



# Release Notes

---

## SUSE Linux Enterprise Desktop 12

These release notes are generic for all SUSE Linux Enterprise Desktop 12 based products. Some parts may not apply to particular architectures or products. Where this is not the case, the respective architecture is listed explicitly.

Manuals can be found in the docu directory of the installation media, or in the directory /usr/share/doc/ on the installed system (if installed).

Publication date: 2014-10-10 , Version: 12.0.20141010

## Contents

- 1 SUSE Linux Enterprise Desktop 2
- 2 Installation and Upgrade 7
- 3 Infrastructure, Package and Architecture Specific Information 12
- 4 Driver Updates 23
- 5 Packages and Functionality Changes 24
- 6 Technical Information 34
- 7 Legal Notices 35

# 1 SUSE Linux Enterprise Desktop

SUSE Linux Enterprise Desktop is the market's only enterprise-quality Linux desktop ready for routine business use. Developed and backed by SUSE, SUSE Linux Enterprise Desktop provides market-leading usability, seamless interoperability with existing IT systems, and dozens of essential applications—all at a fraction of the price of proprietary operating systems. It comes bundled with the latest versions of leading applications such as LibreOffice office productivity suite, Mozilla Firefox web browser, and Evolution e-mail and calendar suite. In addition, it integrates with Microsoft SharePoint and Novell Teaming for group collaboration and supports a wide range of multimedia file formats, wireless and networking standards, and plug-and-play devices.

Through the latest enhancements in power management and security, SUSE Linux Enterprise Desktop also provides an environmentally friendly IT experience (Green IT) and an error-proof desktop. Finally, SUSE Linux Enterprise Desktop offers unparalleled flexibility. You can deploy it on a wide range of thick client devices (including desktops, notebooks, netbooks, and workstations), on thin client devices, or as a virtual desktop. By leveraging the power of SUSE Linux Enterprise Desktop, your business can dramatically reduce costs, improve end user security and increase workforce productivity.

SUSE Linux Enterprise Desktop 12 has a 7 years life cycle. The current version (GA) will be fully maintained and supported until 6 months after the release of SUSE Linux Enterprise Desktop 12 SP1.

## 1.1 What's New?



## Note: Fix Status of the GNU Bourne Again Shell (bash)

Given the proximity of the SUSE Linux Enterprise 12 release to the publication of the “shellshock” series of vulnerabilities in the GNU Bourne Again Shell (bash), we want to provide customers with information on the fix status of the bash version shipped in the SLE 12 GA release:

- CVE-2014-6271 (original shellshock)
- CVE-2014-7169 (taviso bug)
- CVE-2014-7186 (redir\_stack bug)
- CVE-2014-7187 and
- non-exploitable CVE-2014-6277
- non-exploitable CVE-2014-6278

Up-to-date information is available online: <https://www.suse.com/support/shellshock/> .

SUSE Linux Enterprise Desktop 12 introduces a number of innovative changes. Here are some of the highlights:

- Robustness on administrative errors and improved management capabilities with full system rollback based on btrfs as the default file system for the operating system partition and SUSE's snapper technology.
- An overhaul of the installer introduces a new workflow that allows you to register your system and receive all available maintenance updates as part of the installation.
- New core technologies like systemd, replacing the time honored System V based init process.
- GNOME 3.10, giving users a modern desktop environment with a choice of several different look and feel options, including a special SLE Classic mode for easier migration from earlier SUSE Linux Enterprise desktop environments

- For users wishing to use the full range of productivity applications of a Desktop with SUSE Linux Enterprise Server, we are now offering the SUSE Linux Enterprise Workstation Extension
- Integration with the new SUSE Customer Center, SUSE's central web portal to manage Subscriptions, Entitlements, and provide access to Support.

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

- *Section 1.4, "Support Statement for SUSE Linux Enterprise Desktop"*
- *Section 2.2, "Update-Related Notes"*
- *Section 6, "Technical Information"*


## 1.2 Documentation and Other Information

### 1.2.1 Available on the Product Media

- Read the READMEs on the media.
- Get the detailed changelog information about a particular package from the RPM:

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

<FILENAME>. is the name of the RPM.

- Check the ChangeLog file in the top level of the media for a chronological log of all changes made to the updated packages.
- Find more information in the docu directory of the media of SUSE Linux Enterprise Desktop 12 CDs. This directory includes PDF versions of the SUSE Linux Enterprise Desktop 12 Installation Quick Start and Deployment Guides. Documentation (if installed) is available below the /usr/share/doc/ directory of an installed system.
- These Release Notes are identical across all architectures, and the most recent version is always available online at <http://www.suse.com/releasenotes/> .

## 1.2.2 Additional or Updated Documentation

For SUSE Linux Enterprise Desktop 12 documentation, see <http://www.suse.com/documentation/sled12/>, where you can download PDF documents. For installation with YaST software management or with zypper, packages are available on the product media. Some of these packages are installed by default. These are the package names:

- [sled-installquick\\_en-pdf](#): SLED 12 Installation Quick Start
- [sled-gnomeuser\\_en-pdf](#): SLED 12 GNOME User Guide
- [sled-admin\\_en-pdf](#): SLED 12 Administration Guide
- [sled-deployment\\_en-pdf](#): SLED 12 Deployment Guide
- [sled-security\\_en-pdf](#): SLED 12 Security Guide
- [sled-tuning\\_en-pdf](#): SLED 12 Tuning Guide
- [sled-manuals\\_en](#): the set of all SLED books in HTML format

## 1.3 How to Obtain Source Code

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

## 1.4 Support Statement for SUSE Linux Enterprise Desktop

To receive support, see <http://www.suse.com/products/desktop/>.

### 1.4.1 Technology Previews

Technology Preview features are either not supported or supported in a limited fashion. These features are mainly included for customer convenience and be functionally incomplete, unstable or in other ways not suitable for production use.

### 1.4.2 Software Requiring Specific Contracts

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

## 1.5 Derived and Related Products

### 1.5.1 Software Development Kit (SDK)

SUSE provides a Software Development Kit (SDK) for SUSE Linux Enterprise 12. This SDK contains libraries, development environments, and tools along the following patterns:

- C/C++ Development
- Certification

## 1.6 Security, Standards, and Certification

### 1.6.1 Support GB18030 Standard

SUSE Linux Enterprise conforms with Unicode 3.0 or higher, and thus it will be GB18030 compliant.

Unicode 3.0 has been supported by glibc since version 2.2. and currently SUSE Linux Enterprise uses a much newer version of glibc , so it is GB18030 compliant.

## 2 Installation and Upgrade

This section includes installation related information for this release.

### 2.1 Installation

#### 2.1.1 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).

#### 2.1.2 UEFI 2.3.1 Support

SLE 12 is supporting booting systems following UEFI specification up to version 2.3.1 errata C.

Note: Installing SLE 12 on Apple hardware is not supported.

#### 2.1.3 UEFI Secure Boot

SLES 12 and SLED 12 implement UEFI Secure Boot. Installation media supports Secure Boot. Secure Boot is only supported on new installations, if Secure Boot flag is enabled in the UEFI firmware at installation time.

For more informations, see *Administration Guide* , section *Secure Boot* .

## 2.1.4 Current Features and Limitations in a UEFI Secure Boot Context

*Support for Secure Boot on EFI machines is enabled by default.*

When booting with Secure Boot mode enabled in the firmware, the following features apply:

- Installation to UEFI default boot-loader location with a mechanism to restore boot entries.
- Reboot via UEFI.
- Xen hypervisor can be booted without MSFT signature.
- UEFI IPv6 PXE boot support.
- UEFI get videomode support, the kernel is able to retrieve the video mode from UEFI to configure KMS mode with the same parameters.
- UEFI booting from USB devices is supported

Simultaneously, the following limitations 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, not even as root user.
- Access to I/O port is not possible, not 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.
- The `acpi_rsdp` parameter does not have any effect on the kernel.

When booting with Secure Boot mode disabled in the firmware, the following features apply:

- None of the limitations listed above are active.
- The machine always stays bootable, regardless whether secure boot is later toggled in the firmware.
- The feature to retain EFI boot-manager entries after firmware updates or NVRAM resets is available even on systems without (or with disabled) Secure Boot support.



Simultaneously, the following limitations apply:

- shim.efi is always used in the boot process.

Secure boot on EFI machines can be disabled during installation by deactivating the respective option on the installation settings screen under "Bootloader".

## 2.1.5 Rollback with Snapper on Btrfs

*If an update fails or causes trouble, it is sometimes helpful to be able to go back to the last working state.*

### Requirements to Create Atomic Snapshots

- Root filesystem needs to be btrfs
- Root filesystem needs to be on one device, including /usr

That is needed since snapshots need to be atomic, and that is not possible if the data is stored on different partitions, devices, or subvolumes.

### How to Do the Rollback

During boot, you can select an old snapshot. This snapshot will then be booted in something like a read-only modus. All the snapshot data is read-only, all other filesystems or btrfs subvolumes are in read-write mode and can be modified. To make this snapshot the default for the next reboot and switch it into a read-write mode, use "snapper rollback".

### What Will Not Be Rolled Back

The following directories are excluded from rollback. This means that changes below this sub-directory will not be reverted when an old snapshot is booted, in order to not lose valuable data. On the other hand, this may prevent some third-party services from starting correctly when booting from an old snapshot.

```
/boot/grub2/i386-pc (We cannot rollback bootloader)
/boot/grub2/x86_64-efi (We cannot rollback bootloader)
/boot/grub2/power-ieee1275 (We cannot rollback bootloader)
/home (if not already on an own partition)
/opt (Prevents rollback if addons or packages are installed there)
/srv (web services may not be functional after a rollback anymore)
/tmp
```

```
/usr/local
/var/crash
/var/log (services which move files and/or permissions may not be functional anymore
after a rollback)
/var/mail (if not a symlink to /var/spool/mail)
/var/opt
/var/spool (services which move/convert files and/or permissions may not be
functional anymore after rollback)
/var/tmp
```

### Known Issues or Limitations

In general, roolback can result in inconsistencies between the data on the root partition (which has been rolled back to an earlier state) and data on other subvolumes or partitions. These inconsistencies may include the use of different file paths, formats and permissions.

- Add-ons and third party software installed in separate subvolumes or partitions, such as `/opt`, can be completely broken after a rollback of a Service Pack.
- Newly created users will vanish from `/etc/passwd` during a rollback, but the data is still in `/home` , `/var/spool` , `/var/log` and similar directories. If a new user is created later, it may be given the same user id, making it the owner of these files. This can be a security and privacy problem.
- If a package update changes permissions/ownership of files/directories inside of a subvolume (like `/var/log` , `/srv` , ...), the service may be broken after a rollback, because it is no longer able to write/access/read the files/data.
- General: if there are subvolumes like `/srv` , containing a mix of code and data, rollback may lead to loss of data or broken/non-functional code.
- General: if an update to a service introduces a new data format, rolling back to an old snapshot may render the service non-functional, if the older version is unable to handle the new data format.
- Rollback of the boot loader is not possible, since all "stages" of the boot loader must match. However, as there is only one MBR (Master Boot Record) per disk, there cannot be different snapshots of the other stages.

### 2.1.6 Installing from a USB Flash Disk

The ISO installation images can be directly dumped to a USB device such as a flash disk. This way you can install the system without the need of a DVD drive.

Several tools for dumping are listed at [http://en.opensuse.org/SDB:Live\\_USB\\_stick](http://en.opensuse.org/SDB:Live_USB_stick) .

### 2.1.7 UEFI Secure Boot

When booting the installer from the DVD product media on a secure boot enabled system, the installation process is validated by the secure boot signature.

For more information about UEFI and secure boot, see the *Administration Guide* .

## 2.2 Update-Related Notes

This section includes update-related information for this release.

### 2.2.1 File System Layout

For general information about the file system layout, see the *Administration Guide*, Chapter *Snap-per*.

#### **Additional Information**

/run/media/<user\_name> is now used as top directory for removable media mount points. It replaces /media , which is not longer available.

### 2.2.2 dhcpd Replaced by wicked and dhcp-client

dhcpd package was replaced by wicked and dhcp-client packages.

### 2.2.3 /tmp Cleanup from sysconfig Automatically Migrated into systemd Configuration

*By default, systemd cleans tmp directories daily, and systemd does not honor sysconfig settings in /etc/sysconfig/cron such as TMP\_DIRS\_TO\_CLEAR. Thus it is needed to transform sysconfig settings to avoid potential data loss or unwanted misbehavior.*

When updating to SLE 12, the variables in `/etc/sysconfig/cron` will be automatically migrated into an appropriate systemd configuration (see `/etc/tmpfiles.d/tmp.conf`). The following variables are affected:

```
MAX_DAYS_IN_TMP
MAX_DAYS_IN_LONG_TMP
TMP_DIRS_TO_CLEAR
LONG_TMP_DIRS_TO_CLEAR
CLEAR_TMP_DIRS_AT_BOOTUP
OWNER_TO_KEEP_IN_TMP
```

### 2.2.4 Migrating to SUSE Linux Enterprise 12

Migration is supported from SUSE Linux Enterprise 11 SP3 (or higher) using the following methods:

- Booting from an installation medium (ISO image)
- Automated migration from SLE 11 SP3 to 12

For more information, see the *Deployment Guide* coming with SUSE Linux Enterprise.

## 3 Infrastructure, Package and Architecture Specific Information

### 3.1 Architecture Independent Information

#### 3.1.1 Kernel

### 3.1.1.1 Ext4: Experimental Features

Ext4 has some features that are under development and still experimental. Thus, using these features poses a significant risk to data. To clearly indicate such features, the Ext4 driver in SUSE Linux Enterprise 12 refuses to mount (or mount read-write) file systems with such features. To mount such file systems set the `allow_unsupported` module parameter (either when loading the module or via `/sys/module/ext4/parameters/allow_unsupported` ). However setting this option will render your kernel, and thus your system unsupported.

Features which are treated this way are: bigalloc, metadata checksumming, and journal checksumming.

### 3.1.1.2 Enabling Full Heap Randomization

[All architectures] `CONFIG_COMPAT_BRK` has been disabled to allow randomisation of the start address of the userspace heap. This can break old binaries based on libc5. To revert to the old behavior, set the `kernel.randomize_va_space` `sysctl` to 2.

[x86\_64 only] `CONFIG_COMPAT_VDSO` has been disabled to enforce randomization of the VDSO address of 32bit binaries on x86\_64. This can break 32bit binaries using glibc older than 2.3.3. To revert to the old behavior, specify `vdso=2` on the kernel command line.

### 3.1.1.3 Format of the 'microcode' Field in /proc/cpuinfo Changed

*Due to a missing backport, the SLE 11 SP3 kernel is displaying the microcode revision in /proc/cpuinfo as a decimal number.*

The SLE 12 kernel changed the format to a hexadecimal number. Now it is compatible with the mainline kernel.

### 3.1.1.4 Initrd File Compression Format

By default, the initrd file is now compressed with:

```
xz -0 --check=crc32 --memlimit-compress=50%
```

Previously, it was compressed with gzip.

### 3.1.1.5 Blacklisting iTCO\_wdt in toshiba-aipsvcp

*If iTCO\_wdt driver is enabled, the sensor driver shows that the service processor is reporting a constant temperature in spite of heavy CPU load or the CPU fan is stopped.*

To disable the Intel watchdog functionality, we blacklist the iTCO\_wdt driver for SLES, SLED, and SLEPOS installations.

### 3.1.1.6 SDIO 3.0 Support

*Linux Kernel version 3.3 started supporting SD/SDIO version 3.0 that provides faster read/write speed and enhanced security.*

A SDIO (Secure Digital Input Output) card is an extension of the SD specification to cover I/O functions.

Host devices that support SDIO can use the SD slot to support Wi-Fi, Bluetooth, Ethernet, IrDA, etc.

SDIO 3.0 cards and hosts add support for UHS-I bus speed mode, which can be as fast as 104MB/s.

## 3.1.2 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.

For more information about the handling of kernel modules, see the SUSE Linux Enterprise Administration Guide.

## 3.1.3 Systems Management

### 3.1.3.1 New XFS On-disk Format

SUSE Linux Enterprise 12 supports the new on-disk format (v5) of the XFS file system. XFS file systems created by YaST will use this new format. The main advantages of this format are automatic checksumming of all XFS metadata, file type support, and support for a larger number of access control lists for a file.

Caveat: Pre SLE 12 kernels, xfsprogs before version 3.2.0, and the grub2 bootloader before the one released in SLE 12 do not understand the new file system format and thus refuse to work with it. This can be problematic if the file system should also be used from older or other distribution.

If you require interoperability of the XFS file system with older or other distributions, format the filesystem manually using the mkfs.xfs command. That will create a filesystem in the old format unless you use the "-m crc=1" option.

### 3.1.3.2 Systemd Daemon

SLE12 has moved to Systemd, a new way of managing services. For more information, see the *SUSE Linux Enterprise Admin Guide* , Section *The Systemd Daemon* .

## 3.1.4 Storage

### 3.1.4.1 /dev/disk/by-path/ Links for virtio Disks No Longer Available

Because virtio numbers are not stable, by-path links for virtio disks are no longer available. These names are not persistent.

### 3.1.4.2 Support for the Btrfs File System

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

#### *Support*

With SUSE Linux Enterprise 12, Btrfs is the default file system for the operating system, xfs is the default for all other use cases. We also continue to support the Ext-family of file systems, Reiserfs and ocfs2. Each file system offers distinct advantages. Customers are advised to use the YaST partitioner (or AutoYaST) to build their systems: YaST will prepare the Btrfs file system for use with subvolumes and snapshots. Snapshots will be automatically enabled for the root file system using SUSE's snapper infrastructure. For more information about snapper, its integration into ZYpp and YaST, and the YaST snapper module, see the SUSE Linux Enterprise documentation.

### *Migration from "Ext" and Reiserfs File Systems to Btrfs*

Migration from existing "Ext" file systems (Ext2, Ext3, ext4) and Reiserfs is supported "offline" and "in place", if the original filesystem has been created with a 4k block size (this is the case for most file systems on the x86-64 and System z architectures). Calling "btrfs-convert <device>" will convert the file system. This is an offline process, which needs at least 15% free space on the device, but is applied in place. Roll back: calling "btrfs-convert -r <device>" will roll back. **Caveat:** when rolling back, all data will be lost that has been added after the conversion into Btrfs; in other words: the roll back is complete, not partial.

### *RAID*

Btrfs is supported on top of MD (multiple devices) and DM (device mapper) configurations. Use the YaST partitioner to achieve a proper setup. Multivolume Btrfs is supported in RAID0, RAID1, and RAID10 profiles in SUSE Linux Enterprise 12, higher RAID levels are not yet supported, but might be enabled with a future service pack.

### *SWAP files*

Using swap files on top of Btrfs is not supported. In general, we are advising to use partitions for swapping, and not swap files on top of any file system for performance reasons.

### *Future Plans*

- Compression functionality for Btrfs is currently under development and will be supported once the development has matured.
- We are committed to actively work on the Btrfs file system with the community, and we keep customers and partners informed about progress and experience in terms of scalability and performance. This may also apply to cloud and cloud storage infrastructures.

### *Filesystem Maintenance, Online Check, and Repair Functionality*

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

We recommend to apply regular "maintenance" to the Btrfs file system to optimize performance and disk usage. Specifically we recommend to "balance" and "defrag" the file system on a regular basis. Check the "btrfs-maintenance" package and see the SUSE Linux Enterprise documentation for more information.

### *Capacity Planning*



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

#### *Backward compatibility - Hard Link Limitation*

Previous products had a limitation on low hard link count per file in a directory. This has been fixed and is 65535 now. It requires a file system created with "-O extref", which is done by default. **Caveat:** Such a file system might not be mountable on older products.

#### *Backward compatibility - Enhanced metadata*

The file systems are by default created with a more space efficient format of metadata, the feature is called "skinny-metadata" for mkfs. **Caveat:** Such a file system will not be mountable on previous products.

#### *Backward compatibility - metadata block size is 16k*

The default metadata block size has changed to 16 kilobytes, reducing metadata fragmentation. **Caveat:** Such a file system will not be mountable on older products.

#### *Other Limitations*

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

#### *For More Information*

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

### 3.1.4.3 Default File System

With SUSE Linux Enterprise 12, the default file system in new installations was changed from Ext3 to Btrfs for the root system partition. XFS is the default file system for the /home partition and other data partitions.

In the expert partitioner, the default file system is Btrfs. The user can change it if another file system is more suitable to accomplish the intended work load.

#### **POWER Architecture**

On POWER, the pagesize is 64K. Due to the assumption made by Btrfs regarding data blocksize (i.e. data blocksize being equal to the page size), a Btrfs installation on POWER will use a blocksize of 64K. This means that a Btrfs created on x86 will not be mountable and readable via Btrfs on POWER, and vice versa.

If data sharing in mixed architecture environments is a major concern, make sure to use XFS on POWER for data partitions.

## 3.1.5 Security

### 3.1.5.1 Installing CA Certificates

*For legacy reasons, `/etc/ssl/certs` may only contain CA certificates in PEM format. Because this format does not transport usage information `/etc/ssl/certs` may only contain CA certificates that are intended for server authentication.*

OpenSSL understands a different format that transports the usage information, therefore OpenSSL internally uses a different location, which contains certificates of all kinds of usage type ( `/var/lib/ca-certificates/openssl` ). If you put a certificate in plain PEM format in `/etc/pki/trust/anchors/` and call `update-ca-certificates` it should end up in both `/var/lib/ca-certificates/pem` (i.e., `/etc/ssl/certs` ) and `/var/lib/ca-certificates/openssl` [as well as other locations like the cert bundle or the Java keyring].

### 3.1.5.2 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 framebuffer drivers.

### 3.1.5.3 Linux Filesystem Capabilities

Our kernel is compiled with support for Linux Filesystem Capabilities. Since SLE 12, it is enabled by default.

Disable it by adding `file_caps=0` as a kernel boot option.

### 3.1.5.4 Increased dmesg Restrictions

*`dmesg` was providing all kinds of system internal information to any users. It includes kernel addresses, crashes of services, and similar things that could be used by local attackers.*

The use of `dmesg` is now restricted to the root user.

### 3.1.5.5 Restricting Access to Removable Media

Use `udisks2` to restrict access to removable media. For more information, see the *Security and Hardening Guide* .

## 3.1.6 Networking

### 3.1.6.1 systemd: Activating a network.service Implementation

By default, you use the YaST Network Settings dialog (yast2 network) to activate or deactivate NetworkManager. For manual configuration without YaST, proceed as follows.

In the past, the `NETWORKMANAGER` sysconfig variable in `/etc/sysconfig/network/config` was used to activate and deactivate NetworkManager. This variable is gone and replaced with a proper systemd `network.service` alias link, which points to the currently enabled network service.

The alias link will be created by the

```
systemctl enable NetworkManager.service
```

or

```
systemctl enable wicked.service
```

commands.

Further, the `/etc/init.d/network` script has been removed in favor of native systemd services. The `rcnetwork` shortcut executes action of `network.service`.

The command

```
systemctl -p Id show network.service
```

allows to query the currently selected network service, the

```
systemctl status network.service
```

shows the user readable details about currently used network service.

**Procedure to enable NetworkManager manually:**

1) First, stop the running network (wicked) service to get a clean state (configuration may differ):

```
systemctl is-active network.service && systemctl stop network.service
```

2) Then, stop the wicked-daemon services as well:

```
systemctl is-active wickedd.service && \
```

```
systemctl      stop      wickedd.service
```

3) Disable wicked, enable NetworkManager.service (creates alias link):

```
systemctl disable wicked.service  
systemctl --force enable NetworkManager.service
```

4) Start the NetworkManager service via the alias link:

```
systemctl      start      network.service
```

or directly:

```
systemctl start NetworkManager.service
```

**Procedure to disable NetworkManager and switch to wicked.service manually:**

1) Stop the running NetworkManager.service:

```
systemctl      is-active NetworkManager.service && \  
systemctl --kill-who=all kill NetworkManager.service
```

Note: The normal NetworkManager.service stop action stops NetworkManager, but leaves processes such as dhcp clients running to not break network connectivity when it is restarted on update or there is a remote fs mounted while shutdown. The --kill-who = all kill action ensures to stop them too as they conflict with the wicked service using a different implementation.

2) Disable NetworkManager, enable wicked.service (creates alias link):

```
systemctl disable NetworkManager.service  
systemctl --force enable wicked.service
```

3) Start the new network.service, which now is wicked.service:

```
systemctl start wicked.service
```

or via the alias link:

```
systemctl start network.service
```

The wickedd daemon service are started automatically via dependencies.

**To query the currently selected service, use:**

```
systemctl -p Id show network.service
```

It returns "Id=NetworkManager.service" if the NetworkManager service is enabled, otherwise "Id=network.service" and /etc/init.d/network is acting as the network service.

### 3.1.6.2 Remote Login with XDMCP

Depending on your XDMCP client, the following configurations are supported:

- If GLX is available from your X client (such as Xephyr), the default settings for the display manager (gdm) and for the window manager (GNOME3/sle-classic) should be used.
- If GLX is not available from your X client to connect to the XDMCP server (such as XNest), XDM as the display manager should be used ( DISPLAYMANAGER="xdm" in /etc/sysconfig/displaymanager ) and icewm should be used as the window manager ( DEFAULT\_WM="icewm" in /etc/sysconfig/windowmanager ).

If both Xephyr and Xnest are available as the X client, Xephyr is the preferred client to use.

### 3.1.6.3 How to enable the wicked "nanny" framework

*Within the wicked family of tools, the nanny daemon is a policy engine that is responsible for asynchronous or unsolicited scenarios such as hotplugging devices.*

The nanny framework is not enabled by default in SUSE Linux Enterprise 12. To enable it either temporarily specify "nanny = 1" on the boot prompt or activate it in /etc/wicked/common.xml :

```
<config>
...
<use-nanny>true</use-nanny>
</config>
```

After a change at runtime, restart the network:

```
systemctl restart wickedd.service
wicked ifup all
```

For more information, see the *SUSE Linux Enterprise Admin Guide* , Section *The wicked Network Configuration* .

### 3.1.6.4 Passing Options to `/etc/resolv.conf`

With `NETCONFIG_DNS_RESOLVER_OPTIONS` in `/etc/sysconfig/network/config` you can specify arbitrary options that netconfig will write to `/etc/resolv.conf` .

For more information about available options, see the `resolv.conf` man page.

## 3.1.7 Performance

### 3.1.7.1 Enabling VEBOX on Haswell in the `drm/i915` Kernel Driver

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

## 3.1.8 Virtualization

### 3.1.8.1 Others

#### 3.1.8.1.1 `open-vm-tools` Now Included

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

SUSE Linux Enterprise 12 includes the `open-vm-tools` package. These tools are pre-selected when installing on a VMware platform.

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

## 3.2 AMD64/Intel64 64-Bit (x86\_64) Specific Information

### 3.2.1 Trackpoint or Pointing Stick Configuration

*In the past, the default settings of trackpoint or pointing stick devices were different on various machines, and thus the behavior of these devices was not consistent.*

These days people prefer to use the combination of trackpoint or pointing stick and middle button for scrolling. This means pressing the middle button while moving the trackpoint or pointing stick emulates a mouse wheel.

To make it work reliably, the following options are set by default:

EmulateWheel	= on
EmulateWheelButton	= 2
Emulate3Buttons	= on

Commenting these three options with the '#' character at the beginning of the lines in `/etc/X11/xorg.conf.d/11-evdev.conf` will restore the upstream defaults to have a real middle button and the scrollwheel emulation disabled again.

### 3.2.2 System and Vendor Specific Information

#### 3.2.2.1 Installation on Native 4KiB Sector Drives (4kn) Supported with UEFI

*For the last 20 years, hard disk with 512 byte sectors have been in use. Since some years there are drives providing a 4KiB sector size internally, but showing 512 byte sectors externally as a backward compatibility layer (512 byte emulation / 512e). These devices are fully supported in SUSE Linux Enterprise.*

The installation on native 4KiB sector drives (4kn) in x86\_64 systems with UEFI is supported, as is the use of 4 KiB sector drives as non-boot disks. Legacy (non UEFI) installations on x86\_64 systems are not supported on 4KiB drives for technical reasons.

## 4 Driver Updates

### 4.1 Storage Drivers


#### 4.1.1 Driver for IMSM and DDF

For IMSM and DDF RAIDs the mdadm driver is used unconditionally.

### 4.2 Network Drivers

#### 4.2.1 Myricom 10-Gigabit Ethernet Driver and Firmware

SUSE Linux Enterprise 12 (x86\_64) 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 12.

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

## 5 Packages and Functionality Changes


### 5.1 New Packages

#### 5.1.1 New Package: Scribus

Scribus is a powerful desktop publishing software that helps with creating documents of all kinds. Scribus is now available on SLED 12.

### 5.2 Updated Packages

#### 5.2.1 Samba: Changing "winbind expand groups" to "0"

*Forthcoming Samba 4.2.0 provided by <http://www.samba.org> (<http://www.samba.org>)  will come with "winbind expand groups" set to "0" by default.*



Samba post 4.1.10 provided by SUSE anticipates the new default.

The new default makes winbindd more reliable because it does not require SAMR access to domain controllers of trusted domains.

Note: Some legacy applications calculate the group memberships of users by traversing groups; such applications will require winbind expand groups = 1 .

## 5.2.2 GNOME 3.10

We ship GNOME 3.10 with SUSE Linux Enterprise 12.

GNOME on SUSE Linux Enterprise is available in three different setups, which are modifying desktop user experience:

- SLE Classic: this setup uses a single bottom panel, similar to GNOME desktop as available on SUSE Linux Enterprise 11. This setup is default on SUSE Linux Enterprise 12.
- GNOME: this is GNOME 3 upstream user experience, also sometime called "GNOME Shell". This setup might be more adequate with touchscreen.
- GNOME Classic: this setup uses two panels (one top panel, one bottom panel) similar to upstream GNOME 2 desktop

The setup can be changed at login time, in GDM, using the gear icon in the password prompt screen. It can also be modified using YaST, systemwide.

### **Caveats:**

With SLE 11 after joining a Microsoft domain, GDM displayed the available domain names as a drop-down box below the user name and password fields. This behavior has changed.

With SLE 12, you must prefix the domain and the winbind separator manually to login. As soon as you click the 'Not listed?' text, GDM will display a hint such as '(e.g., domain\user)'.

## 5.2.3 Support for Qt5

*We received requests to support QML as part of the Qt framework.*

While Qt4 (minimum ver 4.8.2-260.1) would have been possible to use, directly upgrading to and supporting Qt5 (QML supported) is the better and more future proof solution.

## 5.2.4 Bluetooth Implementation BlueZ 5

*BlueZ 4 is no longer maintained upstream. Thus upgrading to BlueZ 5 ensures that you will get all the latest upstream bug fixes and enhancements.*

BlueZ 5 comes with numerous new features, API simplification and other improvements such as Low Energy support. It is new major version of the Bluetooth handling daemon and utilities.

Note: The new major version indicates that the API is not backwards compatible with BlueZ 4, which means that all applications, agents, etc. must be updated.

## 5.2.5 MOK List Manipulation Tools

A Machine Owner Key (MOK) is a type of key that a user generates and uses to sign an EFI binary. This is a way for the machine owner to have ownership over the platform's boot process.

Suitable tools are coming with the `mokutil` package.

## 5.2.6 Kernel and Toolchain

- GCC 4.8
- glibc 2.19
- Linux kernel 3.12

## 5.2.7 Desktop

- GNOME 3.10
- X.org 7.7

## 5.2.8 Other Changes and Version Updates

- Samba 4.1.3
- UEFI Enablement on AMD64
- SWAP over NFS

- Python 2.7
- Perl 5.18.2
- Ruby 2.0

## 5.3 Deprecated Functionality

### 5.3.1 PCMCIA is deprecated

The old PCMCIA based on ISA and 16-bit only will no more be supported under SLE12. Latest modern laptop uses CardBus (based on PCI), which continues to be supported.

### 5.3.2 Command Line Interface for Managing Packages

YaST as a command line tool for managing packages is deprecated. Instead of yast with the command line switches -i , --install , --update , or --remove for installing, updating, or removing packages, use zypper .

For more information, see the zypper man page.

### 5.3.3 libsysfs obsoleted by libudev

libsysfs has been deprecated and has been replaced by libudev. If you have self-compiled applications using libsysfs previously, you have to recompile using libudev .

### 5.3.4 dhcpd Replaced by wicked and dhcp-client

dhcpd package was replaced by wicked and dhcp-client packages.

### 5.3.5 Raw Devices Are Deprecated

Raw devices are deprecated.

### 5.3.6 Packages Removed with SUSE Linux Enterprise Desktop 12

The following packages were removed with the major release of SUSE Linux Enterprise Desktop 12:

#### 5.3.6.1 Libreoffice Language Tools Removed

Libreoffice language tools, which is a collection of grammar and common errors for a number of languages, is no longer provided as part of SLED. Those tools are still available from Libreoffice.org Web site, as extensions. Spellcheckers for a number of languages are still part of SLED.

#### 5.3.6.2 `scsirastools` is deprecated

`scsirastools` was designed to work with now obsolete SCSI parallel enclosure. This package is not more available in SLE12.

#### 5.3.6.3 Adobe Discontinues Support for Adobe Reader on Linux

*Adobe has discontinued support for Adobe Reader 9 on Linux ([http://www.adobe.com/support/products/enterprise/eol/eol\\_matrix.html#863](http://www.adobe.com/support/products/enterprise/eol/eol_matrix.html#863)) and is no longer providing security updates.*

In order to not lose functionality Adobe Acrobat Reader will be kept on released products, but to avoid security issues with accessing PDFs online the PDF viewer browser plugin will however be removed. In order to maintain functionality the latest Firefox ESR releases include a feature to display PDF documents, which receives maintenance and security updates via Firefox updates.

#### 5.3.6.4 LPRng Discontinued

As announced on SLE 11, `LPRng` is discontinued with SLE 12.

### 5.3.6.5 The Number of Kernel Modules in the kernel-extra Package Reduced

The following unsupported kernel modules have been dropped from the `kernel-extra` package:

- Staging drivers
- IDE drivers on POWER
- Open Sound System on x86\_64
- WAN drivers on x86\_64
- 1-Wire drivers
- File systems: adfs, affs, befs, bfs, efs, freevxfs, hpfs, qnx4, jffs2, jfs, logfs, nilfs2, ubifs

### 5.3.6.6 Unsupported Graphical Chipsets

The following X11 drivers are no longer provided in SLE 12:

- xf86-video-ark
- xf86-video-chips
- xf86-video-geode
- xf86-video-glint
- xf86-video-i128
- xf86-video-neomagic
- xf86-video-newport
- xf86-video-r128
- xf86-video-savage
- xf86-video-siliconmotion
- xf86-video-tdfx
- xf86-video-tga

- xf86-video-trident
- xf86-video-voidoo
- xf86-video-sis
- xf86-video-sisusb
- xf86-video-openchrome
- xf86-video-unichrome
- xf86-video-mach64

#### 5.3.6.7 **suseRegister replaced by SUSEConnect**

s useRegister was replaced by SUSEConnect .

#### 5.3.6.8 **Mono Platform and Programs No Longer Provided**

*Starting with SLE 12, the Mono platform and Mono based programs are no longer supported.*

These are the replacement applications:

- gnote (instead of Tomboy)
- shotwell (instead of F-Spot)
- rhythmbox (instead of Banshee)

#### 5.3.6.9 **YaST No Longer Supports Configuring Modem Devices**

YaST ( yast2-ntework ) no longer offers modem configuration dialogs.

It is still possible to configure modems manually.

#### 5.3.6.10 **YaST No Longer Supports Configuring ISDN Devices**

YaST ( yast2-ntework ) no longer supports configuring ISDN devices. If needed, NetworkManager supports such devices.

### 5.3.6.11 YaST No Longer Supports Configuring DSL Devices

YaST ( [yast2-network](#) ) no longer supports configuring DSL devices. If needed, NetworkManager supports such devices (e.g., DSL cable modems).

## 5.3.7 Packages and Features to Be Removed in the Future

The following packages are deprecated and will be removed with SUSE Linux Enterprise Desktop 13:

- [...](#)

### 5.3.7.1 Support for Qt4

SLE 12 features the Qt4 toolkit. Qt4 will be supported at least until the release of SLE 12 Service Pack 3. Hence it is recommended to migrate applications to Qt5 and start new projects using Qt5.

### 5.3.7.2 Use /etc/os-release Instead of /etc/SuSE-release

*Starting with SLE 12, /etc/SuSE-release file is deprecated. It should not be used to identify a SUSE Linux Enterprise system. This file will be removed in a future Service Pack or release.*

The file [/etc/os-release](#) now is decisive. This file is a cross-distribution standard to identify a Linux system. For more information about the syntax, see the os-release man page ( [man os-release](#) ).

## 5.4 Changes in Packaging and Delivery

### 5.4.1 module-init-tools Replaced by kmod

module-init-tools is replaced by kmod.

Caveat: With the replacement, the `modprobe list` command ( [-l](#) ) is no longer available. As a workaround you can make use of [find](#) or [grep](#) ; for example, if you are looking for modules starting with [xt](#) :

```
grep '/xt[^\]*\.ko:' /lib/modules/$(uname -r)/modules.dep
```

## 5.4.2 AppArmor: Normalized Command Names

AppArmor now offers normalized command names:

- aa-notify instead of aa-apparmor\_notify or apparmor\_notify
- aa-status instead of aa-apparmor\_status ( apparmor\_status is still supported)

## 5.4.3 Legacy module-init-tools Replaced with kmod

Kmod package is a replacement of the former module-init-tools . In addition to the well known tools like lsmod , modprobe , and modinfo , the package offers a shared library for use by system management services which need to query and manipulate Linux kernel modules.


## 5.4.4 Replacing syslog-ng and syslog With rsyslog

On new installations, rsyslog will get installed instead of the former syslog-ng and syslog .

## 5.4.5 Printing System: Improvements and Incompatible Changes

### CUPS Version Upgrade to 1.7

CUPS >= 1.6 has major incompatible changes compared to CUPS up to version 1.5.4 in particular when printing via network:

The IPP protocol default version increased from 1.1 to 2.0. Older IPP servers like CUPS 1.3.x (e.g. in SLE11) reject IPP 2.0 requests with "Bad Request" (see <http://www.cups.org/str.php?L4231>  ). By adding '/version=1.1' to ServerName in client.conf (e.g., ServerName older.server.example.com/version=1.1) or to the CUPS\_SERVER environment variable value or by adding it to the server name value of the '-h' option (e.g., lpstat -h older.server.example.com/version=1.1 -p) the older IPP protocol version for older servers must be specified explicitly.

CUPS Browsing is dropped in CUPS but the new package cups-filters provides the cups-browsed that provides basic CUPS Browsing and Polling functionality. The native protocol in CUPS for automatic client discovery of printers is now DNS-SD. Start cups-browsed on the local host to receive traditional CUPS Browsing information from traditional remote CUPS servers. To broadcast traditional CUPS Browsing information into the network so that traditional remote CUPS clients can receive it, set "BrowseLocalProtocols CUPS" in /etc/cups/cups-browsed.conf and start cups-browsed.



Some printing filters and back-ends are dropped in CUPS but the new package cups-filters provides them. So cups-filters is usually needed (recommended by RPM) but cups-filters is not strictly required.

The cupsd configuration directives are split into two files: cupsd.conf (can also be modified via HTTP PUT e.g. via cupsctl) and cups-files.conf (can only be modified manually by root) to have better default protection against misuse of privileges by normal users who have been specifically allowed by root to do cupsd configuration changes (see <http://www.cups.org/str.php?L4223> , CVE-2012-5519, and SUSE Bugzilla bnc#789566).

CUPS banners and the CUPS test page are no longer supported since CUPS  $\geq$  1.6. The banners and the test page from cups-filters must be used. The CUPS banner files in `/usr/share/cups/banners/` and the CUPS testpage `/usr/share/cups/data/testprint` (which is also a CUPS banner file type) are no longer provided in the cups RPM because they do no longer work since CUPS  $\geq$  1.6 (see <http://www.cups.org/str.php?L4120>) because there is no longer a filter that can convert the CUPS banner files. Since CUPS  $\geq$  1.6 only the banner files and testpage in the cups-filters package work via the cups-filters PDF workflow and the cups-filters package also provides the matching bannertopdf filter.

For details, see the SUSE Bugzilla bnc#735404 issue.

### **Traditional CUPS version 1.5.4 Provided in the Legacy Module**

We provide the last traditional CUPS version 1.5.4 as "cups154" RPMs in the "legacy" module. If CUPS version 1.7 does not support particular needs, you can still use CUPS 1.5.4 (under the conditions of the "legacy" module). This could be important, if you need a traditional CUPS server with original CUPS Browsing features.

For those users any (semi)-automated CUPS version upgrade must be prohibited because CUPS  $>$  1.5.4 has major incompatible changes compared to CUPS  $\leq$  1.5.4. Therefore the CUPS 1.5.4 RPM package name contains the version and it conflicts with higher versions. This way we avoid that an installed CUPS 1.5.4 gets accidentally replaced with a higher version. It is not possible to have different CUPS libraries versions installed at the same time.

The API in CUPS 1.7 is compatible with the CUPS 1.5.4 API (existing functions are not changed) but newer CUPS libraries provide some new functions. There could be applications that might use newer CUPS library functions so that such applications would require the current CUPS 1.7 libraries. It is not possible to use CUPS 1.5.4 together with applications that require the current CUPS 1.7 libraries.

### **PDF Now Common Printing Data Format**


There is a general move away from PostScript to PDF as the standard print job format. This change is advocated by the OpenPrinting workgroup of the Linux Foundation and the CUPS author.

This means that application programs usually no longer produce PostScript output by default when printing but instead PDF.

As a consequence the default processing how application programs printing output is converted into the "language" that the particular printer accepts (the so called "CUPS filter chain") has fundamentally changed from a PostScript-centric workflow to a PDF-centric workflow.

Accordingly the upstream standard for CUPS under Linux (using CUPS plus the cups-filters package) is now PDF-based job processing, letting every non-PDF input be converted to PDF first, page management options being applied by a pdftopdf filter and Ghostscript being called with PDF as input.

With PDF as the standard print job format traditional PostScript printers can no longer print application's printing output directly so that a conversion step in the printing workflow is required that converts PDF into PostScript. But there are also PostScript + PDF printers that can print both PostScript and PDF directly.

For details, see the section "Common printing data formats" in the SUSE wiki article "Concepts printing" at [http://en.opensuse.org/Concepts\\_printing](http://en.opensuse.org/Concepts_printing) .

## 6 Technical Information

This section contains a number of technical changes and enhancements for the experienced user.

### 6.1 Kernel Limits

### 6.1.1 The Number of Kernel Modules in the kernel-extra Package Reduced

The following unsupported kernel modules have been dropped from the `kernel-extra` package:

- Staging drivers
- IDE drivers on POWER
- Open Sound System on x86\_64
- WAN drivers on x86\_64
- 1-Wire drivers
- File systems: adfs, affs, befs, bfs, efs, freevxfs, hpfs, qnx4, jffs2, jfs, logfs, nilfs2, ubifs

## 6.2 File Systems

### 6.2.1 File System Layout

For general information about the file system layout, see the *Administration Guide*, Chapter *Snapshots*.


#### **Additional Information**


`/run/media/<user_name>` is now used as top directory for removable media mount points. It replaces `/media`, which is no longer available.



## 7 Legal Notices

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

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

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

Copyright © 2010, 2011, 2012, 2013, 2014 SUSE LLC. This release notes document is licensed under a Creative Commons Attribution-NoDerivs 3.0 United States License (CC-BY-ND-3.0 US, <http://creativecommons.org/licenses/by-nd/3.0/us/> .

SUSE has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.suse.com/company/legal/>  and one or more additional patents or pending patent applications in the U.S. and other countries. For SUSE trademarks, see SUSE Trademark and Service Mark list (<http://www.suse.com/company/legal/> ). All third-party trademarks are the property of their respective owners.