will only support the usage of original (e.g., unchanged or un-recompiled) packages.

4.1.1 Supportconfig include dmidecode output

The supportconfig tool is now include the dmidecode output.

4.1.2 Support for the btrfs File System

Btrfs is a copy-on-write (CoW) general purpose file system. Based on the CoW functionality, btrfs provides snapshotting. Beyond that data and metadata checksums improve the reliability of the file system. btrfs is highly scalable, but also supports online shrinking to adopt to real-life environments. On appropriate storage devices btrfs also supports the TRIM command.

Support

With SUSE Linux Enterprise 11 SP2, the btrfs file system joins ext3, reiserfs, xfs and ocfs2 as commercially supported file systems. Each file system offers distinct advantages. While the installation default is ext3, we recommend xfs when maximizing data performance is desired, and btrfs as a root file system when snapshotting and rollback capabilities are required. Btrfs is supported as a root file system (i.e. the file system for the operating system) across all architectures of SUSE Linux Enterprise 11 SP2. Customers are advised to use the YaST partitioner (or AutoYaST) to build their systems: YaST will prepare the btrfs file system for use with subvolumes and snapshots. Snapshots will be automatically enabled for the root file system using SUSE's snapper infrastructure. For more information about snapper, its integration into ZYpp and YaST, and the YaST snapper module, see the SUSE Linux Enterprise documentation.

Migration from "ext" File Systems to btrfs

Migration from existing "ext" file systems (ext2, ext3, ext4) is supported "offline" and "in place". Calling "btrfs-convert [device]" will convert the file system. This is an offline process, which needs at least 15% free space on the device, but is applied in place. Roll back: calling "btrfs-convert -r [device]" will roll back. Caveat: when rolling back, all data will be lost that has been added after the conversion into btrfs; in other words: the roll back is complete, not partial.

RAID

Btrfs is supported on top of MD (multiple devices) and DM (device mapper) configurations. Please use the YaST partitioner to achieve a proper setup. Multivolume/RAID with btrfs is not supported yet and will be enabled with a future maintenance update.

Future Plans

- We are planning to announce support for btrfs' built-in multi volume handling and RAID in a later version of SUSE Linux Enterprise.

- Starting with SUSE Linux Enterprise 12, we are planning to implement bootloader support for /boot on btrfs.
- Compression and Encryption functionality for btrfs is currently under development and will be supported once the development has matured.
- We are committed to actively work on the btrfs file system with the community, and we keep customers and partners informed about progress and experience in terms of scalability and performance. This may also apply to cloud and cloud storage infrastructures.

Online Check and Repair Functionality

Check and repair functionality ("scrub") is available as part of the btrfs command line tools. "Scrub" is aimed to verify data and metadata assuming the tree structures are fine. "Scrub" can (and should) be run periodically on a mounted file system: it runs as a background process during normal operation.

The "fsck.btrfs" tool is available in the SUSE Linux Enterprise update repositories.

Capacity Planning

If you are planning to use btrfs with its snapshot capability, it is advisable to reserve twice as much disk space than the standard storage proposal. This is automatically done by the YaST2 partitioner for the root file system.

Hard Link Limitation

In order to provide a more robust file system, btrfs incorporates back references for all file names, eliminating the classic "lost+found" directory added during recovery. A temporary limitation of this approach affects the number of hard links in a single directory that link to the same file. The limitation is dynamic based on the length of the file names used. A realistic average is approximately 150 hard links. When using 255 character file names, the limit is 14 links. We intend to raise the limitation to a more usable limit of 65535 links in a future maintenance update.



Note

With SLE 11 SP3 you can now raise this limitation. The so-called "extended inode refs" are not turned on by default in the SUSE kernels. This is because enabling them involves turning on an incompat bit in the file system which would make it unmountable by old versions of SLE.

If you want extended inode refs on though use 'btrfstune' to turn them on. There is no way to turn them off so it is a 1-way conversion. The command is (replace /PATH/TO/DEVICE with your device):

```
btrfsctl -r /PATH/TO/DEVICE
```

Other Limitations

At the moment, btrfs is not supported as a seed device.

For More Information

For more information about btrfs, see the SUSE Linux Enterprise 11 documentation.

4.1.3 Tomcat6 and Related Packages

Tomcat6 and related packages are fully supported on the Intel/AMD x86 (32bit), AMD64/Intel64, IBM POWER, and IBM System z architectures.

4.1.4 SELinux

The SELinux subsystem is supported. Arbitrary SELinux policies running on SLES are not supported, though. Customers and Partners who have an interest in using SELinux in their solutions, are encouraged to contact SUSE to evaluate the level of support that is needed, and how support and services for the specific SELinux policies will be granted.

4.2 Software Requiring Specific Contracts

The following packages require additional support contracts to be obtained by the customer in order to receive full support:

- BEA Java (Itanium only)
- MySQL Database
- PostgreSQL Database

4.2.1 openMPI Support

openMPI is used in HPC as a standard for communication. It is now supported in SLE 11 SP4.

The libraries are now in a separate RPM package (`openmpi-libs`). The library name has been changed from `libmpi.so.0` to `libmpi.so.1` .

4.3 Technology Previews

Technology previews are packages, stacks, or features delivered by SUSE. These features are not supported. They may be functionally incomplete, unstable or in other ways not suitable for production use. They are mainly included for customer convenience and give customers a chance to test new technologies within an enterprise environment.

Whether a technical preview will be moved to a fully supported package later, depends on customer and market feedback. A technical preview does not automatically result in support at a later point in time. Technical previews could be dropped at any time and SUSE is not committed to provide a technical preview later in the product cycle.

Please, give your SUSE representative feedback, including your experience and use case.

4.3.1 Support for the z Systems 10GbE RoCE Express

Support for the z Systems 10GbE RoCE Express feature can be used on zEC12, zBC12, z13 via the TCP/IP layer without restrictions. SLES 11 SP4 includes RDMA enablement and DAPL/OFED for IBM z Systems as a technology preview but these can only be used on LPAR when running on IBM z Systems zEC12, zBC12 and cannot be used on IBM z Systems z13.

4.3.2 XEN: VMCS Shadowing support for Xen


VMCS Shadowing is a VT-x feature that allows software in VMX non-root operation to execute the VMREAD and VMWRITE instructions. Such executions do not read from the current VMCS (the one supporting VMX non-root operation) but instead from a shadow VMCS. This feature improve nested virtualization performance.

4.3.3 Libvirt: VMware vpx and esx drivers

Libvirt will now be shipped with the VMware ESX hypervisor driver, allowing limited management of ESX hosts and vCenter environments. See <http://libvirt.org/drvesx.html>  for more details.

4.3.4 Technology Preview: Hot-Add Memory

Hot-add memory is currently only supported on the following hardware:

- IBM x3800, x3850, single node x3950, x3850 M2, single node x3950 M2,
- certified systems based on recent Intel Xeon Architecture,
- certified systems based on recent Intel IPF Architecture,
- all IBM servers and blades with POWER5, POWER6, POWER7, or POWER7+ processors and recent firmware. (This requires the Power Linux service and productivity tools available at <http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/yum.html> )

If your specific machine is not listed, please call SUSE support to confirm whether or not your machine has been successfully tested. Also, regularly check our maintenance update information, which will explicitly mention the general availability of this feature.

Restriction on using IBM eHCA InfiniBand adapters in conjunction with hot-add memory on IBM Power: The current eHCA Device Driver will prevent dynamic memory operations on a partition as long as the driver is loaded. If the driver is unloaded prior to the operation and then loaded again afterwards, adapter initialization may fail. A Partition Shutdown / Activate sequence on the HMC may be needed to recover from this situation.

4.3.5 Technology Preview: Internet Storage Naming Service (iSNS)

The Internet Storage Naming Service (iSNS) package is by design suitable for secure internal networks only. SUSE will continue to work with the community on improving security.

4.3.6 Technology Preview: Read-Only Root File System


It is possible to run SUSE Linux Enterprise Server 11 on a shared read-only root file system. A read-only root setup consists of the read-only root file system, a scratch and a state file system. The [/etc/rwtab](#) file defines which files and directories on the read-only root file system are replaced by which files on the state and scratch file systems for each system instance.

The readonlyroot kernel command line option enables read-only root mode; the state= and scratch= kernel command line options determine the devices on which the state and scratch file systems are located.

In order to set up a system with a read-only root file system, set up a scratch file system, set up a file system to use for storing persistent per-instance state, adjust /etc/rwtab as needed, add the appropriate kernel command line options to your boot loader configuration, replace /etc/mtab with a symlink to /proc/mounts as described below, and (re)boot the system.

To replace /etc/mtab with the appropriate symlinks, call:

```
ln -sf /proc/mounts /etc/mtab
```

See the `rwtab(5)` manual page for further details and <http://www.redbooks.ibm.com/abstracts/redp4322.html>  for limitations on System z.

5 Installation

5.1 Detecting FCoE Storage During Installation or Booting

An FCoE storage device is not detected or not properly coming up during installation or booting.

If an FCoE storage device is not detected or not properly coming up during installation or booting, it is recommended to add the following parameter to the kernel boot parameters.

```
with_fcoe=1
```

5.2 Standard Installation with DHCPv4 and DHCPv6

If you configure both DHCPv4 and DHCPv6 during the second stage of a standard installation, only DHCPv4 is enabled when calling 'rcnetwork restart'.

Make sure that the "dhcp-client" package is installed. Then call 'rcnetwork restart' again.

5.3 Top Level Domain ".site" No Longer Available for Private Use

Until SLE 11 SP4, when no hostname was provided by the user or DHCP, the installer was generating a hostname ending with .site. Since 2015, the top level domain (tld) ".site" is officially registered and should no longer be used for private purposes.

We recommend to rename the system using a proper fully qualified resolveable domain name. If impossible, use .test (or .invalid) as the domain name instead of .site (for more information, see RFC 6761). A new installation done with the SLE 11 SP4 installer will default to "linux.suse" instead of "linux.site", when none is provided.

5.4 Running SMT on SLES 11 SP4

SMT 11 SP3 can run on top of SLES 11 SP4. In order to make it possible, it is necessary to install the latest maintenance updates of SMT 11 SP3. This needs to be done before SLES is updated from version 11 SP3 to version 11 SP4. After the update of SLES, SMT keeps running as-is.

In order to install SMT 11 SP3 on top of SLES 11 SP4, use the latest media available from [download.suse.com](https://download.suse.com/index.jsp) (<https://download.suse.com/index.jsp>)  .

5.5 Booting i586 Machines

The provided ISO image is able to boot i586 machines if burnt on a DVD medium. It does not work to dump it on a USB device and use it for booting.

The x86_64 architecture is not affected by this limitation. On x86_64 booting from a USB device is supported.

5.6 AutoYaST Installation and Multipath

During AutoYaST installations, it is now possible to enable multipath.

5.7 Deployment

SUSE Linux Enterprise Server can be deployed in three ways:

- Physical Machine,
- Virtual Host,
- Virtual Machine in paravirtualized environments.

5.8 CJK Languages Support in Text-mode Installation

CJK (Chinese, Japanese, and Korean) languages do not work properly during text-mode installation if the framebuffer is not used (Text Mode selected in boot loader).

There are three alternatives to resolve this issue:

1. Use English or some other non-CJK language for installation then switch to the CJK language later on a running system using *YaST > System > Language*.
2. Use your CJK language during installation, but do not choose *Text Mode* in the boot loader using *F3 Video Mode*. Select one of the other VGA modes instead. Select the CJK language of your choice using *F2 Language*, add **textmode=1** to the boot loader command-line and start the installation.
3. Use graphical installation (or install remotely via SSH or VNC).

5.9 Booting from Harddisks larger than 2 TiB in Non-UEFI Mode

Booting from harddisks larger than 2 TiB in non-UEFI mode (but with GPT partition table) fails.

To successfully use harddisks larger than 2 TiB in non-UEFI mode, but with GPT partition table (i.e., grub bootloader), consider one of the following options:

- Use a 4k sector harddisk in 4k mode (in this case, the 2 TiB limit will become a 16 TiB limit).
- Use a separate /boot partition. This partition must be one of the first 3 partitions and end below the 2 TiB limit.
- Switch from legacy mode to UEFI mode, if this is an option for you.

5.10 Installation Using Persistent Device Names

The installer uses persistent device names by default. If you plan to add storage devices to your system after the installation, we strongly recommend you use persistent device names for all storage devices.

To switch to persistent device names on a system that has already been installed, start the YaST2 partitioner. For each partition, select *Edit* and go to the *Fstab Options* dialog. Any mount option except *Device name* provides you persistent device names. In addition, rerun the Boot Loader module in YaST and select *Propose New Config* to switch the boot loader to using the persistent device name, or manually adjust all boot loader sections. Then select *Finish* to write the new proposed configuration to disk. Alternatively, edit /boot/grub/menu.lst and /boot/grub/device.map according to your needs.

This needs to be done before adding new storage devices.

For further information, see the “Storage Administration Guide” about "Device Name Persistence".

5.11 iSCSI Booting with iBFT in UEFI Mode

If booting over iSCSI, iBFT information cannot be parsed when booting via native UEFI. The system should be configured to boot in legacy mode if iSCSI booting using iBFT is required.

5.12 Using iSCSI Disks when Installing

To use iSCSI disks during installation, passing the withiscsi boot parameter is no longer needed.

During installation, an additional screen provides the option to attach iSCSI disks to the system and use them in the installation process.

Bootting from an iSCSI server on i386, x86_64 and ppc64 is supported if iSCSI-enabled firmware is used.

5.13 Using qla3xxx and qla4xxx Drivers at the Same Time

QLogic iSCSI Expansion Card for IBM BladeCenter provides both Ethernet and iSCSI functions. Some parts on the card are shared by both functions. The current qla3xxx (Ethernet) and qla4xxx (iSCSI) drivers support Ethernet and iSCSI function individually. In contrast to previous SLES releases, using both functions at the same time is now supported.

If you happen to use brokenmodules=qla3xxx or brokenmodules=qla4xxx before upgrading to SLES 11 SP2, these options can be removed.

5.14 Using EDD Information for Storage Device Identification

EDD information (in /sys/firmware/edd/<device>) is used by default to identify your storage devices.

EDD Requirements:

- BIOS provides full EDD information (found in /sys/firmware/edd/<device>)
- Disks are signed with a unique MBR signature (found in /sys/firmware/edd/<device>/mbr_signature).

Add edd=off to the kernel parameters to disable EDD.

5.15 Automatic Installation with AutoYaST in an LPAR (System z)

For automatic installation with AutoYaST in an LPAR, the parmfile used for such an installation must have blank characters at the beginning and at the end of each line (the first line does not need to start with a blank). The number of characters in one line should not exceed 80.

5.16 Adding DASD or zFCP Disks During Installation (System z)

Adding of DASD or zFCP disks is not only possible during the installation workflow, but also when the installation proposal is shown. To add disks at this stage, please click on the *Expert* tab and scroll down. There the DASD and/or zFCP entry is shown. These added disks are not displayed in the partitioner automatically. To make the disks visible in the partitioner, you have to click on *Expert* and select *reread partition table*. This may reset any previously entered information.

5.17 Network Installation via eHEA on POWER

If you want to carry out a network installation via the IBM eHEA Ethernet Adapter on POWER systems, no huge (16GB) pages may be assigned to the partition during installation.

5.18 For More Information

For more information, see *Chapter 12, Infrastructure, Package and Architecture Specific Information*.

6 Features and Versions

6.1 Linux Kernel and Toolchain

6.1.1 qeth: Display Switch Port Mode

A new `sysfs` attribute allows to display the port mode settings of an adjacent switch, for example to see if it is set for hairpin mode

6.1.2 New Kernel Taint Flag - Unsigned Module

In the past modules without proper cryptographic signature loaded into kernel caused the kernel's "forced module load" flag to be set. As a result, tracing was disabled for such modules.

A flag has been introduced to indicate modules with invalid signature. Its internal kernel name is `TAINT_UNSIGNED_MODULE`, and is represented with letter 'E' in kernel debugging output.

6.1.3 Enabling VEBOX on Haswell in the drm/i915 Kernel Driver

Linux Cloud Video Transcode is an Intel GEN based hardware solution to support high quality and performance video transcoding on a server. With enabling VEBOX on Haswell for some video pre and post process features like DN/ADI SUSE Linux Enterprise features improved transcode quality.

6.1.4 General Version Information

- GCC 4.3.4
- glibc 2.11.3
- Linux kernel 3.0
- perl 5.10
- php 5.3

- python 2.6.9
- ruby 1.8.7

6.1.5 SUSE Linux Enterprise Real Time Extension

To take advantage of the Real Time extension the extension must be at the same version as the base SUSE Linux Enterprise Server. An updated version for SUSE Linux Enterprise Real Time extension is provided later after the release of SUSE Linux Enterprise Server.

6.2 Server



Note

Note: in the following text version numbers do not necessarily give the final patch- and security-status of an application, as SUSE may have added additional patches to the specific version of an application.

6.2.1 numactl and libnuma

numactl and libnuma have been updated to the latest version.

This update comes with many bug fixes and some new features that are especially important for large NUMA systems, e.g.:

- IO affinity support
- New option to memhog to disable transparent huge pages
- Show distances on machines without a node 0

6.3 Desktop

- GNOME 2.28

GNOME was updated with SP2 and uses PulseAudio for sound.

- KDE 4.3.5

KDE was updated with SP2.

- X.org 7.4

6.3.1 fbdev Driver Needs Reboot After Resolution Changes

SaX2 offers to change the resolution even for the fbdev driver. Because this is controlled via a VGA kernel option, rebooting is needed after resolution changes. In other words: Modifications will take effect the next time the graphics system is restarted; in some cases a reboot of the machine is needed.

6.3.2 GNOME: Primary Monitor for the Greeter

With GNOME designating the primary monitor for the greeter now is possible.

6.4 Security

6.4.1 chpasswd with Support for SHA256 and SHA512

Various parts of the password checking frameworks in PAM and pwutils already had support for SHA256 and SHA512 based password hashing functions.

Support was missing in chpasswd, a program usable for scripting password setting.

chpasswd was made able to also set SHA256 and SHA512 based passwords.

6.4.2 openSSH update to 6.6p1

OpenSSH is constantly improving and gaining new and more secure cipher suites. Backporting them is occasionally not possible.

OpenSSH has been updated to version 6.6p1, same version as used in SUSE Linux Enterprise 12.

The ciphers now includes modern elliptic curve based on the elliptic curve Curve25519, resulting in public key types Ed25519.

Also the new transport cipher "chacha20-poly1305@openssh.com" was added, using the ChaCha20 stream cipher and Poly1305 MAC developed by Dan Bernstein.

Various other improvements and bugfixes are also included.

6.4.3 Switch repomd from sha to sha26

The update repository integrity used by SUSE is ensured by a GPG signature and the checksums of the YUM repomd XML metadata.

SUSE Linux Enterprise 11 so far used sha1 as intermediate checksum, which should no longer be used.

With SUSE Linux Enterprise 12 and SUSE Linux Enterprise 11 SP4 we start to use sha256 for the XML integrity handling and so get rid of the old sha1 hashing methods.

If you have tools parsing the XML metadata yourself, please verify they can handle also the newer sha256 hashes.

6.4.4 zypper: search aware of --cve

zypper list-patches and install are aware of the CVE metadata inside the update information. Zypper search was however not able to search for CVEs yet.

Zypper search has been improved to also be able to search for CVE names using e.g. `zypper se --cve="CVE-2014-0001"`

6.4.5 PAM Configuration

The common PAM configuration files (`/etc/pam.d/common-*`) are now created and managed with `pam-config`.

6.4.6 SELinux Enablement

In addition to AppArmor, SELinux capabilities have been added to SUSE Linux Enterprise Server. While these capabilities are not enabled by default, customers can run SELinux with SUSE Linux Enterprise Server if they choose to.

What does SELinux enablement mean?

- The kernel ships with SELinux support.
- We will apply SELinux patches to all “common” userland packages.
- The libraries required for SELinux (`libselinux`, `libsepol`, `libsemanage`, etc.) have been added to openSUSE and SUSE Linux Enterprise.
- Quality Assurance is performed with SELinux disabled—to make sure that SELinux patches do not break the default delivery and the majority of packages.
- The SELinux specific tools are shipped as part of the default distribution delivery.
- Arbitrary SELinux policies running on SLES are not supported, though, and we will not be shipping any SELinux policies in the distribution. Reference and minimal policies may be available from the repositories at some future point.
- Customers and Partners who have an interest in using SELinux in their solutions, are encouraged to contact SUSE to evaluate the level of support that is needed, and how support and services for the specific SELinux policies will be granted.

By enabling SELinux in our codebase, we add community code to offer customers the option to use SELinux without replacing significant parts of the distribution.

6.4.7 Enablement for TPM/Trusted Computing

SUSE Linux Enterprise Server 11 comes with support for Trusted Computing technology. To enable your system's TPM chip, make sure that the "security chip" option in your BIOS is selected. TPM support is entirely passive, meaning that measurements are being performed, but no action is taken based on any TPM-related activity. TPM chips manufactured by Infineon, NSC and Atmel are supported, in addition to the virtual TPM device for Xen.

The corresponding kernel drivers are not loaded automatically. To do so, enter:

```
find /lib/modules -type f -name "tpm*.ko"
```

and load the kernel modules for your system manually or via **`MODULES_LOADED_ON_BOOT`** in `/etc/sysconfig/kernel`.

If your TPM chip with taken ownership is configured in Linux and available for use, you may read PCRs from `/sys/devices/*/*/pcrs`.

The `tpm-tools` package contains utilities to administer your TPM chip, and the `trousers` package provides `tcstd`—the daemon that allows userland programs to communicate with the TPM driver in the Linux kernel. `tcstd` can be enabled as a service for the runlevels of your choice.

To implement a trusted ("measured") boot path, use the package `trustedgrub` instead of the `grub` package as your bootloader. The trustedgrub bootloader does not display any graphical representation of a boot menu for informational reasons.

6.4.8 Linux File System Capabilities

Our kernel is compiled with support for Linux File System Capabilities. This is disabled by default. The feature can be enabled by adding `file_caps=1` as kernel boot option.

6.5 Network

IPv6 Improvements

SUSE Linux Enterprise Server has successfully completed the USGv6 test program designated by NIST that provides a proof of compliance to IPv6 specifications outlined in current industry standards for common network products.

Being IPv6 Consortium Member and Contributor Novell/SUSE have worked successfully with University of New Hampshire InterOperability Laboratory (UNH-IOL) to verify compliance to IPv6 specifications. The UNH-IOL offers ISO/IEC 17025 accredited testing designed specifically for the USGv6 test program. The devices that have successfully completed the USGv6 testing at the UNH-IOL by December 2012 are SUSE Linux Enterprise Server 11 SP2. Testing for subsequent releases of SUSE Linux Enterprise Server is in progress, and current and future results will be listed at <http://www.iol.unh.edu/services/testing/ipv6/usgv6tested.php?company=105&type=#eqplist>.

SUSE Linux Enterprise Server can be installed in an IPv6 environment and run IPv6 applications. When installing via network, do not forget to boot with "`ipv6=1`" (accept v4 and v6) or "`ipv6only=1`" (only v6) on the kernel command line. For more information, see the Deployment Guide and also *Section 14.6, "IPv6 Implementation and Compliance"*.

10G Networking Capabilities

OFED 1.5

traceroute 1.2

Support for traceroute over TCP.

FCoE

FCoE is an implementation of the Fibre Channel over Ethernet working draft. Fibre Channel over Ethernet is the encapsulation of Fibre Channel frames in Ethernet packets. It allows users with a FCF (Fibre Channel over Ethernet Forwarder) to access their existing Fibre Channel storage using an Ethernet adapter. When leveraging DCB's PFC technology to provide a loss-less environment, FCoE can run SAN and LAN traffic over the same link.

Data Center Bridging (DCB)

Data Center Bridging (DCB) is a collection of Ethernet enhancements designed to allow network traffic with differing requirements (e.g., highly reliable, no drops vs. best effort vs. low latency) to operate and coexist on Ethernet. Current DCB features are:

- *Enhanced Transmission Selection* (aka *Priority Grouping*) to provide a framework for assigning bandwidth guarantees to traffic classes.
- *Priority-based Flow Control (PFC)* provides a flow control mechanism which can work independently for each 802.1p priority.
- *Congestion Notification* provides a mechanism for end-to-end congestion control for protocols, which do not have built-in congestion management.

6.5.1 iSCSI Booting Using HBA Mode

SLES 11 SP4 now allows iSCSI booting from some adapters using HBA mode, even if the boot target is on a different subnetwork. It does this by gathering and using three new iBFT boot parameters: boot_root , boot_nic , and boot_target . The kernel as well as open-iscsi was changed to provide this feature.

6.5.2 sshd: ipv6 configuration and X11 Forwarding

If ipv6 is disabled sshd automatically starts with the -4 option. This way enabling X11 forwarding is possible.

6.5.3 Linuxrc: Option for AutoYaST to Define a vlan 802.1q

Linuxrc now writes `VlanID: XXX` to `install.inf`, but only if a vlan ID was set.

6.6 Resource Management

6.6.1 LXC Requires Correct Network Configuration

LXC now comes with support for network gateway detection. This feature will prevent a container from starting, if the network configuration setup of the container is incorrect. For instance, you must make sure that the network address of the container is within the host ip range, if it was set up as bridged on host. You might need to specify the netmask of the container network address (using the syntax "`lxc.network.ipv4 = X.Y.Z.T / cidr`") if the netmask is not the network class default netmask).

When using DHCP to assign a container network address, ensure "`lxc.network.ipv4 = 0.0.0.0`" is used in your configuration template.

Previously a container would have been started but the network would not have been working properly. Now a container will refuse to start, and print an error message stating that the gateway could not be set up. For containers created before this update we recommend running `rcnetwork restart` to reestablish a container network connection.



Tip: LXC Maintenance Update

After installing LXC maintenance update, we recommend clearing the LXC SLES cache template (stored by default in `/var/cache/lxc/sles/rootfs-*`) to ensure changes in the SLES template are available in newly created containers.

For containers created before the update, we recommend to install the packages "supportconfig", "sysconfig", and "iputils" using zypper.

6.7 Systems Management

Improved Update Stack

SUSE Linux Enterprise Server 11 provides an improved update stack and the new command line tool zypper to manage the repositories and install or update packages.

Enhanced YaST Partitioner

Extended Built-in Management Infrastructure

SUSE Linux Enterprise Server provides CIM/WBEM enablement with the SFCB CIMOM.

The following CIM providers are available:



- `cmpi-pywbem-base`
- `cmpi-pywbem-power-management` (DSP1027)
- `cmpi-pywbem-software` (DSP1023)
- `libvirt-cim` (DSP1041, DSP1043, DSP1045, DSP1057, DSP1059, DSP1076, DSP1081)
- `sblim-cmpi-base`
- `sblim-cmpi-dhcp`
- `sblim-cmpi-ethport_profile` (DSP1014)
- `sblim-cmpi-fsvol`
- `sblim-cmpi-network`
- `sblim-cmpi-nfsv3`
- `sblim-cmpi-nfsv4`
- `sblim-cmpi-sysfs`
- `sblim-gather-provider`
- `smis-providers`
- `sblim-cmpi-dns`
- `sblim-cmpi-samba`
- `sblim-cmpi-smbios`

Support for Web Services for Management (WS Management)

The WS-Management protocol is supported via Openwsman, providing client (package: openwsman-client) and server (package: openwsman-server) implementations.

This allows for interoperable management with the Windows 'winrm' stack.

WebYaST — Web-Based Remote Management

WebYaST is an easy to use, web-based administration tool targeted at casual Linux administrators. WebYaST is an add-on product. To deploy it, download the WebYaST media from <http://download.novell.com>  (strings search or direct link: <https://download.suse.com/Download?buildid=uVizILaPtzg~> ) and install the add-on product e.g., via the YaST add-on module. After installation, follow these steps:

- Open firewall port (note port number change!):

```
SuSEfirewall2 open EXT TCP 4984
SuSEfirewall2 restart
```

- Start services:

```
rccollected start
rcwebyast start
```

The last command will display the URL to connect to with a Web browser.

For information about migrating to SP3, see *Section 10.4.3, “Migrating SUSE Linux Enterprise Server 11 SP2 with WebYaST Installed via wagon”*.

6.7.1 YaST iSCSI Client Keeps Startup Mode of Connected Targets

There was no possibility to keep the startup mode ('automatic', 'manual' or 'on boot') of already connected targets. Using either 'Discovery' on 'Discovered Targets' or 'Add' of the 'Connected Targets' dialog used to reset the startup mode to the default mode 'manual'.

Now it is possible when using the Add button of the Connected Targets dialog to detect additional targets. Then the startup mode of already connected targets will not change.

6.7.2 IPMI Update

IPMI tools have been updated to version 1.8.15 to support newer hardware.

6.7.3 snmpd Improvements

snmpd has been enhanced to allow the monitoring of tmpfs file systems with SUSE Linux Enterprise 11 SP4.

6.7.4 Tomcat: Support for renaming JSESSIONID

When Tomcat is used together with other application servers the JSESSIONID session cookie can be overwritten by each other.

The solution is to allow the JSESSIONID to be renamed. This is possible with Tomcat 6.0.19 and higher.

6.8 Other

EVMS2 Replaced with LVM2

Default File System

With SUSE Linux Enterprise Server 11, the default file system in new installations has been changed from ReiserFS to ext3. A public statement can be found at <http://www.suse.com/products/server/technical-information/#FileSystem>.


UEFI Enablement on AMD64/Intel64

SWAP over NFS

Linux Foundation's Carrier Grade Linux (CGL)

SUSE supports the Linux Foundation's Carrier Grade Linux (CGL) specification. SUSE Linux Enterprise 11 meets the latest CGL 4.0 standard, and is CGL registered. For more information, see <http://www.suse.com/products/server/cgl/>.

Hot-Add Memory and CPU with vSphere 4.1 or Newer

Hot-add memory and CPU is supported and tested for both 32-bit and 64-bit systems when running vSphere 4.1 or newer. For more information, see the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php?deviceCategory=software&partner=465&virtualHardware=23> .

6.8.1 `/etc/os-release`

In addition to the `/etc/SuSE-release` file the file `/etc/os-release` is now available.

`/etc/os-release` is a cross-distribution standard to identify a Linux system. For more information about the syntax, see the os-release man page (`man os-release`).

7 Driver Updates

7.1 Network Drivers

- Updated bnx driver to version 2.0.4
- Updated bnx2x driver to version 1.52.1-7
- Updated e100 driver to version 3.5.24-k2
- Updated tg3 driver to version 3.106
- Added bna driver for Brocade 10Gbit LAN card in version 2.1.2.1
- Updated bfa driver to version 2.1.2.1
- Updated qla3xxx driver to version 2.03.00-k5
- Updated sky2 driver to version 1.25

7.1.1 Network Modules

The following modules got updated (selection):

- ixgbe - 3.19.1-k
- ixgbevf - 2.12.1-k
- igb - 5.2.15-k
- i40evf - 1.0.5
- e1000e - 2.3.2-k
- e1000 - 7.3.21-k8-NAPI
- e100 - 3.5.24-k2-NAPI
- i40e - 1.0.11-k

- igbvf - 2.0.2-k
- tg3 - 3.137
- cnic - 2.5.20
- bnx2 - 2.2.5
- bnx2i - 2.7.10.1
- bnx2fc - 2.8.1
- qla2xxx - 8.07.00.18-k
- qla4xxx - 5.03.00.00.11.4-k0
- qlcnlc - 5.3.62
- qlge - 1.00.00.34
- netxen_nic - 4.0.82
- cxgb4 - 2.0.0-ko
- cxgb3 - 1.1.5-ko
- cxgb4i - 0.9.1
- cxgb3i - 2.0.0
- cxgb4 - 2.0.0-ko
- cxgb3 - 1.1.5-ko

7.1.2 PSM Library for the Intel Infiniband Solution (OFED)

The PSM library for the Intel Infiniband solution (OFED) is now available.

7.1.3 Updating Firmware for QLogic 82XX based CNA

For QLogic 82XX based CNA, update the firmware to the latest from the QLogic website or whatever is recommended by the OEM in case you are running 4.7.x FW version.

7.2 Storage Drivers

- Updated qla2xxx to version 8.04.00.13.11.3-k
- Updated qla4xxx to version v5.03.00.06.11.3-k0
- Updated megaraid_mbox driver to version 2.20.5.1
- Updated megaraid_sas to version 4.27
- Updated MPT Fusion to version 4.22.00.00
- Updated mpt2sas driver to version 04.100.01.02
- Updated lpfc driver to version 8.3.5.7
- Added bnx2i driver for Broadcom NetXtreme II in version 2.1.1
- Updated bfa driver to version 2.1.2.1
- The enic driver was updated to version 1.4.2 to support newer Cisco UCS systems. This update also replaces LRO (Large Receive Offload) to GRO (Generic Receive Offload).

7.2.1 Storage Modules

The following modules got updated (selection):

- 3w-sas - 3.26.02.000
- 3w-9xxx - 2.26.02.014
- megaraid_sas - 06.806.08.00-rc1
- mptsas - 4.28.00.01suse
- mpt2sas - 20.100.00.00
- mpt3sas - 09.100.00.00
- pmcraid - 1.0.3
- aacraid - 1.2-0[30300]-ms

- mtip32xx - 1.3.1
- mlx4_core - 2.2-1
- mlx5_core - 2.2-1
- be2net - 10.4s
- lpfc - 0:10.4.8000.0.

7.2.2 Brocade FCoE Switch Does Not Accept Fabric Logins from Initiator

1. Once link is up, LLDP query QoS to get the new PFC, send FCoE incapable right away, which is right.
2. After negotiating with neighbor, we got lldp frame with un-recognized ieee dcbb, so we declare link is CEE incapable, and send out FCoE Capable event with PFC = 0 to fcoe kernel.
3. Then neighbor adjusts its version to match our CEE version, now we find right DCBX tlv in incoming LLDP frame, we declare link CEE capable. At this time we did not send FCoE capable again since we already sent it in step 2.

To solve this, upgrade the switch firmware to v6.4.3 or above.

7.3 Other Drivers

- Updated CIFS to version 1.74
- Updated intel-i810 driver
- Added X11 driver for AMD Geode LX 2D (xorg-x11-driver-video-amd)
- Updated X11 driver for Radeon cards
- Updated XFS and DMAPI driver
- Updated Wacom driver to version 1.46

7.3.1 SaX2: Changing Video Resolution

With the update to SLE 11 SP4, SaX2 no longer lets you select a video resolution when KMS is active. With KMS and the native or the modesetting driver RandR > 1.1 is available, which lets you change the resolution on the fly. The Gnome desktop provides a tool to do this and save the settings persistently across sessions.

For any UMS (and RandR 1.1) drivers you will still get the full list of video modes. If you select an unsupported mode, it will be ignored and a monitor preferred default mode will be used instead.

7.3.2 Intel Processor support

This Service Pack adds support for the following Intel Processors.

- Intel® Xeon® processor E7-4800 v3 product family
- Intel® Xeon® processor E7-8800 v3 product family
- 5th Gen Intel® Core™ processors
- Intel® Core™ M Processor
- Intel® Xeon® processor E3-1200 v4 product family

7.3.3 Hyper-v: Update drivers to latest upstream version

The updated drivers provide the following features:

A userland daemon to handle the file copy service is included.

The VMBUS driver utilizes all virtual CPUs (vCPUs) to communicate with the host, this will improve performance.

Support for Generation2 VMs is included. 'Secure Boot' must be disabled in the VM settings on the host side, otherwise the VM will not start.

The network driver was updated to remove the warning about outdated 'Integration Services' that was shown in the Hyper-V Manager GUI.

The network driver was updated to handle hotplug, which is a feature of Windows Server 2016.

8 Other Updates

8.1 Update of PostgreSQL to Version 9.4

The upstream end-of-life for version 9.1 is announced for September 2016. Customers need to switch to a newer supported version until then.

PostgreSQL was updated to version 9.4, prolonging the timeframe during which PostgreSQL is supported. Thus there is enough time for switching.

8.2 List of Updated Packages (SLES)

- Updated autoyast2 to version 2.17.76
- Updated binutils to version 2.24
- Updated biosdevname to version 0.6.1
- Updated brocade-firmware to version 3.2.3.0
- Updated btrfsprogs to version 3.18.2
- Updated cpupower to version 3.19
- Updated crash to version 7.0.9
- Updated cryptsetup to version 1.1.3
- Updated dosfstools to version 3.0.26
- Updated efibootmgr to version 0.6.0
- Updated gdb to version 7.7
- Updated gfxboot to version 4.1.34
- Updated hwdm to version 15.54
- Updated hyper-v to version 6

- Updated installation-images to version 11.221
- Updated ipmitool to version 1.8.15
- Updated iproute2 to version 3.0
- Updated iprutils to version 2.4.1
- Updated iptables to version 1.4.16.3
- Updated ledmon to version 0.79
- Updated libcgroupl to version 0.41.rc1
- Updated libdrm to version 2.4.52
- Updated libguestfs to version 1.20.12
- Updated libHBAAPI2 to version 2.2.9
- Updated libvirt to version 1.2.5
- Updated libvdp2 to version 2.2.4
- Updated libzypp to version 9.38.8
- Updated linux-kernel-headers to version 3.0
- Updated linuxrc to version 3.3.107
- Updated lldpad to version 0.9.46
- Updated makedumpfile to version 1.5.6
- Updated mcelog to version 1.0.2014.12.20
- Updated mdadm to version 3.3.1
- Updated mstflint to version 3.6.0
- Updated ntp to version 4.2.8p2
- Updated numactl to version 2.0.10
- Updated openCryptoki to version 3.2

- Updated open-fcoe to version 1.0.29
- Updated openssh to version 6.6p1
- Updated pciutils-ids to version 2014.11.18
- Updated perf to version 3.0.101
- Updated perl-Bootloader to version 0.4.89.70
- Updated perl-Sys-Virt to version 1.2.5
- Updated postgresql-init to version 9.4
- Updated release-notes-sles to version 11.4.10
- Updated sg3_utils to version 1.40
- Updated sles-installquick_cs to version 11.3
- Updated sles-manuals_en to version 11.4
- Updated smartmontools to version 6.3
- Updated src_vipa to version 2.1.0
- Updated suse-sam to version 0.8.5
- Updated suse-sam-data to version 0.8.5
- Updated virt-manager to version 0.9.5
- Updated virt-viewer to version 0.5.7
- Updated vm-install to version 0.6.37
- Updated xen to version 4.4.2_06
- Updated xrdp to version 0.6.1
- Updated yast2 to version 2.17.138
- Updated yast2-bootloader to version 2.17.98
- Updated yast2-fcoe-client to version 2.17.26

- Updated yast2-http-server to version 2.17.17
- Updated yast2-installation to version 2.17.114
- Updated yast2-iscsi-client to version 2.17.41
- Updated yast2-iscsi-lio-server to version 2.17.14
- Updated yast2-iscsi-server to version 2.17.11
- Updated yast2-ncurses to version 2.17.23
- Updated yast2-network to version 2.17.207
- Updated yast2-nfs-client to version 2.17.19
- Updated yast2-ntp-client to version 2.17.16
- Updated yast2-online-update to version 2.17.24
- Updated yast2-packager to version 2.17.113
- Updated yast2-repair to version 2.17.13
- Updated yast2-samba-server to version 2.18.0
- Updated yast2-storage to version 2.17.157
- Updated yast2-support to version 2.17.21
- Updated yast2-trans-ar to version 2.17.40
- Updated yast2-trans-cs to version 2.17.49
- Updated yast2-trans-da to version 2.17.36
- Updated yast2-trans-de to version 2.17.61
- Updated yast2-trans-el to version 2.17.19
- Updated yast2-trans-en_GB to version 2.17.30
- Updated yast2-trans-en_US to version 2.17.35
- Updated yast2-trans-es to version 2.17.56

- Updated yast2-trans-fi to version 2.17.40
- Updated yast2-trans-fr to version 2.17.60
- Updated yast2-trans-hu to version 2.17.60
- Updated yast2-trans-it to version 2.17.58
- Updated yast2-trans-ja to version 2.17.50
- Updated yast2-trans-ko to version 2.17.56
- Updated yast2-trans-nb to version 2.17.33
- Updated yast2-trans-nl to version 2.17.55
- Updated yast2-trans-pl to version 2.17.52
- Updated yast2-trans-pt to version 2.17.11
- Updated yast2-trans-pt_BR to version 2.17.55
- Updated yast2-trans-ru to version 2.17.50
- Updated yast2-trans-sv to version 2.17.40
- Updated yast2-trans-tr to version 2.17.17
- Updated yast2-trans-uk to version 2.17.29
- Updated yast2-trans-zh_CN to version 2.17.45
- Updated yast2-trans-zh_TW to version 2.17.39
- Updated yast2-update to version 2.17.26
- Updated yast2-users to version 2.17.55
- Updated yast2-vm to version 2.17.17
- Updated zypper to version 1.6.323

9 Software Development Kit

SUSE provides a Software Development Kit (SDK) for SUSE Linux Enterprise 11 Service Pack 4. This SDK contains libraries, development environments and tools along the following patterns:

- C/C++ Development
- Certification
- Documentation Tools
- GNOME Development
- Java Development
- KDE Development
- Linux Kernel Development
- Programming Libraries
- .NET Development
- Miscellaneous
- Perl Development
- Python Development
- Qt 4 Development
- Ruby on Rails Development
- Ruby Development
- Version Control Systems
- Web Development
- YaST Development

9.1 samba-test subpackage with smbtorture and other binaries

The samba-test subpackage was added, offering smbtorture and other binaries for testing.

10 Update-Related Notes

This section includes update-related information for this release.

10.1 General Notes

10.1.1 Lower Version Numbers in SUSE Linux Enterprise 11 SP4 Than in Version 11 SP3

When upgrading from SUSE Linux Enterprise Server or Desktop 11 SP3 to version 11 SP4, you may encounter a version downgrade of specific software packages, including the Linux Kernel.

SLE 11 SP4 has all its software packages and updates in the SLE 11 SP4 repositories. No packages from SLE 11 SP3 repositories are needed for installation or upgrade, not even from the SLE 11 SP3 update repositories.

Note

It is important to remember that the version number is not sufficient to determine which bugfixes are applied to a software package.

In case you add SLE 11 SP3 update repositories, be aware of one characteristic of the repository concept: Version numbers in the SP3 update repository can be higher than those in the SP4 repository. Thus, if you update with the SP3 repositories enabled, you may get the SP3 version of a package instead of the SP4 version. This is admittedly unfortunate.

It is recommended to avoid using the version from a lower product or SP, because using the SLE 11 SP3 package instead of the SP4 package can result in unexpected side effects. Thus we advise to switch off all the SLE 11 SP3 repositories, if you do not really need them. Keep old repositories only, if your system depends on a specific older package version. If you need a package from a lower product or SP though, and thus have SLE 11 SP3 repositories enabled, make sure that the packages you intended to upgrade have actually been upgraded.

Summarizing: If you have an SLE 11 SP3 installation with all patches and updates applied, and then migrate off-line to SLE 11 SP4, you will see a downgrade of some packages. This is expected behavior.

10.1.2 Upgrading from SLES 10 (GA and Service Packs) or SLES 11 GA

There are supported ways to upgrade from SLES 10 GA and SPx or SLES 11 GA and SP1 to SLES 11 SP4, which may require intermediate upgrade steps:

- SLES 10 GA -> SLES 10 SP1 -> SLES 10 SP2 -> SLES 10 SP3 -> SLES 10 SP4 -> SLES 11 SP4, or
- SLES 11 GA -> SLES 11 SP1 -> SLES 11 SP2 -> SLES 11 SP3 -> SLES 11 SP4

10.1.3 Online Migration from SP3 to SP4 via "YaST wagon"

The online migration from SP3 to SP4 is supported via the "YaST wagon" module.

10.1.4 Migrating to SLE 11 SP4 Using Zypper

To migrate the system to the Service Pack 4 level with zypper, proceed as follows:

- Open a root shell.
- Run zypper ref -s to refresh all services and repositories.
- Run zypper patch to install package management updates.
- Now it is possible to install all available updates for SLES/SLED 11 SP3; run zypper patch again.
- Now the installed products contain information about distribution upgrades and which migration products should be installed to perform the migration. Read the migration product information from /etc/products.d/*.prod and install them.
- Enter the following command:

```
grep '<product' /etc/products.d/*.prod
```

A sample output could be as follows:

```
<product>sle-sdk-SP4-migration</product>  
<product>SUSE_SLES-SP4-migration</product>
```

- Install these migration products (example):

```
zypper in -t product sle-sdk-SP4-migration SUSE_SLES-SP4-migration
```

- Run **`suse_register -d 2 -L /root/.suse_register.log`** to register the products in order to get the corresponding SP4 Update repositories.
- Run **`zypper ref -s`** to refresh services and repositories.
- Check the repositories using **`zypper lr`**. Disable SP3 repositories after the registration and enable the new SP4 repositories (such as SP4-Pool, SP4-Updates):

```
zypper mr --disable <repo-alias>  
zypper mr --enable <repo-alias>
```

Also disable repositories you do not want to update from.

- Then perform a distribution upgrade by entering the following command:

```
zypper dup --from SLES11-SP4-Pool --from SLES11-SP4-Updates \  
--from SLE11-SP2-WebYaST-1.3-Pool --from SLE11-SP2-WebYaST-1.3-Updates
```

Add more SP4 repositories here if needed, e.g. in case add-on products are installed. For WebYaST, it is actually **`SLE11-SP2-*`**, because there is one WebYaST release that runs on three SP code bases.



Note

If you make sure that only repositories, which you migrate from, are enabled, you can omit the **`--from`** parameters.

- zypper will report that it will delete the migration product and update the main products. Confirm the message to continue updating the RPM packages.
- To do a full update, run **`zypper patch`**.
- After the upgrade is finished, register the new products again:

```
suse_register -d 2 -L /root/.suse_register.log
```

- Run **zypper patch** after re-registering. Some products donot use the update repositories during the migration and they are not active at this point of time.
- Reboot the system.

10.1.5 Migration from SUSE Linux Enterprise Server 10 SP4 via Bootable Media

Migration is supported from SUSE Linux Enterprise Server 10 SP4 via bootable media (incl. PXE boot).

10.1.6 Migrating Hosts Running SMT 11 SP2 to SMT 11 SP3

CHECKIT:11_3

As part of the release of the SLE 11 SP3 product family, SUSE will also release Subscription Management Tool 11 SP3 (SMT 11 SP3). We expect to release SMT 11 SP3 within a month after the release of SLES 11 SP3.

Do not migrate hosts running SMT 11 SP2 to SLES 11 SP3 before SMT 11 SP3 is available.

You can update SLE 11 SP3 hosts via SMT 11 SP2 without any limitations until SMT 11 SP3 is released.

10.1.7 Online Migration with Debuginfo Packages Not Supported

Online migration from SP3 to SP4 is not supported if debuginfo packages are installed.

10.1.8 Upgrading to SLES 11 SP3 with Root File System on iSCSI

CHECKIT:11_3

The upgrade or the automated migration from SLES 10 to SLES 11 SP3 may fail if the root file system of the machine is located on iSCSI because of missing boot options.

There are two approaches to solve it, if you are using AutoYaST (adjust IP addresses and hostnames according to your environment!):

With Manual Intervention:

Use as boot options:

withiscsi=1 autoupgrade=1 autoyast=http://myserver/autoupgrade.xml

Then, in the dialog of the iSCSI initiator, configure the iSCSI device.

After successful configuration of the iSCSI device, YaST will find the installed system for the upgrade.

Fully Automated Upgrade:

Add or modify the `<iscsi-client>` section in your autoupgrade.xml as follows:

```
<iscsi-client>
  <initiatorname>iqn.2012-01.com.example:initiator-example</initiatorname>
  <targets config:type="list">
    <listentry>
      <authmethod>None</authmethod>
      <iface>default</iface>
      <portal>10.10.42.84:3260</portal>
      <startup>onboot</startup>
      <target>iqn.2000-05.com.example:disk01-example</target>
    </listentry>
  </targets>
  <version>1.0</version>
</iscsi-client>
```

Then, run the automated upgrade with these boot options:

autoupgrade=1 autoyast=http://myserver/autoupgrade.xml

10.1.9 Kernel Split in Different Packages

With SUSE Linux Enterprise Server 11 the kernel RPMs are split in different parts:

- kernel-flavor-base
Very reduced hardware support, intended to be used in virtual machine images.
- kernel-flavor
Extends the base package; contains all supported kernel modules.
- kernel-flavor-extra

All other kernel modules which may be useful but are not supported. This package will not be installed by default.

10.1.10 Tickless Idle

SUSE Linux Enterprise Server uses tickless timers. This can be disabled by adding `nohz=off` as a boot option.

10.1.11 Development Packages

SUSE Linux Enterprise Server will no longer contain any development packages, with the exception of some core development packages necessary to compile kernel modules. Development packages are available in the SUSE Linux Enterprise Software Development Kit.

10.1.12 Displaying Manual Pages with the Same Name

The `man` command now asks which manual page the user wants to see if manual pages with the same name exist in different sections. The user is expected to type the section number to make this manual page visible.

If you want to revert back to the previously used method, please set `MAN_POSIXLY_CORRECT=1` in a shell initialization file such as `~/.bashrc`.

10.1.13 YaST LDAP Server No Longer Uses /etc/openldap/slapd.conf

The YaST LDAP Server module no longer stores the configuration of the LDAP Server in the file /etc/openldap/slapd.conf. It uses OpenLDAP's dynamic configuration backend, which stores the configuration in an LDAP database itself. That database consists of a set of `.ldif` files in the directory /etc/openldap/slapd.d. You should - usually - not need to access those files directly. To access the configuration you can either use the `yast2-ldap-server` module or any capable LDAP client (e.g., `ldapmodify`, `ldapsearch`, etc.). For details on the dynamic configuration of OpenLDAP, refer to the OpenLDAP Administration Guide.

10.1.14 AppArmor

This release of SUSE Linux Enterprise Server ships with AppArmor. The AppArmor intrusion prevention framework builds a firewall around your applications by limiting the access to files, directories, and POSIX capabilities to the minimum required for normal operation. AppArmor protection can be enabled via the AppArmor control panel, located in YaST under Security and Users. For detailed information about using AppArmor, see the documentation in [/usr/share/doc/packages/apparmor-docs](#).

The AppArmor profiles included with SUSE Linux have been developed with our best efforts to reproduce how most users use their software. The profiles provided work unmodified for many users, but some users may find our profiles too restrictive for their environments.

If you discover that some of your applications do not function as you expected, you may need to use the AppArmor Update Profile Wizard in YaST (or use the `aa-logprof(8)` command line utility) to update your AppArmor profiles. Place all your profiles into learning mode with the following: **`aa-complain /etc/apparmor.d/*`**

When a program generates many complaints, the system's performance is degraded. To mitigate this, we recommend periodically running the Update Profile Wizard (or `aa-logprof(8)`) to update your profiles even if you choose to leave them in learning mode. This reduces the number of learning events logged to disk, which improves the performance of the system.

10.1.15 Updating with Alternative Boot Loader (Non-Linux) or Multiple Boot Loader Programs



Note

Before updating, check the configuration of your boot loader to assure that it is not configured to modify any system areas (MBR, settings active partition or similar). This will reduce the amount of system areas that you need to restore after update.

Updating a system where an alternative boot loader (not grub) or an additional boot loader is installed in the MBR (Master Boot Record) might override the MBR and place grub as the primary boot loader into the system.

In this case, we recommend the following: First backup your data. Then either do a fresh installation and restore your data, or run the update nevertheless and restore the affected system areas (in particular, the MBR). It is always recommended to keep data separated from the system software. In other words, /home, /srv, and other volumes containing data should be on separate partitions, volume groups or logical volumes. The YaST partitioning module will propose doing this.

Other update strategies (except booting the install media) are safe if the boot loader is configured properly. But the other strategies are not available, if you update from SUSE Linux Enterprise Server 10.

10.1.16 Upgrading MySQL to SUSE Linux Enterprise Server 11

CHECKIT:11_3

During the upgrade to SUSE Linux Enterprise Server 11 MySQL is also upgraded to the latest version. To complete this migration you may have to upgrade your data as described in the MySQL documentation.

10.1.17 Fine-Tuning Firewall Settings

SuSEfirewall2 is enabled by default, which means you cannot log in from remote systems. This also interferes with network browsing and multicast applications, such as SLP and Samba ("Network Neighborhood"). You can fine-tune the firewall settings using YaST.

10.1.18 Upgrading from SUSE Linux Enterprise Server 10 SP4 with the Xen Hypervisor May Have Incorrect Network Configuration

CHECKIT:11_3

We have improved the network configuration: If you install SUSE Linux Enterprise Server 11 SP3 and configure Xen, you get a bridged setup through YaST.

However, if you upgrade from SUSE Linux Enterprise Server 10 SP4 to SUSE Linux Enterprise Server 11 SP3, the upgrade does not configure the bridged setup automatically.

To start the bridge proposal for networking, start the "YaST Control Center", choose "Virtualization", then "Install Hypervisor and Tools". Alternatively, call **yast2 xen** on the commandline.

10.1.19 LILO Configuration Via YaST or AutoYaST

The configuration of the LILO boot loader on the x86 and x86_64 architecture via YaST or AutoYaST is deprecated, and not supported anymore. For more information, see Novell TID 7003226 <http://www.novell.com/support/documentLink.do?externalID=7003226> ↗.

10.2 Update from SUSE Linux Enterprise Server 11

10.2.1 Changed Routing Behavior

SUSE Linux Enterprise Server 10 and SUSE Linux Enterprise Server 11 set `net.ipv4.conf.all.rp_filter = 1` in `/etc/sysctl.conf` with the intention of enabling route path filtering. However, the kernel fails to enable routing path filtering, as intended, by default in these products.

Since SLES 11 SP1, this bug is fixed and most simple single-homed unicast server setups will not notice a change. But it may cause issues for applications that relied on reverse path filtering being disabled (e.g., multicast routing or multi-homed servers).

10.2.2 Kernel Devel Packages

Starting with SUSE Linux Enterprise Server 11 Service Pack 1 the configuration files for recompiling the kernel were moved into their own sub-package:

kernel-flavor-devel

This package contains only the configuration for one kernel type (“flavor”), such as `default` or `desktop`.

10.3 Update from SUSE Linux Enterprise Server 11 SP1

The direct update from SUSE Linux Enterprise Server 11 SP1 to SP4 is not supported. For more information, see *Section 10.1.2, “Upgrading from SLES 10 (GA and Service Packs) or SLES 11 GA”*.

10.4 Update from SUSE Linux Enterprise Server 11 SP2

10.4.1 Support for 46bit memory addressing in makedumpfile and crash

The makedumpfile and crash utilities can now analyze memory dumps taken on systems with 46bit addresses.

10.4.2 Update from SUSE Linux Enterprise Server 11 SP2

Updating from SUSE Linux Enterprise Server 11 SP2 with AutoYaST is supported.

10.4.3 Migrating SUSE Linux Enterprise Server 11 SP2 with WebYaST Installed via wagon

For migrating SLES 11 SP2 with WebYaST installed to SP3 via wagon, it is necessary to install the WebYaST product metadata before starting the migration. To do so, make sure the packages "sle-11-SP2-WebYaST-release" and "sle-11-SP2-WebYaST-release-cd" are installed. You can ignore, if wagon reports an unknown registration status of WebYaST at the beginning of the migration.



Note

Without the WebYaST product metadata installed, WebYaST will not be migrated.

The product metadata are not needed when upgrading SLES via booting the installation media.

10.5 Update from SUSE Linux Enterprise Server 11 SP3

10.5.1 Update from SUSE Linux Enterprise Server 11 SP3

Updating from SUSE Linux Enterprise Server 11 SP3 with AutoYaST is supported.

11 Deprecated Functionality

11.1 Packages Removed with SUSE Linux Enterprise Server 11 SP4

The following packages were removed with the release of SUSE Linux Enterprise Server 11 SP4:

11.1.1 IBM Java 6

IBM Java 6 is no longer available on the SUSE Linux Enterprise SDK.

11.2 Packages Removed with SUSE Linux Enterprise Server 11 SP3

The following packages were removed with the release of SUSE Linux Enterprise Server 11 SP3:

N/A

11.3 Packages Removed with SUSE Linux Enterprise Server 11 Service Pack 2

The following packages were removed with the release of SUSE Linux Enterprise Server 11 Service Pack 2:

hyper-v-kmp

hyper-v-kmp has been removed.

32-bit Xen Hypervisor as a Virtualization Host

The 32-bit Xen hypervisor as a virtualization host is not supported anymore. 32-bit virtual guests are not affected and fully supported with the provided 64-bit hypervisor.

11.4 Packages Removed with SUSE Linux Enterprise Server 11 Service Pack 1

The following packages were removed with the release of SUSE Linux Enterprise Server 11 Service Pack 1:

brocade-bfa

The brocade-bfa kernel module is now part of the main kernel package.

enic-kmp

The enic kernel module is now part of the main kernel package.

fnic-kmp

The fnic kernel module is now part of the main kernel package.

kvm-kmp

The KVM kernel modules are now part of the main kernel package.

java-1_6_0-ibm-x86

11.5 Packages Removed with SUSE Linux Enterprise Server 11

The following packages were removed with the major release of SUSE Linux Enterprise Server 11:

dante

JFS

The JFS file system is no longer supported and the utilities have been removed from the distribution.

EVMS

Replaced with LVM2.

ippl

powertweak

SUN Java

uw-imapd

mapped-base Functionality

The mapped-base functionality, which is used by 32-bit applications that need a larger dynamic data space (such as database management systems), has been replaced with flexmap.

zmd

11.6 Packages and Features to Be Removed in the Future

The following packages and features are deprecated and will be removed with the next Service Pack or major release of SUSE Linux Enterprise Server:

- The reiserfs file system is fully supported for the lifetime of SUSE Linux Enterprise Server 11 specifically for migration purposes. We will however remove support for creating new reiserfs file systems starting with SUSE Linux Enterprise Server 12.
- The sendmail package is deprecated and might be discontinued with SUSE Linux Enterprise Server 12.
- The lprng package is deprecated and will be discontinued with SUSE Linux Enterprise Server 12.
- The dhcpcv6 package is deprecated and will be discontinued with SUSE Linux Enterprise Server 12.
- The qt3 package is deprecated and will be discontinued with SUSE Linux Enterprise Server 12.
- syslog-ng will be replaced with rsyslog.
- The smpppd package is deprecated and will be discontinued with one of the next Service Packs or SUSE Linux Enterprise Server 12.
- The raw block devices (major 162) are deprecated and will be discontinued with one of the next Service Packs or SUSE Linux Enterprise Server 12.

12 Infrastructure, Package and Architecture Specific Information

12.1 Systems Management

12.1.1 YaST: IPoIB Mode Configuration

In YaST, an IPoIB mode configuration is available.

12.1.2 Zypper: new option that shows enabled repo

A new option `lr` is available for `zypper` to show enabled repositories.

12.1.3 Modified Operation against Novell Customer Center

Effective on 2009-01-13, provisional registrations have been disabled in the Novell Customer Center. Registering an instance of SUSE Linux Enterprise Server or Open Enterprise Server (OES) products now requires a valid, entitled activation code. Evaluation codes for reviews or proofs of concept can be obtained from the product pages and from the download pages on novell.com.

If a device is registered without a code at setup time, a provisional code is assigned to it by Novell Customer Center (NCC), and it will be entered in your NCC list of devices. No update repositories are assigned to the device at this time.

Once you are ready to assign a code to the device, start the YaST Novell Customer Center registration module and replace the un-entitled provisional code that NCC generated with the appropriate one to fully entitle the device and activate the related update repositories.

12.1.4 Operation against Subscription Management Tool

Operation under the Subscription Management Tool (SMT) package and registration proxy is not affected. Registration against SMT will assign codes automatically from your default pool in NCC until all entitlements have been assigned. Registering additional devices once the pool is depleted will result in the new device being assigned a provisional code (with local access to updates) The SMT server will notify the administrator that these new devices need to be entitled.

12.1.5 Minimal Pattern

The minimal pattern provided in YaST's Software Selection dialog targets experienced customers and should be used as a base for your own specific software selections.

Do not expect a minimal pattern to provide a useful basis for your business needs without installing additional software.

This pattern does not include any dump or logging tools. To fully support your configuration, Novell Technical Services (NTS) will request installation of all tools needed for further analysis in case of a support request.

12.2 Performance Related Information

12.2.1 Linux Completely Fair Scheduler Affects Java Performance

Problem (Abstract)

Java applications that use synchronization extensively might perform poorly on Linux systems that include the Completely Fair Scheduler. If you encounter this problem, there are two possible workarounds.

Symptom

You may observe extremely high CPU usage by your Java application and very slow progress through synchronized blocks. The application may appear to hang due to the slow progress.

Cause

The Completely Fair Scheduler (CFS) was adopted into the mainline Linux kernel as of release 2.6.23. The CFS algorithm is different from previous Linux releases. It might change the performance properties of some applications. In particular, CFS implements `sched_yield()` differently, making it more likely that a thread that yields will be given CPU time regardless.

The new behavior of `sched_yield()` might adversely affect the performance of synchronization in the IBM JVM.

Environment

This problem may affect IBM JDK 5.0 and 6.0 (all versions) running on Linux kernels that include the Completely Fair Scheduler, including Linux kernel 2.6.27 in SUSE Linux Enterprise Server 11.

Resolving the Problem

If you observe poor performance of your Java application, there are two possible workarounds:

- Either invoke the JVM with the additional argument `"-Xthr:minimizeUserCPU"`.
- Or configure the Linux kernel to use the more backward-compatible heuristic for `sched_yield()` by setting the `sched_compat_yield` tunable kernel property to `1`. For example:

```
echo "1" > /proc/sys/kernel/sched_compat_yield
```

You should not use these workarounds unless you are experiencing poor performance.

12.2.2 Tuning Performance of Simple Database Engines

Simple database engines like Berkeley DB use memory mappings (`mmap(2)`) to manipulate database files. When the mapped memory is modified, those changes need to be written back to disk. In SUSE Linux Enterprise 11, the kernel includes modified mapped memory in its calculations for deciding when to start background writeback and when to throttle processes which modify additional memory. (In previous versions, mapped dirty pages were not accounted for and the amount of modified memory could exceed the overall limit defined.) This can lead to a decrease in performance; the fix is to increase the overall limit. The maximum amount of dirty memory is 40% in SUSE Linux Enterprise 11 by default. This value is chosen for average workloads, so that enough memory remains available for other uses. The following settings may be relevant when tuning for database workloads:

- `vm.dirty_ratio`
Maximum percentage of dirty system memory (default 40).
- `vm.dirty_background_ratio`

Percentage of dirty system memory at which background writeback will start (default 10).

- `vm.dirty_expire_centisecs`

Duration after which dirty system memory is considered old enough to be eligible for background writeback (in centiseconds).

These limits can be observed or modified with the `sysctl` utility (see `sysctl(1)` and `sysctl.conf(5)`).

12.3 Storage

12.3.1 SUSE Enterprise Storage (Powered by Ceph) Client

SUSE Linux Enterprise Server 11 SP3 and SP4 now provides the functionality to act as a client for SUSE Enterprise Storage. `qemu` can now use storage provided by the SUSE Enterprise Storage Ceph cluster via the RADOS Block Device (`rbd`) backend. Applications can now be enhanced to directly incorporate object or block storage backed by the SUSE Enterprise Storage cluster, by linking with the `librados` and `librbd` client libraries.

Also included is the `rbd` tool to manage RADOS block devices mapped via the `rbd` kernel module, for use as a standard generic block device.

12.3.2 Enable `set_4k_mtu` in Mellanox for All ConnectX Ports

The `mlx4_ib` Infiniband module allows for dynamic `mtu` configuration through a new, `per-port sysfs` entry.

The `mtu` may be set per port via the following `sysfs` entry:

- `/sys/class/infiniband/mlx4_{x}/device/mlx4_port{y}_mtu`

Here `{x}` is the HCA number, and `{y}` is the port number on the HCA

When writing a value to this `sysfs` entry, five values are allowed:

- 256
- 512
- 1024

- 2048
- 4096

All other values will result in an error.

12.3.3 Multipath Configuration Change

With the update to version 0.4.9 on SLES 11 SP2, rr_min_io is replaced by rr_min_io_rq in multipath.conf. The old option is now ignored. Check this setting, if you encounter performance issues.

For more information, see the “Storage Administration Guide” shipped with SLES 11 SP3.

12.3.4 Capturing kdump on a Target using Devicemapper (Incl. Multipath)

If the root device is *not* using devicemapper (multipath), as a workaround add additional parameters to KDUMP_COMMANDLINE_APPEND in /etc/sysconfig/kdump, to capture kdump on a target that is using devicemapper (multipath):

```
KDUMP_COMMANDLINE_APPEND="root_no_dm=1 root_no_mpath=1"
```

Then start the kdump service.

If you use multipath for both root and kdump, these options must not be added.

An example use case with System z could be a kdump target on multipath zfcp-attached SCSI devices and a root file system on DASD.

12.3.5 Multipathing: SCSI Hardware Handler

Some storage devices, e.g. IBM DS4K, require special handling for path failover and failback. In SUSE Linux Enterprise Server 10 SP2, dm layer served as hardware handler.

One drawback of this implementation was that the underlying SCSI layer did not know about the existence of the hardware handler. Hence, during device probing, SCSI would send I/O on the passive path, which would fail after a timeout and also print extraneous error messages in the console.

In SUSE Linux Enterprise Server 11, this problem is resolved by moving the hardware handler to the SCSI layer, hence the term SCSI Hardware Handler. These handlers are modules created under the SCSI directory in the Linux Kernel.

In SUSE Linux Enterprise Server 11, there are four SCSI Hardware Handlers: scsi_dh_alua, scsi_dh_rdac, scsi_dh_hp_sw, scsi_dh_emc.

These modules need to be included in the initrd image so that SCSI knows about the special handling during probe time itself.

To do so, carry out the following steps:

- Add the device handler modules to the INITRD_MODULES variable in /etc/sysconfig/kernel
- Create a new initrd with:

```
mkinitrd -k /boot/vmlinux-<flavour> \  
-i /boot/initrd-<flavour>-scsi_dh \  
-M /boot/System.map-<flavour>
```

- Update the grub.conf/lilo.conf/yaboot.conf file with the newly built initrd.
- Reboot.

12.3.6 Local Mounts of iSCSI Shares

An iSCSI shared device should never be mounted directly on the local machine. In an OCFS2 environment, doing so causes all hardware to hard hang.

12.4 Hyper-V

12.4.1 Hyper-V: Time Synchronization

The system time of a guest will drift several seconds per day.

To maintain an accurate system time it is recommended to run ntpd in a guest. The ntpd daemon can be configured with the YaST "NTP Client" module. In addition to such a configuration, the following two variables must be set manually to "yes" in /etc/sysconfig/ntp:


```
NTPD_FORCE_SYNC_ON_STARTUP="yes"
NTPD_FORCE_SYNC_HWCLOCK_ON_STARTUP="yes"
```

12.4.2 Change of Kernel Device Names in Hyper-V Guests

Starting with SP2, SLES 11 has a newer block device driver, which presents all configured virtual disks as SCSI devices. Disks, which used to appear as /dev/hda in SLES 11 SP1 will from now on appear as /dev/sda.

12.4.3 Using the "Virtual Machine Snapshot" Feature

The Windows Server Manager GUI allows to take snapshots of a Hyper-V guest. After a snapshot is taken the guest will fail to reboot. By default, the guest's root file system is referenced by the serial number of the virtual disk. This serial number changes with each snapshot. Since the guest expects the initial serial number, booting will fail.

The solution is to either delete all snapshots using the Windows GUI, or configure the guest to mount partitions by file system UUID. This change can be made with the YaST partitioner and boot loader configurator.

12.5 Architecture Independent Information

12.5.1 Changes in Packaging and Delivery

12.5.1.1 Updating tcsh

tcsh 6.15 has a locking issue when used concurrently.

On SLE 11 SP3, SUSE updated tcsh to version 6.18 to solve a locking issue when used concurrently.

12.5.1.2 New Ruby Packaging Scheme with the Update to Ruby 1.8

The different Ruby package versions cannot clearly be handled on one system with the old packaging scheme. To help packagers with the new scheme introduced with SLE 12, two new scripts in the ruby package helps to find the correct version suffix for new packages.

This improvement is now available as a backport in SLE 11 SP4, too.

12.5.1.3 New Package: vhostmd / vm-dump-metrics

vhostmd (Virtual Host Metrics Daemon) allows virtual machines to see limited information about the host they are running on. vm-dump-metrics runs inside the guest to dump the host metrics.

12.5.1.4 New Package: cloud-init

cloud-init is the defacto package that handles early initialization of a cloud instance.

To install it, enable the repository of the Public Cloud Module.

12.5.1.5 Update to IBM Java 7 Release 1

IBM Java 7 Release 1 provides performance improvements through IBM POWER8 and IBM zEnterprise EC12 exploitation.

12.5.1.6 SUSE Linux Enterprise High Availability Extension 11

With the *SUSE Linux Enterprise High Availability Extension 11*, SUSE offers the most modern open source High Availability Stack for Mission Critical environments.

12.5.1.7 Kernel Has Memory Cgroup Support Enabled By Default

While this functionality is welcomed in most environments, it requires about 1% of memory. Memory allocation is done at boot time and is using 40 Bytes per 4 KiB page which results in 1% of memory.

In virtualized environments, specifically but not exclusively on s390x systems, this may lead to a higher basic memory consumption: e.g., a 20GiB host with 200 x 1GiB guests consumes 10% of the real memory.

This memory is not swappable by Linux itself, but the guest cgroup memory is pageable by a z/VM host on an s390x system and might be swappable on other hypervisors as well.

Cgroup memory support is activated by default but it can be deactivated by adding the Kernel Parameter `cgroup_disable=memory`

A reboot is required to deactivate or activate this setting.

12.5.1.8 Kernel Development Files Moved to Individual kernel-\$flavor-devel Packages

Up to SLE 11 GA, the kernel development files (`.config`, `Module.symvers`, etc.) for all flavors were packaged in a single kernel-syms package. Starting with SLE 11 SP1, these files are packaged in individual kernel-\$flavor-devel packages, allowing to build KMPs for only the required kernel flavors. For compatibility with existing spec files, the kernel-syms package still exists and depends on the individual kernel-\$flavor-devel packages.

12.5.1.9 Live Migration of KVM Guest with Device Hot-Plugging

Hot-plugging a device (network, disk) works fine for a KVM guest on a SLES 11 host since SP1. However, migrating the same guest with the hotplugged device (available on the destination host) fails.

Since SLES 11 SP1, supports the hotplugging of the device to the KVM guest, but migrating the guest with the hot-plugged device is not supported and expected to fail.

12.5.2 Security

12.5.2.1 openldap2-client 2.4: New Options

These new options are especially noteworthy:

1. Specify the handshake protocol and the strength of minimally acceptable SSL/TLS ciphers for the operation of OpenLDAP server.
2. Specify the handshake protocol and the strength of proposed SSL/TLS ciphers for the operation of OpenLDAP client.

General information:

The parameter "TlsParameterMin" helps both use cases. The parameter value controls both handshake protocol and cipher strength. The interpretation of the value by server and client is identical, however the parameter name appears differently in server's and client's configuration files.

The value format is "X.Y" where X and Y are single digits:

- If X is 2, handshake is SSLv2, the usable ciphers are SSLv2 and up.
- If X is 3, handshake is TLSv1.0 (SLES 11) or TLSv1.2 (SLES 12), the usable ciphers are TLSv1.(Y-1) and up.

Examples:

- 2.0 - Handshake is SSLv2, usable ciphers are SSLv2, SSLv3, and TLSv1.x
- 2.1 - Same as above
- 3.1 - Handshake is TLSv1.0 (SLES 11), usable ciphers are SSLv3 and up.
- 3.2 - Handshake is TLSv1.0 (SLES 11), usable ciphers are TLSv1.1 and up.

Important: OpenSSL identifies TLSv1.0 ciphers as "SSLv3", if the parameter value prohibits SSLv3 operation, then TLSv1.0 ciphers will be rejected too, and vice versa.

Use case 1:

Supported by SLES 12 only. SLES 11 is too old to support this use case. Add parameter TLSProtocolMin to slapd.conf and restart server.

Example - reject SSLv2 handshake, accept TLSv1.0 handshake and TLSv1.x ciphers:

```
TLSProtocolMin 3.1
```

Use case 2:

Supported by both SLE 12 and SLE 11 server and desktop products. Add parameter TLS_PROTOCOL_MIN to either /etc/openldap/ldap.conf or ~/.ldaprc.

Example - do not use SSLv2 handshake, use TLSv1.0 handshake, and propose SSLv3 and TLSv1.x ciphers:

```
TLS_PROTOCOL_MIN 3.1
```

Debug tips for Client operation:

Run ldap client programs with debug level 5 (-d 5) will trace TLS operations. Be aware that OpenSSL will misleadingly print this message:

```
SSL_connect:SSLv2/v3 write client hello A
```

which apparently suggests the usage of SSLv2, but in fact OpenSSL has not decided on the handshake protocol yet!

References:

- Original feature commit by OpenLDAP developers: <http://www.openldap.org/its/index.cgi/Software%20Enhancements?id=5655>
- OpenLDAP client configuration manual: <http://man7.org/linux/man-pages/man5/ldap.conf.5.html>
- OpenLDAP server configuration manual (note the lack of TlsProtocolMin usage instruction): <http://www.openldap.org/doc/admin24/tls.html>

12.5.2.2 Kdump Over Network Via SSH

For improved security, the following steps are now required:

1. Add the root user's public key to the dump user's authorized_keys on the target machine.
2. If the target machine's host key is not in machine's known_hosts file, the following additional steps are required:
 1. Add the target machine's host key to the dump machine's known_hosts file.
 2. Re-generate the kdump initrd (mkdumprd -f).
 3. Restart kdump (rckdump restart).

12.5.2.3 ECDSA Support for kdump Over SSH

When saving kernel dumps over SSH, kdump has relied on an external library (libssh2). This library only supports RSA and DSA keys. The default SSH key type in SLE 11 SP4 is ECDSA, which is not supported by libssh2. Consequently, if a SUSE Linux Enterprise product is installed on the target machine, the admin must configure it with one of the supported ciphers.

Kernel dumping now uses the "ssh" binary to establish the connection. This automatically enables all ciphers and key exchange protocols that are supported by the openssh2 package, including forthcoming additions.

12.5.2.4 New Symlink and Hardlink Security Restrictions

A number of security breaches is based on symlink and hardlink exploitation.

To increase the system robustness, restrictions on link creation have been implemented in the operating system kernel. The mechanism can be switched off by the system administrator by setting the fs.protected_hardlinks and fs.protected_symlinks sysctl parameters to 0.

12.5.2.5 Removable Media

To allow a specific user ("joe") to mount removable media, run the following command as root:

```
polkit-auth --user joe \  
--grant org.freedesktop.hal.storage.mount-removable
```

To allow all locally logged in users on the active console to mount removable media, run the following commands as root:

```
echo 'org.freedesktop.hal.storage.mount-removable no:no:yes' \  
>> /etc/polkit-default-privs.local  
/sbin/set_polkit_default_privs
```

12.5.2.6 Verbose Audit Records for System User Management Tools

Install the package "pwdutils-plugin-audit". To enable this plugin, add "audit" to /etc/pwdutils/logging. See the "Security Guide" for more information.

12.5.3 Networking

12.5.3.1 openssl1 Enablement

Customers require TLS 1.2 support in the openssl1 library, partially for their own programs, but also for selected SUSE ones.

We provide openssl1 enablement packages in a separate repository.

12.5.3.2 Improved Samba libsmbcclient Support for Microsoft Distributed File System (DFS)

libsmbcclient previously resolved DFS referrals on every API call, always using the first entry in the referral response. With random DFS referral ordering, libsmbcclient would often open a new server connection, rather than reusing an existing (cached) connection established from a previous DFS referred API call.

libsmbcclient now checks the connection cache for any of the DFS referral response entries before creating a new connection.

12.5.3.3 Providing TLS 1.2 Support for Apache2 Via mod_nss

The Apache Web server offers HTTPS protocol support via mod_ssl, which in turn uses the openssl shared libraries. SUSE Linux Enterprise Server 11 SP2 and SP3 come with openssl version 0.9.8j. This openssl version supports TLS version up to and including TLSv1.0, support for newer TLS versions like 1.1 or 1.2 is missing.

Recent recommendations encourage the use of TLSv1.2, specifically to support Perfect Forward Secrecy. To overcome this limitation, the SUSE Linux Enterprise Server 11 SP2, SP3, and SP4 are supplied with upgrades to recent versions of the mozilla-nss package and with the package apache2-mod_nss, which makes use of mozilla-nss for TLSv1.2 support for the Apache Web server.

An additional mod_nss module is supplied for apache2, which can coexist with all existing libraries and apache2 modules. This module uses the mozilla netscape security services library, which supports TLS 1.1 and TLS 1.2 protocols. It is not a drop-in replacement; configuration and certificate storages are different. It can coexist with mod_ssl if necessary.

The package includes a sample configuration and a README-SUSE.txt for setup guidance.

12.5.3.4 Bind Update to Version 9.9

The DNS Server Bind has been updated to the long term supported version 9.9 for longer stability going forward. In version 9.9, the commands 'dnsssec-makekeyset' and 'dnsssec-signkey' are not available anymore.

DNSSEC tools provided by Bind 9.2.4 are not compatible with Bind 9.9 and later and have been replaced where applicable. Specifically, DNSSEC-bis functionality removes the need for dnsssec-signkey(1M) and dnsssec-makekeyset(1M); dnsssec-keygen(1M) and dnsssec-signzone(1M) now provide alternative functionality.

For more information, see [TID 7012684](https://www.suse.com/support/kb/doc.php?id=7012684) (<https://www.suse.com/support/kb/doc.php?id=7012684>)  (<https://www.suse.com/support/kb/doc.php?id=7012684>).

12.5.3.5 Enabling NFS 4.1 for nfsd

Support for NFS 4.1 is now available.

The parameter `NFS4_SERVER_MINOR_VERSION` is now available in `/etc/nfs/syconfig` for setting the supported minor version of NFS 4.

12.5.3.6 Mounting NFS Volumes Locally on the Exporting Server

Mounting NFS volumes locally on the exporting server is not supported on SUSE Linux Enterprise systems, as it is the case on all Enterprise class Linux systems.

12.5.3.7 Loading the mlx4_en Adapter Driver with the Mellanox ConnectX2 Ethernet Adapter

There is a reported problem that the Mellanox ConnectX2 Ethernet adapter does not trigger the automatic load of the `mlx4_en` adapter driver. If you experience problems with the `mlx4_en` driver not automatically loading when a Mellanox ConnectX2 interface is available, create the file `mlx4.conf` in the directory `/etc/modprobe.d` with the following command:

```
install mlx4_core /sbin/modprobe --ignore-install mlx4_core \  
&& /sbin/modprobe mlx4_en
```



12.5.3.8 Using the System as a Router

As long as the firewall is active, the option `ip_forwarding` will be reset by the firewall module. To activate the system as a router, the variable `FW_ROUTE` has to be set, too. This can be done through **yast2** firewall or manually.

12.5.4 Cross Architecture Information

12.5.4.1 Myricom 10-Gigabit Ethernet Driver and Firmware

SUSE Linux Enterprise 11 (x86, x86_64 and IA64) is using the Myri10GE driver from mainline Linux kernel. The driver requires a firmware file to be present, which is not being delivered with SUSE Linux Enterprise 11.

Download the required firmware at <http://www.myricom.com> .

12.6 AMD64/Intel64 64-Bit (x86_64) and Intel/AMD 32-Bit (x86) Specific Information

12.6.1 System and Vendor Specific Information

12.6.1.1 Installation on 4KB Sector Drives Not Supported

Legacy installations are not supported on 4KB sector drives that are installed in x86/x86_64 servers. (UEFI installations and the use of the 4KB sector disks as non-boot disks are supported).

12.6.1.2 Insecurity with XEN on Some AMD Processors

This hardware flaw ("AMD Erratum #121") is described in "Revision Guide for AMD Athlon 64 and AMD Opteron Processors" (<http://support.amd.com/TechDocs/25759.pdf> 

The following 130nm and 90nm (DDR1-only) AMD processors are subject to this erratum:

- First-generation AMD-Opteron(tm) single and dual core processors in either 939 or 940 packages:
 - AMD Opteron(tm) 100-Series Processors
 - AMD Opteron(tm) 200-Series Processors
 - AMD Opteron(tm) 800-Series Processors
 - AMD Athlon(tm) processors in either 754, 939 or 940 packages
 - AMD Sempron(tm) processor in either 754 or 939 packages
 - AMD Turion(tm) Mobile Technology in 754 package
- This issue does not affect Intel processors.

(End quoted text.)

As this is a hardware flaw. It is not fixable except by upgrading your hardware to a newer revision, or not allowing untrusted 64-bit guest systems, or accepting that someone stops your machine. The impact of this flaw is that a malicious PV guest user can halt the host system.

The SUSE XEN updates will fix it via disabling the boot of XEN GUEST systems. The HOST will boot, just not start guests. In other words: If the update is installed on the above listed AMD64 hardware, the guests will no longer boot by default.

To reenale booting, the "allow_unsafe" option needs to be added to XEN_APPEND in /etc/sysconfig/bootloader as follows:

```
XEN_APPEND="allow_unsafe"
```

12.6.1.3 Boot Device Larger than 2 TiB

Due to limitations in the legacy x86/x86_64 BIOS implementations, booting from devices larger than 2 TiB is technically not possible using legacy partition tables (DOS MBR).

Since SUSE Linux Enterprise Server 11 Service Pack 1 we support installation and boot using uEFI on the x86_64 architecture and certified hardware.

12.6.1.4 i586 and i686 Machines with More than 16 GB of Memory

Depending on the workload, i586 and i686 machines with 16GB-48GB of memory can run into instabilities. Machines with more than 48GB of memory are not supported at all. Lower the memory with the `mem=` kernel boot option.

In such memory scenarios, we strongly recommend using a x86-64 system with 64-bit SUSE Linux Enterprise Server, and run the (32-bit) x86 applications on it.

12.6.1.5 Directly Addressable Memory on x86 Machines

When running SLES on an x86 machine, the kernel can only address 896MB of memory directly. In some cases, the pressure on this memory zone increases linearly according to hardware resources such as number of CPUs, amount of physical memory, number of LUNs and disks, use of multipath, etc.

To workaround this issue, we recommend running an x86_64 kernel on such large server machines.

12.6.1.6 NetXen 10G Ethernet Expansion Card on IBM BladeCenter HS12 System

When installing SUSE Linux Enterprise Server 11 on a HS12 system with a "NetXen Incorporated BladeCenter-H 10 Gigabit Ethernet High Speed Daughter Card", the boot parameter `pcie_aspm=off` should be added.

12.6.1.7 NIC Enumeration

Ethernet interfaces on some hardware do not get enumerated in a way that matches the marking on the chassis.

12.6.1.8 Service Pack for HP Linux ProLiant

The hpilo driver is included in SUSE Linux Enterprise Server 11. Therefore, no hp-ilo package will be provided in the Linux ProLiant Service Pack for SUSE Linux Enterprise Server 11.

For more details, see Novell TID 7002735 <http://www.novell.com/support/documentLink.do?externalID=7002735> .

12.6.1.9 HP High Performance Mouse for iLO Remote Console.

The desktop in SUSE Linux Enterprise Server 11 now recognizes the HP High Performance Mouse for iLO Remote Console and is configured to accept and process events from it. For the desktop mouse and the HP High Performance Mouse to stay synchronized, it is necessary to turn off mouse acceleration. As a result, the HP iLO2 High-Performance mouse (hpmouse) package is no longer needed with SUSE Linux Enterprise Server 11 once one of the following three options are implemented.

1. In a terminal run `xset m 1` — this setting will not survive a reset of the desktop.
2. (Gnome) In a terminal run `gconf-editor` and go to desktop->gnome->peripherals->mouse. Edit the "motion acceleration" field to be 1.
(KDE) Open "Personal Settings (Configure Desktop)" in the menu and go to "Computer Administration->Keyboard&Mouse->Mouse->Advanced" and change "Pointer Acceleration" to 1.
3. (Gnome) In a terminal run "gnome-mouse-properties" and adjust the "Pointer Speed" slide scale until the HP High Performance Mouse and the desktop mouse run at the same speed across the screen. The recommended adjustment is close to the middle, slightly on the "Slow" side.

After acceleration is turned off, sync the desktop mouse and the ILO mouse by moving to the edges and top of the desktop to line them up in the vertical and horizontal directions. Also if the HP High Performance Mouse is disabled, pressing the <Ctrl> key will stop the desktop mouse and allow easier synching of the two pointers.

For more details, see Novell TID 7002735 <http://www.novell.com/support/documentLink.do?externalID=7002735>.

12.6.1.10 Missing 32-Bit Compatibility Libraries for libstdc++ and libg++ on 64-Bit Systems (x86_64)

32-bit (x86) compatibility libraries like "libstdc++-libc6.2-2.so.3" have been available on x86_64 in the package "compat-32-bit" with SUSE Linux Enterprise Server 9, SUSE Linux Enterprise Server 10, and are also available on the SUSE Linux Enterprise Desktop 11 medium (compat-32-bit-2009.1.19), but are not included in SUSE Linux Enterprise Server_11.

Background

The respective libraries have been deprecated back in 2001 and shipped in the compatibility package with the release of SUSE Linux Enterprise Server 9 in 2004. The package was still shipped with SUSE Linux Enterprise Server 10 to provide a longer transition period for applications requiring the package.

With the release of SUSE Linux Enterprise Server 11 the compatibility package is no longer supported.

Solution

In an effort to enable a longer transition period for applications still requiring this package, it has been moved to the unsupported "Extras" channel. This channel is visible on every SUSE Linux Enterprise Server 11 system, which has been registered with the Novell Customer Center. It is also mirrored via SMT alongside the supported and maintained SUSE Linux Enterprise Server 11 channels.

Packages in the "Extras" channel are not supported or maintained.

The compatibility package is part of SUSE Linux Enterprise Desktop 11 due to a policy difference with respect to deprecation and deprecated packages as compared to SUSE Linux Enterprise Server 11.

We encourage customers to work with SUSE and SUSE's partners to resolve dependencies on these old libraries.

12.6.1.11 32-Bit Devel-Packages Missing from the Software Development Kit (x86_64)

Example: libcap0-devel-32-bit package was available in Software Development Kit 10, but is missing from Software Development Kit 11

Background

SUSE supports running 32-bit applications on 64-bit architectures; respective runtime libraries are provided with SUSE Linux Enterprise Server 11 and fully supported. With SUSE Linux Enterprise 10 we also provided 32-bit devel packages on the 64-bit Software Development Kit. Having 32-bit devel packages and 64-bit devel packages installed in parallel may lead to side-effects during the build process. Thus with SUSE Linux Enterprise 11 we started to remove some (but not yet all) of the 32-bit devel packages from the 64-bit Software Development Kit.

Solution

With the development tools provided in the Software Development Kit 11, customers and partners have two options to build 32-bit packages in a 64-bit environment (see below). Beyond that, SUSE's appliance offerings provide powerful environments for software building, packaging and delivery.

- Use the "build" tool, which creates a chroot environment for building packages.
- The Software Development Kit contains the software used for the Open Build Service. Here the abstraction is provided by virtualization.

12.6.2 Virtualization

12.6.2.1 XEN: Watchdog Usage

Multiple XEN watchdog instances are not supported. Enabling more than one instance can cause system crashes.


12.6.2.2 Inclusion of the virt-top tools

virt-top is a top-like utility for showing stats of virtualized domains. Many keys and command line options are the same as for ordinary top .

12.6.2.3 open-vm-tools Now Included

In the past, it was necessary to install VMware tools separately, because they had not been shipped with the distribution.

SUSE Linux Enterprise 11 SP4 includes the open-vm-tools package. These tools are pre-selected when installing on a VMware platform.

Partnering with VMware, SUSE provides full support for these tools. For more information, see <http://kb.vmware.com/kb/2073803>  .

12.6.2.4 Xen: Kernel Dom0 and Raw Hardware Characteristics

Because the kernel dom0 is running virtualized, tools such as irqbalance or lscpu will not reflect the raw hardware characteristics.

12.6.2.5 VMware: Enabling X2APIC

For improved performance, X2APIC is now supported.

12.6.2.6 KVM

Since SUSE Linux Enterprise Server 11 SP1, KVM is fully supported on the x86_64 architecture. KVM is designed around hardware virtualization features included in both AMD (AMD-V) and Intel (VT-x) CPUs produced within the past few years, as well as other virtualization features in even more recent PC chipsets and PCI devices. For example, device assignment using IOMMU and SR-IOV.

The following website identifies processors, which support hardware virtualization:

- http://en.wikipedia.org/wiki/X86_virtualization 

The KVM kernel modules will not load if the basic hardware virtualization features are not present and enabled in the BIOS. If KVM does not start, please check the BIOS settings.

KVM allows for memory overcommit and disk space overcommit. It is up to the user to understand the impact of doing so. Hard errors resulting from exceeding available resources will result in guest failures. CPU overcommit is supported but carries performance implications.

KVM supports a number of storage caching strategies which may be employed when configuring a guest VM. There are important data integrity and performance implications when choosing a caching mode. As an example, cache=writeback is not as safe as cache=none. See the online "SUSE Linux Enterprise Server Virtualization with KVM" documentation for details.

The following guest operating systems are supported:

- Starting with SLES 11 SP2, Windows guest operating systems are fully supported on the KVM hypervisor, in addition to Xen. For the best experience, we recommend using WHQL-certified virtio drivers, which are part of SLE VMDP.
SUSE Linux Enterprise Server 11 SP2 and SP3 as fully virtualized. The following virtualization aware drivers are available: kvm-clock, virtio-net, virtio-block, virtio-balloon
- SUSE Linux Enterprise Server 10 SP3 and SP4 as fully virtualized. The following virtualization aware drivers are available: kvm-clock, virtio-net, virtio-block, virtio-balloon
- SUSE Linux Enterprise Server 9 SP4 as fully virtualized. For 32-bit kernel, specify clock=pmtmr on the Linux boot line; for 64-bit kernel, specify ignore_lost_ticks on the Linux boot line.

For more information, see </usr/share/doc/packages/kvm/kvm-supported.txt>.

12.6.2.7 VMI Kernel (x86, 32-bit only)

VMware, SUSE and the community improved the kernel infrastructure in a way that VMI is no longer necessary. Starting with SUSE Linux Enterprise Server 11 SP1, the separate VMI kernel flavor is obsolete and therefore has been dropped from the media. When upgrading the system, it will be automatically replaced by the PAE kernel flavor. The PAE kernel provides all features, which were included in the separate VMI kernel flavor.

12.6.2.8 CPU Overcommit and Fully Virtualized Guest

Unless the hardware supports Pause Loop Exiting (Intel) or Pause Intercept Filter (AMD) there might be issues with fully virtualized guests with CPU overcommit in place becoming unresponsive or hang under heavy load.

Paravirtualized guests work flawlessly with CPU overcommit under heavy load.

This issue is currently being worked on.

12.6.2.9 IBM System x x3850/x3950 with ATI Radeon 7000/VE Video Cards and Xen Hypervisor

When installing SUSE Linux Enterprise Server 11 on IBM System x x3850/x3950 with ATI Radeon 7000/VE video cards, the boot parameter 'vga=0x317' needs to be added to avoid video corruption during the installation process.

Graphical environment (X11) in Xen is not supported on IBM System x x3850/x3950 with ATI Radeon 7000/VE video cards.

12.6.2.10 Video Mode Selection for Xen Kernels

In a few cases, following the installation of Xen, the hypervisor does not boot into the graphical environment. To work around this issue, modify /boot/grub/menu.lst and replace vga=<number> with vga=mode-<number>. For example, if the setting for your native kernel is vga=0x317, then for Xen you will need to use vga=mode-0x317.

12.6.2.11 Time Synchronization in virtualized Domains with NTP

Paravirtualized (PV) DomUs usually receive the time from the hypervisor. If you want to run "ntp" in PV DomUs, the DomU must be decoupled from the Dom0's time. At runtime, this is done with:

```
echo 1 > /proc/sys/xen/independent_wallclock
```

To set this at boot time:

1. either append "independent_wallclock=1" to kernel cmd line in DomU's grub configuration file
2. or append "xen.independent_wallclock = 1" to /etc/sysctl.conf in the DomU.

If you encounter time synchronization issues with Paravirtualized Domains, we encourage you to use NTP.

12.7 Intel Itanium (ia64) Specific Information

12.7.1 Installation on Systems with Many LUNs (Storage)

While the number of LUNs for a running system is virtually unlimited, we suggest not having more than 64 LUNs online while installing the system, to reduce the time to initialize and scan the devices and thus reduce the time to install the system in general.

12.8 POWER (ppc64) Specific Information

12.8.1 ppc64-diag

ppc64-diag is a RAS package used to retrieve platform error logs and take some of the actions (like DLPAR, EPOW actions, etc). It is not part of the base installation pattern for Power on SLES 11 SP4 but should be installed manually on any Power system.

12.8.2 GPT (GUID Partition Tables) Support

Support for installation on GPT (GUID Partition Tables) is now available.

12.8.3 Supported Hardware and Systems

All POWER3, POWER4, PPC970 and RS64-based models that were supported by SUSE Linux Enterprise Server 9 are no longer supported.

12.8.4 Using btrfs as /root File System on IBM Power Systems

Configure a minimum of 32MB for the PReP partition when using btrfs as the /root file system.

12.8.5 Loading the Installation Kernel via Network on POWER

With SUSE Linux Enterprise Server 11 the bootfile DVD1/suseboot/inst64 can not be booted directly via network anymore, because its size is larger than 12MB. To load the installation kernel via network, copy the files yaboot.ibm, yaboot.cnf and inst64 from the DVD1/suseboot directory to the TFTP server. Rename the yaboot.cnf file to yaboot.conf. yaboot can also load config files for specific Ethernet MAC addresses. Use a name like yaboot.conf-01-23-45-ab-cd-ef to match a MAC address. An example yaboot.conf for TFTP booting looks like this:

```
default=sles11
timeout=100
image[64-bit]=inst64
    label=sles11
    append="quiet install=nfs://hostname/exported/sles11dir"
```

12.8.6 Huge Page Memory Support on POWER

Huge Page Memory (16GB pages, enabled via HMC) is supported by the Linux kernel, but special kernel parameters must be used to enable this support. Boot with the parameters "hugepagesz=16G hugepages=N" in order to use the 16GB huge pages, where N is the number of 16GB pages assigned to the partition via the HMC. The number of 16GB huge pages available can not be changed once the partition is booted. Also, there are some restrictions if huge pages are assigned to a partition in combination with eHEA / eHCA adapters:

IBM eHEA Ethernet Adapter:

The eHEA module will fail to initialize any eHEA ports if huge pages are assigned to the partition and Huge Page kernel parameters are missing. Thus, no huge pages should be assigned to the partition during a network installation. To support huge pages after installation, the huge page kernel parameters need to be added to the boot loader configuration before huge pages are assigned to the partition.

IBM eHCA InfiniBand Adapter:


The current eHCA device driver is not compatible with huge pages. If huge pages are assigned to a partition, the device driver will fail to initialize any eHCA adapters assigned to the partition.

12.8.7 Installation on POWER onto IBM VSCSI Target


The installation on a vscsi client will fail with old versions of the AIX VIO server.

Solution: Upgrade the AIX VIO server to version 1.5.2.1-FP-11.1 or later.


12.8.8 IBM Linux VSCSI Server Support in SUSE Linux Enterprise Server 11


Customers using SLES 9 or SLES 10 to serve Virtual SCSI to other LPARs, using the ibmvscsis driver, who wish to migrate from these releases, should consider migrating to the IBM Virtual I/O server. The IBM Virtual I/O server supports all the IBM PowerVM virtual I/O features and also provides integration with the Virtual I/O management capabilities of the HMC. It can be downloaded from: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/download/home.html> 

12.8.9 Virtual Fibre Channel Devices

When using IBM Power Virtual Fibre Channel devices utilizing N-Port ID Virtualization, the Virtual I/O Server may need to be updated in order to function correctly. Linux requires VIOS 2.1, Fixpack 20.1, and the LinuxNPIV I-Fix for this feature to work properly. These updates can be downloaded from: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html> 

12.8.10 Virtual Tape Devices

When using virtual tape devices served by an AIX VIO server, the Virtual I/O Server may need to be updated in order to function correctly. The latest updates can be downloaded from: <http://www14.software.ibm.com/webapp/set2/sas/f/vios/home.html> 

For more information about IBM Virtual I/O Server, see <http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/home.html> .

12.8.11 Chelsio cxgb3 iSCSI Offload Engine

The Chelsio hardware supports ~16K packet size (the exact value depends on the system configuration). It is recommended that you set the parameter `MaxRecvDataSegmentLength` in `/etc/iscsid.conf` to 8192.

For the cxgb3i driver to work properly, this parameter needs to be set to 8192.

In order to use the cxgb3i offload engine, the cxgb3i module needs to be loaded manually after open-scsi has been started.

For additional information, refer to `/usr/src/linux/Documentation/scsi/cxgb3i.txt` in the kernel source tree.

12.8.12 Known TFTP Issues with Yaboot

When attempting to netboot yaboot, users may see the following error message:

```
Can't claim memory for TFTP download (01800000 @ 01800000-04200000)
```

and the netboot will stop and immediately display the yaboot "boot:" prompt. Use the following steps to work around the problem.

- Reboot the system and at the IBM splash screen select '8' to get to an Open Firmware prompt "0>"
- At the Open Firmware prompt, type the following commands:

```
setenv load-base 4000
setenv real-base c00000
dev /packages/gui obe
```

- The second command will take the system back to the IBM splash screen and the netboot can be attempted again.

12.8.13 Graphical Administration of Remotely Installed Hardware

If you do a remote installation in text mode, but want to connect to the machine later in graphical mode, be sure to set the default runlevel to 5 via YaST. Otherwise xdm/kdm/gdm might not be started.

12.8.14 InfiniBand - SDP Protocol Not Supported on IBM Hardware

To disable SDP on IBM hardware set SDP=no in `openib.conf` so that by default SDP is not loaded. After you have set this setting in `openib.conf` to 'no' run **openibd restart** or reboot the system for this setting to take effect.

12.8.15 RDMA NFS Server May Hang During Shutdown (OFED)

If your system is configured as an NFS over RDMA server, the system may hang during a shutdown if a remote system has an active NFS over RDMA mount. To avoid this problem, prior to shutting down the system, run "openibd stop"; run it in the background, because the command will hang and otherwise block the console:

```
/etc/init.d/openibd stop &
```

A shutdown can now be run cleanly.

The steps to configure and start NFS over RDMA are as follows:

- On the server system:

1. Add an entry to the file `/etc/exports`, for example:

```
/home  
192.168.0.34/255.255.255.0(fsid=0,rw,async,insecure,no_root_squash)
```

2. As the root user run the commands:

```
/etc/init.d/nfsserver start  
echo rdma 20049 > /proc/fs/nfsd/portlist
```

- On the client system:

1. Run the command: **modprobe xprtrdma**.
2. Mount the remote file system using the command **/sbin/mount.nfs**. Specify the ip address of the ip-over-ib network interface (ib0, ib1...) of the server and the options: **proto=rdma,port=20049**, for example:


```
/sbin/mount.nfs 192.168.0.64:/home /mnt \  
-o proto=rdma,port=20049,nolock
```

12.8.16 XFS Stack Overflow

Under heavy IO load on a fragmented filesystem, XFS can overflow the stack on ppc64 architecture leading to system crash.

This problem is fixed with the first SLE 11 SP3 maintenance update. The released kernel version is 3.0.82-0.7.9

12.9 System z (s390x) Specific Information

Look at http://www.ibm.com/developerworks/linux/linux390/documentation_novell_suse.html  for more information.

IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114) further on referred to as z196 and z114.

12.9.1 Hardware

12.9.1.1 IBM System z Architecture Level Set (ALS) Preparation

To exploit new IBM System z architecture capabilities during the lifecycle of SUSE Linux Enterprise Server 11, support for machines of the types z900, z990, z800, z890 is deprecated in this release. SUSE plans to introduce an ALS earliest with SUSE Linux Enterprise Server 11 Service Pack 1 (SP1), latest with SP2. After ALS, SUSE Linux Enterprise Server 11 only executes on z9 or newer processors.

With SUSE Linux Enterprise Server 11 GA, only machines of type z9 or newer are supported.

When developing software, we recommend to switch gcc to z9/z10 optimization:

- install gcc
- install gcc-z9 package (change gcc options to -march=z9-109 -mtune=z10)

12.9.1.2 Minimum Storage Firmware Level for LUN Scanning

For LUN Scanning to work properly, the minimum storage firmware level should be:

- DS8000 Code Bundle Level 64.0.175.0
- DS6000 Code Bundle Level 6.2.2.108

12.9.1.3 Issue with SLES 11 and NSS under z/VM

Starting SLES 11 under z/VM with NSS sometimes causes a guest to logoff by itself.

Solution: IBM addresses this issue with APAR VM64578.

12.9.2 Virtualization

12.9.2.1 Support of Live Guest Relocation (LGR) with z/VM 6.2

Live guest relocation (LGR) with z/VM 6.2 requires z/VM service applied, especially with Collaborative Memory Management (CMMA) active (cmma=on).

Apply z/VM APAR VM65134.

12.9.2.2 Linux Guests Running on z/VM 5.4 and 6.1 Require z/VM Service Applied

Linux guests using dedicated devices may experience a loop, if an available path to the device goes offline prior to the IPL of Linux.

Apply recommended z/VM service APARs VM65017 and VM64847.

12.9.3 Storage

12.9.3.1 QSAM Access Method for Data sharing with z/OS - Stage 1

This feature introduces a new interface that enables Linux applications like Data Stage to access and process (read only) data in z/OS owned physical sequential data sets without interfering with z/OS. By avoiding FTP or NFS transfer of data from z/OS the turnaround time for batch processing is significantly reduced.

12.9.4 Network

12.9.4.1 10GbE RoCE Express

SLES 11 SP4 supports the 10GbE RoCE Express feature on zEC12, zBC12, z13 via the TCP/IP layer without restrictions. SLES 11 SP4 includes RDMA enablement and DAPL/OFED for s390x as a technology preview but these can only be used on LPAR when running on IBM z Systems zEC12, zBC12 and cannot be used on IBM z Systems z13.

12.9.4.2 `src_vipa`: IPv6 Enablement

Adds support for IPv6 addresses to the `src_vipa` tool, that only supported IPv4 up to now

12.9.4.3 YaST May Fail to Activate Hipersocket Devices in Layer 2 Mode

In rare occasions Hipersocket devices in layer 2 mode may remain in softsetup state when configured via YaST.

Perform **ifup** manually.

12.9.4.4 YaST Sets an Invalid Default MAC Address for OSA Devices in Layer 2 Mode

OSA devices in layer 2 mode remain in softsetup state when "Set default MAC address" is used in YaST. Do not select "Set default MAC address" in YaST. If default MAC address got selected in YaST remove the line LLADDR='00:00:00:00:00:00' from the ifcfg file in /etc/sysconfig/network.

12.9.4.5 Limitations with the "qetharp" Utility

qetharp -d

Deleting: An ARP entry, which is part of Shared OSA should not get deleted from the arp cache.

Current Behavior: An ARP entry, which is part of shared OSA is getting deleted from the arp cache.

qetharp -p

Purging: It should remove all the remote entries, which are not part of shared OSA.

Current Behavior: It is only flushing out the remote entries, which are not part of shared OSA for first time. Then, if the user pings any of the purged ip address, the entry gets added back to the arp cache. Later, if the user runs purge for a second time, that particular entry is not getting removed from the arp cache.

12.9.5 RAS

12.9.5.1 Keywords for ipl and Console Device for Use in cio_ignore

Enable the use of keywords, "IPLDEV" for the IPL device and "CONDEV" for the console devices to ease installation when a system uses cio_ignore to blacklist all devices at install time and does not have a default CCW console device number, has no devices other than the IPL device as a base to clone Linux guests, or with ramdisk based installations with no devices other than the CCW console.

12.9.5.2 libica 2.4.2 Available since SLES 11 SP4 for s390x

The libica package contains the interface library routines used by IBM modules to interface with IBM Cryptographic Hardware (ICA). Starting with SLES 11 SP1, libica is provided in the s390x distribution in three flavors of packages: libica-1_3_9, libica-2_0_2, libica-2_1_0 and libica-2_4_2, providing libica versions 1.3.9, 2.0.2, 2.1.0 and 2.4.2 respectively.

libica 1.3.9 is provided for compatibility reasons with legacy hardware. For s390x users it is always recommended to use the new libica 2.4.2 library since it supports all newer s390x hardware, improved crypto usage statistics and is backwards compatible with any crypto device driver in the s390x architecture. You may choose to continue using libica 1.3.9, 2.0.2 or 2.1.0 if you do not have newer Cryptographic hardware to exploit or wish continue using custom applications that do not support the libica 2.4.2 library yet. Both openCryptoki and openssl-ibmca, the two main exploiters for the libica interface, are provided starting with SLES 11 SP4 to support the newer libica 2.4.2 library.

12.9.5.3 Support of Enterprise PKCS#11 (EP11)

Exploitation of Enterprise wide PKCS#11 (EP11) in CryptoExpress4 device driver and openCryptoki token for access to the Enterprise PKCS#11 (EP11) features of the CEX4S crypto adapter that implements certified PKCS#11 mechanism.

12.9.6 Performance

12.9.6.1 Add support for hardware sampling to the perf tool

With support for the CPU-measurement sampling facility available with IBM System z10 (z10) and later hardware the perf program can be used to capture performance data for processes, shared libraries, the kernel and device drivers.

12.9.6.2 snIPL Interface to Control Dynamic CPU Capacity

Remote control of the capacity of target systems in HA setups allows to maintain the bandwidth during failure situations and removes the need for keeping unused capacity activated during normal operation

13 Resolved Issues

- Bugfixes

This Service Pack contains all the latest bugfixes for each package released via the maintenance Web since the GA version. For details, see <https://bugzilla.suse.com> .

- Security Fixes

This Service Pack contains all the latest security fixes for each package released via the maintenance Web since the GA version.

- Program Temporary Fixes

This Service Pack contains all the PTFs (Program Temporary Fix) for each package released via the maintenance Web since the GA version which were suitable for integration into the maintained common codebase.

14 Technical Information

This section contains information about system limits, a number of technical changes and enhancements for the experienced user.

When talking about CPUs we are following this terminology:

CPU Socket

The visible physical entity, as it is typically mounted to a motherboard or an equivalent.

CPU Core

The (usually not visible) physical entity as reported by the CPU vendor.

On System z this is equivalent to an IFL.

Logical CPU

This is what the Linux Kernel recognizes as a "CPU".

We avoid the word "thread" (which is sometimes used), as the word "thread" would also become ambiguous subsequently.

Virtual CPU

A logical CPU as seen from within a Virtual Machine.

14.1 Kernel Limits

<http://www.suse.com/products/server/technical-information/#Kernel> 

This table summarizes the various limits which exist in our recent kernels and utilities (if related) for SUSE Linux Enterprise Server 11.

<i>SLES 11 (3.0)</i>	<i>x86</i>	<i>ia64</i>	<i>x86_64</i>	<i>s390x</i>	<i>ppc64</i>
CPU bits	32	64	64	64	64
max. # Logical CPUs	32	4096	4096	64	1024
max. RAM (theoretical / certified)	64/16 GiB	1 PiB/8+ TiB	64 TiB/16 TiB	4 TiB/256 GiB	1 PiB/512 GiB

<i>SLES 11 (3.0)</i>	<i>x86</i>	<i>ia64</i>	<i>x86_64</i>	<i>s390x</i>	<i>ppc64</i>
max. user-/ kernelspace	3/1 GiB	2 EiB/#	128 TiB/128 TiB	##	2 TiB/2 EiB
max. swap space	up to 29 * 64 GB (i386 and x86_64) or 30 * 64 GB (other architectures)				
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			
FD_SETSIZE	1024				

14.2 KVM Limits

Guest RAM size	2 TB
Virtual CPUs per guest	160
Maximum number of NICs per guest	8
Block devices per guest	4 emulated, 20 para-virtual
Maximum number of guests	Limit is defined as the total number of vCPUs in all guests being no greater than eight times the number of CPU cores in the host.

14.2.1 Virtualization: Supported Live Migration Scenarios

The following KVM host operating system combinations will be fully supported (L3) for live migrating guests from one host to another:

- VM from a SLES 12 host to SLES 12 host
- VM from a SLES 11 SP4 host to SLES 12 host

The following KVM host operating system combinations will be fully supported (L3) for live migrating guests from one host to another, later when released:

- VM from a SLES 12 host to SLES 12 SP1 host

All guests as outlined in the *Virtualization Guide* , chapter *Supported VM Guests* , are supported.

Backward migration is not supported:

- VM from a SLES 12 host to SLES 11 SP4 host
- VM from a SLES 11 SP4 host to SP3/SP3 host

14.2.2 KVM: QXL Video Driver

The QXL video driver is now available. With the QXL video virtual GPU you will get para-virtualized performance. Thus SLES 11 SP4 will run better as a guest under SLES 12.

14.2.3 KVM: QEMU Guest Agent

SLES11 SP4 now provide a QEMU Guest Agent to enable better controls and interactions with a SLES 11 SP4 guest.

14.2.4 KVM: Online disk resizing

KVM Guest are able to see the new size of it's disk after a resize on the host using the [virsh blockresize](#) command.

14.2.5 TLS Support for QEMU Websockets

Since SLE 11 SP3 we ship QEMU with TLS encryption support for QEMU Websockets. This feature allows every modern browser to create a secure VNC connection to QEMU without any additional plugins or configuration on the user side.

14.3 Xen Limits

<i>SLES 11 SP3</i>	<i>x86</i>
CPU bits	64
Logical CPUs (Xen Hypervisor)	256
Virtual CPUs per VM	64
Maximum supported memory (Xen Hypervisor)	2 TB
Maximum supported memory (Dom0)	500 GiB
Virtual memory per VM	16 GB (32-bit), 512 GB (64-bit)
Total virtual devices per host	2048
Maximum number of NICs per host	8
Maximum number of vNICs per guest	8
Maximum number of guests per host	64

In Xen 4.2, the hypervisor bundled with SUSE Linux Enterprise Server 11 SP3, dom0 is able to see and handle a maximum of 512 logical CPUs. The hypervisor itself, however, can access up to logical 256 logical CPUs and schedule those for the VMs.

With SUSE Linux Enterprise Server 11 SP2, we removed the 32-bit hypervisor as a virtualization host. 32-bit virtual guests are not affected and are fully supported with the provided 64-bit hypervisor.

14.3.1 Virtualization: Supported Disks Formats and Protocols

Currently, the disk formats raw , qed (only KVM), qcow (only Xen) and qcow2 has read-write (rw) support. The disk formats vmdk , vpc , and vhd/vhdx are only supported in read-only (ro) mode. The http , https , ftp , ftps , tftp protocols are supported for read-only access to images.

Under Xen the qed format will not be displayed as a selectable storage under virt-manager .

14.3.2 XEN: Update Xen to Version 4.4

Xen updated to Version 4.4.

14.3.3 Libvirt: Enhancement of virsh/libvirt "send-key" command

The `xm` and `xl` management tools have long supported sending SysRq keys to a domain using the `sysrq` subcommand. `virsh` now supports sending SysRq keys to a domain using the `send-key` subcommand.

14.3.4 Libvirt: enhancement of the virsh/libvirt "migrate" command

Similar to the `xm` tool, `virsh migrate` now supports options to control the process of migration. Large memory virtual machines running busy workloads pose migration challenges. Memory pages can be dirtied at a higher rate than they are transferred to the migration destination, resulting in prolonged migration time. The default migration algorithm will detect stalls due to high dirty page rate, suspend the virtual machine, and transfer the remaining dirty pages. In virtual machines running busy guest workloads, the final memory transfer can take considerable time, which could affect guest workloads sensitive to the time jump incurred when resuming the virtual machine. These new options allow fine-tuning the migration process

- `--max_iters <number>` : Number of transfer iterations before suspend (default: 30)
- `--max_factor <factor>` : Maximum amount of memory to transfer before suspend (default: 3*RAM)
- `--min_remaining <number>` : Minimum number of dirty pages remaining to be transferred before suspend (default: 50)
- `--abort_if_busy` : Abort migration instead of doing suspend and final memory transfer

14.4 File Systems

<https://www.suse.com/products/server/technical-information/#FileSystem> 

SUSE Linux Enterprise was the first enterprise Linux distribution to support journaling file systems and logical volume managers back in 2000. Today, we have customers running XFS and ReiserFS with more than 8TiB in one file system, and our own SUSE Linux Enterprise engineering team is using all 3 major Linux journaling file systems for all its servers.

We are excited to add the OCFS2 cluster file system to the range of supported file systems in SUSE Linux Enterprise.

We propose to use XFS for large-scale file systems, on systems with heavy load and multiple parallel read- and write-operations (e.g., for file serving with Samba, NFS, etc.). XFS has been developed for such conditions, while typical desktop use (single write or read) will not necessarily benefit from its capabilities.

Due to technical limitations (of the bootloader), we do not support XFS to be used for /boot.

<i>Feature</i>	<i>Ext 3</i>	<i>Reiserfs 3.6</i>	<i>XFS</i>	<i>Btrfs *</i>	<i>OCFS 2 **</i>
Data/Metadata Journaling	••	#/•	#/•	n/a *	#/•
Journal internal/external	••	••	••	n/a *	•/#
Offline extend/shrink	••	••	##	##	•/#
Online extend/shrink	•/#	•/#	•/#	••	•/#
Sparse Files	•	•	•	•	•
Tail Packing	#	•	#	•	#
Defrag	#	#	•	•	#
Extended Attributes/Access Control Lists	••	••	••	••	••
Quotas	•	•	•	^	•

<i>Feature</i>	<i>Ext 3</i>	<i>Reiserfs 3.6</i>	<i>XFS</i>	<i>Btrfs *</i>	<i>OCFS 2 **</i>
Dump/Restore	•	#	•	#	#
Blocksize default	4 KiB	4 KiB	4 KiB	4/64 KiB	4 KiB
max. File System Size	16 TiB	16 TiB	8 EiB	16 EiB	16 TiB
max. Filesize	2 TiB	1 EiB	8 EiB	16 EiB	1 EiB
	<p>* Btrfs is supported in SUSE Linux Enterprise Server 11 Service Pack3; Btrfs is a copy-on-write logging-style file system. Rather than journaling changes before writing them in-place, it writes them to a new location, then links it in. Until the last write, the new changes are not "committed". Due to the nature of the filesystem, quotas will be implemented based on subvolumes in a future release. The blocksize default varies with different host architectures. 64KiB is used on ppc64 and IA64, 4KiB on most other systems. The actual size used can be checked with the command "getconf PAGE_SIZE".</p>				
	<p>** OCFS2 is fully supported as part of the SUSE Linux Enterprise High Availability Extension.</p>				

The maximum file size above can be larger than the file system's actual size due to usage of sparse blocks. Note that unless a file system 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 file systems (including ext3 and ReiserFS) have LFS, which gives a maximum file size of 2^{63} bytes in theory. The numbers in the above tables assume that the file systems are using 4 KiB block size. When using different block sizes, the results are different, but 4 KiB reflects the most common standard.

In this document: 1024 Bytes = 1 KiB; 1024 KiB = 1 MiB; 1024 MiB = 1 GiB; 1024 GiB = 1 TiB; 1024 TiB = 1 PiB; 1024 PiB = 1 EiB. See also <http://physics.nist.gov/cuu/Units/binary.html>.

NFSv4 with IPv6 is only supported for the client side. A NFSv4 server with IPv6 is not supported.

This version of Samba delivers integration with Windows 7 Active Directory Domains. In addition we provide the clustered version of Samba as part of SUSE Linux Enterprise High Availability 11 SP3.

14.4.1 ext4: Runtime Switch for Write Support

The SUSE Linux Enterprise 11 kernel contains a fully supported ext4 file system module, which provides read-only access to the file system. A separate package is not required.

Read-write access to an ext4 file system can be enabled by using the `rw=1` module parameter. The parameter can be passed while loading the ext4 module manually, by adding it for automatic use by creating `/etc/modprobe.d/ext4` with the contents `options ext4 rw=1`, or after loading the module by writing `1` to `/sys/module/ext4/parameters/rw`. Note that read-write ext4 file systems are still officially unsupported by SUSE Technical Services.

ext4 is not supported for the installation of the SUSE Linux Enterprise operating system.

Since SLE 11 SP2 we support offline migration from ext4 to the supported btrfs file system.

The ext4-writeable package is still available for compatibility with systems with kernels from both the SLE 11 SP2 and SLE 11 SP3 releases installed.

14.5 Kernel Modules

An important requirement for every Enterprise operating system is the level of support a customer receives for his environment. Kernel modules are the most relevant connector between hardware ("controllers") and the operating system. Every kernel module in SUSE Linux Enterprise Server 11 has a flag 'supported' with three possible values: `"yes"`, `"external"`, `" "` (empty, not set, "unsupported").

The following rules apply:

- All modules of a self-recompiled kernel are by default marked as unsupported.
- Kernel Modules supported by SUSE partners and delivered using SUSE's Partner Linux Driver process are marked "external".
- If the `"supported"` flag is not set, loading this module will taint the kernel. Tainted kernels are not supported. To avoid this, not supported Kernel modules are included in an extra RPM (kernel-`<flavor>-extra`) and will not be loaded by default (`"flavor"=default|smp|xen|...`). In addition, these unsupported modules are not available in the installer, and the package `kernel-$flavor-extra` is not on the SUSE Linux Enterprise Server media.
- Kernel Modules not provided under a license compatible to the license of the Linux kernel will also taint the kernel; see `/usr/src/linux/Documentation/sysctl/kernel.txt` and the state of `/proc/sys/kernel/tainted`.

- Linux Kernel

The value of `/proc/sys/kernel/unsupported` defaults to 2 on SUSE Linux Enterprise Server 11 ("do not warn in syslog when loading unsupported modules"). This is the default used in the installer as well as in the installed system. See [`/usr/src/linux/Documentation/sysctl/kernel.txt`](#) for more information.

- `modprobe`

The **`modprobe`** utility for checking module dependencies and loading modules appropriately checks for the value of the "supported" flag. If the value is `"yes"` or `"external"` the module will be loaded, otherwise it will not. See below, for information on how to override this behavior.

Note: SUSE does not generally support removing of storage modules via **`modprobe -r`**.

Working with Unsupported Modules

While the general supportability is important, there might occur situations where loading an unsupported module is required (e.g., for testing or debugging purposes, or if your hardware vendor provides a hotfix):

- You can override the default by changing the variable `allow_unsupported_modules` in `/etc/modprobe.d/unsupported-modules` and set the value to `"1"`.

If you only want to try loading a module once, the `--allow-unsupported-modules` command-line switch can be used with `modprobe`. (For more information, see **`man modprobe`**).

- During installation, unsupported modules may be added through driver update disks, and they will be loaded.

To enforce loading of unsupported modules during boot and afterwards, please use the kernel command line option `oem-modules`.

While installing and initializing the `module-init-tools` package, the kernel flag `"TAINT_NO_SUPPORT"` (`/proc/sys/kernel/tainted`) will be evaluated. If the kernel is already tainted, `allow_unsupported_modules` will be enabled. This will prevent unsupported modules from failing in the system being installed. (If no unsupported modules are present during installation and the other special kernel command line option (`oem-modules=1`) is not used, the default will still be to disallow unsupported modules.)

- If you install unsupported modules after the initial installation and want to enable those modules to be loaded during system boot, please do not forget to run **`depmod`** and **`mkinitrd`**.

Remember that loading and running unsupported modules will make the kernel and the whole system unsupported by SUSE.

14.6 IPv6 Implementation and Compliance

SUSE Linux Enterprise Server 11 is compliant to IPv6 Logo Phase 2. However, when running the respective tests, you may see some tests failing. For various reasons, we cannot enable all the configuration options by default, which are necessary to pass all the tests. For details, see below.

- Section 3: RFC 4862 - IPv6 Stateless Address Autoconfiguration

Some tests fail because of the default DAD handling in Linux; disabling the complete interface is possible, but not the default behavior (because security-wise, this might open a DoS attack vector, a malicious node on a network could shutdown the complete segment) this is still conforming to RFC 4862: the shutdown of the interface is a "should", not a mandatory ("must") rule.

The Linux kernel allows you to change the default behavior with a sysctl parameter. To do this on SUSE Linux Enterprise Server 11, you need to make the following changes in configuration:

- Add ipv6 to the modules load early on boot

Edit `/etc/sysconfig/kernel` and add ipv6 to `MODULES_LOADED_ON_BOOT` e.g. `MODULES_LOADED_ON_BOOT="ipv6"`. This is needed for the second change to work, if ipv6 is not loaded early enough, setting the sysctl fails.

- Add the following lines to `/etc/sysctl.conf`

```
## shutdown IPV6 on MAC based duplicate address detection
net.ipv6.conf.default.accept_dad = 2
net.ipv6.conf.all.accept_dad = 2
net.ipv6.conf.eth0.accept_dad = 2
net.ipv6.conf.eth1.accept_dad = 2
```

Note: if you use other interfaces (e.g., eth2), modify the lines. With these changes, all tests for RFC 4862 should pass.

- Section 4: RFC 1981 - Path MTU Discovery for IPv6

- Test v6LC.4.1.10: Multicast Destination - One Router

- Test v6LC.4.1.11: Multicast Destination - Two Routers

On these two tests ping6 needs to be told to allow defragmentation of multicast packets. Newer ping6 versions have this disabled by default. Use: `ping6 -M want <other parameters>`. See `man ping6` for more information.

- Enable IPv6 in YaST for SCTP Support

SCTP is dependent on IPv6, so in order to successfully insert the SCTP module, IPv6 must be enabled in YaST. This allows for the IPv6 module to be automatically inserted when **modprobe sctp** is called.

14.6.1 YaST: IPv6 open-iscsi support

YaST has been extended to support installation using IPv6 iSCSI target as root device.

14.7 Other Technical Information

14.7.1 YaST Support for Layer 2 Devices

YaST writes the MAC address for layer 2 devices only if they are of the card_types:

1. OSD_100
2. OSD_1000
3. OSD_10GIG
4. OSD_FE_LANE
5. OSD_GbE_LANE
6. OSD_Express

Per intent YaST does not write the MAC address for devices of the types:

1. HiperSockets
2. GuestLAN/VSWITCH QDIO
3. OSM
4. OSX

14.7.2 Changes to Network Setup

The script `modify_resolvconf` is removed in favor of a more versatile script called `netconfig`. This new script handles specific network settings from multiple sources more flexibly and transparently. See the documentation and man-page of `netconfig` for more information.

14.7.3 Memory cgroups

Memory cgroups are now disabled for machines where they cause memory exhaustion and crashes. Namely, X86 32-bit systems with PAE support and more than 8G in any memory node have this feature disabled.

14.7.4 MCELog

The `mcelog` package logs and parses/translates Machine Check Exceptions (MCE) on hardware errors (also including memory errors). Formerly this has been done by a cron job executed hourly. Now hardware errors are immediately processed by an `mcelog` daemon.

However, the `mcelog` service is not enabled by default resulting in memory and CPU errors also not being logged by default. In addition, `mcelog` has a new feature to also handle predictive bad page offlining and automatic core offlining when cache errors happen.

The service can either be enabled via the YaST runlevel editor or via commandline with:

```
chkconfig mcelog on
rcmcelog start
```

14.7.5 Locale Settings in `~/.i18n`

If you are not satisfied with locale system defaults, change the settings in `~/.i18n`. Entries in `~/.i18n` override system defaults from `/etc/sysconfig/language`. Use the same variable names but without the `RC_` namespace prefixes; for example, use `LANG` instead of `RC_LANG`. For more information about locales in general, see "Language and Country-Specific Settings" in the Administration Guide.

14.7.6 Configuration of kdump

kdump is useful, if the kernel is crashing or otherwise misbehaving and a kernel core dump needs to be captured for analysis.

Use YaST (*System > Kernel Kdump*) to configure your environment.

14.7.7 Configuring Authentication for kdump through YaST with ssh/scp as Target

When kdump is configured through YaST with ssh/scp as target and the target system is SUSE Linux Enterprise, then enable authentication using either of the following ways:

1. Copy the public keys to the target system:

```
ssh-copy-id -i ~/.ssh/id_*.pub <username>@<target system IP>
```

or

2. Change the PasswordAuthentication setting in /etc/ssh/sshd_config of the target system from:

```
PasswordAuthentication no
```


to:

```
PasswordAuthentication yes
```

3. After changing PasswordAuthentication in /etc/ssh/sshd_config restart the sshd service on the target system with:

```
rcsshd restart
```

14.7.8 JPackage Standard for Java Packages

Java packages are changed to follow the JPackage Standard (<http://www.jpackage.org/> ) . For more information, see the documentation in /usr/share/doc/packages/jpackage-utils/.

14.7.9 Stopping Cron Status Messages


To avoid the mail-flood caused by cron status messages, the default value of SEND_MAIL_ON_NO_ERROR in /etc/sysconfig/cron is set to "no" for new installations. Even with this setting to "no", cron data output will still be send to the MAILTO address, as documented in the cron manpage.

In the update case it is recommended to set these values according to your needs.



15 Documentation and Other Information

- Read the READMEs on the DVDs.
- Get the detailed changelog information about a particular package from the RPM (with filename <FILENAME>):



```
rpm --changelog -qp <FILENAME>.rpm
```

- Check the ChangeLog file in the top level of DVD1 for a chronological log of all changes made to the updated packages.
- Find more information in the docu directory of DVD1 of the SUSE Linux Enterprise Server 11 Service Pack 4 DVDs. This directory includes PDF versions of the SUSE Linux Enterprise Server 11 Installation Quick Start and Deployment Guides.
- These Release Notes are identical across all architectures, and are available online at <http://www.suse.com/releasenotes/> .

15.1 Additional or Updated Documentation

- <http://www.suse.com/documentation/sles11/>  contains additional or updated documentation for SUSE Linux Enterprise Server 11 Service Pack 4.
- Find a collection of White Papers in the SUSE Linux Enterprise Server Ressource Library at <https://www.suse.com/products/server/resource-library/?ref=b#WhitePapers> .

15.2 Product and Source Code Information


Visit <http://www.suse.com/products/>  for the latest product news from SUSE and <http://www.suse.com/download-linux/source-code.html>  for additional information on the source code of SUSE Linux Enterprise products.

16 Miscellaneous


17 Legal Notices


SUSE makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, SUSE reserves the right to revise this publication and to make changes to its content, at any time, without the obligation to notify any person or entity of such revisions or changes.

Further, SUSE makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, SUSE reserves the right to make changes to any and all parts of SUSE software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classifications to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical/biological weaponry end uses. Please refer to <http://www.suse.com/company/legal/>  for more information on exporting SUSE software. SUSE assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2010, 2011, 2012, 2013, 2014, 2015 SUSE LLC. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

SUSE has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.suse.com/company/legal/>  and one or more additional patents or pending patent applications in the U.S. and other countries.

For SUSE trademarks, see the Trademark and Service Mark List (<http://www.suse.com/company/legal/> ). All third-party trademarks are the property of their respective owners.

Colophon

Thanks for using SUSE Linux Enterprise Server in your business.

The SUSE Linux Enterprise Server Team.