



SUSE Linux Enterprise Server 15 SP4

Release Notes

Abstract

SUSE Linux Enterprise Server is a modern, modular operating system for both multimodal and traditional IT. This document provides a high-level overview of features, capabilities, and limitations of SUSE Linux Enterprise Server 15 SP4 and highlights important product updates.

These release notes are updated periodically. The latest version of these release notes is always available at <https://www.suse.com/releasesnotes>. General documentation can be found at <https://documentation.suse.com/sles/15-SP4>.

Publication Date: 2024-01-19, **Version:** 15.4.20240119

1 About the release notes

2 SUSE Linux Enterprise Server

2.1 Interoperability and hardware support

2.2 What is new?

2.3 Important sections of this document

2.4 Security, standards, and certification

2.5 Documentation and other information

2.6 Support and life cycle

2.7 Support statement for SUSE Linux Enterprise Server

2.8 Technology previews

3 Modules, extensions, and related products

3.1 Modules in the SLE 15 SP4 product line

3.2 SLE extensions

3.3 Derived and related products

4 Installation and upgrade

4.1 Installation

4.2 Upgrade-related notes

4.3 Minimal-VM and Minimal-Image

4.4 JeOS renamed Minimal-VM and Minimal-Image

4.5 Minimal-VM VMware flavor with cloud-init support

4.6 Alternative Python 3 development interpreter moved to a separate module

4.7 For more information

5 Changes affecting all architectures

5.1 Authentication

5.2 Basic utilities

5.3 Containers

5.4 Databases

5.5 Desktop

5.6 Development

5.7 Hardware

5.8 Kernel

5.9 Miscellaneous

5.10 Networking

5.11 Security

5.12 Storage and file systems

5.13 SUSE Package Hub

5.14 System management

5.15 Virtualization

6 AMD64/Intel 64-specific changes (x86-64)

6.1 System-specific and vendor-specific information

7 POWER-specific changes (ppc64le)

7.1 Hardware

7.2 Performance

7.3 Security

7.4 Virtualization

7.5 Miscellaneous

8 IBM Z-specific changes (s390x)

8.1 Hardware

8.2 Networking

8.3 Performance

8.4 Security

8.5 Storage

8.6 Virtualization

8.7 Miscellaneous

9 Arm 64-bit-specific changes (AArch64)

9.1 System-on-Chip driver enablement

9.2 New features

9.3 Known limitations

9.4 Removal of NXP Layerscape LX2160A rev. 1 silicon support

10 Removed and deprecated features and packages

10.1 Removed features and packages

10.2 Deprecated features and packages

11 Obtaining source code

12 Legal notices

A Changelog for 15 SP4

A.1 2024-01-19

A.1.1 New

A.2 2023-05-22

A.2.1 Updated

A.3 2023-05-11

A.3.1 New

A.4 2023-05-10

A.4.1 New

A.5 [2023-03-01](#)

A.5.1 [New](#)

A.5.2 [Updated](#)

A.6 [2022-11-30](#)

A.6.1 [New](#)

A.7 [2022-09-30](#)

A.7.1 [New](#)

A.7.2 [Changed](#)

A.8 [2022-08-31](#)

A.8.1 [New](#)

A.8.2 [Changed](#)

A.9 [2022-07-14](#)

A.9.1 [New](#)

A.9.2 [Updated](#)

A.9.3 [Removed](#)

A.10 [2022-05-11](#)

A.10.1 [New](#)

A.10.2 [Updated](#)

A.11 [2022-04-20](#)

A.11.1 [New](#)

A.12 [2022-03-23](#)

A.12.1 [New](#)

A.12.2 [Updated](#)

A.12.3 [Removed](#)

A.13 [2022-02-16](#)

A.14 [2022-01-19](#)

A.15 [2021-12-08](#)


A.16 [2021-11-17](#)

A.17 [2021-11-03](#)

B [Kernel parameter changes](#)

B.1 [Changes from SP3 to SP4](#)

1 About the release notes

These Release Notes are identical across all architectures, and the most recent version is always available online at <https://www.suse.com/releasesnotes> .

Entries are only listed once but they can be referenced in several places if they are important and belong to more than one section.

Release notes usually only list changes that happened between two subsequent releases. Certain important entries from the release notes of previous product versions are repeated. To make these entries easier to identify, they contain a note to that effect.

However, repeated entries are provided as a courtesy only. Therefore, if you are skipping one or more service packs, check the release notes of the skipped service packs as well. If you are only reading the release notes of the current release, you could miss important changes.

2 SUSE Linux Enterprise Server

SUSE Linux Enterprise Server 15 SP4 is a multimodal operating system that paves the way for IT transformation in the software-defined era. It is a modern and modular OS that helps simplify multimodal IT, makes traditional IT infrastructure efficient and provides an engaging platform for developers. As a result, you can easily deploy and transition business-critical workloads across on-premises and public cloud environments.

SUSE Linux Enterprise Server 15 SP4, with its multimodal design, helps organizations transform their IT landscape by bridging traditional and software-defined infrastructure.

2.1 Interoperability and hardware support

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

This modular, general-purpose operating system runs on four 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-performance guest on leading hypervisors. This makes SUSE Linux Enterprise Server the perfect guest operating system for virtual computing.

2.2 What is new?

2.2.1 General changes in SLE 15

SUSE Linux Enterprise Server 15 introduces many innovative changes compared to SUSE Linux Enterprise Server 12. The most important changes are listed below.

Migration from openSUSE Leap to SUSE Linux Enterprise Server

SLE 15 SP2 and later support migrating from openSUSE Leap 15 to SUSE Linux Enterprise Server 15. Even if you decide to start out with the free community distribution, you can later easily upgrade to a distribution with enterprise-class support. For more information, see the *Upgrade Guide* at <https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-upgrade-online.html#sec-upgrade-online-opensuse-to-sle> ↗.

Extended package search

Use the new Zypper command `zypper search-packages` to search across all SUSE repositories available for your product, even if they are not yet enabled. For more information see [Section 5.14.13, “Searching packages across all SLE modules”](#).

Software Development Kit

In SLE 15, packages formerly shipped as part of the Software Development Kit are now integrated into the products. Development packages are packaged alongside other packages. In addition, the *Development Tools* module contains tools for development.

RMT replaces SMT

SMT (Subscription Management Tool) has been removed. Instead, RMT (Repository Mirroring Tool) now allows mirroring SUSE repositories and custom repositories. You can then register systems directly with RMT. In environments with tightened security, RMT can also proxy other RMT servers. If you are planning to migrate SLE 12 clients to version 15, RMT is the supported product to handle such migrations. If you still need to use SMT for these migrations, beware that the migrated clients will have *all* installation modules enabled. For more information see [Section 4.2.5, "SMT has been replaced by RMT"](#).

Media changes

The *Unified Installer* and *Packages* media known from SUSE Linux Enterprise Server 15 SP1 have been replaced by the following media:

- **Online Installation Medium:** Allows installing all SUSE Linux Enterprise 15 products. Packages are fetched from online repositories. This type of installation requires a registration key. Available SLE modules are listed in [Section 3.1, "Modules in the SLE 15 SP4 product line"](#).
- **Full Installation Medium:** Allows installing all SUSE Linux Enterprise Server 15 products without a network connection. This medium contains all packages from all SLE modules. SLE modules need to be enabled manually during installation. RMT (Repository Mirroring Tool) and SUSE Manager provide additional options for disconnected or managed installations.

MAJOR UPDATES TO THE SOFTWARE SELECTION:

Salt

SLE 15 SP4 can be managed via Salt, making it integrate better with modern management solutions such as SUSE Manager.

Python 3

As the first enterprise distribution, SLE 15 offers full support for Python 3 development in addition to Python 2.

Directory Server

389 Directory Server replaces OpenLDAP as the LDAP directory service.

2.2.2 Changes in 15 SP4

SUSE Linux Enterprise Server 15 SP4 introduces changes compared to SUSE Linux Enterprise Server 15 SP3. The most important changes are listed below:

2.2.3 Package and module changes in 15 SP4

The full list of changed packages compared to 15 SP3 can be seen at this URL:

- https://documentation.suse.com/package-lists/sle/15-SP4/package-changes_SLE-15-SP3-GA_SLE-15-SP4-GA.txt ↗

The full list of changed modules compared to 15 SP3 can be seen at this URL:

- https://documentation.suse.com/package-lists/sle/15-SP4/module-changes_SLE-15-SP3-GA_SLE-15-SP4-GA.txt ↗

2.3 Important sections of this document

If you are upgrading from a previous SUSE Linux Enterprise Server release, you should review at least the following sections:

- Section 2.7, “Support statement for SUSE Linux Enterprise Server”
- Section 4.2, “Upgrade-related notes”
- Section 5, “Changes affecting all architectures”

2.4 Security, standards, and certification

SUSE Linux Enterprise Server 15 SP4 has been submitted to the certification bodies for:

- Common Criteria Certification, see <https://www.commoncriteriaportal.org/> ↗
- NIST FIPS 140-3 validation, see <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.140-3.pdf> ↗

The status of SUSE submitted modules can be reviewed at:

<https://csrc.nist.gov/Projects/Cryptographic-Module-Validation-Program/Modules-In-Process/Modules-In-Process-List> ↗

For more information about certification, see

<https://www.suse.com/support/security/certifications/> ↗.

2.5 Documentation and other information

2.5.1 Available on the product media

- Read the READMEs on the media.
- Get the detailed change log information about a particular package from the RPM (where FILENAME.rpm is the name of the RPM):

```
rpm --changelog -qp FILENAME.rpm
```

- Check the ChangeLog file in the top level of the installation medium for a chronological log of all changes made to the updated packages.
- Find more information in the docu directory of the installation medium of SUSE Linux Enterprise Server 15 SP4. This directory includes PDF versions of the SUSE Linux Enterprise Server 15 SP4 Installation Quick Start Guide.
- Get list of manual pages with usage information about a particular package from the RPM (where FILENAME.rpm is the name of the RPM):

```
rpm --docfiles -qp FILENAME.rpm | grep man
```

2.5.2 Online documentation

- For the most up-to-date version of the documentation for SUSE Linux Enterprise Server 15 SP4, see <https://documentation.suse.com/sles/15-SP4> ↗.

2.6 Support and life cycle

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.

SUSE Linux Enterprise Server 15 has a 13-year life cycle, with 10 years of General Support and three years of Extended Support. The current version (SP4) will be fully maintained and supported until six months after the release of SUSE Linux Enterprise Server 15 SP5.

If you need additional time to design, validate and test your upgrade plans, Long Term Service Pack Support can extend the support duration. You can buy an additional 12 to 36 months in twelve month increments. This means that you receive a total of three to five years of support per Service Pack.

For more information, see the pages [Support Policy](#) and [Long Term Service Pack Support](#).

2.7 Support statement for SUSE Linux Enterprise Server

To receive support, you need an appropriate subscription with SUSE. For more information, see https://www.suse.com/support/?id=SUSE_Linux_Enterprise_Server.

The following definitions apply:

L1

Problem determination, which means technical support designed to provide compatibility information, usage support, ongoing maintenance, information gathering, and basic troubleshooting using the documentation.

L2

Problem isolation, which means technical support designed to analyze data, reproduce customer problems, isolate the problem area, and provide a resolution for problems not resolved by Level 1 or 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 is delivered with L3 support for all packages, except for the following:

- Technology Previews, see [Section 2.8, “Technology previews”](#)
- Sound, graphics, fonts and artwork
- Packages that require an additional customer contract, see [Section 2.7.2, “Software requiring specific contracts”](#)
- Some packages shipped as part of the module *Workstation Extension* are L2-supported only

- Packages with names ending in `-devel` (containing header files and similar developer resources) will only be supported together with their main packages.

SUSE will only support the usage of original packages. That is, packages that are unchanged and not recompiled.

2.7.1 General support

To learn about supported features and limitations, refer to the following sections in this document:

- [Section 5.8, “Kernel”](#)
- [Section 5.12, “Storage and file systems”](#)
- [Section 5.15, “Virtualization”](#)
- [Section 10, “Removed and deprecated features and packages”](#)

2.7.2 Software requiring specific contracts

Certain software delivered as part of SUSE Linux Enterprise Server may require an external contract. Check the support status of individual packages using the RPM metadata that can be viewed with `rpm`, `zypper`, or YaST.

Major packages and groups of packages affected by this are:

- PostgreSQL (all versions, including all subpackages)

2.7.3 Software under GNU AGPL

SUSE Linux Enterprise Server 15 SP4 (and the SUSE Linux Enterprise modules) includes the following software that is shipped *only* under a GNU AGPL software license:

- Ghostscript (including subpackages)

SUSE Linux Enterprise Server 15 SP4 (and the SUSE Linux Enterprise modules) includes the following software that is shipped under multiple licenses that include a GNU AGPL software license:

- MySpell dictionaries and LightProof

- ArgyllCMS

2.8 Technology previews

Technology previews are packages, stacks, or features delivered by SUSE to provide glimpses into upcoming innovations. Technology previews are included for your convenience to give you a chance to test new technologies within your environment. We would appreciate your feedback! If you test a technology preview, contact your SUSE representative and let them know about your experience and use cases. Your input is helpful for future development.

Technology previews come with the following limitations:

- Technology previews are still in development. Therefore, they may be functionally incomplete, unstable, or in other ways not suitable for production use.
- Technology previews are **not** supported.
- Technology previews may only be available for specific hardware architectures. Details and functionality of technology previews are subject to change. As a result, upgrading to subsequent releases of a technology preview may be impossible and require a fresh installation.
- Technology previews can be removed from a product at any time. This may be the case, for example, if SUSE discovers that a preview does not meet the customer or market needs, or does not comply with enterprise standards.

2.8.1 Technology previews for all architectures

2.8.1.1 schedutil

schedutil is a CPU frequency scaling governor that makes decisions based on the utilization data provided by the scheduler, as opposed to other governors that use CPU idle time, such as ondemand. It was introduced in the Linux kernel version 4.7. However, it is only viable for production use together with an optimization called util_est (short for "utilization estimation") that makes it much more responsive. This optimization is only available in Linux kernel version 4.17 and newer. For this reason it is only offered as technology preview in SLE 15 SP4.

2.8.1.2 Redfish-finder functionality in wicked

The new version of wicked in SLES 15 SP4 has added initial support to decode the SMBIOS Management Controller Host Interface (Type 42) structure. It will expose it as wicked firmware:redfish configuration to setup a Host Network Interface (to the BMC) using the Redfish over IP protocol. This allows access to the Redfish Service (via redfish-localhost in /etc/hosts) used to manage the computer system.

This functionality has been added as a technical preview.

2.8.1.3 Support for Intel's Alderlake graphics platform

SLES 15 SP4 adds support for Intel's Alderlake graphics platform as technology preview. You can enable it by adding the i915.force_probe=<Device-ID> parameter to your kernel options in GRUB configuration. In order to figure out the <Device ID> of the Intel graphics adapter, use the inxi -aG command.

The output should look like this:

```
~> inxi -aG
Graphics:
  Device-1: Intel (R) Graphics vendor: Lenovo driver: i915 v: kernel
  bus ID: 00:02.0 chip ID: 8086:46a6
```

So in this case, use i915.force_probe=46a6 as the kernel option. The command to add the option to the bootloader configuration would then be:

```
pbl --add-option 'i915.force_probe=46a6' --config
```

2.8.1.4 zypper single transaction mode

Traditionally, zypper executes the rpm command separately for each operation in a transaction. This is among other things a lot slower for a large number of packages. Therefore we have implemented a new backend that runs all the operations in a single transaction using librpm.

This feature can be enabled by setting the environmental variable ZYPP_SINGLE_RPMTRANS to 1. Because this feature is offered as a technology preview, enabling it system-wide is known to have issues, thus we recommend enabling this feature per command, for example:

```
env ZYPP_SINGLE_RPMTRANS=1 zypper dup
```

However, it is possible to enable it permanently by editing the `/etc/zypp/zypp.conf` file by adding the following text:

```
[main]
techpreview.ZYPP_SINGLE_RPMTRANS=1
```

2.8.1.5 `zypper` parallel downloads

There is new `zypper` HTTP backend that enables a new asynchronous downloader. While it currently will not have massive impact on performance due to the frontend code not requesting files asynchronously, it will do some additional mirror rating and as soon as we update the frontend code will bring more benefits.

This feature can be enabled by setting the environmental variable `ZYPP_MEDIANETWORK` to `1`. Because this feature is offered as a technology preview, enabling it system-wide is known to have issues, thus we recommend enabling this feature per command, for example:

```
env ZYPP_MEDIANETWORK=1 zypper dup
```

However, it is possible to enable it permanently by editing the `/etc/zypp/zypp.conf` file by adding the following text:

```
[main]
techpreview.ZYPP_MEDIANETWORK=1
```

2.8.2 Technology previews for Arm 64-Bit (AArch64)

2.8.2.1 64K page size kernel flavor is available

SUSE Linux Enterprise Server for Arm 12 SP2 and later kernels have used a page size of 4K. This offers the widest compatibility also for small systems with little RAM, allowing to use Transparent Huge Pages (THP) where large pages make sense.

As a technology preview, SUSE Linux Enterprise Server for Arm 15 SP3 added a kernel flavor `64kb`, offering a page size of 64 KiB and physical/virtual address size of 52 bits. Same as the `default` kernel flavor, it does not use preemption.

Main purpose at this time is to allow for side-by-side benchmarking for High Performance Computing, Machine Learning and other Big Data use cases. Contact your SUSE representative if you notice performance gains for your specific workloads.

**Note: Default file system no longer needs to be changed**

SUSE Linux Enterprise Server for Arm 15 SP4 newly allows the use of Btrfs based file systems with 4 KiB block size also with 64 KiB page size kernels.

See [Section 5.8.12, “Btrfs sub-page block size support”](#) for details and known limitations.

**Important: Swap needs to be re-initialized**

After booting the 64K kernel, any swap partitions need to be re-initialized to be usable. To do this, run the `swapon` command with the `--fixpgsz` parameter on the swap partition. Note that this process deletes data present in the swap partition (for example, suspend data). In this example, the swap partition is on `/dev/sdc1`:

```
swapon --fixpgsz /dev/sdc1
```

**Warning: RAID 5 uses page size as stripe size**

It is currently possible to configure stripe size by setting the following kernel parameter:

```
echo 16384 > /sys/block/md1/md/stripe_size
```

Keep in mind that `stripe_size` must be in multiples of 4KB and not bigger than `PAGE_SIZE`. Also, it is only supported on systems where `PAGE_SIZE` is not 4096, such as arm64.

Avoid RAID 5 volumes when benchmarking 64K vs. 4K page size kernels.

See the *Storage Guide* for more information on software RAID.

**Note: Cross-architecture compatibility considerations**

The SUSE Linux Enterprise Server 15 SP4 kernels on x86-64 use 4K page size.

The SUSE Linux Enterprise Server for POWER 15 SP4 kernel uses 64K page size.

2.8.2.2 Driver enablement for NVIDIA BlueField-2 DPU as host platform

SUSE Linux Enterprise Server for Arm 15 SP1 and later kernels include drivers for installing on NVIDIA* BlueField* Data Processing Unit (DPU) based server platforms and SmartNIC (Network Interface Controller) cards.

As a technology preview, the SUSE Linux Enterprise Server for Arm 15 SP3 and SP4 kernels include drivers for running on NVIDIA BlueField-2 DPU.

Should you wish to use SUSE Linux Enterprise Server for Arm on NVIDIA BlueField-2 or BlueField-2X (or BlueField-3) in production, contact your SUSE representative.



Note: Host drivers and tools for NVIDIA BlueField-2 SmartNICs

This Technology Preview status applies only to installing SUSE Linux Enterprise Server for Arm 15 SP4 **on** NVIDIA BlueField-2 DPUs.

For an NVIDIA BlueField-2 DPU PCIe card inserted as **SmartNIC** into a SUSE Linux Enterprise Server 15 SP4 or SUSE Linux Enterprise Server for Arm 15 SP4 based server, check [Section 2.8, “Technology previews”](#) and [Section 5.8, “Kernel”](#) for support status or known limitations of NVIDIA ConnectX* network drivers for BlueField-2 DPUs ([mlx5_core](#) and others).

The [rshim](#) tool is available from SUSE Package Hub ([Section 5.13, “SUSE Package Hub”](#)).

2.8.2.3 etnaviv drivers for Vivante GPUs are available

The NXP* Layerscape* LS1028A/LS1018A System-on-Chip (SoC) contains a Vivante GC7000UL Graphics Processor Unit (GPU), and the NXP i.MX 8M SoC contains a Vivante GC7000L GPU.

As a technology preview, the SUSE Linux Enterprise Server for Arm 15 SP4 kernel includes [etnaviv](#), a Display Rendering Infrastructure (DRI) driver for Vivante GPUs, and the [Mesa-dri](#) package contains a matching [etnaviv_dri](#) graphics driver library. Together they can avoid the need for third-party drivers and libraries.

**Note**

To use them, the Device Tree passed by the bootloader to the kernel needs to include a description of the Vivante GPU for the kernel driver to get loaded. You may need to contact your hardware vendor for a bootloader firmware upgrade.

2.8.2.4 **lima driver for Arm Mali Utgard GPUs available**

The Xilinx* Zynq* UltraScale*+ MPSoC contains an Arm* Mali*-400 Graphics Processor Unit (GPU).

Prior to SUSE Linux Enterprise Server for Arm 15 SP2, this GPU needed third-party drivers and libraries from your hardware vendor.

As a technology preview, the SUSE Linux Enterprise Server for Arm 15 SP2 kernel added lima, a Display Rendering Infrastructure (DRI) driver for Mali *Utgard* microarchitecture GPUs, such as Mali-400, and the Mesa-dri package contains a matching lima_dri graphics driver library.

**Note**

To use them, the Device Tree passed by the bootloader to the kernel needs to include a description of the Mali GPU for the kernel driver to get loaded. You may need to contact your hardware vendor for a bootloader firmware upgrade.

**Note**

The panfrost driver for Mali *Midgard* microarchitecture GPUs is supported since SUSE Linux Enterprise Server for Arm 15 SP2.

2.8.2.5 **mali-dp driver for Arm Mali Display Processors available**

The NXP* Layerscape* LS1028A/LS1018 System-on-Chip contains an Arm* Mali*-DP500 Display Processor.

As a technology preview, the SUSE Linux Enterprise Server for Arm 15 SP2 kernel added `mali-dp`, a Display Rendering Manager (DRM) driver for Mali Display Processors. It has undergone only limited testing because it requires an accompanying physical-layer driver for DisplayPort* output (see [Section 9.3.1, “No DisplayPort graphics output on NXP LS1028A and LS1018A”](#)).

2.8.2.6 Btrfs file system is enabled in U-Boot bootloader

For Raspberry Pi* devices, SUSE Linux Enterprise Server for Arm 12 SP3 and later include *Das U-Boot* as bootloader, in order to align the boot process with other platforms. By default, it loads GRUB as UEFI application from a FAT-formatted partition, and GRUB then loads Linux kernel and ramdisk from a file system such as Btrfs.

As a technology preview, SUSE Linux Enterprise Server for Arm 15 SP2 added a Btrfs driver to U-Boot for the Raspberry Pi (package `u-boot-rpiarm64`). This allows its commands `ls` and `load` to access files on Btrfs-formatted partitions on supported boot media, such as microSD and USB.

The U-Boot command `btrfsbvol` lists Btrfs subvolumes.

2.8.3 Technology previews for Intel 64/AMD64 (x86-64)

2.8.3.1 LUKS2 support in the installer

LUKS2 is supported in the YaST Partitioner as a tech preview. This means that currently it has to be explicitly enabled. This can be done in the following ways:

- set the `YAST_LUKS2_AVAILABLE` environmental variable
- use a checkbox in the YaST Expert Console (Ctrl+Alt+Shift+C in graphical interface, Ctrl+D Shift+C in text interface)

Use the Help button in the installer to see more information about configuring LUKS2.

2.8.3.2 Wayland now works with the latest NVIDIA proprietary driver

The NVIDIA proprietary display driver on Linux has been updated to version 470.57.02 as technology preview. This release provides an enhanced support of Wayland as well as providing X applications on Wayland (via XWayland) with 3D acceleration.

See the [full changelog](#) ↗ for more details.

2.8.3.3 virt-tuner

virt-tuner is a tool for optimizing libvirt XML definitions of a virtual machine for specific use cases. It is shipped as technology preview.

2.8.3.4 virt-manager SEV detection

virt-manager is now providing a simple check box that is only enabled on SEV-compatible hardware and only when an installation is being configured. The GUI will not allow the feature to be enabled for a guest that is already installed nor can it be turned off for a guest that was installed with the feature enabled.

2.8.4 Technology previews for POWER (ppc64le)

2.8.4.1 keylime has been added

The keylime package provides an end-to-end solution for utilizing TPM technology to provide remote trust. See <https://github.com/keylime/keylime> ↗ for more information. The keylime package is offered as a technical preview.

3 Modules, extensions, and related products

This section comprises information about modules and extensions for SUSE Linux Enterprise Server 15 SP4. Modules and extensions add functionality to the system.

**Note: Package and module changes in 15 SP4**

For more information about all package and module changes since the last version, see [Section 2.2.3, “Package and module changes in 15 SP4”](#).

3.1 Modules in the SLE 15 SP4 product line

The SLE 15 SP4 product line is made up of modules that contain software packages. Each module has a clearly defined scope. Modules differ in their life cycles and update timelines.

The modules available within the product line based on SUSE Linux Enterprise 15 SP4 at the release of SUSE Linux Enterprise Server 15 SP4 are listed in the *Modules and Extensions Quick Start* at

<https://documentation.suse.com/sles/15-SP4/html/SLES-all/article-modules.html> ↗.

Not all SLE modules are available with a subscription for SUSE Linux Enterprise Server 15 SP4 itself (see the column *Available for*).

For information about the availability of individual packages within modules, see <https://scc.suse.com/packages> ↗.

3.2 SLE extensions

SLE Extensions add extra functionality to the system and require their own registration key, usually at additional cost. Most extensions have their own release notes documents that are available from

<https://www.suse.com/releasesnotes> ↗.

The following extensions are available for SUSE Linux Enterprise Server 15 SP4:

- SUSE Linux Enterprise Live Patching:
<https://www.suse.com/products/live-patching> ↗
- SUSE Linux Enterprise High Availability Extension:
<https://www.suse.com/products/highavailability> ↗
- SUSE Linux Enterprise Workstation Extension:
<https://www.suse.com/products/workstation-extension> ↗

The following extension is not covered by SUSE support agreements, available at no additional cost and without an extra registration key:

- SUSE Package Hub: <https://packagehub.suse.com/> ↗ (see Section 5.13, “SUSE Package Hub”)

3.3 Derived and related products

This sections lists derived and related products. Usually, these products have their own release notes documents that are available from <https://www.suse.com/releasesnotes> ↗.

- SUSE Linux Enterprise JeOS:
<https://www.suse.com/products/server/jeos> ↗ (see Section 4.3, “Minimal-VM and Minimal-Image”)
- SUSE Linux Enterprise Desktop:
<https://www.suse.com/products/desktop> ↗
- SUSE Linux Enterprise Server for SAP Applications:
<https://www.suse.com/products/sles-for-sap> ↗
- SUSE Linux Enterprise for High-Performance Computing:
<https://www.suse.com/products/server/hpc> ↗
- SUSE Linux Enterprise Real Time:
<https://www.suse.com/products/realtime> ↗
- SUSE Manager: <https://www.suse.com/products/suse-manager> ↗

4 Installation and upgrade

SUSE Linux Enterprise Server can be deployed in several ways:

- Physical machine
- Virtual host
- Virtual machine
- System containers
- Application containers

4.1 Installation

This section includes information related to the initial installation of SUSE Linux Enterprise Server 15 SP4.



Important: Installation documentation

The following release notes contain additional notes regarding the installation of SUSE Linux Enterprise Server. However, they do not document the installation procedure itself.

For installation documentation, see the *Deployment Guide* at <https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-deployment.html> ↗.

Also see the following additional notes:

- Section 5.15.1.2, “Native graphical installer with [virtio-gpu](#)”

4.1.1 Secure Boot shim update blacklists some older bootloaders

Due to an update to the secure booting process, some older bootloaders will be blacklisted permanently in the UEFI flash memory after an updated shim has been installed. As a result, older installation media containing those blacklisted bootloaders can not be used anymore on Secure Boot-enabled systems. That includes any ISOs created for SLES 12 or SLES 15 except for the upcoming SP5 and SP4 Quarterly Update 3.

Registered installations done using the SLES 15 SP4 Quarterly Update 3 ISO need to apply updates during installation otherwise they will become unbootable.

To summarize, these are the supported scenarios for installation:

- Use full medium, Quarterly Update, that contains the new package, but don't register.
- If you register, or if you use online media (and register), use update repositories.

- Use DUD (Driver Update Disk) that also contain the updated RPM for the system being installed.
- If installed by AutoYaST, custom scripts are necessary to perform additional actions.

4.1.2 New media layout

The set of media has changed with 15 SP2. There still are two different installation media, but the way they can be used has changed:

- You can install with registration using either the online-installation medium (as with SUSE Linux Enterprise Server 15 SP1) or the full medium.
- You can install without registration using the full medium. The installer has been added to the full medium and the full medium can now be used universally for all types of installations.
- You can install without registration using the online-installation medium. Point the installer at the required SLE repositories, combining the `install=` and `instsys=` boot parameters:
 - With the `install=` parameter, select a path that contains either just the product repository or the full content of the media.
 - With the `inst-sys=` parameter, point at the installer itself, that is, `/boot/ARCHITECTURE/root` on the medium.

For more information about the parameters, see https://en.opensuse.org/SDB:Linuxrc#p_install.

4.2 Upgrade-related notes

This section includes upgrade-related information for SUSE Linux Enterprise Server 15 SP4.



Important: Upgrade documentation

The following release notes contain additional notes regarding the upgrade of SUSE Linux Enterprise Server. However, they do not document the upgrade procedure itself.

For upgrade documentation, see the *Upgrade Guide* at <https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-upgrade-online.html>.

4.2.1 Migration from SLES 11 to SLES 15 SP4 no longer supported

Upgrading from SLES 11 directly is not supported. See [the upgrade guide](#) ↗ for more information.

4.2.2 Hibernation requires manual intervention

Previously, it was possible for data loss to occur due to the system not hibernating correctly.

In 15 SP4, a sanity check was introduced to prevent this. It works by removing the kernel `resume` parameter if it points to a non-existent device. However, that means a system would not use the hibernation data. To fix it, do the following:

1. Edit `/etc/default/grub` and correct the `resume` parameter to point to an existing device.
2. Regenerate `initrd`.
3. Reboot.

4.2.3 Make sure the current system is up-to-date before upgrading

Upgrading the system is only supported from the most recent patch level. Make sure the latest system updates are installed by either running `zypper patch` or by starting the YaST module *Online Update*. An upgrade on a system that is not fully patched may fail.

4.2.4 Skipping service packs requires LTSS

Skipping service packs during an upgrade is only supported if you have a Long Term Service Pack Support contract. Otherwise, you need to first upgrade to SLE 15 SP3 before upgrading to SLE 15 SP4.

4.2.5 SMT has been replaced by RMT

SLE 12 is the last codestream that SMT (Subscription Management Tool) is available for.

When upgrading your OS installation to SLE 15, we recommend also upgrading from SMT to its replacement RMT (Repository Mirroring Tool). RMT provides the following functionality:

- Mirroring of SUSE-originated repositories for the SLE 12-based and SLE 15-based products your organization has valid subscriptions for.
- Synchronization of subscriptions from SUSE Customer Center using your organization's mirroring credentials. (These credentials can be found in SCC under *Select Organization, Organization, Organization Credentials*)
- Selecting repositories to be mirrored locally via `rmt-cli` tool.
- Registering systems directly to RMT to get required updates.
- Adding custom repositories from external sources and distributing them via RMT to target systems.
- Improved security with proxying: If you have strict security requirements, an RMT instance with direct Internet access can proxy to another RMT instance without direct Internet access.
- Nginx as Web server: The default Web server of RMT is Nginx which has a smaller memory footprint and comparable performance than that used for SMT.

Note that unlike SMT, RMT does not support installations of SLE 11 and earlier.

For more feature comparison between RMT and SMT, see

https://github.com/SUSE/rmt/blob/master/docs/smt_and_rmt.md ↗.

For more information about RMT, also see the new RMT Guide at

<https://documentation.suse.com/sles/15-SP3/html/SLES-all/book-rmt.html> ↗.

.

4.3 Minimal-VM and Minimal-Image

SUSE Linux Enterprise Server Minimal-VM and Minimal-Image is a slimmed-down form factor of SUSE Linux Enterprise Server that is ready to run in virtualization environments and the cloud. With SUSE Linux Enterprise Server Minimal-VM and Minimal-Image, you can choose the right-sized SUSE Linux Enterprise Server option to fit your needs.

SUSE provides virtual disk images for Minimal-VM and Minimal-Image in the file formats `.qcow2`, `.vhd`, and `.vmdk`, compatible with KVM, Xen, OpenStack, Hyper-V, and VMware environments. All Minimal-VM and Minimal-Image images set up the same disk size (24 GB) for the system. Due to the properties of different file formats, the size of Minimal-VM and Minimal-Image image downloads differs between formats.

4.4 JeOS renamed Minimal-VM and Minimal-Image

We have received feedback from users confused by the name JeOS, as a matter of fact the acronym JeOS, which meant Just enough Operating System, was not well understood and could be confused with other images provided by SUSE or openSUSE.

We have decided to go with simplicity and rename JeOS by "Minimal-VM" for all our Virtual Machine Images and "Minimal-Image" for the Raspberry Pi Image. We have also removed a few other characters, in the full images name to make it more simple and clear:

- [SLES15-SP4-Minimal-VM.x86_64-kvm-and-xen-GM.qcow2](#)
- [SLES15-SP4-Minimal-VM.x86_64-OpenStack-Cloud-GM.qcow2](#)
- [SLES15-SP4-Minimal-VM.x86_64-MS-HyperV-GM.vhdx.xz](#)
- [SLES15-SP4-Minimal-VM.x86_64-VMware-GM.vmdk.xz](#)
- [SLES15-SP4-Minimal-VM.aarch64-kvm-GM.qcow2](#)
- [SLES15-SP4-Minimal-Image.aarch64-RaspberryPi-GM.raw.xz](#)

4.5 Minimal-VM VMware flavor with cloud-init support

Rancher and private cloud users requested a VMware image that would include `cloud-init` to allow them to deploy their virtual machines with ease in their environment.

To increase the flexibility of deployment of our Minimal-VM image on VMware, we have added `cloud-init` along with `jeos-firstboot`. `jeos-firstboot` is still used by default except if the user has provided `cloud-init` datasource file.

We have also disabled `firewalld` by default to comply with the cloud environment recommendation.

4.6 Alternative Python 3 development interpreter moved

to a separate module

SLE 15 SP4 introduces a new *Python 3* Module, which includes the alternatively available development Python interpreter, formerly included in the *Basesystem* Module. This new module will allow for more flexibility for the life-cycle of the packages provided within it and a clean separation between the system and development interpreter.

As the `python39` package was part of the *Basesystem* Module on SLE 15 SP3, the introduction of this new module will require some changes when migrating to SLE 15 SP4. If you are using `python39` and migrate from SLE 15 SP3, you will have to add the *Python 3* module after migration via SUSEConnect to receive updates for this alternative interpreter. Otherwise the package will remain orphaned and without security updates.

Packages inside this module can have differing support level and support life-cycle. For more information, see documentation.

4.7 For more information

For more information, see [Section 5, “Changes affecting all architectures”](#) and the sections relating to your respective hardware architecture.

5 Changes affecting all architectures

Information in this section applies to all architectures supported by SUSE Linux Enterprise Server 15 SP4.

5.1 Authentication

5.1.1 User negation in `sudoers.ldap` now works

Previously, the `sudoUser` attribute in `sudoers.ldap` did not accept negation (that is, every user except the specified user).

This has now been enabled and requires `sudo` version 1.9.9 or higher. See `man 5 sudoers.ldap` for more information.

5.1.2 389 Directory Server is the primary LDAP server, the

OpenLDAP server has been removed

The OpenLDAP server (package `openldap2`, part of the *Legacy* SLE module) has been removed from SUSE Linux Enterprise Server 15 SP4. The OpenLDAP client libraries are widely used for LDAP integrations and are compatible with 389 Directory Server. Hence, the OpenLDAP client libraries and command-line tools will continue to be supported on SLES 15 to provide an easier transition for customers that currently use the OpenLDAP Server.

To replace OpenLDAP server, SLES includes 389 Directory Server. 389 Directory Server (package `389-ds`) is a fully-featured LDAPv3-compliant server suited for modern environments and for very large LDAP deployments. 389 Directory Server also comes with command-line tools of its own.

For information about setting up and upgrading to 389 Directory Server, see the [SLES 15 SP3 Security Guide, chapter LDAP—A Directory Service](#) ↗.

5.2 Basic utilities

5.2.1 `util-linux` has been updated

The `util-linux` package has been updated to version 2.37.2.

The deprecated `raw` utility has been removed. Applications have to be ported to `open(2)` device files, such as `/dev/sda1`, with the `O_DIRECT` flag.

5.2.2 `fish` has been updated and moved to SUSE Package Hub

The `fish` package has been updated to version 3. At the same time, it is no longer part of SLE but has been moved to SUSE Package Hub.

5.2.3 Some RPM 4.15 macros have been added

The following RPM 4.15 macros have been added:

- `set_build_flags`
- `smp_build_ncpus`
- `vpath_srcdir`
- `vpath_builddir`

5.3 Containers

5.3.1 Samba size improved for container usage

Previously, installing the Samba package always also installed some large dependencies.

In SLES 15 SP4, we have made some of those components optional so that when installing the package on its own, for example in container environment, these can be omitted, reducing the final footprint of the whole container.

5.3.2 SLE BCI language container images

These are container images providing language SDKs and runtimes. The language container contains and is updated with the same version of the particular language that is in the respective Service Pack of SLES. The following containers are now available:

- Rust BCI
- Ruby BCI

See the [SUSE registry](#) ↗ for more information.

5.3.3 SLE BCI minimal container image

The current SLE container images were not small enough for cloud-native applications. Even though they had fewer packages compared to a regular SLE system, they still included many that were not required. These extra packages increased the size of the image and, most importantly, its attack surface.

As a solution, a minimal container image based on the SUSE BCI (Base Container Image) has been made available. See the [SUSE registry](#) ↗ for more information.

**Note**

The container does not include the `zypper` package but it includes the `rpm` package. That means:

- applications can be deployed into the container in the RPM format
- there is no simple way to install dependencies in the container except for manually copying all the RPM packages and installing them

5.3.4 Busybox SLE BCI (Base Container Image)

Starting with SLES 15 SP4, we will be shipping a new and even smaller variant as part of our BCI portfolio: the Busybox container. This container image ships Busybox as a replacement for Bash and the GNU Coreutils, thereby drastically decreasing its footprint. Additionally, we have included the standard set of CA certificates and the `rpm` database in the image. Note that neither `rpm` nor `zypper` are included in this image as it is only intended for shipping prebuilt applications which include all their dependencies. As this image contains neither Bash nor GNU Coreutils, it is completely free of GPLv3 code. This eases legal requirements in certain cases.

Additional changes to SLE

We have adjusted SLE itself to ensure that the Busybox BCI is built from the same baseline as the rest of the distribution so that it can meet our quality standards. This resulted in the following changes to SLE:

1. Busybox has been updated to version 1.34.1
2. The new package `busybox-links` has been added to SLE. This is a helper package that provides links in `PATH` to `/bin/busybox` for every function that Busybox provides. Thereby it is possible to use the Busybox-provided core utilities instead of the GNU coreutils without having to change the script (assuming it is compatible with Busybox).
3. Bash now no longer provides `/bin/sh` by default; instead this capability has been moved into the `bash-sh` subpackage.
4. `sysuser-tools` has been updated to version 3.1 which includes support for `busybox-adduser` additionally to `adduser` from the GNU coreutils.

5.3.5 RPM Repository Mirroring Tool (RMT) container has been

added

RMT is a tool that allows you to mirror RPM repositories in your own private network.

In a container-native world, running a separate (physical or virtual) host as an RMT server is violating the expectations of a fully containerized experience. That is why to make SUSE Linux Enterprise software updates available in such an environment, we now provide a container with a pre-configured RMT.

The RMT Helm chart provides an easy way to deploy an RMT server on top of a Kubernetes installation. It needs customization to fit your needs:

- list of repositories (modules) you want to mirror
- SUSE Customer Center secrets
- a decent volume size depending on the number of repositories you want to mirror

Once deployed, it will take care of updating the repository mirror daily via a `cron` job in Kubernetes.

**Note: Technical details**

This is an attempt to deliver a software using a containerized architecture. Every component of the stack is defined in its own container, and Helm is used to ease deployment on top of Kubernetes.

RMT server

A containerized version of the RMT application, with the ability to pass its configuration via Helm values. Storage is done on a volume, thus you need to adapt its size depending on the number of repositories you need to mirror.

Mariadb

MariaDB is the database backend for RMT. RMT does create the database and tables at startup if needed so no specific post-installation task is required for it to be usable. Passwords are self-generated unless explicitly specified in the values file.

Nginx

The web server with proper configuration for RMT routes. Having a properly configured webserver out of the box allows you to target your ingress traffic (for RMT) to it directly. You do not have to configure ingress for RMT-specific path handling, as Nginx is configured to do so.

5.3.6 Supported 389 Directory Server has been added

A container for the 389 Directory Server has been added. The pull URL is registry.suse.com/bci/389-ds:latest.

5.3.7 Podman upgrade from 3.4.x to 4.3.1

Podman 4.x is a major release with 60 new features and more than 50 bug fixes compared to Podman 3. It also includes a complete rewrite of the network stack.

Podman 4.x brings a new container network stack based on [Netavark ↗](#), the new container network stack and [Aardvark DNS server ↗](#) in addition to the existing container network interface (CNI) stack used by Podman 3.x. The new stack brings 3 important improvement:

- Better support for containers in multiple networks
- Better IPv6 support
- Better performance

To ensure that nothing break with this major change, the old CNI stack will remain the default on existing installations, while new installs will use Netavark.

New installations can opt to use CNI by explicitly specifying it via the `containers.conf` configuration file, using the `network_backend` field.

If you have run Podman 3.x before upgrading to Podman 4, Podman will continue to use CNI plugins as it had before. There is a marker in Podman's local storage that indicates this. In order to begin using Podman 4, you need to destroy that marker with `podman system reset`. This will destroy the marker, all of the images, all of the networks, and all of the containers.



Warning

Before testing Podman 4 and the new network stack, you will have to destroy all your current containers, images, and networks. You must export/save any import containers or images on a private registry, or make sure that your Dockerfiles are available for rebuilding and scripts/playbooks/states to reapply any settings, regenerate secrets, etc.

Last but not least CNI will be deprecated from upstream at a future date:

<https://github.com/containers/podman/tree/main/cni> ↗

For a complete overview of the changes, please check out the upstream [4.0.0](#) ↗ but also [4.1.1](#) ↗, [4.2.0](#) ↗ and [4.3.0](#) ↗ to be informed about all the new features and changes.

5.3.8 LXC containers have been removed

System containers using LXC have been removed in SUSE Linux Enterprise Server 15 SP4. This includes the following packages:

- `libvirt-lxc`
- `virt-sandbox`

As a replacement, we recommend commonly used alternatives like Docker or Podman.

5.3.9 [suse/sle15](#) container uses NDB as the database back-end for RPM

Starting with SUSE Linux Enterprise 15 SP3, the [rpm](#) package in the [suse/sle15](#) container image no longer supports the BDB back-end (based on Berkeley DB) and switches to the NDB back-end. Tools for scanning, diffing, and building container image using the [rpm](#) binary of the host for introspection can fail or return incorrect results if the host's version of [rpm](#) does not recognize the NDB format.

To use such tools, make sure that the host supports reading NDB databases, such as hosts with SUSE Linux Enterprise 15 SP2 and later.

5.4 Databases

5.4.1 MariaDB 10.6 has been added

The [mariadb](#) package has been updated to version 10.6. See [the full changelog](#) ↗ for more information.

5.4.2 [unixODBC](#) package drivers not for production

Drivers in the [unixODBC](#) package are not suitable for production use. The drivers are provided for test purposes only. We have added a reference to the package's README file with information about third-party [unixODBC](#) drivers that are suitable for production use (<http://www.unixodbc.org/drivers.html> ↗).

5.4.3 The ODBC driver location has changed

Previously in SLES 12, the unixODBC driver for PostgreSQL was included in the [postgresql10-odbc](#) package and was located in [/usr/pgsql-10/lib/psqlodbcw.so](#). In SLES 15 SP4, this driver is part of the [psqlODBC-<version>](#) package and it is located in [/usr/lib64/psqlodbcw.so](#).

For some more information, see: https://bugzilla.suse.com/show_bug.cgi?id=1169697 ↗.

5.4.4 PostgreSQL 14 has been added

PostgreSQL 14 has been added to SUSE Linux Enterprise Server. For information about changes between PostgreSQL 14 and 13, see the [upstream release notes](#) ↗.

At the same time, PostgreSQL 13 has been deprecated and has been moved to the Legacy module. PostgreSQL 12 has been removed.

5.4.5 PostgreSQL REINDEX is required when migrating

If you migrate a PostgreSQL server from an earlier version than SLES 15 SP3, a REINDEX is required before using the database productively again to avoid database corruptions. See <https://www.suse.com/support/kb/doc/?id=000020305> ↗ for details.

5.5 Desktop

Also see the following notes:

- [Section 2.8.3.2, “Wayland now works with the latest NVIDIA proprietary driver”](#)

5.5.1 SLES-specific Windows Terminal shortcuts in WSL

WSL-DistroLauncher will now install branded shortcuts for Windows Terminal upon installation.

5.5.2 Graphical application support in Windows Subsystem for Linux

Users who plan to run Linux graphical applications under Windows using WSL are recommended to install the wsl_gui pattern. This pattern includes packages that enable running of graphical applications without issues like missing fonts etc.

5.5.3 Pipewire has been added

Both pulseaudio and pipewire have been updated to their latest versions.

Right now, `pipewire` is mainly used to provide support for screen sharing in the Wayland session. In the default installation, `pipewire` doesn't have sound support because it is still currently provided by `pulseaudio`.

To enable `pipewire` for audio, install the `pipewire-pulseaudio` package which will remove all `pulseaudio`-related packages and install `wireplumber-audio`, enabling audio support in `pipewire` and `pulseaudio` emulation so that most applications will keep working with `pipewire`.

5.5.4 Printing in GNOME

With GNOME we provide a fully-featured printing stack, which includes `cups`, GNOME itself, and `avahi`. We encourage users to use GNOME settings to manage their printers as it is the most complete solution.

Additionally:

- the relevant GNOME components (`gnome-shell`, `gnome-control-center`, `gnome-settings-daemon`) have been updated to version 41
- `avahi` has been updated to version 0.8
- CUPS has been updated to version 2.2.7

5.5.5 GNOME has been updated

The GNOME desktop has been updated to version 41. Among others, the changes include:

- power profiles
- updated app store appearance
- new multitasking options
- new connections app

See the [full changelog](#) ↗ for more information.

5.5.6 High-quality Bluetooth codecs are now supported

In 15 SP4, the `pulseaudio` package has been updated to version 15, which among other changes brings support for the LDAC, AptX and SBC XQ codecs. See the [full changelog](#) ↗ for more information.

5.5.7 Qt 5 has been updated

The Qt 5 stack has been updated to version 5.15.2. This service pack update also contains KDE's Qt 5 Patch Collection. See

<https://dot.kde.org/2021/04/06/announcing-kdes-qt-5-patch-collection> ↗ for more information.

5.5.8 GTK has been updated

The GTK toolkit has been updated to version 4.0.

This is a major release with many notable changes. Some of the areas that have seen work are the following:

- Data transfers
- Event controllers
- Layout managers
- Render nodes
- Media playback
- Scalable lists
- Shaders
- Accessibility

See [the full changelog](#) ↗ for more information.

5.6 Development

5.6.1 Python

5.6.1.1 Python 3.10 modules, Reduced Python Stack

- Add Python 3.10 interpreter and modules

- We have added an additional long-term supported Python 3.10 interpreter and modules needed to be able to get python modules from PyPI (`python-setuptools`, `python-pip`, `python-virtualenv`) via the *Python 3* Module. It will be regularly updated to the latest patch version.
- Add Python Reduced Stack
 - A reduced subset of what we provide today in SLES 15 SP4 has been added as modules.
- Keep Python 3.6
 - The current Python 3.6 interpreter and packages stay intact.
 - Renaming of Python packages would be done only to avoid source name conflicts.
- Life-cycle & support
 - Python Reduced Stack will be supported until 2026, which is the upstream end-of-life date for Python 3.10. The interpreter will be updated regularly to the latest patch version but modules will stay stable as much as possible.
 - We will continue delivering new interpreters (along with the respective `setuptools` / `wheel` / `pip`) with each new service pack. These are short-term supported for each Service Pack.

5.6.1.2 Python 3.10 has been added, replaces Python 3.9

Python 3.9 that had been available in SLE 15 SP3 has been replaced with Python 3.10 in SLE 15 SP4.

5.6.1.3 All Python packages have been updated

All `python-*` packages have been updated to their most recent versions. Combined with the removal of Python 2 described in [Section 5.6.1.4, “Python 2 has been removed”](#), using external packages from the Python Package Index (PyPI) should now be easier due to less compatibility problems.

5.6.1.4 Python 2 has been removed

With SUSE Linux Enterprise Server 15 SP1, SUSE has started to phase out support for Python 2 in SLE.

In SUSE Linux Enterprise Server 15 SP4, standard Python 2 (executable names `python2` and `python`), and the temporarily available *Python 2* module have been removed. Only Python 3 (executable name `python3`) is now available.

Python scripts usually expect the `python` executable (without a version number) to refer to the Python 2.x interpreter. If the Python 3 interpreter is started instead, this can lead to applications failing or misbehaving. For this reason, SUSE has decided *not* to ship a symbolic link `/usr/bin/python` pointing to the Python 3 executable.

5.6.1.5 Alternative Python 3 development interpreter moved to a separate module

SLE 15 SP4 introduces a new *Python 3* Module, which includes the alternatively available development Python interpreter, formerly included in the *Basesystem* Module. This new module will allow for more flexibility for the life-cycle of the packages provided within it and a clean separation between the system and development interpreter.

As the `python39` package was part of the *Basesystem* Module on SLE 15 SP3, the introduction of this new module will require some changes when migrating to SLE 15 SP4. If you are using `python39` and migrate from SLE 15 SP3, you will have to add the *Python 3* module after migration via SUSEConnect to receive updates for this alternative interpreter. Otherwise the package will remain orphaned and without security updates.

Packages inside this module can have differing support level and support life-cycle. For more information, see documentation.

5.6.2 Node

5.6.2.1 `nodejs16` has been added

The `nodejs16` package has been added. This is the new default NodeJS version.

5.6.2.2 nodejs-common has been changed

Previously, nodejs-common pointed to nodejs14 as this was the default version. With the addition of nodejs16, nodejs-common now points to nodejs-default instead.

5.6.3 Other

5.6.3.1 debuginfod packages have been added

The following packages were added to the *Basesystem Module*:

- debuginfod-client
- libdebuginfod1

We do not provide any debuginfod service but the packages enable creation of one.

See <https://sourceware.org/elfutils/Debuginfod.html> ↗ for more information.

5.6.3.2 Prometheus has been updated

The golang-github-prometheus-prometheus package has been updated to version 2.32.1. See the [full changelog](#) ↗ for more information.

5.6.3.3 tcl has been updated

The tcl package has been updated to version 8.6.12. See the [full changelog](#) ↗ for more information.

5.6.3.4 bzr has been replaced with breezy

The bzr package has been removed from SLES because it requires the removed Python 2. As a replacement, the breezy package has been added. breezy is a Python 3 implementation of the Bazaar VCS.

5.6.3.5 'subversion' has been updated

The `subversion` package has been updated to version 1.14.1.

Among others, this version includes:

- Python 3.x support
- breaking change for the experimental shelving feature

See [the full changelog](#) ↗ for more information.

5.6.3.6 `sccache` and `rustup` have been added

`sccache` is a compiler caching tool for Rust, C, and C++, with optional cloud storage. `rustup` is a tool for managing user Rust toolchains. These two tools have been added in an effort to improve Rust developer tools.

5.6.3.7 Squid has been updated

The `squid` package has been updated from version 4.17 to version 5.2.

See the [full changelog](#) ↗ for more information.

5.6.3.8 TCK compliance testing in SUSE Linux Enterprise

We run the TCK test suite provided by Oracle to ensure that our version of OpenJDK is in compliance with the Java specification.


5.6.3.9 PHP 8 has been added, replacing PHP 7

PHP version 8.0.10 has been added. There are many improvements in this version, some of which are:

- Union Types
- Named Arguments
- Match Expressions
- Attributes
- Constructor Property Promotion

- Nullsafe Operator
- Weak Maps
- Just In Time Compilation

For the full changelog, see <https://www.php.net/ChangeLog-8.php#8.0.0> ↗ and <https://www.php.net/ChangeLog-8.php#8.0.10> ↗.

**Note: Removal of PHP 7**

PHP 8 has replaced PHP 7, which has moved to the *Legacy Module*. After upgrading, make sure that all needed packages, including extension modules, are installed for PHP 8. Additionally, the current version of the package `apparmor-profiles` includes the `php-fpm` profile that was absent before. This profile effectively blocks PHP-FPM processes from accessing the socket file (if its name differs from the template in this file) as well as to directories with PHP scripts (for example, the standard directory `/srv/www/htdocs` and its subdirectories).

5.6.4 Supported Java versions

The following Java implementations are available in SUSE Linux Enterprise Server 15 SP4:

Name (Package Name)	Version	Module	Support
OpenJDK (<code>java-11-openjdk</code>)	11	Base System	SUSE, L3, until 2026-12-31
OpenJDK (<code>java-17-openjdk</code>)	17	Base System	SUSE, L3, until 2027-12-31
OpenJDK (<code>java-1_8_0-openjdk</code>)	1.8.0	Legacy	SUSE, L3, until 2026-12-31
IBM Java (<code>java-1_8_0-ibm</code>)	1.8.0	Legacy	External, until 2025-04-30

5.7 Hardware

5.7.1 Realtek RTL8821CE support

Support for the Realtek RTL8821CE WiFi chip has been added. For more information, see <https://www.realtek.com/en/products/communications-network-ics/item/rtl8821ce> ↗.

5.8 Kernel

Also see the following notes:

- Section 5.14.9, “[zram](#) is now officially supported”
- Section 10.2.1, “[sched_](#) parameters are deprecated”

5.8.1 Advanced Matrix Extensions (AMX) instruction support

SLES 15 SP4 now enabled support for Intel's AMX in their new Sapphire Rapids line of CPUs.

The main use cases for AMX is deep learning inference and training (CNN, DNN), and other data analytics and machine learning applications. In practical terms, AMX can be more than 3x as performant as VNNI/AVX-512 using MKL-DNN and similar low-level libraries.

5.8.2 Tomoyo is not supported

The Tomoyo kernel module is not supported. The primary confinement technology is AppArmor. For more information about the module see <https://www.kernel.org/doc/html/v4.16/admin-guide/LSM/tomoyo.html> ↗.

5.8.3 New functionality in the SUSE kernel module tools package

The SUSE kernel module tools have been updated to better comply with the file system hierarchy standards and also clearly indicate that certain kernel modules will be disabled in a future SUSE Linux Enterprise release.

Distribution-provided configuration files previously placed in the `/etc` directory are now located in the `/lib` directory. The tools continue to recognize the user-supplied configuration files in the `/etc` directory. The `modprobe(8)` tool now presents an interactive dialog in case the user attempts to load one of the obsolete kernel modules. The dialog offers to abort the load operation, load the kernel module once, or override the blacklisting status.

See the package documentation in

</usr/share/doc/packages/suse-module-tools/README.md> for more information.

5.8.4 `zstd` compression of kernel modules

The `zstd` algorithm achieves much higher compression and decompression speed compared to `xz`, at the cost of somewhat lesser compression ratio. As a result, some reading operations during boot and installation are much faster. The module file extension has changed from `.ko.xz` to `.ko.zst` and the content is `zstd`-compressed. All SLE components that manipulate the kernel modules have been adapted. Third-party software that does in-depth examination of kernel modules may require adjustments.

5.8.5 Unified cgroups hierarchy support

The kernel cgroups API comes in two variants: v1 and v2. Additionally, there can be multiple cgroups hierarchies, exposing different APIs. The main two that are relevant in this case are:

- hybrid: v2 hierarchy without controllers, controllers on v1 hierarchies
- unified: v2 hierarchy with controllers

The kernel cgroups v2 is now supported in unified mode. However, the default is still hybrid mode.

See the [kernel documentation](#) ↗ for more information about cgroups.

5.8.6 SEV instance live migration in GCE

Support for live migration in SEV-based Confidential VM images on Google Compute Engine is now supported.

5.8.7 The `kernel-preempt` kernel variant has been replaced with

a boot-time option

In SLE SP2 we have introduced the `kernel-preempt` package for latency-sensitive workloads on x86-64 and AArch64 hardware architectures. The settings of `kernel-preempt` support timely reaction to external events and precise timing at the cost of overall system throughput.

In SLE 15 SP4, the functionality embedded in the `kernel-preempt` package can be activated by adding the boot-time `preempt=full` parameter to the default SLE kernel. The specialized `kernel-preempt` package has been consequently removed from the distribution.

5.8.8 Loading `lpfc` driver in INTx mode

Due to limitations in legacy interrupt routing setup by the firmware/hardware and a change in the kernel, loading the `lpfc` driver in INTx mode does not work.

As a workaround, use the kernel parameter `pci=noioapicquirk` to successfully boot the `lpfc` driver in INTx mode.

For more information see the relevant [kernel commit](#) and the [kernel documentation on boot interrupts](#).

5.8.9 `zstd` compression of `initramfs`

`dracut` supports compression of the `initramfs` image file with `zstd`. `zstd` is superior to `xz` both in terms of speed and compression ration. However, the kernel did not support decompressing a `zstd`-compressed `initramfs` image before.

The feature has now been enabled in the kernel but the default compression of `dracut` is still `xz` for now.

5.8.10 Kernel firmware files are now compressed

In addition to the firmware files being compressed, the packaging scheme has also been changed. Previously, all firmware files were shipped in the `kernel-firmware` package. Now, the files are split into sub-packages, and the `kernel-firmware-all` package will pull all the sub-packages into the system using the `kernel-firmware` `provides` symbol.

5.8.11 BTF has been enabled

BTF (BPF Type Format) has been enabled in the kernel in SLES 15 SP4.

It has not been enabled for kernel modules (`DEBUG_INFO_BT_MODULES=n`). This is because it introduced a new kind of binary compatibility check, which is currently not compatible with the kernel in 15 SP4. It may also prevent loading modules in unexpected ways. However, we still keep BTF of `vmlinux` (`DEBUG_INFO_BT=y`). This way there will be no BTF information on the modules but the Compile-Once-Run-Everywhere feature is still available to BPF programs that only trace kernel functions found within `vmlinux`.

5.8.12 Btrfs sub-page block size support

In previous SLES versions, the Btrfs file system implementation could not work with file systems formatted with a block size smaller than the configured kernel page size. That means a file system formatted with 4-kilobyte block size could be mounted by the kernel using 4-kilobyte page size but not on another system that uses 64-kilobyte pages.

Starting with SLES 15 SP4, kernel with 64-kilobyte page size can use Btrfs file systems formatted with the smaller block size smaller than the kernel page size.

However, writing to compressed files on such a volume is not yet supported.

5.8.13 BPF tooling has been updated

In SLES 15 SP4 the (e)BPF tooling has been updated to the latest version.

`bpftool` is a high-level tracing language for Linux enhanced Berkeley Packet Filter (eBPF) available in the Linux kernel. `bpftool` uses LLVM as a backend to compile scripts to BPF bytecode and makes use of BCC for interacting with the Linux BPF system, as well as existing Linux tracing capabilities: kernel dynamic tracing (`kprobes`), user-level dynamic tracing (`uprobes`), and tracepoints.

The existing packages (`libbpf`, `bcc`, and `bpftool`) have been updated and a new package (`cereal`, the build-time dependency of `bpftool`) has been added.

5.8.14 BlueZ has been updated to version 5.62

In SLES 15 SP4, BlueZ has been upgraded from version 5.55 to version 5.62.

In 5.62 some of the changes were the following:

- API to add new properties for GATT and Adapter.

- For MESH, it updates the configuration client and adds a new API to export the keys.

For the full changelog, see

<https://github.com/bluez/bluez/blob/master/ChangeLog> ↗.

5.8.15 Unprivileged eBPF usage has been disabled

A large amount of security issues was found and fixed in the Extended Berkeley Packet Filter (eBPF) code. To reduce the attack surface, its usage has been restricted to privileged users only.

Privileged users include `root`. Programs with the `CAP_BPF` capability in the newer versions of the Linux kernel can still use eBPF as-is.

To check the privileged state, you can check the value of the `/proc/sys/kernel/unprivileged_bpf_disabled` parameter. Value of 0 means "unprivileged enable", and value of 2 means "only privileged users enabled".

This setting can be changed by the `root` user:

- to enable it temporarily for all users by running the command `sysctl kernel.unprivileged_bpf_disabled=0`
- to enable it permanently by adding `kernel.unprivileged_bpf_disabled=0` to the `/etc/sysctl.conf` file.

5.8.16 Kernel limits

This table summarizes the various limits which exist in our recent kernels and utilities (if related) for SUSE Linux Enterprise Server 15 SP4.

SLES 15 SP4 (Linux 5.14)	AMD64/Intel 64 (x86_64)	IBM Z (s390x)	POWER (ppc64le)	ARMv8 (AArch64)
CPU bits	64	64	64	64
Maximum number of logical CPUs	8192	256	2048	768
Maximum amount of RAM (theoretical/certified)	>1 PiB/64 TiB	10 TiB/256 GiB	1 PiB/64 TiB	256 TiB/n.a.
Maximum amount of user space/kernel space	128 TiB/128 TiB	n.a.	512 TiB ¹ /2 EiB	256 TiB/256 TiB
Maximum amount of swap space	Up to 29 * 64 GB	Up to 30 * 64 GB		
Maximum number of processes	1,048,576			
Maximum number of threads per process	Upper limit depends on memory and other parameters (tested with more than 120,000) ² .			
Maximum size per block device	Up to 8 EiB on all 64-bit architectures			
FD_SETSIZE	1024			

¹ By default, the user space memory limit on the POWER architecture is 128 TiB. However, you can explicitly request mmmaps up to 512 TiB.

² The total number of all processes and all threads on a system may not be higher than the "maximum number of processes".

5.8.17 AMD SEV-ES host support

With QEMU 6.1, the Linux kernel in SLES 15 SP4 now provides SEV-ES (Secure Encrypted Virtualization Encrypted State) host support on AMD EPYC processors. SEV-ES builds off the base AMD SEV to also encrypt CPU register contents when exiting a virtual machine to ensure there is no register information leakage to the hypervisor. In addition, SEV-ES can detect malicious modifications to the CPU register state.

5.8.18 tmon has been updated

tmon is a monitoring and testing tool for the Linux kernel thermal subsystem. Although the version number is still the same in SLES 15 SP4, there have been added some patches.

5.8.19 Shared Virtual Addressing support

The Linux kernel of SLES 15 SP4 now supports Shared Virtual Addressing (SVA), also known as Shared Virtual Memory (SVM). This feature allows sharing of CPU address spaces with devices, and simplifies I/O memory management for device drivers and userspace processes.

Sharing address spaces of processes with devices makes it possible to rely on core kernel memory management for DMA, removing some complexity from application and device drivers. After binding to a device, applications can instruct it to perform DMA on buffers obtained with malloc.

SVA mostly aims at simplifying DMA management but also improves security by isolating address spaces in devices.

5.9 Miscellaneous

5.9.1 Change of SUSEConnect handling of licenses

Some third party repositories available as SLE extension modules come with their own EULAs. Previously, SUSEConnect silently accepted these licenses when registering such modules.

Now SUSEConnect will display the license text and explicitly ask user for acceptance in interactive mode.



Note

This can break some existing scripts which relied on automatic acceptance of licenses. Users who want to use SUSEConnect with third party licenses in an automatic way can use the --auto-agree-with-licenses CLI option.

5.9.2 Enriched system visibility in the SUSE Customer Center (SCC)

SUSE is committed to helping provide better insights into the consumption of SUSE subscriptions regardless of where they are running or how they are managed; physical or virtual, on-prem or in the cloud, connected to SCC or Repository Mirroring Tool (RMT), or managed by SUSE Manager. To help you identify or filter out systems in SCC that are no longer running or decommissioned, SUSEConnect now features a daily “ping”, which will update system information automatically.

For more details see the documentation at <https://documentation.suse.com/subscription/suseconnect/single-html/SLE-suseconnect-visibility/>.

5.9.3 `modinfo` does not show driver versions

Outputting the driver version when using `modinfo` has been removed upstream. This has been done mainly because drivers developed as part of the Linux kernel tree are already versioned along with the kernel, and having a separate version was optional and confusing.

5.9.4 Use `/dev/mapper` instead of UUID in `fstab` for LUKS-back-up devices

During installation, the entries generated for LUKS devices in `/etc/fstab` used UUID. This meant that tools such as systemd generators could not know which LUKS device to activate to make a filesystem appear, unless all volumes were set up at boot.

To fix this, entries in `/etc/fstab` now use the name of the resulting encrypted block device (`/dev/mapper/cr_xxx`) because it identifies the LUKS-backed device without ambiguity.

5.9.5 `adcli` now supports setting password expiry

The `adcli` command now supports the `--dont-expire-password` parameter.

This parameter sets or unsets the `DONT_EXPIRE_PASSWORD` flag in the `userAccountControl` attribute to indicate if the machine account password should expire or not. By default `adcli` will set this flag while joining the domain which corresponds to the default behavior of Windows clients.

5.9.6 NTLM support in the Unified Installer

The online SLES media require that customers register with SUSE Customer Center at installation time. However, previously the Unified Installer proxy configuration did not support NTLM authentication. NTLM is a common form of authentication in enterprise environments with Microsoft Active Directory.

In SLES 15 SP4, support for NTLM authentication in the Unified Installer has been added.

5.9.7 chrony Network Time Security (NTS) support

This option enables authentication using the Network Time Security (NTS) mechanism. Unlike with the key option, the server and client do not need to share a key in a key file. NTS has a Key Establishment (NTS-KE) protocol using the Transport Layer Security (TLS) protocol to get the keys and cookies required by NTS for authentication of NTP packets.

5.9.8 New version of SUSEConnect eliminates Ruby requirements

Previously, SUSEConnect was written in Ruby and therefore required the Ruby stack to be present in the installed system. This conflicted with the increasing demand for minimal product footprint, especially for products that were targeted for edge and embedded use cases.

In SLES 15 SP4, SUSEConnect has been replaced by the new version written in Go called suseconnect-ng. This new version also obsoletes the previously separate plugins zypper-migration-plugin and zypper-search-packages-plugin, which have been removed.



Note: Abbreviated options

Abbreviated options not mentioned in --help are not supported. Previously, some abbreviated options worked due to the way Ruby parses options but they were not officially supported nor documented.

5.9.9 Boot-time graphics DRM enablement for UEFI and VESA

framebuffers

On system start-up, the graphics console is first serviced by the framebuffer drivers. Later in the process, the framebuffer driver hands over the graphics-card memory to the Direct Rendering Manager (DRM). In some scenarios, the handover can fail and the system graphics console can appear frozen. 15 SP4 provides a DRM native boot-time graphics driver, called simplifiedrm, as a replacement to the framebuffer drivers.

To use the new graphics driver, simplifiedrm, the module has to be loaded during boot. As root, on the console, type:

```
echo "simplifiedrm" > /etc/modules-load.d/simplifiedrm.conf
```

systemd will automatically load the simplifiedrm driver on the next startup. To avoid this, simply remove the file. To use the driver, pass the kernel parameter enable_sysfb on the next boot. This can be done from within the GRUB boot menu.

There should be no difference from regular boot. Everything should look as before. To verify that the simplifiedrm driver has been used, in the console type:

```
dmesg | grep drm
```

The output should mention simplifiedrm.

By default, the hardware's native driver replaces simplifiedrm during boot. To disable native drivers, pass the kernel parameters enable_sysfb and nomodeset to the kernel on the next boot. The former parameter enables simplifiedrm and the latter disables the native driver. Afterwards, all the graphic output will be done by simplifiedrm.

5.9.10 Adding a new welcome screen for jeos-firstboot to all consoles

Finding the right console for the jeos-firstboot wizard can be tricky for the user and nothing was in place before to introduce the jeos-firstboot wizard to the user.

This features addressed these two issues:

- It adds a welcome screen to greet the user and tell them about which distribution is about to be started and configured.
- It shows the welcome screen on all the consoles. This solves the issue where the user might not know which console needs to be used for the jeos-firstboot wizard.

5.10 Networking

Also see the following notes:

- [Section 2.8.1.2, “Redfish-finder functionality in wicked”](#)

5.10.1 `set-hostname` has been deprecated

The `set-hostname` command has been deprecated. Use `hostnamectl hostname` instead.

5.10.2 `bind-chrootenv` has been removed, functionality replaced with `systemd`

The `bind-chrootenv` package has been removed. Instead of protecting the system using a chroot jail, the built-in `systemd` methods to provide chroot like containment are now used.

To fix potential issues arising from this change, in `/etc/sysconfig/named` remove any mention of `run_chrooted`. The log file paths also need to be adjusted because they are different from the chroot setup.

5.10.3 Samba

The version of Samba shipped with SUSE Linux Enterprise Server 15 SP4 delivers integration with Windows Active Directory domains. In addition, we provide the clustered version of Samba as part of SUSE Linux Enterprise High Availability Extension 15 SP4.

5.10.3.1 Samba has been updated to 4.15

The `samba` package has been updated to version 4.15.

Some of the changes in this version are the following:

- File server
 - The following SMB (development) dialects are no longer supported: `SMB2_22`, `SMB2_24`, and `SMB3_10`
 - Modernized VFS interface, basing all access to the server's filesystem on file handles and not on paths

- "server multi channel support" no longer experimental, enabled by default
- `samba-tool` available without `ad-dc`
- Improved command line user experience
- Winbind
 - Scanning of trusted domains disabled by default
 - Enterprise principals enabled by default
 - The `net` utility supports Offline Domain Join

5.10.3.2 Samba Active Directory Domain Controller has been deprecated

The Samba Active Directory Domain Controller (package `ad-dc`) has been deprecated. It had previously been available only as a technical preview.

5.10.3.3 SMB1 support has been deprecated

With Samba 4.17 it is planned to disable the SMB1 protocol. We therefore deprecated SMB1 for a possible future update of Samba. This affects the Samba file server, its libraries and clients, as well as the kernel CIFS client (`cifs.ko` and `mount.cifs`). This version of the protocol is insecure and usage of version 2.02 or later is recommended.

5.10.3.4 NT Domain Controller has been deprecated

The NT Domain Controller support has been deprecated since Samba 4.13.

5.10.3.5 Samba NIS support has been deprecated

NIS (netgroups) support has been deprecated. It was removed upstream in Samba 4.15 but we have retained it for this release in a deprecated state.

5.10.4 NFS

5.10.4.1 NFSv4

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

5.10.4.2 NFS readahead size has been decreased

Previously, the default readahead size was set to a 15× the maximum read size reported by the server. This artificial increase has been dropped in 15 SP4.

Networks and storage are now faster and having a large amount of reads in flight is usually not beneficial anymore. However, the reduced readahead size can result in worse performance in case of multiple parallel sequential IO streams on slow networks or with slow workloads. To increase the maximum readahead size, change the value of

`/sys/devices/virtual/bdi/$BDI_NUM/read_ahead_kb`. The `$BDI_NUM` of the mount can be found, for example, in the third column of `/proc/self/mountinfo`.

5.11 Security

5.11.1 TLS 1.1 and 1.0 are no longer recommended for use

The TLS 1.0 and 1.1 standards have been superseded by TLS 1.2 and TLS 1.3. TLS 1.2 has been available for considerable time now.

SUSE Linux Enterprise Server packages using OpenSSL, GnuTLS, or Mozilla NSS already support TLS 1.3. We recommend no longer using TLS 1.0 and TLS 1.1, as SUSE plans to disable these protocols in a future service pack. However, not all packages, for example, Python, are TLS 1.3-enabled yet as this is an ongoing process.

5.11.2 `p11-kit-server` has been added

The `p11-kit-server` package has been added. It provides command line tools that enable exporting of PKCS#11 modules through a Unix domain socket.

5.11.3 Global crypto-policies have been added

Previously, every cryptography-using tool and library had been using its own cryptographic configuration. Some packages used DEFAULT_SUSE SSL cipherset in older releases but this was a hardcoded set of ciphers.

In SLES 15 SP4, we use global cryptographic policies. The default setting will already block some algorithms that are considered insecure. To enable all legacy ciphers use:

```
update-crypto-policies --set LEGACY
```

Please see the man page for crypto-policies for details on further configuration.

5.11.4 fail2ban has been added

The fail2ban package has been added to the *Basesystem Module*. It is used to ban addresses that make too many authentication failures, based on scanning various log files.

5.11.5 OpenSSL 3.0 availability

Starting with 15 SP4, SLES includes the OpenSSL 3.0 library in addition to the system OpenSSL 1.1.1 library.

The OpenSSL 3 library is currently not used by system applications but can be used by third-party libraries and applications already. It can be used in processes together with the 1.1.1 library.

5.11.6 /dev is not mounted noexec anymore

Since systemd v248, /dev is not mounted noexec anymore. This did not provide any significant security benefits and conflicted with the executable mappings used with /dev/sgx device nodes. The previous behavior can be restored for individual services with NoExecPaths=/dev (or by allow-listing and excluding /dev from ExecPaths=).

5.11.7 Certificate Auto Enrollment

Certificate Auto Enrollment allows devices to enroll for certificates from Active Directory Certificate Services. It is enabled by Group Policy using Samba's samba-gpupdate command.

5.11.8 Unlocking LUKS volumes with TPM 2.0 or FIDO2


The unlocking of fully-encrypted devices using TPM 2.0 or FIDO2 is now supported.

There are at least 2 common use cases for this:

- laptops and similar devices: unlocking encrypted disk only with an external, secure factor
- server or edge: automated encryption of server disks at boot, especially in remote locations, that are made unusable if the disk is physically stolen

5.11.9 FIPS mode now available

SLES now supports enabling FIPS mode. The Federal Information Processing Standard 140-2 (FIPS 140-2) is a security standard for cryptographic modules. It is frequently needed when doing work for the United States federal government.

See the [Enabling compliance with FIPS 140-2](#)  section in the *Security and Hardening Guide* for more information.

5.11.10 [sigstore](#) support has been added

[sigstore](#) is a project that aims to improve the open source software supply chain by easing the adoption of cryptographic software signing, backed by transparency log technologies.

As part of adding support for [sigstore](#), the following were added:

- [rekor](#) – a global log, includes server and client
- [cosign](#) – container signing and verification

For more information see <https://sigstore.github.io/> .

5.12 Storage and file systems

Also see the following release notes:

- [Section 5.8.12, “Btrfs sub-page block size support”](#)

5.12.1 cryptsetup has been updated

The cryptsetup package has been updated to version 2.4.3. Among the various improvements, it enables the use of FIDO and TPM tokens for unlocking LUKS filesystems.

See the following changelogs for more information:

- [2.4.0](#) ↗
- [2.4.3](#) ↗

5.12.2 DFS failover support

Support has been added for DFS target failover not only when the original connection is lost but also when refreshing DFS-cached referrals by either forcing it through mount -o remount or cached entries expired:

- DFS target hostname resolved to a different IP address
- Refreshed DFS referral no longer matches the currently connected DFS share:
 - Try to connect to all new DFS targets and then select the working one as target hint.
 - If the reconnected target is a regular share, then tree connect to it. Otherwise, resolve any nested DFS links at reconnect time and avoid having to umount and mount it again. There is a maximum of 8 nested DFS links per mount (2) and failover to avoid looping forever on broken DFS referrals.

5.12.3 iotop support

iotop does not display values for SWAPIN and IO %.

Since Linux kernel 5.14, either kernel boot parameter delayacct needs to be specified or kernel.task_delayacct sysctl needs to be enabled.

5.12.4 Improved booting from remote disks

Systems with mount points located in network-based disks can fail to boot after installation unless the `_netdev` option is set in `/etc/fstab`. However previously, the installer did not consider all the scenarios and thus might not have set the flag correctly.

In SUSE Linux Enterprise Server 15 SP4, YaST will now:

- only add `_netdev` in the last step of the so-called Guided Proposal
- will no longer add `_netdev` to the list of default mount options
- will never remove any `_netdev` previously added by the user

YaST will add the `_netdev` option in these cases:

- the mount point is not `/` or `/var` and it is also not on the same device as `/` or `/var`
- the mount point does not have the mount option `x-initrd.mount` and is not on the same device as any other mount point with this option

YaST will also show a warning in the Expert Partitioner if it thinks `_netdev` should be added but the user omitted it, though it is possible to ignore it.

5.12.5 NVMe-oF-TCP CDC support

In SLES 15 SP4, in order to support new features of NVMe such as Centralized Discovery Controller (CDC), the package `nvme-cli` has been updated to v2.0, and two new packages have been added: `libnvme` v1.0 and `nvme-stas` v1.0.

NVMe-oF suffers from a well-known discovery problem that fundamentally limits the size of realistic deployments. To address this discovery problem, thanks to the newly added and updated packages in 15 SP4, it is now possible to manage NVMe-oF via a “network-centric” (Centralized Discovery Controller) provisioning process instead of an “end node-centric” (Direct Discovery Controller) one by using the following approaches:

1. Automated Discovery of NVMe-oF Centralized Discovery Controllers in an IP Network and preventing the user from manually configuring the IP Address of Discovery Controllers.

2. The Centralized Discovery Controller (CDC) allows users to manage connectivity from a single point of management on an IP Fabric by IP Fabric basis. Keep in mind that the user is still able to perform explicit registration with CDCs and DDCs.

5.12.6 `/etc/fstab` option to disable `fstrim` has been added

Previously, file systems that supported `fstrim` were always trimmed if the device supported the TRIM command.

In 15 SP4, the `X-fstrim.nostrim` option has been added. Adding this option to a device in `/etc/fstab` will opt it out of the `fstrim` functionality without disabling the `fstrim` service.

5.12.7 XFS V4 format file systems have been deprecated

Customers who have created XFS file system on SLE 11 or prior will see the following message:

```
Deprecated V4 format (crc=0) will not be supported after September 2030
```

While the file system will work and be supported until the date mentioned, it is best to re-create the file system:

1. Backup all the data to another drive or partition
2. Create the file system on the device
3. Restore the data from the backup

5.12.8 Comparison of supported file systems

SUSE Linux Enterprise was the first enterprise Linux distribution to support journaling file systems and logical volume managers in 2000. Later, we introduced XFS to Linux, which allows for reliable large-scale file systems, systems with heavy load, and multiple parallel reading and writing operations. With SUSE Linux Enterprise 12, we started using the copy-on-write file system Btrfs as the default for the operating system, to support system snapshots and rollback.

The following table lists the file systems supported by SUSE Linux Enterprise.

Support status: + supported / - unsupported

Feature	Btrfs	XFS	Ext4	OCFS 2 ¹
Supported in product	SLE	SLE	SLE	SLE HA
Data/metadata journaling	N/A ²	- / +	+ / +	- / +
Journal internal/external	N/A ²	+ / +	+ / +	+ / -
Journal checksumming	N/A ²	+	+	+
Subvolumes	+	-	-	-
Offline extend/shrink	+ / +	- / -	+ / +	+ / - ³
Inode allocation map	B-tree	B+-tree	Table	B-tree
Sparse files	+	+	+	+
Tail packing	-	-	-	-
Small files stored inline	+ (in metadata)	-	+ (in inode)	+ (in inode)
Defragmentation	+	+	+	-
Extended file attributes/ACLs	+ / +	+ / +	+ / +	+ / +
User/group quotas	- / -	+ / +	+ / +	+ / +
Project quotas	-	+	+	-
Subvolume quotas	+	N/A	N/A	N/A
Data dump/restore	-	+	-	-
Block size default	4 KiB ⁴			
Maximum file system size	16 EiB	8 EiB	1 EiB	4 PiB
Maximum file size	16 EiB	8 EiB	1 EiB	4 PiB

¹ OCFS 2 is fully supported as part of the SUSE Linux Enterprise High Availability Extension.

² Btrfs is a copy-on-write file system. Instead of journaling changes before writing them in-place, it writes them to a new location and then links the new location in. Until the last write, the changes are not "committed". Because of the nature of the file system, quotas are implemented based on subvolumes (qgroups).

³ To extend an OCFS 2 file system, the cluster must be online but the file system itself must be unmounted.

⁴ The block size default varies with different host architectures. 64 KiB is used on POWER, 4 KiB on other systems. The actual size used can be checked with the command `getconf PAGE_SIZE`.

Additional notes

Maximum file size above can be larger than the file system's actual size because of the use of sparse blocks. All standard file systems on SUSE Linux Enterprise Server have LFS, which gives a maximum file size of 2^{63} bytes in theory.

The numbers in the table above assume that the file systems are using a 4 KiB block size which is the most common standard. When using different block sizes, the results are different.

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> ↗.

Some file system features are available in SUSE Linux Enterprise Server 15 SP4 but are not supported by SUSE. By default, the file system drivers in SUSE Linux Enterprise Server 15 SP4 will refuse mounting file systems that use unsupported features (in particular, in read-write mode). To enable unsupported features, set the module parameter `allow_unsupported=1` in `/etc/modprobe.d` or write the value `1` to `/sys/module/MODULE_NAME/parameters/allow_unsupported`. However, note that setting this option will render your kernel and thus your system unsupported.

5.12.9 Supported Btrfs features

The following table lists supported and unsupported Btrfs features across multiple SLES versions.

Support status: + supported / - unsupported

Feature	SLES 11 SP4	SLES 12 SP5	SLES 15 GA	SLES 15 SP1	SLES 15 SP2	SLES 15 SP3
Copy on write	+	+	+	+	+	+
Free space tree (Free Space Cache v2)	-	-	-	+	+	+
Snapshots/subvolumes	+	+	+	+	+	+
Swap files	-	-	-	+	+	+
Metadata integrity	+	+	+	+	+	+
Data integrity	+	+	+	+	+	+
Online metadata scrubbing	+	+	+	+	+	+
Automatic defragmentation	-	-	-	-	-	-
Manual defragmentation	+	+	+	+	+	+
In-band deduplication	-	-	-	-	-	-
Out-of-band deduplication	+	+	+	+	+	+
Quota groups	+	+	+	+	+	+
Metadata duplication	+	+	+	+	+	+
Changing metadata UUID	-	-	-	+	+	+
Multiple devices	-	+	+	+	+	+
RAID 0	-	+	+	+	+	+
RAID 1	-	+	+	+	+	+
RAID 5	-	-	-	-	-	-
RAID 6	-	-	-	-	-	-
RAID 10	-	+	+	+	+	+
Hot add/remove	-	+	+	+	+	+
Device replace	-	-	-	-	-	-
Seeding devices	-	-	-	-	-	-
Compression	-	+	+	+	+	+
Big metadata blocks	-	+	+	+	+	+
Skinny metadata	-	+	+	+	+	+
Send without file data	-	+	+	+	+	+
Send/receive	-	+	+	+	+	+
Inode cache	-	-	-	-	-	-
Fallocate with hole punch	-	+	+	+	+	+

5.13 SUSE Package Hub

SUSE Package Hub brings open-source software packages from openSUSE to SUSE Linux Enterprise Server and SUSE Linux Enterprise Desktop.

Usage of software from SUSE Package Hub is not covered by SUSE support agreements. At the same time, usage of software from SUSE Package Hub does not affect the support status of your SUSE Linux Enterprise systems. SUSE Package Hub is available at no additional cost and without an extra registration key.

5.13.1 Important package additions to SUSE Package Hub

Among others, the following packages have been added to SUSE Package Hub:

- Section 5.2.2, “[fish](#)” has been updated and moved to SUSE Package Hub”

5.14 System management

Also see the following notes:

- Section 6.1.1, “User Space Live Patching (ULP) infrastructure and live patches for shared libraries”

5.14.1 Silence `KillMode=None` messages

The log level of the deprecation warnings regarding `killmode=None` have been reduced. Instead of `warning`, they are now logged at the `debug` log level.

5.14.2 systemd updated to 249

systemd has been updated to version 249. Find a summary of changes below. See the [full changelog](#) ↗ for more information.

New features

- Cryptography

- A new `systemd-cryptenroll` tool has been added to enroll TPM 2.0, FIDO2 and PKCS#11 security tokens to LUKS volumes, list and destroy them. It also supports enrolling "recovery keys" and regular passphrases.
- Support has been added to `systemd-cryptsetup` for extracting the PKCS#11 token URI and encrypted key from the LUKS2 JSON embedded metadata header.
- `systemd-cryptsetup` gained support for unlocking LUKS2 volumes using TPM 2.0 hardware, as well as FIDO2 security tokens.
- The `ConditionSecurity=tpm2` unit file setting may be used to check if the system has at least one TPM 2.0 (`tpmrm` class) device.
- A new credentials logic has been added to system services. This is a simple mechanism to pass privileged data to services in a safe and secure way.
- A concept of system extension images is introduced. Such images may be used to extend the `/usr/` and `/opt/` directory hierarchies at runtime with additional files (even if the file system is read-only). When a system extension image is activated, its `/usr/` and `/opt/` hierarchies and `os-release` information are combined via `overlayfs` with the file system hierarchy of the host OS. A new `systemd-sysext` tool can be used to merge, un-merge, list, and refresh system extension hierarchies.
- udev rules may now set `log_level=` option. This allows debug logs to be enabled for select events, for example, just for a specific subsystem or even a single device.
- A new udev hardware database has been added for FireWire devices (IEEE 1394).

Deprecation warnings

- Builds with support for separate `/` and `/usr/` hierarchies (so-called "split-usr" builds, "non-merged-usr" builds) are now officially deprecated. A warning is emitted during build.
- Systems with the legacy cgroup v1 hierarchy are now marked as "tainted", to make it clearer that using the legacy hierarchy is not recommended.

- `systemctl --check-inhibitors=true` may now be used to obey inhibitors even when invoked non-interactively. The old `--ignore-inhibitors` switch is now deprecated and has been replaced with `--check-inhibitors=false`.

Incompatible changes

- The "net_id" built-in of udev has been updated with three backwards-incompatible changes:
 - PCI hotplug slot names on s390 systems are now parsed as hexadecimal numbers. They were incorrectly parsed as decimal previously, or ignored if the name was not a valid decimal number.
 - PCI onboard indices up to 65535 are allowed. Previously, numbers above 16383 were rejected. This primarily impacts s390 systems, where values up to 65535 are used.
 - Invalid characters in interface names are replaced with the character "_".
- Kernel API incompatibility: Linux 4.14 introduced two new `uevents` to the Linux device model: `bind` and `unbind`. The introduction of these new `uevents` (which are typically generated for USB devices and devices needing a firmware upload before being functional) resulted in a number of issues. To minimize issues resulting from this kernel change starting with `systemd-udev` 247, the udev tags concept (which is a concept for marking and filtering devices during enumeration and monitoring) has been reworked: udev tags are now "sticky", meaning that once a tag is assigned to a device it will not be removed from the device again until the device itself is removed (that is, unplugged).
- Units using `ConditionNeedsUpdate=` will no longer be activated in the `initrd`.
- `systemd-hostnamed` will now respect hostname being explicitly set to `localhost` instead of silently ignoring it.
- PAM configuration in `/etc/pam.d` will take precedence before `/usr/lib/pam.d/`.
- Support for the `ConditionNull=` unit file condition has been removed.

5.14.3 AutoYaST per-product schema

AutoYaST provides a scheme package, which can be used to manually validate a created AutoYaST profile. However, there are AutoYaST modules, which are only available in some products.

Now there are different versions of the `yast2-schema` package, which only include the modules relevant for the particular product.

5.14.4 YaST now offers several visual themes

YaST now makes it possible to select from several different visual themes. This includes a dark or a high-contrast mode, and several others.

5.14.5 YaST now assigns subuids/subgids

Previously, users added using YaST did not have subuids/subgids assigned. This is required, for example, for running rootless containers.

In 15 SP4, users created using YaST are always assigned subuids/subgids.

5.14.6 Dropped support in YaST for groups password

Previously, it was possible to set a group password in YaST. However, group passwords are an inherent security problem. This even more true in SUSE Linux Enterprise because, for historical reasons, a separate `/etc/gshadow` file is not used.

Thus this features has been removed from both the user interface and AutoYaST. When cloning a system with AutoYaST, the group description does not include the `<group_password>` or `<encrypted>` tags anymore. Those elements are also ignored when importing a group from an existing AutoYaST profile.

5.14.7 Changes in the section `<user_defaults>` of the AutoYaST profile

The `<user_defaults>` section of the AutoYaST profile has been updated to only include relevant settings.

As a result, the entries `<groups>`, `<no_groups>`, and `<skel>` will not longer be exported when cloning a system and they will be ignored when importing an existing AutoYaST profile during installation.

5.14.8 AutoYaST GRUB2 password protection

AutoYaST now supports setting password protection in GRUB2 either in plain text or encrypted/hashed form. See the [password](#) option in the [AutoYaST Guide](#)  for more information.

5.14.9 [zram](#) is now officially supported

[zram](#) is a Linux kernel feature that provides a form of virtual memory compression. Previously, it has only been available in SUSE Package Hub.

In 15 SP4, the `systemd-zram-service` package has been moved from SUSE Package Hub and is thus now officially supported.

See the [package's official website](#)  and the [kernel documentation](#)  for more information.

5.14.10 AutoYaST UEFI detection

AutoYaST can now detect whether the system was booted in UEFI mode. This is exposed via the `boot_efi` ERB helper and the `efi` predefined system attribute.

See the *AutoYaST Guide* at <https://documentation.suse.com/sles/15-SP4/html/SLES-autoyast/>  for more information.

5.14.11 Hibernation proposal in installer

The installer proposes hibernation (including adding the `resume` kernel option) only if these conditions are met:

- Architecture is `x86_64`
- There must be a swap partition

In other cases, hibernation is not proposed but you can change it manually.

5.14.12 Support for System V `init.d` scripts is deprecated

`systemd` in SUSE Linux Enterprise Server 15 SP4 automatically converts System V `init.d` scripts to service files. Support for System V `init.d` scripts is deprecated and will be removed with the next major version of SUSE Linux Enterprise Server. In the next major version of SUSE Linux Enterprise Server, `systemd` will also stop converting System V `init.d` scripts to `systemd` service files.

To prepare for this change, use the automatically generated systemd service files directly instead of using System V `init.d` scripts. To do so, copy the generated service files to `/etc/systemd/system`. To then control the associated services, use `systemctl`.

The automatic conversion provided by systemd (specifically, `systemd-sysv-generator`) is only meant to ensure backward compatibility with System V `init.d` scripts. To take full advantage of systemd features, it can be beneficial to manually rewrite the service files.

This deprecation also causes the following changes:

- The `/etc/init.d/halt.local` initscript is deprecated. Use systemd service files instead.
- `rcSERVICE` controls of systemd services are deprecated. Use systemd service files instead.
- `insserv.conf` is deprecated.

5.14.13 Searching packages across all SLE modules

In SLE 15 SP4 you can search for packages both within and outside of currently enabled SLE modules using the following command:

```
zypper search-packages -d SEARCH_TERM
```

This command contacts the SCC and searches all modules for matching packages. This functionality makes it easier for administrators and system architects to find the software packages needed.

5.15 Virtualization

For more information about acronyms used below, see <https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html>.



Important: Virtualization limits and supported hosts/guests

These release notes only document changes in virtualization support compared to the immediate previous service pack of SUSE Linux Enterprise Server. Full information regarding virtualization limits for KVM and Xen as well as supported guest and host systems is now available as part of the SUSE Linux Enterprise Server documentation.

See the *Virtualization Guide* at <https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-virt-support.html> ↗.

5.15.1 KVM

5.15.1.1 Virtualized TPM (vTPM) support for Windows Server 2022

The new Windows Server Virtualization Validation Program (SVVP) now requires TPM.

For this reason, in SLE 15 SP4 virtualized TPM (vTPM) now works with KVM.

5.15.1.2 Native graphical installer with [virtio-gpu](#)

Support for native graphical installer has been added if [virtio-gpu](#) is used.

To that effect, the display type dialog shown during installation has been changed:

- the [X11](#) option has been renamed to [Remote X11](#)
- the [ASCII Console](#) option has been renamed to [Text-based UI](#)
- a [Graphical UI](#) option has been added, which is a graphical Qt-based UI

You can also explicitly display the dialog by adding the [netsetup=display](#) parameter to boot options.

For additional information see the [Connecting to the SUSE Linux Enterprise Server installation system](#) ↗ section in the *Deployment Guide*.

5.15.1.3 Support for AMD SEV-ES

Support for AMD Secure Encrypted Virtualization-Encrypted State (SEV-ES) has been added. The main use case is preventing access by third parties to data hosted in a public cloud. For more information see <https://developer.amd.com/sev/> ↗.

5.15.2 Xen

5.15.2.1 Dom0 suspend/resume not supported

The suspend/resume feature is not supported with Xen Dom0.

5.15.2.2 Automatic virtual firmware selection

Before, firmware such as OMVF had to be specified by an explicit path to the firmware.

With this change, the upstream communities now define metadata that describe the firmware. This allows firmware to be automatically selected based on user-friendly configuration. For example, the user can now simply specify EFI and the appropriate firmware will be selected.

5.15.2.3 Xen has been updated to 4.16.0

Xen has been updated to version 4.16.0. Some of the changes in this version are the following: * Miscellaneous fixes to the TPM manager software in preparation for TPM 2.0 support. * Increased reliance on the PV shim as 32-bit PV guests will only be supported in shim mode going forward. This change reduces the attack surface in the hypervisor. * Increased hardware support by allowing Xen to boot on Intel devices that lack a Programmable Interval Timer. * Cleanup of legacy components by no longer building QEMU Traditional or PV-Grub by default. Note both projects have upstream Xen support merged now, so it is no longer recommended to use the Xen specific forks. * Xen can now export Intel Processor Trace (IPT) data from guests to tools in dom0. * Xenstored and oxenstored both now support LiveUpdate (tech preview). * Switched x86 MSR accesses to deny by default policy. * Named PCI devices for xl/libxl and improved documentation for xl PCI configuration format. * x86: Allow domains to use AVX-VNNI instructions. * Added `XEN_SCRIPT_DIR` configuration option to specify location for Xen scripts. * Increase the maximum number of guests which can share a single IRQ from 7 to 16, and make this configurable with `irq-max-guests`.

5.15.3 QEMU

5.15.3.1 QEMU has been updated to 6.2

QEMU has been updated to version 6.2. For the full changelog see <https://wiki.qemu.org/ChangeLog/6.2>.



Note: Deprecation notice

In previous versions, if no explicit image format was provided, some QEMU tools tried to guess the format of the image, and then process it accordingly. Because this feature is a potential source of security issues, it has been deprecated and removed. It is now necessary to explicitly specify the image format. For more information, see <https://qemu-project.gitlab.io/qemu/about/removed-features.html#qemu-img-backing-file-without-format-removed-in-6-1>.

5.15.4 libvirt

5.15.4.1 libvirt has been updated to 0.8.0

libvirt has been updated to version 0.8.0. For more information see <https://libvirt.org/news.html> ↗.

5.15.5 Others

5.15.5.1 cloud-init has been updated

The cloud-init package has been updated to version 21.3.

Additionally, the update also provides the VMware data source.

5.15.5.2 apparmor-parser is now installed by default in Minimal-VM images

Enforcing good behavior and preventing both known and unknown security flaws from being exploited is highly recommended in the Linux world.

For this reason, our Minimal-VM images now have AppArmor packages installed by default. This allows a user to configure AppArmor policies at will right after the first boot of our Minimal-VM images. It also makes it easier to install Rancher Kubernetes Engine (RKE)/K3s on our images.

5.15.5.3 KubeVirt

KubeVirt is a technology which enables container-native virtualization. A specific documentation about KubeVirt can be found at

<https://documentation.suse.com/en-us/sbp/all/html/SBP-KubeVirt-SLES15SP3/> ↗.

5.15.5.4 virt-manager

virt-manager has been updated to version 4.0.0

- It is now possible to prefer UEFI when creating new virtual machines. Add an option to allow users to default to UEFI when creating a new VM. libvirt decides which firmware file to use.
- Add `virtiofs` filesystem driver UI option
- Enable a TPM by default when UEFI is used
- Use `virtio-gpu` video for most modern distros
- Default to extra PCIe root ports for q35
- Set `discard=unmap` by default for sparse disks and block devices
- `virt-install`: missing `--os-variant` / `--osinfo` is now a hard error

5.15.5.5 `virt-viewer` has been updated

`virt-viewer` has been updated to version 11.0. Some of the changes in this version are the following:

- Remove clashing `-r` command line shortcut for `resize` that clashed with existing `reconnect` shortcut
- Support modifier-only hotkeys for cursor release
- Add USB device reset hotkey support
- Second display support
- Remapping keys using the `--keymap` argument
- Bash completion for the client

For more information see <https://gitlab.com/virt-viewer/virt-viewer/-/tree/v11.0> ↗.

5.15.5.6 `libguestfs` has been updated to 1.44.2

`libguestfs` has been updated to version 1.44.2. `virt-v2v` and `virt-p2v` have been separated from `libguestfs` into their own packages.

5.15.5.7 VM installer of YaST can no longer install LXC containers

The YaST module for installing VMs (`yast2-vm`) has the following changes:

- As support for libvirt LXC containers has been removed with SUSE Linux Enterprise Server 15 SP4, the option to install the libvirt-daemon-lxc package has been removed.
- As Xen is only supported on x86-64, Xen-related options have been disabled for AArch64.

6 AMD64/Intel 64-specific changes (x86-64)

Information in this section applies to SUSE Linux Enterprise Server 15 SP4 for the AMD64/Intel 64 architectures.

6.1 System-specific and vendor-specific information

6.1.1 User Space Live Patching (ULP) infrastructure and live patches for shared libraries

Complementing the Kernel Live Patching (KLP), SUSE now offers an infrastructure for live patching user-space applications. SUSE has enabled the shared libraries from the system packages glibc and openssl for live patching.

The technology targets patching shared libraries at runtime and is part of the SUSE Linux Enterprise Live Patching extension. The respective packages are libpulp0, the live patching core that must be pre-loaded into the application on start, and libpulp-tools containing the essential tools for building and deployment of patches. Next, there are containers for the future live patches for each library, for example glibc-livepatches for glibc, that will receive the fixes through future maintenance updates.

ULP is currently offered for the x86-64 platform.

See the *Administration Guide* at <https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-ulp.html> for more detailed information.

7 POWER-specific changes (ppc64le)

Information in this section applies to SUSE Linux Enterprise Server for POWER 15 SP4.

7.1 Hardware

7.1.1 IBM Power10 support

On SLES 15 SP4, the Power10 CPU is supported in default mode, which includes performance counters, prefixed instructions, new idle state timings, and MMA unit. Previous SLES releases that support the POWER9 CPU can work on Power10 (POWER9 Compatibility mode). However, new features and performance counters are not supported and the use of idle states might not be optimal.

7.2 Performance

There were the following performance-related changes:

- OpenBlas package for POWER provides Power10 capabilities and includes optimization for matrix multiplication GEMM kernels
- enablement of GZIP engines on PowerVM to send GZIP compression requests directly to NX without kernel involvement

7.2.1 Enhanced libgcrypt cryptography performance

Community contributions improved the performance of GHASH and SHA2 for POWER9 and Power10.

7.2.2 Enhanced OpenSSL cryptography performance

ECC improvements were added for Power10 in OpenSSL 3.0. These are backported to OpenSSL 1.1.1.

7.2.3 Enhanced NSS FreeBL cryptography performance

Community contributions improved the performance of Chacha20 for POWER9 and Power10.

7.3 Security

Also see the following notes:

- Section 2.8.4.1, “[keylime](#) has been added”
- Enable Strict RWX support for IBM Power, improving kernel protection and security.

7.3.1 POWER guest secure boot with static keys

PowerVM LPAR guest secure boot with static keys with verification to extend the chain of trust from partition firmware to the OS kernel and includes key management.

7.3.2 The LPAR security flavor in human-readable format

The LPAR security flavor is available in a human-readable format from inside the LPAR via the `lparstat -x` option.

7.3.3 Key Management Interoperability Protocol (KMIP)

The Key Management Interoperability Protocol (KMIP) C client `libkmip` package from OpenKMIP has been added. KMIP provides a standard protocol for managing keys over the network to automate many key management tasks.

7.4 Virtualization

The following new features are supported in SLES 15 SP4 under PowerVM:

- Provide dump capture support to HNV based dump target.
- Linux Hybrid Network Virtualization (HNV) support for Wicked
- Support IBM vNIC as backend device for Hybrid Network Virtualization (HNV)

Additionally, there are the following virtualization-related notes:

7.4.1 Logical partition migration when using NX coprocessor

Each Power10+ chip has NX coprocessor to support hardware compression. Logical partitions can access to the NX coprocessor with Virtual Accelerator Switchboard (VAS) windows without going through the kernel. Once the VAS window is established, the userspace may use `copy` and `paste` instruction pairs to issue compression requests directly to NX coprocessor. So all VAS windows opened on a coprocessor belongs to a specific PowerPC chip. For the logical partition migration, the hypervisor expects the partition to close all active windows on the sources system and reopen them after migration on the destination machine.

The partition migration support with the NX coprocessor is not included in SLES 15 SP4 but is expected in 15 SP5. That means the logical partition migration can not be used in 15 SP4 if NX is used by applications.

The following command can be used to determine which process is currently using the NX coprocessor:

```
fuser /dev/crypto/nx-gzip
```



Note

Make sure no workload that uses hardware compression is running at the time of partition migration because it is possible that a workload might open VAS windows after the migration is initiated.

7.4.2 Multiqueue support for `ibmvfc` SCSI driver (NPIV)

The `ibmvfc` client can negotiate with the VIOS server adapter the use of multiple queues such that those queues can be exploited by the `blk-mq` / `scsi-mq` in Linux.

7.5 Miscellaneous

There were the following miscellaneous changes:

- Improved management of cached writes for persistent memory devices like NVDIMMs.

- Improve robustness of `fadump` further by isolating `initrd` to capture the vmcore inside the default `initrd` and activating it only during dump capture. Improved the numa locality of workload allocations.

7.5.1 LPM and DLPAR cannot be used with Secure Boot

When Secure Boot is enabled for a logical partitioning (LPAR), the Linux kernel enables lockdown which disables access to kernel memory from userspace. Some Run-Time Abstraction Service (RTAS) services are not available when access to the kernel memory is disabled. Without these RTAS services, RMC connection to LPAR cannot be established. As a result, LPM and dynamic logical partitioning (DLPAR) operation is not possible.

It is expected that in the future a new interface to RTAS that does not require kernel memory access will be provided.

7.5.2 Enhanced mechanism to handle the installer errors

Enhanced mechanism to handle the installer errors and summarize the errors in the installer (a single popup message for everything and a page listing all the details).

7.5.3 Transactional memory is deprecated and disabled

On POWER9, transactional memory is partially emulated by the hypervisor, but this does not give the expected performance.

Therefore, transactional memory is now disabled by default in the kernel. For legacy applications on platforms that still support transactional memory, it can be enabled with the `ppc_tm=on` kernel parameter.

8 IBM Z-specific changes (s390x)

Information in this section applies to SUSE Linux Enterprise Server for IBM Z and LinuxONE 15 SP4. For more information, see

<https://www.ibm.com/docs/en/linux-on-systems?topic=distributions-suse-linux-enterprise-server> ↗

8.1 Hardware

There were the following hardware-related changes:

- support has been added for IBM z15 instructions in Valgrind
- support has been added for IBM z16 instructions in [glibc](#), [gdb](#), and [binutils](#)
- support has been added for IBM z16 in kernel
- added the zDNN library that provides a user space API for exploitation of the Neural Network Processing Assist Facility of the IBM z16

8.2 Networking

8.2.1 [zdsfs](#): transparent dataset conversion

Enabled zdsfs to read and write EBCDIC-encoded data sets as ASCII and read data sets in the same format as resulting from an FTP transfer from z/OS to Linux (including record translations).

8.2.2 [zipl](#): implemented environment block

Introduces new tool zipl-editenv that allows a Linux on Z user to specify persistent configuration information that is evaluated during boot without the need to rewrite IPL records.

8.2.3 PCI auto-activate for Dynamic Partition Manager

Allows a Linux on Z user to automatically use any PCI function defined for an LPAR on Dynamic Partition Manager without the need to manually configure the PCI function online.

8.2.4 SMC-Rv2 support

Lifts the restriction of traffic limited to be within a single IP subnet only.

8.2.5 SMC: statistics support

Adds statistics for traffic run across RoCE (RDMA) and ISM devices.

8.2.6 SMC: user-defined EID (Enterprise ID) support

Adds a tool to display and set EIDs (SMC Enterprise IDs).

8.2.7 wireshark: updated to include SMC-D support

Provides support for SMC-R, SMC-D and SMC-Dv2 in wireshark.

8.2.8 HSCI (HiperSockets Converged Interface): multi-MAC support

Enhances HSCI to support multiple MAC Addresses as required by Open vSwitch, as well as the corresponding tool for exploitation.

8.2.9 RoCE: predictable interface names

Up to SLES 15 SP3:

- Interface names for RoCE Express adapters were very hard to predict
- Interface names could change between re-boots, invalidating any previously stored network card configuration To fix this, changes were made in the Linux kernel to indicate whether UIDs are unique to have systemd generate easy to predict interface names on (preferably) UID or FID.

8.3 Performance

- Nettle-Community contributions improved the performance of the AES-GCM algorithm for POWER9 and Power10.

8.3.1 Improved performance on RoCE ConnectX-4 hardware

Fixed performance problem for which the workaround was described in the Release Notes of earlier SLES 15 service packs.

8.4 Security

8.4.1 zcrypt

There were the following zcrypt-related changes:

- provide indications to early exploiters of crypto adapters (e.g. dm-crypt root devices using the PAES cipher) that ap bus initialization and DD bindings are complete
- AP bus and zcrypt device driver uevent extensions that provide uevents for the following AP bus events: online state change, config state change, add crypto mode events
- toleration for new IBM Z crypto hardware

8.4.2 openCryptoki

There were the following openCryptoki-related changes:

- ep11 token: support generation of attribute bound keys and operations with attribute bound keys
- ep11 token: protected key support
- event notification support
- cca token: support the interchange of CCA secure key objects (as generated by the CCA library) between code using the CCA library and openCryptoki
- p11sak: extended to display vendor specific boolean attributes in the long listing and enablement for configuration to learn about additional (boolean) key attributes

8.4.3 zkey

There were the following zkey-related changes (s390-tools): - extended LUKS2 functionality - integration of the zkey repository into an enterprise key mangement system with a KMIP interface

8.4.4 libica

Eliminated implementations of software fallback functions and replaced them by calls to openssl/libcrypto.

8.4.5 [openssl-ibmca](#)

Made [openssl-ibmca](#) engine call libica w/o software fallbacks. Only register [openssl-ibmca](#) functions if [libica](#) signals the existence of a hardware function.

8.4.6 [pkey](#)

Add protected key support for private ECDSA/EdDSA keys.

8.4.7 [libzpc](#)

Added new library to support protected key cryptography: [libzpc](#) – IBM Z Protected-key Cryptography

8.5 Storage

8.5.1 [zfcp](#): handling of firmware update notifications

Enhanced user information of the FCP device driver about HBA firmware version to improve handling of firmware update notifications.

8.5.2 Multi-path re-IPL

List-Directed IPL (for FCP etc.) was restricted to a single FCP-WWPN-LUN path. If this path is unavailable, (re)-IPL fails. This change implements a solution to keep the path to re-IPL up to date, and therefore work around transient path failures in many cases.

8.6 Virtualization

The following new features are supported in SUSE Linux Enterprise Server 15 SP4 under KVM:

8.6.1 [Fix for "Host does not support any virtualization options"](#)

When attempting to run [virt-install](#) on SUSE Linux Enterprise Server 15 SP4 for IBM Z the command may fail with the error "Host does not support any virtualization options". Install the latest qemu package from the update repository to fix this problem.

8.6.2 Provide persistent vfio-ccw device assignments

Establish persistent information about CCW devices intended to be passed through to KVM guests.

8.6.3 Added CPU model for IBM z16

Enable architectural features of the IBM z16 for KVM guests.

8.6.4 Change Secure Execution header defaults for Plaintext Control Flags (PCF)

To improve usability the default SE header PCF settings are now set to allow all PCKMO types. An explicit option has been added to enable/disable PCKMO, so that clients have no need to use the "experimental/expert" flags.

8.6.5 Secure guest indication

Provides an indication in the guest that it is running securely. Cannot replace a real attestation and does not really provide additional security (or could even create the false impression of security), but has been frequently requested by customers.

8.6.6 Enabled vfio-ccw and vfio-ap in virt-* tools in virt-manager

The tools in the virt-manager package, most prominently virt-install and virt-xml, are now aware of the IBM Z specific virtio types. Therefore, it's now possible to install a VM with passed-through DASDs or APQNs.

8.7 Miscellaneous

8.7.1 SCLP (Service-Call Logical Processor) extended length SCCBs

Enable support for machines with more than 256 CPUs.

8.7.2 Improved CPU-MF counter set extraction performance

Performance improvement through reading out complete counter sets with a single instruction and export them to user space without sampling involved.

9 Arm 64-bit-specific changes (AArch64)

Information in this section applies to SUSE Linux Enterprise Server for Arm 15 SP4.

9.1 System-on-Chip driver enablement

SUSE Linux Enterprise Server for Arm 15 SP4 includes driver enablement for the following System-on-Chip (SoC) chipsets:

- AMD* Opteron* A1100
- Ampere* X-Gene*, eMAG*, Altra*, *Altra Max**
- AWS* Graviton, Graviton2, **Graviton3**
- Broadcom* BCM2837/BCM2710, BCM2711
- Fujitsu* A64FX
- Huawei* Kunpeng* 916, Kunpeng 920
- Marvell* ThunderX*, ThunderX2*, ThunderX3*; OCTEON TX*; Armada* 7040, Armada 8040
- NVIDIA* Tegra* X1, Tegra X2, Xavier*; BlueField*, BlueField-2
- NXP* i.MX 8M, 8M Mini; Layerscape* LS1012A, LS1027A/LS1017A, LS1028A/LS1018A, LS1043A, LS1046A, LS1088A, LS2080A/LS2040A, LS2088A, LX2160A
- Qualcomm* Centriq* 2400
- Rockchip RK3399
- Socionext* SynQuacer* SC2A11
- Xilinx* Zynq* UltraScale*+ MPSoC

**Note**

Driver enablement is done as far as available and requested. Refer to the following sections for any known limitations.

Some systems might need additional drivers for external chips, such as a Power Management Integrated Chip (PMIC), which may differ between systems with the same SoC chipset.

For booting, systems need to fulfill either the Server Base Boot Requirements (SBBR) or the Embedded Base Boot Requirements (EBBR), that is, the Unified Extensible Firmware Interface (UEFI) either implementing the Advanced Configuration and Power Interface (ACPI) or providing a Flat Device Tree (FDT) table. If both are implemented, the kernel will default to the Device Tree; the kernel command line argument `acpi=force` can override this default behavior.

Check for SUSE YES! certified systems, which have undergone compatibility testing.

9.2 New features

9.2.1 Uacce support

Uacce (Unified/User-space-access-intended Accelerator Framework) aims to provide Shared Virtual Addressing (SVA) between accelerators and processes.

There are more and more heterogeneous processors, such as encryption/decryption accelerators, TPUs, or EDGE processors. The intention of Uacce is to make sure the accelerator and process can share the same address space, so the accelerator ISA can directly address any data structure of the main CPU. This differs from the data sharing between CPU and IO device, which share data content rather than address.

9.2.2 Support execute-only permissions with Enhanced PAN on ARMv8.7

Enhanced Privileged Access Never (EPAN) allows Privileged Access Never to be used with Execute-only mappings. The feature is detected at runtime, and will remain disabled if the CPU does not implement the feature.

9.2.3 OpenSSL 3 improvements backported to OpenSSL 1.1.1k

OpenSSL 3 contains performance improvements that are beneficial to Arm architectures.

This patchset includes:

- Optimize RSA on Armv8 (A72 and N1) [\[1\]](#)
- Optimize AES-XTS mode in OpenSSL for AArch64 [\[2\]](#)
- Optimize AES-GCM for microarchitectures with unroll and new instructions [\[3\]](#)

9.3 Known limitations

9.3.1 No DisplayPort graphics output on NXP LS1028A and LS1018A

The NXP* Layerscape* LS1028A/LS1018A System-on-Chip contains an Arm* Mali*-DP500 Display Processor, whose output is connected to a DisplayPort* TX Controller (HDP-TX) based on Cadence* High Definition (HD) Display Intellectual Property (IP).

A Display Rendering Manager (DRM) driver for the Arm Mali-DP500 Display Processor is available as technology preview ([Section 2.8.2.5, “mali-dp driver for Arm Mali Display Processors available”](#)).

However, there was no HDP-TX physical-layer (PHY) controller driver ready yet. Therefore no graphics output will be available, for example, on the DisplayPort* connector of the NXP LS1028A Reference Design Board (RDB).

Contact the chip vendor NXP for whether third-party graphics drivers are available for SUSE Linux Enterprise Server for Arm 15 SP4.

Alternatively, contact your hardware vendor for whether a bootloader update is available that implements graphics output, allowing to instead use efifb framebuffer graphics in SUSE Linux Enterprise Server for Arm 15 SP4.



Note

The Vivante GC7000UL GPU driver ([etnaviv](#)) is available as a technology preview ([Section 2.8.2.3, “etnaviv drivers for Vivante GPUs are available”](#)).

9.4 Removal of NXP Layerscape LX2160A rev. 1 silicon support

NXP* Layerscape* LX2160A System-on-Chip silicon revision 1.0 differs from revision 2.0 in the PCIe controller (Mobiveil based vs. Synopsys DesignWare* based respectively).

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel supported the PCIe controllers in both silicon revisions of NXP* Layerscape* LX2160A SoC.



Note

The bootloader of the system may need to detect the chip revision and to patch the Device Tree to pass the right `compatible` string to the kernel:

- `fsl, lx2160a-pcie` for rev. 1.0 silicon,
- `fsl, ls2088a-pcie` for rev. 2.0 silicon.

To verify which one has been passed to the kernel, you can check the DT nodes:

```
cat /sys/firmware/devicetree/base/soc/pcie@3400000/compatible
```

Deprecated with SUSE Linux Enterprise Server for Arm 15 SP3, SP4 removes the support for rev. 1.0 silicon by dropping patches from the kernel. This will now result in failure to boot on rev. 1.0 silicon due to a kernel panic (`SError` interrupt request).

This affects among others the original NXP Layerscape LX2160A Reference Design Board; the RDB revision B uses rev. 2.0 silicon.



Note

To check whether an LX2160A SoC-based machine will be affected by this, read the chip revision from its kernel:

```
cat /sys/bus/soc/devices/soc0/revision
```

If this prints `1.0`, your system is affected; if it prints `2.0`, it is not.

10 Removed and deprecated features and packages

This section lists features and packages that were removed from SUSE Linux Enterprise Server or will be removed in upcoming versions.



Note: Package and module changes in 15 SP4

For more information about all package and module changes since the last version, see [Section 2.2.3, "Package and module changes in 15 SP4"](#).

10.1 Removed features and packages

The following features and packages have been removed in this release.

- Vagrant boxes of SUSE Linux Enterprise Server have been discontinued.
- `bind-chroot` has been removed. See [Section 5.10.2, "`bind-chrootenv` has been removed, functionality replaced with `systemd`"](#).
- `bzr` has been removed. See [Section 5.6.3.4, "`bzr` has been replaced with `breezy`"](#).
- MariaDB 10.5 has been removed.
- PostgreSQL 12 and earlier have been removed.
- The `raw` application of the `util-linux` package has been removed. See [Section 5.2.1, "`util-linux` has been updated"](#).
- `nodejs12` has been removed.
- The `udev` package has been removed as a dependency of the `rpm` package.
- The `imgen` package, containing Mellanox firmware generator, has been removed.
- The OpenLDAP server has been removed. For more information, see [Section 5.1.2, "389 Directory Server is the primary LDAP server, the OpenLDAP server has been removed"](#).

- Python 2 has been removed entirely from SLE with SLE 15 SP4 and is no longer be available via the *Python 2* SLE module. For more information, see [Section 5.6.1.4, “Python 2 has been removed”](#).
- NXP LX2160A revision 1 silicon quirks have been removed. For more information, see [Section 9.4, “Removal of NXP Layerscape LX2160A rev. 1 silicon support”](#).
- Support for libvirt LXC containers has been removed. For more information, see [Section 5.15.5.7, “VM installer of YaST can no longer install LXC containers”](#).
- System containers using LXC have been removed. For more information, see [Section 5.3.8, “LXC containers have been removed”](#).
- `pam_ldap` and `nss_ldap` have been removed. Use SSSD instead.
- `libmodman` has been removed. `libproxy` is now shipped with an internal `modman`.
- Eclipse has been removed.
- `systemd-journal-remote` has been removed.

10.2 Deprecated features and packages

The following features and packages are deprecated and will be removed in a future version of SUSE Linux Enterprise Server.

- PostgreSQL 13 has been deprecated and moved to the Legacy module.
- TLS 1.0 and 1.1 are deprecated and will be removed in a future service pack of SUSE Linux Enterprise Server 15. For more information, see [Section 5.11.1, “TLS 1.1 and 1.0 are no longer recommended for use”](#).
- OSN support on IBM Z has been deprecated.
- The `mkinitrd` wrapper has been replaced with `dracut` everywhere and will be removed in the next major version of SUSE Linux Enterprise Server.
- The `lftp_wrapper` package has been deprecated and will be removed in the near future. It is still available as an `update-alternative` for `ftp`, but it is no longer used by default. The default implementation of `ftp` is now the `lftp` executable.

- Support for System V `init.d` scripts is deprecated and will be removed with the next major version of SUSE Linux Enterprise Server. In consequence, the `/etc/init.d/halt.local` `initscript`, `rcSERVICE` controls, and `insserv.conf` are also deprecated. For more information, see [Section 5.14.12, “Support for System V `init.d` scripts is deprecated”](#).
- `lftp_wrapper` is deprecated. Use `lftp` directly instead.
- On the POWER architecture, transactional memory is deprecated. For more information, see [Section 7.5.3, “Transactional memory is deprecated and disabled”](#).
- The `opa-fmgui` package is not maintained upstream anymore. It has been deprecated, moved to the Legacy module, and will be removed in a future service pack.
- The `thunderbolt-user-space` package does not work properly with a later revision of the TBT hardware. For this reason, SLES 15 SP4 now includes the `bolt-tools` package which can work with both new and old TBT hardware. The `thunderbolt-user-space` package will be removed in SLES 15 SP5 to allow time for customers to adapt.
- NIS is deprecated and will be removed with the next major version of SUSE Linux Enterprise Server. This includes packages implementing NIS, like `ypserv`. NIS code will be removed from SUSE tools and all NIS client code will be dropped with the next major version of SUSE Linux Enterprise Server.

Also see the following release notes elsewhere:

- [Section 5.10.1, “`set-hostname` has been deprecated”](#)
- [Section 5.6.3.9, “PHP 8 has been added, replacing PHP 7”](#)
- [Section 5.12.7, “XFS V4 format file systems have been deprecated”](#)
- [Section 5.10.3.2, “Samba Active Directory Domain Controller has been deprecated”](#)
- [Section 5.10.3.3, “SMB1 support has been deprecated”](#)
- [Section 5.10.3.4, “NT Domain Controller has been deprecated”](#)
- [Section 5.10.3.5, “Samba NIS support has been deprecated”](#)

10.2.1 sched_ parameters are deprecated

The following scheduler parameters have been moved to debugfs and will be removed in the future. They have been moved from /proc/sys/kernel/sched_* to /sys/kernel/debug/sched/:

- sched_latency_ns
- sched_migration_cost_ns
- sched_min_granularity_ns
- sched_nr_migrate
- sched_tunable_scaling
- sched_wakeup_granularity_ns

The sysctls still exist but a deprecation warning will be logged and there is no guarantee that either the sysctls or debugfs options will exist in a future SLE release due to changes in the CPU scheduler implementation.

10.2.2 Berkeley DB removed from packages

Berkeley DB, used as a database in certain packages, is dual-licensed under GNU AGPLv3/Sleepycat licenses. Because service vendors that redistribute our packages could find packages with these licenses potentially detrimental to their solutions, we have decided to remove Berkeley DB as a dependency from these packages. In the long term, SUSE aims to provide a solution without Berkeley DB.

This change affects the following packages:

- apr-util
- cyrus-sasl
- iproute2
- perl
- php7
- postfix
- rpm

11 Obtaining 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 <https://www.suse.com/products/server/download/> on Medium 2. For up to three years after distribution of the SUSE product, upon request, SUSE will mail a copy of the source code. Send requests by e-mail to sle_source_request@suse.com. SUSE may charge a reasonable fee to recover distribution costs.

12 Legal notices

SUSE makes no representations or warranties with regard 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 regard 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. Refer to <https://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 © 2010–2024 SUSE LLC.

This release notes document is licensed under a Creative Commons Attribution-NoDerivatives 4.0 International License (CC-BY-ND-4.0). You should have received a copy of the license along with this document. If not, see <https://creativecommons.org/licenses/by-nd/4.0/>.

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 <https://www.suse.com/company/legal/> and one or more additional patents or pending patent applications in the U.S. and other countries.

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

[1] <https://git.openssl.org/?p=openssl.git;a=commit;h=10646160125ac1328d892f1dd27f2847892d33c5>

[2] <https://git.openssl.org/?p=openssl.git;a=commit;h=9ce8e0d17e608de4f85f7543c52b146e3c6a2291>

[3] <https://github.com/openssl/openssl/pull/15916>

Appendix A Changelog for 15 SP4