Release Notes for SUSE Linux Enterprise Server 11 Service Pack 3 (SP3)

Release Notes for SUSE Linux Enterprise Server 11 Service Pack 3 (SP3)

Version 11.3.12 (2013-04-18)

Abstract

These release notes are generic for all products of our SUSE Linux Enterprise Server 11 product line. Some parts may not apply to a particular architecture or product. Where this is not obvious, the specific architectures or products are explicitly listed.

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.

Chapter 1. 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.

***CHECKIT With the release of SUSE Linux Enterprise Server 11 Service Pack 3 the former SUSE Linux Enterprise Server 11 Service Pack 2 enters the 6 month migration window, during which time SUSE will continue to provide security updates and full support. At the end of the six-month parallel support period, on 2013-MM-DD, support for SUSE Linux Enterprise Server 11 Service Pack 2 will be discontinued. Long Term Service Pack Support (LTSS) for SUSE Linux Enterprise Server 11 Service Pack 1 is available as a separate option.

Chapter 2. Read Me First

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

- Chapter 3, *Support Statement for SUSE Linux Enterprise Server*
- Chapter 9, *Update-Related Notes*
- Chapter 13, *Technical Information*

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.

Chapter 3. 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/.

3.1. Erasing All Registration Data

Sometimes you may want to remove all data that was created during the registration of a SUSE Linux Enterprise system, so you can cleanly re-register it with different credentials.

This can now be accomplished with suse_register by using the new option "--erase-local-regdata". Note that this does not free the subscription that the system may have consumed in the Customer Center. This needs to be done from the Customer Center's Web UI.

3.2. 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.

3.2.1. 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.

3.2.2. 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.

3.3. 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
- WebSphere CE Application Server

3.4. 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. Alternatively, use the Novell Requirements Portal at http://www.novell.com/rms.

3.4.1. Technology Preview: libguestFS

Libguestfs is a set of tools for accessing and modifying virtual machine disk images. It can be used for many virtual image managements tasks such as viewing and editing files inside guests (only Linux one are enable), scripting changes to VMs, monitoring disk used/free statistics, performing partial backups, and cloning VMs. See http://libguestfs.org/ for more information.

3.4.2. Hot-Add Memory

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

- IBM x3800, x3850, single node x3950, x3850 M2, single node x3850 M2, 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 System p:

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.

3.4.3. 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.

3.4.4. 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.

Chapter 4. Installation

4.1. Current Limitations in a UEFI Secure Boot Context

When booting in Secure Boot mode, the following restrictions apply:

- bootloader, kernel and kernel modules must be signed
- kexec and kdump are disabled
- hibernation (suspend on disk) is disabled
- access to /dev/kmem and /dev/mem is not possible, even as root user
- access to IO port is not possible, even as root user. All X11 graphical drivers must use a kernel driver
- PCI BAR access through sysfs is not possible
- 'custom_method' in ACPI is not available
- debugfs for asus-wmi module is not available
- acpi_rsdp parameter doesn't have any effect on kernel

4.2. Support for 4 KB/Sector Hard Disk Drives

Support for 4 KB/sector hard disk drives requires support from all code that directly accesses the hard disk drives.

SUSE Linux Enterprise fully supports 4 KB/sector drives in all conditions and architectures with one exception. The 4KB/sector hard disk drives are not supported as a boot drive on x86_64 systems booting with a legacy BIOS.

4.3. XEN: Pygrub Improvement

Pygrub is used to boot a virtual Xen machine according to a certain menu.lst entry. Pygrub now accepts a new flag [-l|--list_entries] to show grub entries in the guest.

4.4. Installation via USB

FATE 312662 RN missing

4.5. Deployment

SUSE Linux Enterprise Server can be deployed in three ways:

- Physical Machine,

- Virtual Host,
- Virtual Machine in paravirtualized environments.

4.6. 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).

4.7. 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.

4.8. 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".

4.9. 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.

4.10. Using iSCSI Disks when Installing

To use iSCSI disks during installation, add the following parameter to the boot option line: `withiscsi=1`.

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

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

4.11. 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.

4.12. 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.

4.13. 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.

4.14. 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.

4.15. 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.

4.16. For More Information

For more information, see Chapter 11, *Infrastructure, Package and Architecture Specific Information*.

Chapter 5. Features and Versions

5.1. Linux Kernel and Toolchain

5.1.1. Lustre 2.1 Kernel Modules Preparation

Lustre 2.1 builds of kernel modules by 3rd parties needed kernel modifications of the previous shipped SUSE Kernel, and thus breaking the support chain.

To allow the build of kernel modules for Lustre 2.1 by 3rd parties without breaking the support chain for the SUSE Kernel, the needed hooks for Lustre were added to the shipped kernel.

This change does not include Lustre modules or packages, nor support.

5.1.2. Kernel Dumps with LZO Compression

yast2-kdump and crash also support LZO compression as a new target format.

5.1.3. Add Option to mpstat to Only Display Stats for Online CPUs

mpstat added the "-P ON" option to limit statistics displayed to only online CPUs.

5.1.4. Support for Failopen Mode When Using Netfilter's NFQUEUE Target

Adds support for a new failopen mode when using netfilter's NFQUEUE target. This mode allows users to temporarily disable packet inspection and maintain connectivity under heavy network traffic.

5.1.5. libvirt Support for QEMU seccomp Sandboxing

QEMU guests spawned by libvirt are exposed to a large number of system calls that go unused for the entire lifetime of the process.

libvirt's qemu.conf file is updated with a seccomp_sandbox option that can be used to enable use of QEMU's seccomp sandboxing support. This allows execution of QEMU guests with reduced exposure to kernel system calls.

5.1.6. Makedumpfile: Enhanced Elimination of Sensitive Data from Dumps

Enhances the current makedumpfile filtering to eliminate complex data structures and cryptographic keys.

5.1.7. Support for Latest Intel Active Management Technology (AMT)

This Servicepack adds support for Intel AMT version 7 and later by providing the Intel MEI kernel driver.

In order to use Intel AMT, you also must download the Intel LMS and ACUConfig components from Intels website:

http://software.intel.com/en-us/articles/download-the-latest-intel-amt-open-source-drivers

For more information on AMT on Linux, please follow the URLs in the "Additional Information" found on the above mentioned Intel website.

5.1.8. General Version Information

- GCC 4.3.4
- glibc 2.11.3
- Linux kernel 3.0
- perl 5.10
- php 5.3
- python 2.6.8
- ruby 1.8.7

5.1.9. 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.

5.2. 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.

5.2.1. Upgrading MySQL to Version 5.5

Replacing an unmaintained version of MySQL.

SLES11-SP3 introduces the upgrade of the MySQL database to version 5.5. This upgrade involves a change of the database format and the database needs to be converted before MySQL can run again. Therefore MySQL is not running directly after the upgrade.

To mgirate the MySQL database, run following commands as root:

```
touch /var/lib/mysql/.force_upgrade
rcmysql restart
```

To verify failures during the server start check the log files under /var/log/mysql/.

We strongly recommend to back up the database before migrating it (mostly /var/lib/mysql).

5.3. 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

5.4. Security

5.4.1. Support of SHA-256 Hash Algorithm in opencryptoki CCA Token

SLES 11 SP3 includes opencryptoki 2.4.2 which comes with a CCA token that exploits the SHA-256 hash algorithm that is provided by System z crypto hardware.

5.4.2. OpenSCAP Tools and Libraries Added

OpenSCAP is a set of open source libraries providing a path for integration of SCAP (Security Content Automation Protocol). SCAP is a collection of standards managed by NIST with the goal of providing a standard language for the expression of Computer Network Defense related information. For more information about SCAP, see http://nvd.nist.gov .

5.4.3. PAM Configuration

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

5.4.4. 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.

5.4.5. 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 `tcsd`—the daemon that allows userland programs to communicate with the TPM driver in the Linux kernel. `tcsd` 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.

5.4.6. 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.

5.5. 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.

***CHECKIT 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 13.6, “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.

5.5.1. Linux Virtual Server Load Balancer (ipvs) Extends Support for IPv6

The LVS/ipvs load balancing code did not fully support RFC2460 and fragmented IPv6 packets which could lead to lost packets and interrupted connections when IPv6 traffic was fragmented.

The load balancer has been enhanced to fully support IPv6 fragmented extension headers and is now RFC2460 compliant.

5.6. Resource Management

5.6.1. libseccomp

Seccomp filters are expressed as a Berkeley Packet Filter (BPF) program, which is not a well understood interface for most developers.

The libseccomp library provides an easy to use interface to the Linux Kernel's syscall filtering mechanism, seccomp. The libseccomp API allows an application to specify which syscalls, and optionally which syscall arguments, the application is allowed to execute, all of which are enforced by the Linux Kernel.

5.6.2. XEN: Support for PCI Pass-through Bind and Unbind in libvirt Xen Driver

Virt-manager is now able to set up PCI pass-through for Xen without having to switch to the command line to free the PCI device before assigning it to the VM.

5.6.3. 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 brigded 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.

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.

5.7. 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 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:

    ```
    rccollectd start
    rcwebyast start
    ```

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

5.8. 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/detail.php?device_cat=software&device_id=11287~16&release_id=24.

Chapter 6. Driver Updates

6.1. X.Org: fbdev Used in UEFI Secure Boot Mode (ASpeed Chipset)

The unaccelerated fbdev driver is used as a fallback in UEFI secure boot mode with the ast KMS driver, EFI VGA, and other currently unknown frame buffer drivers.

6.2. X.Org Driver Used in UEFI Secure Boot Mode (Matrox)

The unaccelerated "mgag200"/"modesetting" (generic X.Org) driver combo is used instead of the "mga" X.Org driver if machine is running in UEFI secure boot mode. The driver does not load in other cases with a warning message in the kernel log.

6.3. 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

6.3.1. Broadcom bnx2x Driver Limitation

Only the initial SR-IOV Linux support is available.

6.3.2. Add Support for TIPC (Transparent Inter-Process Communication)

The Transparent Inter-Process Communication protocol (TIPC) allows applications in a cluster environment to communicate quickly and reliably with eachother, regardless of their location within the cluster. TIPC includes a network topology service that lets applications track both functional and physical changes in the network, helping to synchronize startup of distributed applications and their responses to failure conditions. A socket API is used to interact with the topology service and other applications. Address assignment and bearer configuration is managed from a userspace application called tipc-config.

6.3.3. 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.

6.3.4. Broadcom 57712 vNICs/NPAR PCIE Functions Disappearing under SP2

***CHECKIT: SP3?

SP2 scans for the functions on a PCI device in a new way using ARI. This can cause some of the functions on the Broadcom 57712 adapter to be missing after upgrading to SP2.

Contact your system vendor to receive the latest firmware for the 57712 adapter that resolves this issue. Alternatively, upgrading to kernel 3.0.26 and adding the boot parameter 'pci=noari' will allow all the functions on the 57712 adapter to become visible under SLES 11 SP2 and later.

6.4. Storage Drivers

- Updated qla2xxx to version 8.03.01.04.11.1-k8
- Updated qla4xxx to version v5.01.00.00.11.01-k13
- 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).

6.4.1. 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 dcbx, 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.

6.5. 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

6.5.1. New Intel Platform and CPU Support

This service pack adds support for the following new Intel CPUs:

- Next generation Intel® Xeon® processor E7-8800/4800/2800 v2 product families

This covers new support for the following platforms:

- Brickland–EX

6.5.2. Support for the Intel Bordenville Microserver

This Service Pack adds support for Intel's Bordenville Microserver based on the Centerton SoC (System On Chip).

Chapter 7. Other Updates

7.1. Package python-ethtool

The Python bindings for ethtool were updated in SLE11 SP2 to version 0.7. This update introduced several stability bugfixes and support for handling IPv6.

7.2. Update Python to 2.6.8

Python 2.6.7 and 2.6.8 are security only updates to 2.6.6.

Python 2.6 helps with migrating to Python 3.0, which is a major redesign of the language. Whenever possible, Python 2.6 incorporates new features and syntax changes from 3.0 while remaining compatible with existing code. In case of conflict, Python 2.6 adds compatibility functions in a `future_builtins` module and a `-3` switch to warn about usages that will become unsupported in 3.0.

Some significant new packages have been added to the standard library, such as the multiprocessing and json modules.

7.3. Individual Timeout Value for Each Direct AutoFS Mount

If there were two direct mounts with different timeouts configured, the second one was ignored and the first timeout value was used for both mount points.

AutoFS was patched to support individual timeout values for each direct mount.

7.4. List of Updated Packages

- Added support for installation from an NFSv4 server.
- Updated binutils to version 2.21.1
- Updated bluez to version 4.51
- Updated clamav to version 0.97.3
- Updated crash to version 5.1.9
- Updated dhcp to version 4.2.3.P2
- Updated gdb to version 7.3
- Updated hplip to version 3.11.10
- Updated ipsec-tools to version 0.7.3
- Updated IBM Java 1.4.2 (java-1_4_2-ibm) to SR13 FP11
- Updated IBM Java 1.6.0 (java-1_6_0-ibm) to SR9.3

- Updated libcgroup1 to version 0.37.1
- Updated libcmpiutil to version 0.5.6
- Updated libelf to version 0.8.12
- Updated QT4 (libqt4) to version 4.6.3
- Updated libvirt to version 0.9.6
- Updated libvirt-cim to version 0.5.12
- Updated mdadm to version 3.2.2
- Updated module-init-tools to version 3.11.1
- Updated MozillaFirefox to version 10
- Added mt_st version 0.9b
- Added netlabel version 0.19
- Updated numactl to version 2.0.7
- Updated openCryptoki to version 2.4
- Updated openldap2 to version 2.4.26
- Added openvas version 3.0
- Added perf: Performance Counters For Linux
- Added perl-WWW-Curl version 4.09
- Added rng-tools: Support daemon for hardware random device
- Updated sblim-cim-client2 to version 2.1.3
- Updated sblim-cmpi-base to version 1.6.1
- Updated sblim-cmpi-fsvol to version 1.5.0
- Updated sblim-cmpi-network to version 1.4.0
- Updated sblim-cmpi-nfsv3 to version 1.1.0
- Updated sblim-cmpi-nfsv4 to version 1.1.0
- Updated sblim-cmpi-params to version 1.3.0
- Updated sblim-cmpi-sysfs to version 1.2.0
- Updated sblim-gather to version 2.2.0
- Updated sblim-sfcb to version 1.3.11
- Updated sblim-sfcc to version 2.2.1
- Updated sblim-wbemcli to version 1.6.1

- Updated strongswan to version 4.4.0
- Added stunnel version 4.36
- Updated virt-viewer to version 0.4.1
- Updated virt-manager to version 0.9.0
- Updated kvm to version 0.15.1
- Updated Xen (xen) to version 4.1.2
- Updated dcbd to version 0.9.24
- Updated e2fsprogs to version 1.41.9
- Updated iprutils to version 2.3.7
- Updated iscsitarget to version 1.4.20
- Updated nfs-utils to version 1.2.3 for improved IPv6 support
- Added apport, a tool to collect data automatically from crashed processes

Chapter 8. Software Development Kit

SUSE provides a Software Development Kit (SDK) for SUSE Linux Enterprise 11 Service Pack 3. 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

8.1. Optional GCC Compiler Suite on SDK

The optional compiler on the SDK has been updated to GCC 4.7. It brings better standard compliance (ISO C 11, ISO C++ 11), improved optimizations and allows to take benefit of new hardware instructions.

SUSE also added support for the IBM zEnterprise EC12 architecture

For details see http://gcc.gnu.org/gcc-4.7/changes.html

Chapter 9. Update-Related Notes

This section includes update-related information for this release.

9.1. General Notes

9.1.1. Upgrading PostgreSQL Installations from 8.3 to 9.1.

To upgrade a PostgreSQL server installation from version 8.3 to 9.1, the database files need to be converted to the new version.

Newer versions of PostgreSQL come with the pg_upgrade tool that simplifies and speeds up the migration of a PostgreSQL installation to a new version. Formerly dump and restore was needed that was much slower.

pg_upgrade needs to have the server binaries of both versions available. To allow this, we had to change the way PostgreSQL is packaged as well as the naming of the packages, so that two or more versions of PostgreSQL can be installed in parallel.

Starting with version 9.1, PostgreSQL package names contain numbers indicating the major version. In PostgreSQL terms the major version consists of the first two components of the version number, i.e. 8.3, 8.4, 9.0, or 9.1. So, the packages for Postgresql 9.1 are named postgresql91, postgresql91-server, etc. Inside the packages the files were moved from their standard locations to a versioned location such as /usr/lib/postgresql83/bin or /usr/lib/postgresql91/bin to avoid file conflicts if packages are installed in parallel. The update-alternatives mechanism creates and maintains symbolic links that cause one version (by default the highest installed version) to re-appear in the standard locations. By default, database data are stored under /var/lib/pgsql/data on SUSE Linux.

The following preconditions have to be fulfilled before data migration can be started:

1. If not already done, the packages of the old PostgreSQL version must be upgraded to the new packaging scheme through a maintenance update. For SLE11 this means to install the patch that upgrades PostgreSQL from version 8.3.14 to 8.3.19 or higher.

2. The packages of the new PostgreSQL major version need to be installed. For SLE11 this means to install postgresql91-server and all the packages it depends on. As pg_upgrade is contained in postgresql91-contrib, that one has to be installed as well, at least until the migration is done.

3. Unless pg_upgrade is used in link mode, the server must have enough free disk space to temporarily hold a copy of the database files. If the database instance was installed in the default location, the needed space in megabytes can be determined by running the follwing command as root: "du -hs /var/lib/pgsql/data". If space is tight, it might help to run the "VACUUM FULL" SQL command on each database in the instance to be migrated, but be aware that it might take very long.

Upstream documentation about pg_upgrade including step by step instructions for performing a database migration can be found under file:///usr/share/doc/packages/postgresql91/html/pgupgrade.html (if the postgresql91-docs package is installed), or online under http://www.postgresql.org/docs/9.1/static/pgupgrade.html. NOTE: The online documentation starts with explaining how you can install PostgreSQL from the upstream sources (which is not necessary on SLES) and also uses other directory names (/usr/local instead of the update-alternatives based path as described above).

For background information about the inner workings of pg_admin and a performance comparison with the old dump and restore method, see http://momjian.us/main/writings/pgsql/pg_upgrade.pdf .

9.1.2. Online Migration from SP2 to SP3 via "YaST wagon"

The online migration from SP2 to SP3 is supported via the "YaST wagon" module.

9.1.3. Online Migration with Debuginfo Packages Not Supported

***CHECKIT Online migration from SP2 to SP3 is not supported if debuginfo packages are installed.

9.1.4. Migrating to SLE 11 SP3 Using Zypper

To migrate the system to the Service Pack 3 level with **zypper**, proceed as follows:

- Open a root shell.
- Run **zypper ref -s** to refresh all services and repositories.
- Run **zypper up -t patch** to install package management updates.
- Now it is possible to install all available updates for SLES/SLED 11 SP2; run **zypper up -t 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-SP3-migration</product>
  <product>SUSE_SLES-SP3-migration</product>
  ```

- Install these migration products (example):

  ```
  zypper in -t product sle-sdk-SP3-migration SUSE_SLES-SP3-migration
  ```

- Run **suse_register -d 2 -L /root/.suse_register.log** to register the products in order to get the corresponding SP3 Update repositories.
- Run **zypper ref -s** to refresh services and repositores.
- ***CHECKIT Check the repositories using **zypper lr**. Only if needed, disable repositories manually (note that the SP1-Pool, SP1-Updates, SP2-Pool and SP2-Updates repos need to stay enabled!) and enable the new SP3 (SP3-Core, SP3-Updates) repositories:

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

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

  ```
  zypper dup --from SLES11-SP3-Core --from SLES11-SP3-Updates \
  --from SLE11-WebYaST-SP3-Pool --from SLE11-WebYaST-SP3-Updates
  ```

  Add more SP3 catalogs here if needed, e.g. in case add-on products are installed.

- 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
  ```

- Reboot the system.

9.1.5. 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).

9.1.6. 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 SP3, which may require intermediate upgrade steps:

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

- SLES 11 GA -> SLES 11 SP1 -> SLES 11 SP2 -> SLES 11 SP3

9.1.7. Upgrading to SLES 11 SP3 with Root File System on iSCSI

***CHECKIT 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
```

9.1.8. 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.

9.1.9. Tickless Idle

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

9.1.10. 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.

9.1.11. 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`.

9.1.12. 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.

9.1.13. 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.

9.1.14. 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.

9.1.15. Upgrading MySQL to SUSE Linux Enterprise Server 11

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.

9.1.16. 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.

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

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.

9.1.18. 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.

9.2. Update from SUSE Linux Enterprise Server 11

9.2.1. 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).

For more details, see http://ifup.org/2011/02/03/reverse-path-filter-rp_filter-by-example/.

9.2.2. 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`.

9.3. Update from SUSE Linux Enterprise Server 11 SP 1

9.3.1. Update from SUSE Linux Enterprise Server 11 SP 1

***CHECKIT Updating from SUSE Linux Enterprise Server 11 SP 1 with AutoYaST is supported.

9.4. Update from SUSE Linux Enterprise Server 11 SP 2

9.4.1. Update of python-lxml to 2.3.x

python-lxml has been updated to version 2.3.6. It brings several features and numerous bug fixes, as well as one API change:

`Element.findtext()` now returns an empty string instead of `None` for elements without text content; it still returns `None` when there is no element matching the request. This brings the lxml implementation of the ElementTree API in conformance with the ElementTree API specification [http://www.effbot.org/zone/pythondoc-elementtree-ElementTree.htm#elementtree.ElementTree._ElementInterface.findtext-method] .

9.4.2. Augeas Framework Updated to Version 0.9

9.4.3. Postfix: Incompatibility Issues and New Features

To benefit from enhancements and improvements which have been developed in the upstream community, postfix is upgraded from version 2.5.13 to the current version 2.9.4.

Incompatibility Issues:

- The default milter_protocol setting is increased from 2 to 6; this enables all available features up to and including Sendmail 8.14.0.
- When a mailbox file is not owned by its recipient, the local and virtual delivery agents now log a warning and defer delivery. Specify "strict_mailbox_ownership = no" to ignore such ownership discrepancies.
- The Postfix SMTP client(!) no longer tries to use the obsolete SSLv2 protocol by default, as this may prevent the use of modern SSL features. Lack of SSLv2 support should never be a problem, since SSLv3 was defined in 1996, and TLSv1 in 1999. You can undo the change by specifying empty main.cf values for smtp_tls_protocols and lmtp_tls_protocols.

- Postfix SMTP server replies for address verification have changed. unverified_recipient_reject_code and unverified_sender_reject_code now handle "5XX" rejects only. The "4XX" rejects are now controlled with unverified_sender_defer_code and unverified_recipient_defer_code.

- postfix-script, postfix-files and post-install are moved away from /etc/postfix to $daemon_directory.

- Postfix now adds (Resent-) From:, Date:, Message-ID: or To: headers only when clients match $local_header_rewrite_clients. Specify "always_add_missing_headers = yes" for backwards compatibility.

- The verify(8) service now uses a persistent cache by default (address_verify_map = btree: $data_directory/verify_cache). To disable, specify "address_verify_map ="

- The meaning of an empty filter next-hop destination has changed (for example, "content_filter = foo:" or "FILTER foo:"). Postfix now uses the recipient domain, instead of using $myhostname as in Postfix 2.6 and earlier. To restore the old behavior specify "default_filter_nexthop = $myhostname", or specify a non-empty next-hop content filter destination.

- Postfix now requests default delivery status notifications when adding a recipient with the Milter smfi_addrcpt action, instead of "never notify" as with Postfix automatically-added recipients.

- Postfix now reports a temporary delivery error when the result of virtual alias expansion would exceed the virtual_alias_recursion_limit or virtual_alias_expansion_limit.

- To avoid repeated delivery to mailing lists with pathological nested alias configurations, the local(8) delivery agent now keeps the owner-alias attribute of a parent alias, when delivering mail to a child alias that does not have its own owner alias.

- The Postfix SMTP client no longer appends the local domain when looking up a DNS name without ".". Specify "smtp_dns_resolver_options = res_defnames" to get the old behavior, which may produce unexpected results.

- The format of the "postfix/smtpd[pid]: queueid: client=host[addr]" logfile record has changed. When available, the before-filter client information and the before-filter queue ID are now appended to the end of the record.

- Postfix by default no longer adds a "To: undisclosed-recipients:;" header when no recipient specified in the message header. For backwards compatibility, specify: "undisclosed_recipients_header = To: undisclosed-recipients:;"

- The Postfix SMTP server now always re-computes the SASL mechanism list after successful completion of the STARTTLS command. Earlier versions only re-computed the mechanism list when the values of smtp_sasl_tls_security_options and smtp_sasl_security_options differ. This could produce incorrect results, because the Dovecot authentication server may change responses when the SMTP session is encrypted.

- The smtpd_starttls_timeout default value is now stress-dependent. By default, TLS negotiations must now complete under overload in 10s instead of 300s.

- Postfix no longer appends the system-supplied default CA certificates to the lists specified with *_tls_CAfile or with *_tls_CApath. This prevents third-party certificates from getting mail relay permission with the permit_tls_all_clientcerts feature. Unfortunately this change may cause compatibility problems when configurations rely on certificate verification for other purposes. Specify "tls_append_default_CA = yes" for backwards compatibility.

- The VSTREAM error flags are now split into separate read and write error flags. As a result of this change, all programs that use Postfix VSTREAMs MUST be recompiled.

- For consistency with the SMTP standard, the (client-side) smtp_line_length_limit default value was increased from 990 characters to 999 (i.e. 1000 characters including <CR><LF>). Specify "smtp_line_length_limit = 990" to restore historical Postfix behavior.

- To simplify integration with third-party applications, the Postfix sendmail command now always transforms all input lines ending in <CR><LF> into UNIX format (lines ending in <LF>). Specify "sendmail_fix_line_endings = strict" to restore historical Postfix behavior.

- To work around broken remote SMTP servers, the Postfix SMTP client by default no longer appends the "AUTH=<>" option to the MAIL FROM command. Specify "smtp_send_dummy_mail_auth = yes" to restore the old behavior.

- Instead of terminating immediately with a "fatal" message when a database file can't be opened, a Postfix daemon program now logs an "error" message, and continues execution with reduced functionality. Logfile-based alerting systems may need to be updated to look for "error" messages in addition to "fatal" messages. Specify "daemon_table_open_error_is_fatal = yes" to get the historical behavior (immediate termination with "fatal" message).

- Postfix now logs the result of successful TLS negotiation with TLS logging levels of 0.

- The default inet_protocols value is now "all" instead of "ipv4", meaning use both IPv4 and IPv6. To avoid an unexpected loss of performance for sites without global IPv6 connectivity, the commands "make upgrade" and "postfix upgrade-configuration" now append "inet_protocols = ipv4" to main.cf when no explicit inet_protocols setting is already present.

New Features:

- Support for managing multiple Postfix instances. Multi-instance support allows you to do the following and more: - Simplify post-queue content filter configuration by using separate Postfix instances before and after the filter. - Implement per-user content filters (or no filter) via transport map lookups instead of content_filter settings. - Test new configuration settings (on a different server IP address or TCP port) without disturbing production instances.

- check_reverse_client_hostname_access, to make access decisions based on the unverified client hostname.

- With "reject_tempfail_action = defer", the Postfix SMTP server immediately replies with a 4xx status after some temporary error.

- The Postfix SMTP server automatically hangs up after replying with "521". This makes overload handling more effective. See also RFC 1846 for prior art on this topic.

- Stress-dependent behavior is enabled by default. Under conditions of overload, smtpd_timeout is reduced from 300s to 10s, smtpd_hard_error_limit is reduced from 20 to 1, and smtpd_junk_command_limit is reduced from 100 to 1.

- Specify "tcp_windowsize = 65535" (or less) to work around routers with broken TCP window scaling implementations.

- New "lmtp_assume_final = yes" flag to send correct DSN "success" notifications when LMTP delivery is "final" as opposed to delivery into a content filter.

- The Postfix SMTP server's SASL authentication was re-structured. With "smtpd_tls_auth_only = yes", SASL support is now activated only after a successful TLS handshake. Earlier Postfix SMTP server versions could complain about unavailable SASL mechanisms during the plaintext phase of the SMTP protocol.

- Improved before-queue filter performance. With "smtpd_proxy_options = speed_adjust", the Postfix SMTP server receives the entire message before it connects to a before-queue content filter. This means you can run more SMTP server processes with the same number of running content filter processes, and thus, handle more mail. This feature is off by default until it is proven to create no new problems.
- sender_dependent_default_transport_maps, a per-sender override for default_transport.
- milter_header_checks: Support for header checks on Milter-generated message headers. This can be used, for example, to control mail flow with Milter-generated headers that carry indicators for badness or goodness. Currently, all header_checks features are implemented except PREPEND.
- Support to turn off the TLSv1.1 and TLSv1.2 protocols. Introduced with OpenSSL version 1.0.1, these are known to cause inter-operability problems with for example hotmail. The radical workaround is to temporarily turn off problematic protocols globally: smtp_tls_protocols = !SSLv2, !TLSv1.1, !TLSv1.2 smtp_tls_mandatory_protocols = !SSLv2, !TLSv1.1, !TLSv1.2
- Prototype postscreen(8) server that runs a number of time-consuming checks in parallel for all incoming SMTP connections, before clients are allowed to talk to a real Postfix SMTP server. It detects clients that start talking too soon, or clients that appear on DNS blocklists, or clients that hang up without sending any command.
- Support for address patterns in DNS blacklist and whitelist lookup results.
- The Postfix SMTP server now supports DNS-based whitelisting with several safety features: permit_dnswl_client whitelists a client by IP address, and permit_rhswl_client whitelists a client by its hostname. These features use the same syntax as reject_rbl_client and reject_rhsbl_client, respectively. The main difference is that they return PERMIT instead of REJECT.
- The SMTP server now supports contact information that is appended to "reject" responses. This includes SMTP server responses that aren't logged to the maillog file, such as responses to syntax errors, or unsupported commands.
- tls_disable_workarounds parameter specifies a list or bit-mask of OpenSSL bug work-arounds to disable.
- The lower-level code in the TLS engine was simplified by removing an unnecessary layer of data copying. OpenSSL now writes directly to the network.
- enable_long_queue_ids Introduces support for non-repeating queue IDs (also used as queue file names). These names are encoded in a mix of upper case, lower case and decimal digit characters. Long queue IDs are disabled by default to avoid breaking tools that parse logfiles and that expect queue IDs with the smaller [A-F0-9] character set.
- memcache lookup and update support. This provides a way to share postscreen(8) or verify(8) caches between Postfix instances.
- Support for TLS public key fingerprint matching in the Postfix SMTP client (in smtp_tls_policy_maps) and server (in check_ccert access maps).
- Support for external SASL authentication via the XCLIENT command. This is used to accept SASL authentication from an SMTP proxy such as NGINX. This support works even without having to specify "smtpd_sasl_auth_enable = yes".

9.4.4. Binutils Update

Binutils was updated to support newer hardware instructions.

9.4.5. unixODBC Updated to Version 2.3.1

unixODBC 2.3.1 provides the most recent upstream fixes; this helps for seamless population of DB2 data using automated tools and improves interoperability with MS SQL server.

9.4.6. stunnel Update to Version 4.54

The "stunnel" package update adds new service options for sni and tcp socket handling, improves handling in a FIPS setup and contains some performance improvements

9.4.7. IBM Java 1.4.2 End of Life

As announced with SUSE Linux Enterprise Server 11 SP2, IBM Java 1.4.2 reached End of Life, and thus we remove support for this specific Java version with SUSE Linux Enterprise Server 11 SP3. We recommend to upgrade your environments.

9.4.8. Update from SUSE Linux Enterprise Server 11 SP 2

Updating from SUSE Linux Enterprise Server 11 SP 2 with AutoYaST is supported.

Chapter 10. Deprecated Functionality

10.1. X.Org: fbdev Used in UEFI Secure Boot Mode (ASpeed Chipset)

The unaccelerated fbdev driver is used as a fallback in UEFI secure boot mode with the ast KMS driver, EFI VGA, and other currently unknown frame buffer drivers.

10.2. X.Org Driver Used in UEFI Secure Boot Mode (Matrox)

The unaccelerated "mgag200"/"modesetting" (generic X.Org) driver combo is used instead of the "mga" X.Org driver if machine is running in UEFI secure boot mode. The driver does not load in other cases with a warning message in the kernel log.

10.3. Support for the JFS File System

In connection with the change in the JFS support status the corresponding kernel module has been moved to the extra kernel RPM (kernel-flavor-extra).

10.4. Support for Portmap to End with SUSE Linux Enterprise 11 SP3

In SUSE Linux Enterprise we provide "rpcbind", which is compatible with portmap. "rpcbind" provides full IPv6 support. Thus portmap is now deprecated, and support for portmap will end end with the release of SUSE Linux Enterprise 11 SP3.

10.5. L3 Support for Openswan Is Scheduled to Expire

L3 support for Openswan is scheduled to expire. This decision is driven by the fact that Openswan development stalled substantially and there are no tangible signs that this will change in the future.

In contrast to this the strongSwan project is vivid and able to deliver a complete implementation of current standards. Compared to Openswan all relevant features are available by the package strongSwan plus strongSwan is the only complete Open Source implementation of the RFC 5996 IKEv2 standard whereas Openswan only implements a small mandatory subset. For now and the expected future only strongSwan qualifies to be an enterprise-ready solution for encrypted TCP/IP connectivity.

10.6. PHP 5.2 Is Deprecated

Based on significant customer demand, we ship PHP 5.3 parallel to PHP 5.2 with SUSE Linux Enterprise 11 SP2.

PHP 5.2 is deprecated though, and will be removed with SLES 11 SP3.

10.7. Packages Removed with SUSE Linux Enterprise Server 11 SP3

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

10.8. 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.

10.9. 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

10.10. 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

10.11. 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 `dhcp-client` 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.

Chapter 11. Infrastructure, Package and Architecture Specific Information

11.1. Hyper-V: KVP IP Injection

Hyper-V now supports the KVP (Key Value Pair) functionality to implement the mechanism to GET/SET IP addresses in the guest. This functionality is used in Windows Server 2012 to implement VM replication functionality.

11.2. Systems Management

11.2.1. Providing the URL of an Add-on Media at the Command Line during Installation

Add-on media like the Software Development Kit or third party driver media can be added to SUSE Linux Enterprise during installation or later in the running system. Sometimes it's advisable that an add-on media is available from the very beginning, for example to make drivers for new hardware available.

It is now possible to provide one or more URLs that point to the location of add-on media at the installer's command line by providing an "addon=url" parameter. Multiple add-ons need to be provided as a comma-separated list ("addon=url1,url2,...").

11.2.2. Individual Timeout Value for Each Direct AutoFS Mount

If there were two direct mounts with different timeouts configured, the second one was ignored and the first timeout value was used for both mount points.

AutoFS was patched to support individual timeout values for each direct mount.

11.2.3. YaST Repair Tool Limitation

The YaST Repair Tool as available from the boot medium does not detect pseudo devices like `/dev/btrfs` and writes a warning about missing partitions instead.

You should skip the repair of such a device, because for such pseudo devices the availability of a partition is not expected.

11.2.4. 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.

11.2.5. 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.

11.2.6. 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.

11.2.7. SPident

SPident is a tool to identify the Service Pack level of the current installation. On SUSE Linux Enterprise Server 11 GA, this tool has been replaced by the new SAM tool (package "suse-sam").

11.3. Performance Related Information

11.3.1. 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. More information on CFS can be found

here: "Multiprocessing with the Completely Fair Scheduler", http://www.ibm.com/developerworks/linux/library/l-cfs/?ca=dgrlnxw06CFC4Linux

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.

11.3.2. 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)).

11.4. Storage

11.4.1. Improved Support for Intel RSTe

This Service Pack adds improved support for Intel Rapid Storage Technology Enterprise (RSTe). It now supports RAID levels 0,1,4,5,6 and 10.

11.4.2. Define disk order for MD Raid with YaST

This enables to specify the disk order if a RAID device is created. Thus you can influence which data of the RAID is written on which disk.

11.4.3. 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.

11.4.4. 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.

11.5. Hyper-V

11.5.1. 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`.

11.5.2. 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.

11.5.3. Formatting Large Disk Partitions on Windows 8 Server

Installing a guest hosted on Windows 8 Server may fail when a large virtual disk image (larger than 50 GB) in `.vhdx` format is assigned to the guest. To workaround this issue use either virtual disk images with a fixed size, or create the dynamically sized disk image using Powershell.

Technical Background about the Issue

The `.vhd` and `.vhdx` images are sparse files. When a dynamic `.vhdx` is created with a maximum size of 127 GB, the initial size is about 256 KB. Because the default block size for `.vhdx` files is 32 MB, writing one 512 byte sector will result in a 32 MB section of the sparse file being allocated. When `ext3` is allocating the MBR, the super block, the backup super blocks, inodes, directories, etc., space is being allocated in the sparse file. Because of `ext3`'s suboptimal IO, how the data structures are laid out on disk, and the default block size, a large partition of the `.vhdx` file is allocated just by formatting. The workaround is to create a `.vhdx` file with a 1 MB block size rather than the default 32 MB.

Changing the block size in the UI is not implemented. It can only be changed when the VHDx file is created through Powershell. To create a VHD with a modified block size, use this Powershell script (all in one line):

```
New-VHD -Path C:\MyVHDs\test.vhdx -SizeBytes (127GB)
  -Dynamic -BlockSizeBytes (1MB) -VHDFormat vhdx
```

11.6. Architecture Independent Information

11.6.1. Current Limitations in a UEFI Secure Boot Context

When booting in Secure Boot mode, the following restrictions apply:

- bootloader, kernel and kernel modules must be signed
- kexec and kdump are disabled
- hibernation (suspend on disk) is disabled
- access to /dev/kmem and /dev/mem is not possible, even as root user
- access to IO port is not possible, even as root user. All X11 graphical drivers must use a kernel driver

- PCI BAR access through sysfs is not possible
- 'custom_method' in ACPI is not available
- debugfs for asus-wmi module is not available
- acpi_rsdp parameter doesn't have any effect on kernel

11.6.2. Changes in Packaging and Delivery

11.6.2.1. Python Updated to Version 2.6.8 with "collections.OrderedDict" Functionality

The "OrderedDict" functionality ensures that Python dictionaries emitted for conversion into strings maintain their original order. This functionality is important for data analytics applications.

11.6.2.2. Postfix: Incompatibility Issues and New Features

To benefit from enhancements and improvements which have been developed in the upstream community, postfix is upgraded from version 2.5.13 to the current version 2.9.4.

Incompatibility Issues:

- The default milter_protocol setting is increased from 2 to 6; this enables all available features up to and including Sendmail 8.14.0.
- When a mailbox file is not owned by its recipient, the local and virtual delivery agents now log a warning and defer delivery. Specify "strict_mailbox_ownership = no" to ignore such ownership discrepancies.
- The Postfix SMTP client(!) no longer tries to use the obsolete SSLv2 protocol by default, as this may prevent the use of modern SSL features. Lack of SSLv2 support should never be a problem, since SSLv3 was defined in 1996, and TLSv1 in 1999. You can undo the change by specifying empty main.cf values for smtp_tls_protocols and lmtp_tls_protocols.
- Postfix SMTP server replies for address verification have changed. unverified_recipient_reject_code and unverified_sender_reject_code now handle "5XX" rejects only. The "4XX" rejects are now controlled with unverified_sender_defer_code and unverified_recipient_defer_code.
- postfix-script, postfix-files and post-install are moved away from /etc/postfix to $daemon_directory.
- Postfix now adds (Resent-) From:, Date:, Message-ID: or To: headers only when clients match $local_header_rewrite_clients. Specify "always_add_missing_headers = yes" for backwards compatibility.
- The verify(8) service now uses a persistent cache by default (address_verify_map = btree: $data_directory/verify_cache). To disable, specify "address_verify_map ="
- The meaning of an empty filter next-hop destination has changed (for example, "content_filter = foo:" or "FILTER foo:"). Postfix now uses the recipient domain, instead of using $myhostname as in Postfix 2.6 and earlier. To restore the old behavior specify "default_filter_nexthop = $myhostname", or specify a non-empty next-hop content filter destination.
- Postfix now requests default delivery status notifications when adding a recipient with the Milter smfi_addrcpt action, instead of "never notify" as with Postfix automatically-added recipients.

- Postfix now reports a temporary delivery error when the result of virtual alias expansion would exceed the virtual_alias_recursion_limit or virtual_alias_expansion_limit.

- To avoid repeated delivery to mailing lists with pathological nested alias configurations, the local(8) delivery agent now keeps the owner-alias attribute of a parent alias, when delivering mail to a child alias that does not have its own owner alias.

- The Postfix SMTP client no longer appends the local domain when looking up a DNS name without ".". Specify "smtp_dns_resolver_options = res_defnames" to get the old behavior, which may produce unexpected results.

- The format of the "postfix/smtpd[pid]: queueid: client=host[addr]" logfile record has changed. When available, the before-filter client information and the before-filter queue ID are now appended to the end of the record.

- Postfix by default no longer adds a "To: undisclosed-recipients:;" header when no recipient specified in the message header. For backwards compatibility, specify: "undisclosed_recipients_header = To: undisclosed-recipients:;"

- The Postfix SMTP server now always re-computes the SASL mechanism list after successful completion of the STARTTLS command. Earlier versions only re-computed the mechanism list when the values of smtp_sasl_tls_security_options and smtp_sasl_security_options differ. This could produce incorrect results, because the Dovecot authentication server may change responses when the SMTP session is encrypted.

- The smtpd_starttls_timeout default value is now stress-dependent. By default, TLS negotiations must now complete under overload in 10s instead of 300s.

- Postfix no longer appends the system-supplied default CA certificates to the lists specified with *_tls_CAfile or with *_tls_CApath. This prevents third-party certificates from getting mail relay permission with the permit_tls_all_clientcerts feature. Unfortunately this change may cause compatibility problems when configurations rely on certificate verification for other purposes. Specify "tls_append_default_CA = yes" for backwards compatibility.

- The VSTREAM error flags are now split into separate read and write error flags. As a result of this change, all programs that use Postfix VSTREAMs MUST be recompiled.

- For consistency with the SMTP standard, the (client-side) smtp_line_length_limit default value was increased from 990 characters to 999 (i.e. 1000 characters including <CR><LF>). Specify "smtp_line_length_limit = 990" to restore historical Postfix behavior.

- To simplify integration with third-party applications, the Postfix sendmail command now always transforms all input lines ending in <CR><LF> into UNIX format (lines ending in <LF>). Specify "sendmail_fix_line_endings = strict" to restore historical Postfix behavior.

- To work around broken remote SMTP servers, the Postfix SMTP client by default no longer appends the "AUTH=<>" option to the MAIL FROM command. Specify "smtp_send_dummy_mail_auth = yes" to restore the old behavior.

- Instead of terminating immediately with a "fatal" message when a database file can't be opened, a Postfix daemon program now logs an "error" message, and continues execution with reduced functionality. Logfile-based alerting systems may need to be updated to look for "error" messages in addition to "fatal" messages. Specify "daemon_table_open_error_is_fatal = yes" to get the historical behavior (immediate termination with "fatal" message).

- Postfix now logs the result of successful TLS negotiation with TLS logging levels of 0.

- The default inet_protocols value is now "all" instead of "ipv4", meaning use both IPv4 and IPv6. To avoid an unexpected loss of performance for sites without global IPv6 connectivity, the commands "make upgrade" and "postfix upgrade-configuration" now append "inet_protocols = ipv4" to main.cf when no explicit inet_protocols setting is already present.

New Features:

- Support for managing multiple Postfix instances. Multi-instance support allows you to do the following and more: - Simplify post-queue content filter configuration by using separate Postfix instances before and after the filter. - Implement per-user content filters (or no filter) via transport map lookups instead of content_filter settings. - Test new configuration settings (on a different server IP address or TCP port) without disturbing production instances.

- check_reverse_client_hostname_access, to make access decisions based on the unverified client hostname.

- With "reject_tempfail_action = defer", the Postfix SMTP server immediately replies with a 4xx status after some temporary error.

- The Postfix SMTP server automatically hangs up after replying with "521". This makes overload handling more effective. See also RFC 1846 for prior art on this topic.

- Stress-dependent behavior is enabled by default. Under conditions of overload, smtpd_timeout is reduced from 300s to 10s, smtpd_hard_error_limit is reduced from 20 to 1, and smtpd_junk_command_limit is reduced from 100 to 1.

- Specify "tcp_windowsize = 65535" (or less) to work around routers with broken TCP window scaling implementations.

- New "lmtp_assume_final = yes" flag to send correct DSN "success" notifications when LMTP delivery is "final" as opposed to delivery into a content filter.

- The Postfix SMTP server's SASL authentication was re-structured. With "smtpd_tls_auth_only = yes", SASL support is now activated only after a successful TLS handshake. Earlier Postfix SMTP server versions could complain about unavailable SASL mechanisms during the plaintext phase of the SMTP protocol.

- Improved before-queue filter performance. With "smtpd_proxy_options = speed_adjust", the Postfix SMTP server receives the entire message before it connects to a before-queue content filter. This means you can run more SMTP server processes with the same number of running content filter processes, and thus, handle more mail. This feature is off by default until it is proven to create no new problems.

- sender_dependent_default_transport_maps, a per-sender override for default_transport.

- milter_header_checks: Support for header checks on Milter-generated message headers. This can be used, for example, to control mail flow with Milter-generated headers that carry indicators for badness or goodness. Currently, all header_checks features are implemented except PREPEND.

- Support to turn off the TLSv1.1 and TLSv1.2 protocols. Introduced with OpenSSL version 1.0.1, these are known to cause inter-operability problems with for example hotmail. The radical workaround is to temporarily turn off problematic protocols globally: smtp_tls_protocols = !SSLv2, !TLSv1.1, !TLSv1.2 smtp_tls_mandatory_protocols = !SSLv2, !TLSv1.1, !TLSv1.2

- Prototype postscreen(8) server that runs a number of time-consuming checks in parallel for all incoming SMTP connections, before clients are allowed to talk to a real Postfix SMTP server. It detects clients that start talking too soon, or clients that appear on DNS blocklists, or clients that hang up without sending any command.

- Support for address patterns in DNS blacklist and whitelist lookup results.

- The Postfix SMTP server now supports DNS-based whitelisting with several safety features: permit_dnswl_client whitelists a client by IP address, and permit_rhswl_client whitelists a client by its hostname. These features use the same syntax as reject_rbl_client and reject_rhsbl_client, respectively. The main difference is that they return PERMIT instead of REJECT.

- The SMTP server now supports contact information that is appended to "reject" responses. This includes SMTP server responses that aren't logged to the maillog file, such as responses to syntax errors, or unsupported commands.

- tls_disable_workarounds parameter specifies a list or bit-mask of OpenSSL bug work-arounds to disable.

- The lower-level code in the TLS engine was simplified by removing an unnecessary layer of data copying. OpenSSL now writes directly to the network.

- enable_long_queue_ids Introduces support for non-repeating queue IDs (also used as queue file names). These names are encoded in a mix of upper case, lower case and decimal digit characters. Long queue IDs are disabled by default to avoid breaking tools that parse logfiles and that expect queue IDs with the smaller [A-F0-9] character set.

- memcache lookup and update support. This provides a way to share postscreen(8) or verify(8) caches between Postfix instances.

- Support for TLS public key fingerprint matching in the Postfix SMTP client (in smtp_tls_policy_maps) and server (in check_ccert access maps).

- Support for external SASL authentication via the XCLIENT command. This is used to accept SASL authentication from an SMTP proxy such as NGINX. This support works even without having to specify "smtpd_sasl_auth_enable = yes".

11.6.2.3. Postfix Banner Less Verbose

The SMTP MTA banner sent to the client upon connection is less verbose now. It does not print the services name and version number anymore.

The SMTP MTA banner sent to the client upon connection is less verbose now. It does not print the services name and version number anymore.

11.6.2.4. Update RRDTool to 1.4.7

RRDTool 1.4.5 is a drop-in replacement without any incompatible changes.

11.6.2.5. IBM Java 1.4.2 End of Life

As announced with SUSE Linux Enterprise Server 11 SP2, IBM Java 1.4.2 reached End of Life, and thus we remove support for this specific Java version with SUSE Linux Enterprise Server 11 SP3. We recommend to upgrade your environments.

11.6.2.6. 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.

11.6.2.7. 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.

11.6.2.8. 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.

11.6.2.9. 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.

11.6.3. Security

11.6.3.1. 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
```

11.6.3.2. 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.

11.6.4. Networking

11.6.4.1. 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.

11.6.4.2. 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
```

11.6.4.3. 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.

11.6.5. Cross Architecture Information

11.6.5.1. ATI Radeon ES1000 Support

***CHECKIT If upgrading from SP1 to SP2 on a system with an ATI Radeon ES1000 video chip, there may be issues with the color palette when running Xorg. To avoid this issue, regenerate a new `xorg.conf` file after the installation with:

```
sax2 -a -r
```

This will allow the Xorg vesa driver to control the video chip.

11.6.5.2. 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.

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

11.7.1. System and Vendor Specific Information

11.7.1.1. IBM System x Servers: Installation on 4KB Sector Drives Not Supported

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

11.7.1.2. 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/us/Processor_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 reenable booting, the "`allow_unsafe`" option needs to be added to `XEN_APPEND` in `/etc/sysconfig/bootloader` as follows:

```
XEN_APPEND="allow_unsafe"
```

11.7.1.3. 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.

11.7.1.4. 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.

11.7.1.5. 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.

11.7.1.6. 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.

11.7.1.7. NIC Enumeration

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

11.7.1.8. HP Linux ProLiant Support Pack for SUSE Linux Enterprise Server 11

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

For more details, see Novell TID 7002735.

11.7.1.9. 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.

11.7.1.10. 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.

11.7.1.11. 32-Bit Devel-Packages Missing from the Software Development Kit (x86_64)

Example: libpcap0-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.

11.7.2. Virtualization

11.7.2.1. Hyper-V: Memory Ballooning Support

Windows hosts dynamically manage the guest memory allocation via a combination memory hot add and ballooning. Memory hot add is used to grow the guest memory upto the maximum memory that can be allocatted to the guest. Ballooning is used to both shrink as well as expand up to the max memory.

11.7.2.2. Xen Support for Booting the Hypervisor to UEFI X64

The hypervisor is now able to boot to UEFI.

11.7.2.3. 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 websites identify processors, which support hardware virtualization:

- http://wiki.xensource.com/xenwiki/HVM_Compatible_Processors
- 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`.

11.7.2.4. 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.

11.7.2.5. 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.

11.7.2.6. 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.

11.7.2.7. 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.

11.7.2.8. Time Synchronization in titlevirtualized 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.

11.7.3. RAS

11.8. Intel Itanium (ia64) Specific Information

11.8.1. 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.

11.9. POWER (ppc64) Specific Information

11.9.1. Support for the IBM POWER7+ Accelerated Encryption and Random Number Generation

For more information on making use of the IBM POWER7+ crypto and RNG accelerators, please see: https://www.ibm.com/developerworks/mydeveloperworks/files/form/anonymous/api/library/f57fde24-5f30-4295-91fb-e612c6a7a75a/document/4a8d6ce4-6e1f-4203-b9b9-1d7747cec644/media/power7%2B-accelerated-encryption-for-linux-v3.pdf

11.9.2. POWER7+ Random Number Generator

Support the POWER7+ on-chip Random Number Generator.

11.9.3. Add Per-process Data Stream Control Register (DSCR) Support

The current kernel supports setting system-wide DSCR (Data Stream Control Register) value using sysfs interface (/sys/devices/system/cpu/dscr_default). This system-wide DSCR value will be inherited by new processes until user changes this value again. So users cannot modify and/or retrieve DSCR value for each process separately.

The powerpc-utils package shipped in this release provides the modified ppc64_cpu command. This command allows users to set and read DSCR value per process basis.

11.9.4. Check Sample Instruction Address Register (SIAR) Valid Bit before Saving Contents of SIAR

The POWER7 processor has a register, referred to as Sample Instruction Address Register. This register is loaded with the contents of instruction address when a sample of a performance monitoring event is taken. If an instruction that was executed speculatively is rolled back, the event is also rolled back but the contents of SIAR are not cleared and thus invalid. The kernel has no way of detecting that the contents of SIAR are invalid. This can result in a few profiling samples with incorrect instruction addresses.

The POWER7+ processor adds a new bit, referred to as SIAR-Valid bit and sets this bit to indicate when the contents of the SIAR are valid. The new SLES 11 SP3 kernel checks this bit before saving the contents of the SIAR in a sample. This ensures that the instruction addresses saved in profiling samples are correct.

11.9.5. LightPath Diagnostics Framework for IBM Power

IBM Power systems have Service indicators (LED) that help identify components (Guiding Light) and also to indicate a component in error (Light Path). Currently, Linux only has a couple of commands that cater to LightPath services.

Deliver a LightPath framework that will help customers to identify a hardware component in error on IBM Power Systems

11.9.6. PRRN Event Handling

The latest versions of firmware for IBM Power Systems provide customers the opportunity to have the affinity for the resources on their systems dynamically updated. This procedure occurs via a Platform Resource Reassignment Notification (PRRN) Event.

The updates to the ppc64-diag, powerpc-utils, and librtas packages allow Linux systems to handle these PRRN events and update the affinity for system cpu and memory resources.

11.9.7. Increase Number of Partitions per Core on IBM POWER7+

Enable support for 20 partitions per core on IBM POWER7+

11.9.8. Enable Firmware Assisted Dump for IBM Power Systems

Starting from IBM POWER6 and above the Power firmware now has a capability to preserve the partition memory dump during system crash and boot into a fresh copy of the kernel with fully-reset system. This feature adds support to exploit the dump capture capability provided by Power firmware

11.9.9. Kernel cpuidle Framework for POWER7

Enable POWER systems to leverage the generic cpuidle framework by taking advantage of advanced heuristics, tunables and features provided by the cpuidle framework. This enables better power management on the systems and helps tune the system and applications accordingly.

11.9.10. 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.

11.9.11. 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.

11.9.12. 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"
```

11.9.13. 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.

11.9.14. Installation on POWER onto IBM VSCSI Target

The installation on a vscsi client will fail with old versions of the AIX VIO server. Please upgrade the AIX VIO server to version 1.5.2.1-FP-11.1 or later.

11.9.15. iSCSI Installations with Multiple NICs Losing Network Connectivity at the End of Firstboot Stage

***CHECKIT (still valid for SP3?) After installing SLES 11 SP1 on an iSCSI target, the system boots properly, network is up and the iSCSI root device is found as expected. The install completes (firstboot part) as usual. However, at the end of firstboot, the network is shut down before the root file system is unmounted, leading to read failures accessing the root (iSCSI) device; the system hangs.

Solution: reboot the system.

11.9.16. 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

11.9.17. 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

11.9.18. 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.

11.9.19. 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.

11.9.20. 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.

11.9.21. 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.

11.9.22. 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.

11.9.23. 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
```

11.10. 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.

11.10.1. Hardware

11.10.1.1. Leverage Cross Memory Attach Functionality for System z

Cross memory attach reduces the number of data copies needed for intra-node interprocess communication. In particular, MPI libraries engaged in intra-node communication can now perform a single copy of the message to shared memory rather than performing a double copy.

11.10.1.2. CryptoExpress4 - Device Driver Exploitation

With SLES 11 SP3 the z90crypt device driver supports the Crypto Express 4 (CEX4) adapter card.

11.10.1.3. Implement lscpu and chcpu

This feature improves handling of CPU hotplug. The lscpu command now displays detailed information about available CPUs. Using a new command, chcpu, you can change the CPU state, disable and enable CPUs, and configure specified CPUs.

11.10.1.4. CPACF Exploitation (libica Part 2)

This feature extends the libica library with new modes of operation for DES, 3DES and AES. These modes of operation (CBC-CS, CCM, GCM, CMAC) are supported by Message Security Assist (CPACF) extension 4, which can be used with z196 and later System z mainframes.

11.10.1.5. Exploitation of Data Routing for FCP

This feature supports the enhanced mode of the System z FCP adapter card. In this mode, the adapter passes data directly from memory to the SAN when there is no free memory on the adapter card because of large or slow I/O requests.

11.10.2. Virtualization

11.10.2.1. VEPA Mode Support

VEPA mode routes traffic between virtual machines on the same mainframe through an external switch. The switch then becomes a single point of control for security, filtering, and management.

11.10.2.2. Technology preview: KVM support on s390x

KVM is now included on the s390x platform as a technology preview.

11.10.2.3. Support of Live Guest Relocation (LGR) with z/VM 6.2 on SLES 11 SP2

***CHECKIT (still valid for SP3?) Live guest relocation (LGR) with z/VM 6.2 on SLES 11 SP2 requires z/VM service applied, especially with Collaborative Memory Management (CMMA) active (`cmma=on`).

Apply z/VM APAR VM65134.

11.10.2.4. 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

11.10.3. Storage

11.10.3.1. Safe Offline Interface for DASD Devices

Instead of setting a DASD device offline and returning all outstanding I/O requests as failed, with this interface you can set a DASD device offline and write all outstanding data to the device before setting the device offline.

11.10.3.2. Flash Express Support for IBM System z

Flash Express memory is accessed as storage-class memory increments. Storage-class memory for IBM System z is a class of data storage devices that combine properties of both storage and memory. This feature improves the paging rate and access performance for temporary storage, for example, for data warehousing.

11.10.3.3. Detect DASD Path Connection Error

This feature enables the Linux DASD device driver to detect path configuration errors that cannot be detected by hardware or microcode. The device driver then does not use such paths. For example, with this feature, the DASD device driver detects paths that are assigned to a specific subchannel, but lead to different storage servers.

11.10.3.4. SAN Utilities for zFCP, hbaapi Completion

Improves systems manageability by supporting pass-through for generic services and retrieving events in the SAN. Improves SAN setup by retrieving information about the SAN fabric including all involved interconnect elements, such as switches.

11.10.3.5. Enhanced DASD Statistics for PAV and HPF

This feature improves DASD I/O diagnosis, especially for Parallel Access Volume (PAV) and High Performance FICON (HPF) environments, to analyze and tune DASD performance.

11.10.3.6. New Partition Types Added to the fdasd Command

In SLES11 SP2 new partition types were added to the **fdasd** command in the `s390-tools` package. Anyone using YaST in SP3 to create partitions will not see this happening. If **fdasd** is used from the command line, it will work as documented and desired.

11.10.4. Network

11.10.4.1. 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.

11.10.4.2. 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 `LLADR='00:00:00:00:00:00'` from the `ifcfg` file in `/etc/sysconfig/network`.

11.10.4.3. 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.

11.10.5. Security

11.10.5.1. Support of SHA-256 Hash Algorithm in opencryptoki CCA Token

SLES 11 SP3 includes opencryptoki 2.4.2 which comes with a CCA token that exploits the SHA-256 hash algorithm that is provided by System z crypto hardware.

11.10.5.2. CryptoExpress4 - Device Driver Exploitation

With SLES 11 SP3 the z90crypt device driver supports the Crypto Express 4 (CEX4) adapter card.

11.10.5.3. CPACF Exploitation (libica Part 2)

This feature extends the libica library with new modes of operation for DES, 3DES and AES. These modes of operation (CBC-CS, CCM, GCM, CMAC) are supported by Message Security Assist (CPACF) extension 4, which can be used with z196 and later System z mainframes.

11.10.5.4. Existing Data Execution Protection Removed for System z

The existing data execution protection for Linux on System z relies on the System z hardware to distinguish instructions and data through the secondary memory space mode. As of System z10, new load-relative-long instructions do not make this distinction. As a consequence, applications that have been compiled for System z10 or later fail when running with the existing data execution protection.

Therefore, data execution protection for Linux on System z has been removed.

11.10.6. RAS

11.10.6.1. Crypto Adapter Resiliency

This feature provides System z typical RAS for cryptographic adapters through comprehensive failure recovery. For example, this feature handles unexpected failures or changes caused by Linux guest relocation, suspend and resume activities or configuration changes.

11.10.6.2. Fuzzy Live Dump for System z

With this feature kernel dumps from running Linux systems can be created, to allow problem analysis without taking down systems. Because the Linux system continues running while the dump is written, and kernel data structures are changing during the dump process, the resulting dump contains inconsistencies.

11.10.6.3. kdump Support for System z

kdump can be used to create system dumps for instances of SUSE Linux Enterprise Server. kdump reduces dump time and size, facilitates dump disk sharing. A setup GUI is provided by YaST. Any IBM System z system with kdump support and more than 4 GB of memory has kdump enabled by default. When performing an upgrade to SLES 11 SP3, note that kdump reserves approximately 128 MB by default and sufficient disk space must be available for storing the dump.

11.10.6.4. Distinguish Dump System and Boot System

A dump system is not necessarily identical to the system that was booted. Linux guest relocation or suspend and resume activities might introduce problems. To help analyze such problems, a system dump now provides location information about the original Linux system.

11.10.7. Performance

11.10.7.1. Leverage Cross Memory Attach Functionality for System z

Cross memory attach reduces the number of data copies needed for intra-node interprocess communication. In particular, MPI libraries engaged in intra-node communication can now perform a single copy of the message to shared memory rather than performing a double copy.

11.10.7.2. Support of the Transactional Execution Facility and Runtime Instrumentation

With the facility the Linux kernel supports hardware runtime instrumentation, an advanced mechanism that improves analysis of and optimization of the code generated by the new IBM JVM. Software locking overhead is minimized and scalability and parallelism increased.

11.10.7.3. System z Performance Counters in the Linux perf Tool

This feature provides simplified performance analysis for software on Linux on System z. It uses the perf tool to access the hardware performance counters.

11.10.7.4. Optimized Compression Library zlib

This feature provides optimization of and support for the general purpose data compression library zlib. This library improves compression performance on System z.

11.10.7.5. Libhugetlbfs support for System z

Enables the transparent exploitation of large pages in C/C++ programs. Applications and middleware programs can profit from the performance benefits of large pages without changes or recompilation.

11.10.8. Miscellaneous

11.10.8.1. 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)

11.10.8.2. 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

11.10.8.3. Large Page Support in IBM System z

Large Page support allows processes to allocate process memory in chunks of 1 MiB instead of 4 KiB. This works through the hugetlbfs.

11.10.8.4. Collaborative Memory Management Stage II (CMM2) Lite

SLES 11 SP2 supports CMM2 Lite for optimized memory usage and to handle memory overcommitment via memory page state transitions based on "stable" and "unused" memory pages of z/VM guests using the existing `arch_alloc_page` and `arch_free_page` callbacks.

11.10.8.5. 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.

Chapter 12. Resolved Issues

- Bugfixes

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

- 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.

Chapter 13. 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.

13.1. 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.

SLES 11 (3.0)	***x86***	***ia64***	***x86_64***	***s390x***	***ppc64***
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
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				

13.2. KVM Limits

Guest RAM size	512 GiB
Virtual CPUs per guest	64
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.

13.2.1. QEMU: Version 1.4 Master Feature

13.2.2. Technology preview: QEMU: Include virtio-blk-data-plane

The virtio-blk-data-plane is a new experimental performance feature for KVM. It provides a streamlined block IO path which favors performance over functionality.

13.2.3. Technology Preview: KVM Nested Virtualization with Intel VT

The KVM kernel module "kvm_intel" now has the nested parameter available, achieving parity with the "kvm_amd" kernel module with respect to nested virtualization capabilities.

13.2.4. XEN/KVM: virt-manager Can Configure PCI Pass-through Devices at VM Creation

Virt-Manager is now capable to allow the configuration of PCI pass-through devices at VM creation in Xen and KVM.

13.2.5. libseccomp

Seccomp filters are expressed as a Berkeley Packet Filter (BPF) program, which is not a well understood interface for most developers.

The libseccomp library provides an easy to use interface to the Linux Kernel's syscall filtering mechanism, seccomp. The libseccomp API allows an application to specify which syscalls, and optionally which syscall arguments, the application is allowed to execute, all of which are enforced by the Linux Kernel.

13.2.6. libvirt Support for QEMU seccomp Sandboxing

QEMU guests spawned by libvirt are exposed to a large number of system calls that go unused for the entire lifetime of the process.

libvirt's qemu.conf file is updated with a seccomp_sandbox option that can be used to enable use of QEMU's seccomp sandboxing support. This allows execution of QEMU guests with reduced exposure to kernel system calls.

13.2.7. libvirt Bridged Networking for Unprivileged Users

libvirt can already spawn QEMU guests with bridged networking support when running under a privileged user ID, however it cannot do the same when run under an unprivileged user ID.

libvirt is updated to enable QEMU guests to be spawned with bridged networking when libvirt is run under an unprivileged user ID. This benefits installations that connect to the libvirtd instance with the qemu:/// session URI. This was achieved by using the new QEMU network helper support when libvirt is running under an unprivileged user ID.

13.2.8. libvirt DAC Isolation

libvirt spawns all QEMU guests created through the qemu:///system URI under the user ID and group ID defined in /etc/libvirt/qemu.conf. This means all guests are run under the same user ID and group ID, removing all Discretionary Access Control (DAC). While Mandatory Access Control (MAC) may already be isolating guests, it would be nice to also have DAC isolation for an added layer of security.

libvirt has been updated to allow spawning of guests under unique user and group IDs. The libvirt domain XML's <seclabel> tag is updated with model='dac' to provide this support, and libvirt APIs are updated to allow applications to inspect the full list of security labels of a domain.

13.2.9. QEMU Network Helper for Unprivileged Users

QEMU guests previously could not be started with bridged networking support when run under an unprivileged user ID.

Infrastructure is introduced to enable a network helper to be executed by QEMU. This also allows third parties to implement user-visible network backends without having to introduce them into QEMU itself. A default network helper is introduced that implements the same bridged networking functionality as the common qemu-ifup script. It creates a tap file descriptor, attaches it to a bridge, and passes it back to QEMU. This helper runs with higher privileges, allowing QEMU to be invoked with bridged networking support under an unprivileged user.

13.2.10. QEMU: Sandboxing with seccomp

New seccomp kernel functionality is intended to be used to declare the whitelisted syscalls and syscall parameters. This will limit QEMU's syscall footprint, and therefore the potential kernel attack surface. The idea is that if an attacker were to execute arbitrary code, they would only be able to use the whitelisted syscalls.

QEMU has been updated with the '-sandbox' option. When set to 'on', the '-sandbox' option will enable seccomp system call filtering for QEMU, allowing only a subset of system calls to be used.

13.2.11. KVM: Export Platform Power Management Capability through libvirt Framework

Libvirt can now discover and update tags in the capabilities XML field based on power management features supported by the platform.

13.2.12. KVM: Support INVPCID's Haswell Instructions

KVM now support the new Haswell CPU instructions: INVPCID. Process-context identifiers (PCIDs) are a facility by which a logical processor may cache information for multiple linear-address spaces so that

the processor may retain cached information when software switches to a different linear address space. INVPCID instruction is used for fine-grained TLB flush which is benefit for kernel. This features is now exposed to the guest. Modern guest can use this new instructiosn to improve the efficiency of KVM. qemu-kvm is required to selects PCID via -cpu option.

13.2.13. KVM: TSC Deadline Timer Support

TSC deadline timer is a new mode in LAPIC timer, which will generate one-shot timer interrupt based on TSC deadline, in place of current APIC clock count interval. It will provide more precise timer interrupt (less than 1 ticks) to benefit OS scheduler etc.

13.2.14. KVM: TSC Offset Timer

TSC is only writable via MSR 0x10 which is a moving target. TSC offset timer feature will provide a new MSR 0x3b that exposes the "Thread Offset" directly.

13.2.15. KVM: Support for APIC Virtualization

Starting from IvyTown processor, APIC Virtualization provides more supports to improve VMM interrupt handling efficiency. There are two features: - APIC-Register Virtualization: a new VM-execution control is introduced, which eliminates almost all VM exits on reads of APIC registers and provides more information for memory-mapped APIC writes - Virtual-Interrupt Delivery: a new VM-execution control is introduced, which CPU delivers virtual interrupt through guest IDT when appropriate

13.2.16. KVM: Haswell New Instructions Support

KVM now support the Haswell CPU new instructions (ie: FP fused Multiply Add, 256-bit Integer vectors, MOVBE support...). Using some of these new instructions will improve the efficiency of KVM.

13.2.17. KVM: support for Supervisor Mode Execution Protection (SMEP)

KVM now support tje Supervisor mode execution protection (SMEP) wich prevents execution of user mode pages while in supervisor mode and addresses class of exploits for hijacking kernel execution.

13.2.18. XEN/KVM/libvirt: Virtual Machine Lock Manager

The virtual machine lock manager is a daemon which will ensure that a virtual machine's disk image cannot be written to by two QEMU/KVM processes at the same time. It provides protection against starting the same virtual machine twice, or adding the same disk to two different virtual machines.

13.3. Xen Limits

SLES 11 SP2	***x86***
CPU bits	64
Logical CPUs (Xen Hypervisor)	255
Virtual CPUs per VM	32
Maximum supported memory (Xen Hypervisor)	2 TiB

SLES 11 SP2	***x86***
Maximum supported memory (Dom0)	512 GiB
Virtual memory per VM	128 MiB-256 GiB
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	128

In Xen 4.1, the hypervisor bundled with SUSE Linux Enterprise Server 11 SP2, 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.

13.3.1. XEN: Secure Boot

Xen hypervisor is shipped as an EFI application, and signed. It will negotiate with the shim loader to validate the Dom0 kernel signature before booting it. Enabling the alternative kernel image format takes as a prerequisite the bumping of the backward compatibility level from 3.2 to 4.X, so we are not able to boot a SLE11 SP3 PV guest on SLE10 SP4, even if secure boot is not enable.

13.3.2. XEN/KVM: virt-manager Can Configure PCI Pass-through Devices at VM Creation

Virt-Manager is now capable to allow the configuration of PCI pass-through devices at VM creation in Xen and KVM.

13.3.3. XEN: Netconsole Support to Netfront Device

XEN now support netconsole on its netfront device.

13.3.4. XEN: TSC Deadline Timer Support

TSC deadline timer is a new mode in LAPIC timer, which will generate one-shot timer interrupt based on TSC deadline, in place of current APIC clock count interval. It will provide more precise timer interrupt (less than 1 ticks) to benefit OS scheduler etc.

13.3.5. XEN: JKT Core Error Recovery

Xen now support the new MCA type to handle errors in the core (like L1/L2 cache error). Previously only uncore errors (like L3 cache error) was handled.

13.3.6. XEN: TSC Offset Support

TSC is only writable via MSR 0x10 which is a moving target. TSC offset timer feature will provide a new MSR 0x3b that exposes the "Thread Offset" directly.

13.3.7. XEN: Haswell New Instructions Support

XEN now support the Haswell CPU new instructions (ie: FP fused Multiply Add, 256-bit Integer vectors, MOVBE support...). Using some of these new instructions will improve the efficiency of XEN.

13.3.8. APIC Virtuatization in Xen and KVM

This Service Pack adds support for the APIC virtualization feature for Intel's IvyBridge and later CPUs. Both hypervisors - Xen and KVM - support APICv.

13.3.9. XEN: Large VT-d Pages

This is an IOMMU performance enhancement to reduce IOMMU page table and IOTLB footprint.

13.3.10. XEN/KVM/libvirt: Virtual Machine Lock Manager

The virtual machine lock manager is a daemon which will ensure that a virtual machine's disk image cannot be written to by two QEMU/KVM processes at the same time. It provides protection against starting the same virtual machine twice, or adding the same disk to two different virtual machines.

13.3.11. XEN: Bios Information to XEN HVM Guest

Bios information of the physical server can now be passed to XEN HVM guest system.

13.3.12. XEN: Support for PCI Pass-through Bind and Unbind in libvirt Xen Driver

Virt-manager is now able to set up PCI pass-through for Xen without having to switch to the command line to free the PCI device before assigning it to the VM.

13.3.13. XEN: xenstore-chmod Command Now Support 256 Permissions

To be able to manage permission using the xenstore-chmod command on more than 16 domUs at the same time, xenstore-chmod command now support 256 permissions.

13.4. 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`.

Feature	***Ext 3***	***Reiserfs 3.6***	***XFS***	***Btrfs ****	***OCFS 2 *****
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	•	•	•	^	•
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
	* 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".				
	** OCFS2 is fully supported as part of the SUSE Linux Enterprise High Availability Extension.				

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.

13.4.1. XFS Realtime Volumes

XFS Realtime Volumes is an experimental feature, available for testing and experimenting. If you encounter any issues, SUSE is interested in feedback. Please, submit a support request through the usual access methods.

13.4.2. ext4: Runtime Switch for Write Support

The SLE 11 SP3 kernel contains a fully supported ext4 file system module, which provides read-only access to the file system.

Read-write access to an ext4 file system can be acquired by setting the `rw` kernel module parameter to 1, either through module load time options or after module load through the kernel sysctl interface. Be aware that this action will render the kernel module and the kernel as the whole as unsupported upon first read-write mount of an ext4 file system.

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

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

13.5. 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`.

Technical Background

- 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 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.

13.6. 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.

13.6.1. IPv6 Support for NFSv3

Kernel configuration and NFS userland utilities have been updated to fully support NFSv3 over the IPv6 protocol. The same functionality for NFSv4 has already been enabled since SUSE Linux Enterprise 11 SP2.

13.6.2. Add IPv6 support to AutoFS

13.6.3. Linux Virtual Server Load Balancer (ipvs) Extends Support for IPv6

The LVS/ipvs load balancing code did not fully support RFC2460 and fragmented IPv6 packets which could lead to lost packets and interrupted connections when IPv6 traffic was fragmented.

The load balancer has been enhanced to fully support IPv6 fragmented extension headers and is now RFC2460 compliant.

13.7. Other Technical Information

13.7.1. libica 2.1.0 Available in SLES 11 SP2 for s390x

***CHECKIT (adjusting for SP3?) 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, and libica-2_1_0 providing libica versions 1.3.9, 2.0.2, and 2.1.0 respectively.

libica 1.3.9 is provided for compatibility reasons with legacy hardware present e.g. in the ppc64 architecture. For s390x users it is always recommended to use the new libica 2.1.0 library since it supports all newer s390x hardware, larger key sizes and is backwards compatible with any ICA device driver in the s390x architecture.

You may choose to continue using libica 1.3.9 or 2.0.2 if you do not have newer Cryptographic hardware to exploit or wish continue using custom applications that do not support the libica 2.1.0 library yet. Both openCryptoki and openssl-ibmca, the two main exploiters for the libica interface, are provided in SLES 11 SP2 to support the newer libica 2.1.0 library.

13.7.2. 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

13.7.3. 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.

13.7.4. 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.

13.7.5. 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
```

13.7.6. 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.

13.7.7. 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.

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

   ```
   rcsshd restart
   ```

13.7.9. 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/`.

13.7.10. 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 now 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.

Chapter 14. 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 3 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/.

14.1. Additional or Update Documentation

- http://www.suse.com/documentation/sles11/ contains additional or updated documentation for SUSE Linux Enterprise Server 11 Service Pack 3.
- 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.

14.2. 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.

Chapter 15. Miscellaneous

Chapter 16. 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.novell.com/info/exports/ for more information on exporting SUSE software. SUSE assumes no responsibility for your failure to obtain any necessary export approvals.

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.novell.com/company/legal/patents/ and one or more additional patents or pending patent applications in the U.S. and other countries.

For SUSE trademarks, see Novell Trademark ad Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html). 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.