



SUSE Linux Enterprise Server 15 SP3

Release Notes

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 SP3 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/releasenotes> . General documentation can be found at <https://documentation.suse.com/sles/15-SP3> .

Publication Date: 2022-09-30, Version: 15.3.20220930

Contents

- 1 About the release notes 3
- 2 SUSE Linux Enterprise Server 3
- 3 Modules, extensions, and related products 14
- 4 Installation and upgrade 16
- 5 Changes affecting all architectures 20
- 6 AMD64/Intel 64-specific changes (x86-64) 59
- 7 POWER-specific changes (ppc64le) 59
- 8 IBM Z-specific changes (s390x) 61
- 9 Arm 64-bit-specific changes (AArch64) 67
- 10 Removed and deprecated features and packages 74

- 11 Obtaining source code 77
- 12 Legal notices 77
- A Kernel parameter changes 78

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/releasenotes> .

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 SP3 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 SP3, 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. A single subscription for SLES allows for running an unlimited number of SLES virtual machines per physical system. 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-SP3/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. This functionality makes it easier for administrators and system architects to find the software packages needed. To do so, it leverages the SUSE Customer Center.

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.

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 SP3 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.

Vagrant

SLES 15 SP3 and SLED 15 SP3 will be available as a Vagrant boxes. For more information, see [Section 5.13.5.1, “Vagrant boxes for SUSE Linux Enterprise Server”](#).

MAJOR UPDATES TO THE SOFTWARE SELECTION:

Salt

SLE 15 SP3 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 SP3

SUSE Linux Enterprise Server 15 SP3 introduces changes compared to SUSE Linux Enterprise Server SP2. The most important changes are listed below:

- xca (X Certificate and Key Management) has been added as the new Certificate Authority (CA) management tool. For more information, see [Section 5.12.6, “xca has been added”](#).
- You can now use Podman without root privileges for enhanced security. For more information, see [Section 5.3.3, “Rootless containers”](#).

2.2.3 Package and module changes in 15 SP3

The full list of changed packages and modules compared to 15 SP2 can be seen at these two URLs:

- https://documentation.suse.com/package-lists/sle/15-SP3/package-changes_SLE-15-SP2-GA_SLE-15-SP3-GA.txt ↗
- https://documentation.suse.com/package-lists/sle/15-SP3/module-changes_SLE-15-SP2-GA_SLE-15-SP3-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 SP3 has been submitted to the certification bodies for:

- Common Criteria Certification, see <https://www.commoncriteriaportal.org/> ↗
- FIPS 140-2 validation, see <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140InProgress.pdf> ↗

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 SP3. This directory includes PDF versions of the SUSE Linux Enterprise Server 15 SP3 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 SP3, see <https://documentation.suse.com/sles/15-SP3>.
- Find a collection of White Papers in the SUSE Linux Enterprise Server Resource Library at <https://www.suse.com/products/server#resources>.

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 (SP3) will be fully maintained and supported until six months after the release of SUSE Linux Enterprise Server 15 SP4.

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 \(https://www.suse.com/support/policy.html\)](https://www.suse.com/support/policy.html) and [Long Term Service Pack Support \(https://www.suse.com/support/programs/long-term-service-pack-support.html\)](https://www.suse.com/support/programs/long-term-service-pack-support.html).

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 5.7, “Kernel”](#)
- [Section 5.11, “Storage and file systems”](#)
- [Section 5.13, “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 `zypper`.

Major packages and groups of packages affected by this are:

- PostgreSQL (all versions, including all subpackages)
 - Additional contract needed: SUSE is providing the package including its dependencies with support based on upstream projects. For L3 support a separate contact with SUSE or another Third-Party is required.
- `keepalived`

2.7.3 Software under GNU AGPL

SUSE Linux Enterprise Server 15 SP3 (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 SP3 (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 SP3.

2.8.1.2 Technology previews for Arm 64-Bit (AArch64)

2.8.1.2.1 64K page size kernel flavor has been added

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 adds 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.



Important: Btrfs file system uses page size as block size

It is currently not possible to use Btrfs file systems across page sizes. Block sizes below page size are not yet supported and block sizes above page size might never be supported.

During installation, change the default partitioning proposal and choose another file system, such as Ext4 or XFS, to allow rebooting from the default 4K page size kernel of the Installer into `kernel-64kb` and back.

See the *Storage Guide* for a discussion of supported file systems.



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 not yet possible to configure stripe size on volume creation. This will lead to sub-optimal performance if page size and block size differ.

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 SP3 kernels on x86-64 use 4K page size.

The SUSE Linux Enterprise Server for POWER 15 SP3 kernel uses 64K page size.

2.8.1.3 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 kernel includes updated 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 SP3 **on** NVIDIA BlueField-2 DPUs.

For an NVIDIA BlueField-2 DPU PCIe card inserted as **SmartNIC** into a SUSE Linux Enterprise Server 15 SP3 or SUSE Linux Enterprise Server for Arm 15 SP3 based server, check [Section 2.8.1, “Technology previews for all architectures”](#) and [Section 5.7, “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 5.14.2, “Important package additions to SUSE Package Hub”](#)).

2.8.1.4 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 SP3 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.

2.8.1.5 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.

2.8.1.6 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 9.3.4, “No DisplayPort graphics output on NXP LS1028A and LS1018A”](#)).

2.8.1.7 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.

3 Modules, extensions, and related products

This section comprises information about modules and extensions for SUSE Linux Enterprise Server 15 SP3. Modules and extensions add functionality to the system.



Note: Package and module changes in 15 SP3

For more information about all package and module changes since the last version, see [Section 2.2.3, “Package and module changes in 15 SP3”](#).

3.1 Modules in the SLE 15 SP3 product line

The SLE 15 SP3 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 SP3 at the release of SUSE Linux Enterprise Server 15 SP3 are listed in the *Modules and Extensions Quick Start* at <https://documentation.suse.com/sles/15-SP3/html/SLES-all/article-modules.html>.

Not all SLE modules are available with a subscription for SUSE Linux Enterprise Server 15 SP3 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/releasenotes>.


The following extensions are available for SUSE Linux Enterprise Server 15 SP3:







- 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 5.14, “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, “JeOS: Just Enough Operating System”*)
- 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 SP3.



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-SP3/html/SLES-all/book-deployment.html> .

Also see the following additional notes:

- *Section 5.13.6.1, “High video resolutions in VMware ESXi need more VRAM”*

4.1.1 Migration failure due to `insserv-compat`

If `insserv-compat` is not installed prior to migration and other software has already created the `/etc/rc.d/` directory, the migration fails. As a workaround, remove the `/etc/rc.d/` directory and continue with migration.

4.1.2 YaST will warn when the `root` account is set up with an SSH key only but SSH access is unavailable

With its default settings, the SLES installer blocks access via SSH. However, during the installation of SLES, you can enable login via SSH key for the `root` user, either exclusively or as an alternative to a password. Combining the default settings with exclusive SSH key login, you can effectively lock yourself out.

Starting with SLES 15 SP3, the page *Installation Summary* will display a warning if the `root` user will not be able to log in after installation.

4.1.3 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 SP3.



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-SP3/html/SLES-all/cha-upgrade-online.html>.

4.2.1 Migration procedure to openSUSE Leap has changed

The migration procedure between SUSE Linux Enterprise and openSUSE Leap has changed. For more information, see the *Upgrade Guide* at <https://documentation.suse.com/sles/15-SP3/html/SLES-all/cha-upgrade-online.html#sec-upgrade-online-opensuse-to-sle>.

4.2.2 Differences between AutoYaST profiles in SLES 12 and 15

Significant changes in SLES 15 required changes in AutoYaST. If you want to reuse existing SLES 12 profiles with SLES 15, you need to adjust them as documented in <https://documentation.suse.com/sles/15-SP2/html/SLES-all/appendix-ay-12vs15.html>.

4.2.3 Upgrading glibc can cause issues in some software

For more information see *Section 5.5.7, “Package compat-libpthread-nonshared has been added”*.

4.2.4 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.

4.2.5 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 SP2 before upgrading to SLE 15 SP3.

4.3 JeOS: Just Enough Operating System

SUSE Linux Enterprise Server JeOS 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 JeOS, you can choose the right-sized SUSE Linux Enterprise Server option to fit your needs.

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

4.3.1 Removing the locale warning from `jeos-firstboot`

With SLES JeOS 15 SP1, the dialog for choosing the system locale was replaced by a warning dialog. It explained about `en_US` being the only locale available and provided instructions on how to change the locale after the first boot. On SLES JeOS 15 SP3, this dialog has been removed. Instructions on how to change the locale are provided by the [JeOS Quick Start Guide \(https://documentation.suse.com/sles-15/html/SLES-all/article-jeos.html\)](https://documentation.suse.com/sles-15/html/SLES-all/article-jeos.html).

4.3.2 JeOS KVM image is available for aarch64

In addition to the SLES JeOS 15 SP3 for KVM on x86_64, we are now providing the same image for aarch64.

4.4 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 SP3.

5.1 Authentication

5.1.1 389 Directory Server is the primary LDAP server, the OpenLDAP server is deprecated

The OpenLDAP server (package `openldap2`, part of the *Legacy SLE* module) is deprecated and will be 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 \(https://documentation.suse.com/sles-15/html/SLES-all/cha-security-ldap.html\)](https://documentation.suse.com/sles-15/html/SLES-all/cha-security-ldap.html) .

5.2 Basic utilities

5.2.1 `at` user's default shell has been set to `/bin/false`

The default shell of the user used by the job manager application `at` was set to `/bin/bash`.

That is considered to be against security best practices. In SUSE Linux Enterprise Server 15 SP3, its default shell is now set to `/bin/false`.

5.2.2 Bash is now available under `/usr/bin/bash`

The Bash is now available at both of the following paths: `/usr/bin/bash` and `/bin/bash`. This is part of the `/usr` merge initiative and provides compatibility with openSUSE Tumbleweed. For more information, see the [the openSUSE wiki \(https://en.opensuse.org/openSUSE:Usr_merge\)](https://en.opensuse.org/openSUSE:Usr_merge) .


5.3 Containers

Also see the following additional note:

- *Section 5.5.5, “Web and Scripting Module: NodeJS 14 has been added, NodeJS 8 and NodeJS 10 has been removed”*

5.3.1 SLE BCI minimal and micro container images

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 and micro container images based on the SUSE BCI (Base Container Image) have been made available. See the [SUSE registry \(https://registry.suse.com/\)](https://registry.suse.com/)  for more information.



Note

The minimal 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

The micro container does not include the zypper nor the rpm packages.

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:

- Ruby
- Python

- Node
- Go
- Java

See the [SUSE registry \(https://registry.suse.com/\)](https://registry.suse.com/) for more information.

5.3.3 Rootless containers

By default, Podman requires root privileges.

You can use Podman without root privileges for enhanced security. For more information, see <https://susedoc.github.io/doc-sle/main/single-html/SLES-container/#cha-podman-install>.

5.3.4 LXC containers have been deprecated

System containers using LXC have been deprecated and will be 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.

5.3.5 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

Also see the following additional notes:

- *Section 5.5.7, “Package compat-libpthread-nonshared has been added”.*

5.4.1 `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> )

5.4.2 The ODBC driver location has changed

Previously in SLES 12, the `postgresql10-odbc` package was located in `/usr/pgsql-10/lib/psqlodbcw.so`. In SLES 15 SP3, the `psqlODBC-10` package is located in `/usr/lib64/psqlodbcw.so`.

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


5.4.3 PostgreSQL 13 has been added

PostgreSQL 13 has been added to SUSE Linux Enterprise Server. For information about changes between PostgreSQL 13 and 12, see the upstream release notes:

- <https://www.postgresql.org/docs/13/release-13.html> 



Warning: REINDEX is required

If you migrate a PostgreSQL server to 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.

PostgreSQL 10 is deprecated and has been moved to the Legacy module.

5.4.4 PostgreSQL JDBC Driver has been added

The PostgreSQL JDBC Driver has been added. This includes the following packages:

- `jdbc-postgresql-42.2.16`
- `postgres-scram-1.0.0-beta.2`

5.4.5 MariaDB has been updated to version 10.5

The `mariadb` package has been updated to 10.5. For more information about upgrading from 10.4 to 10.5, see <https://mariadb.com/kb/en/upgrading-from-mariadb-104-to-mariadb-105/> .

5.5 Development

5.5.1 nodejs - common has been updated

The `nodejs - common` package has been updated. It provides these sub-packages:

- `nodejs-default`
- `npm-default`
- `nodejs-devel-default`

In SUSE Linux Enterprise Server 15 SP3, the NodeJS version of these subpackages is set to `nodejs14`.

5.5.2 python-kubernetes has been added

The `python-kubernetes` package has been added. It is the official Python client library for Kubernetes.

5.5.3 erlang has been updated to version 22.3

The `erlang` package has been updated to version 22.3.

For more information, see <https://www.erlang.org/news/137> .

5.5.4 `rpcgen` has been moved from `glibc-devel` to its own package

`rpcgen` has been removed from `glibc-devel`.

As a replacement, the `rpcgen` package has been added.

5.5.5 Web and Scripting Module: NodeJS 14 has been added, NodeJS 8 and NodeJS 10 has been removed

NodeJS 8 (package `nodejs8`) and NodeJS 10 (package `nodejs10`) have been removed from the SLE Module *Web and Scripting*. NodeJS 14 (package `nodejs14`) has been added to the module.

5.5.6 New Python modules: `python3-kerberos`, `python-cassandra-driver`, and `python-arrow` have been added

The following new Python modules have been added as packages:

- `python3-kerberos` is a Python Kerberos module that is available in addition to `python-krb5`. Both modules provide the same `.so` objects and cannot coexist.
- `python3-cassandra-driver` can initialize tables in Apache Cassandra
- `python3-arrow` handles timestamps

5.5.7 Package `compat-libpthread-nonshared` has been added

A `glibc` package update in SLES 15 SP3 caused some enterprise software to fail due to the missing `libpthread_nonshared.a` file. This includes the products *Oracle Database* and *Oracle Forms & Reports*.

The newly provided `compat-libpthread-nonshared` package enables applications that directly reference `libpthread_nonshared.a` to work properly.

5.5.8 `librabbitmq` has been added

The package `librabbitmq` v0.10.0 has been added. It is C-language AMQP client library for use with the RabbitMQ broker.

5.5.9 Support for Python 3.9 has been added

Support for Python version 3.9 has been added. Right now, this is only an interpreter, including pip and setuptools.

This is in addition to the system-default Python 3.6 that has already been present and continues to be available. All SLE python3-* packages are only verified to be compatible with the system Python.

5.5.10 glibc has been updated to version 2.31

The glibc package has been updated to version 2.31. For more information about changes see <https://www.gnu.org/software/libc/>.

5.5.11 Python 2 is deprecated

The python executable is only provided via the *Python 2* module, not via the default repositories.

With SUSE Linux Enterprise Server 15 SP1, SUSE has started to phase out support for Python 2 in SLE. Within the standard distribution, only Python 3 (executable name python3) is available. Python 2 (executable names python2 and python) is only provided via the *Python 2* SLE module. This module is disabled by default and will be removed entirely starting with SLE 15 SP4.

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.

To run Python 2 scripts, make sure to enable the Python 2 module and install the package python.

5.5.12 Supported Java versions

The following Java implementations are available in SUSE Linux Enterprise Server 15 SP3:

Name (Package Name)	Version	Module	Support
OpenJDK	11	Base System	SUSE, L3, until 2026-12-31

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

5.6 Hardware

5.6.1 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>.

5.7 Kernel

Also see the following:

- [Appendix A, Kernel parameter changes](#)

5.7.1 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.

5.7.2 No firmware reserved region can cover this RMRR

You can see the above message on systems with BIOS. This is not an OS-specific issue. Currently, we are waiting for the BIOS vendor to provide a fix.

5.7.3 Kernel module compression

Kernel module files are now stored in compressed form. As a result, the kernel package storage footprint is almost halved. The module file extension has changed from `.ko` to `.ko.xz` and the content is LZMA-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.

5.7.4 New scheduler preemption mode switch

Until recently, the process scheduler preemption mode could be selected only in the build configuration. This SUSE Linux Enterprise Server release brings the possibility to choose voluntary preemption mode via a kernel command line option. The exact option is `preempt=<value>` and the value can be either `none` (the default) or `voluntary`. Note that `preempt=voluntary` changes the system performance characteristics and performance degradations observed in this mode may be excluded from SUSE support guarantees.

5.7.5 Pstore block oops/panic logging

Oops/panic logs can now be saved to a block or a non-block device before the system crashes. After a reboot, they can be retrieved from the `pstore` file system. The kernel modules responsible for this are `mtdpstore` and `pstore_blk`. For more information, see the documentation file `/usr/src/linux-KERNEL_VERSION/Documentation/admin-guide/pstore-blk.rst` from the `kernel-source` package.

5.7.6 NVMe-oF TCP Support

Starting from SLES 15 SP2, NVMe-oF TCP is supported and validated by several partners such as EMC and NetApp.

5.7.7 RLIMIT_NOFILE has been increased

The Linux kernel's default `RLIMIT_NOFILE` hard limit, `fs.file-max`, and `fs.nr_open` have been increased by a newer version of `systemd`. The primary reason is to allow to serve more files without an administrator intervention. The `RLIMIT_NOFILE` soft limit has to be increased explicitly to benefit from this change. Controlling the maximum number of file descriptors that can be opened by a process is therefore simplified and only the `RLIMIT_NOFILE` hard and soft limits need to be considered by a process.

Note that `select(2)` is not safe to be used with the increased soft limit. For more information, see <https://github.com/openSUSE/systemd/blob/SLE15-SP3/NEWS#L2084>.

5.7.8 Support for Goya deep learning inference hardware

The Linux kernel in SLES 15 SP3 now supports Habana Labs Goya AI Processor (AIP) PCIe cards that are designed to accelerate deep learning inference and training workloads.

5.7.9 util-linux has been updated

The `util-linux` package has been updated to version 2.36.2. For more information about the changes see <https://www.kernel.org/pub/linux/utils/util-linux/v2.36/v2.36-ReleaseNotes>.

5.7.10 Kernel limits

This table summarizes the various limits which exist in our recent kernels and utilities (if related) for SUSE Linux Enterprise Server 15 SP3.

SLES 15 SP3 (Linux 5.3)	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 Networking

5.8.1 All users can now create ICMP Echo sockets

The Linux kernel parameter `net.ipv4.ping_group_range` now covers all groups. This allows all users of the operating system create ICMP Echo sockets without using `setuid` binaries or needing to have the `CAP_NET_ADMIN` and `CAP_NET_RAW` file capabilities. This improves the overall security by using ICMP instead of raw sockets and simplifies configuration of tools like `ping`, `fping`, `traceroute`, `prometheus-blackbox_exporter`, `collectd-ping`, etc.

5.8.2 `nftables` backend in `firewalld`

`firewalld` now supports `nftables` as a firewall backend. `nftables` is a replacement for `iptables` that brings many advantages, such as built-in sets, faster rule updates, and combined IPv4/IPv6 processing.

For more information, see <https://firewalld.org/2018/07/nftables-backend> ↗.

5.8.3 WireGuard userland tools have been added

The package `wireguard-tools` version 1.0.20200827 has been added. It contains userland tools for the kernel WireGuard module.

WireGuard is a secure, fast, and easy-to-use VPN that uses modern cryptography. For more information, see <https://www.wireguard.com> ↗.

5.8.4 NetworkManager not supported for server workloads

NetworkManager is only supported for desktop workloads with SLED or Workstation Extension. All server certifications are done with `wicked` as the network configuration tool and using NetworkManager might render them invalid. NetworkManager is not supported for server workloads. NetworkManager might be removed from the server products in a future release.

5.8.5 RFC2132 DHCP without MAC address

Certain environments, for example, Microsoft Active Directory, require DHCP requests in the RFC2132 format. `linuxrc`, as shipped with previous versions of SUSE Linux Enterprise Server, required passing MAC address as an argument to get RFC2132-formatted DHCP. This could pose a maintenance issue when managing large numbers of machines.

`linuxrc` can now send RFC2132-formatted DHCP requests without providing MAC address.

5.8.6 Samba

The version of Samba shipped with SUSE Linux Enterprise Server 15 SP3 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 SP3.

5.8.6.1 Samba has been updated

The `samba` package has been updated to version 4.15. This update will require also updating other packages depending on it, namely:

- `apparmor`
- `cacertificates`
- `gnutls`
- `libldb`
- `libnettle`
- `libtalloc`
- `libtdb`
- `libtevent`
- `pll-kit`
- `sssd`

After the update, the `ldb` and `tdb` files created by `samba` (internal databases in `/var/lib/samba/**`) should be managed with the tools installed in `/usr/lib[64]/samba/bin/`. The reason is that we are not updating these libraries because the newer versions required by `samba`

4.15 do not provide Python 2 bindings and we can not remove them with the update, so we built samba with the libraries bundled in. The system's talloc, tdb, tevent, and ldb packages will not be modified by the update so the databases created using them (not by samba) must be administrated with the tools provided by the system's ldb-tools and tdb-tools packages as before the update.

5.8.7 NFSv4

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

5.9 Performance-related information

5.9.1 perf stat allows configuring whether to run used events in kernel space or user space

The perf tool offers a rich set of commands to collect and analyze performance and trace data. perf record supports --all-kernel/--all-user to configure all used events to run in kernel space or run in user space. However, in the version of perf shipped with SUSE Linux Enterprise Server 15 SP2, perf stat does not support these options.

In SUSE Linux Enterprise Server 15 SP3, we have updated perf stat to support the --all-kernel and --all-user options to keep the same semantics available in both commands.

5.10 Security

5.10.1 dm-crypt target supports synchronous encryption for increased performance

By default, dm-crypt performs data encryption and decryption through an asynchronous thread. Starting with SLE 15 SP3, the target supports synchronous operation which is controlled with no-read-workqueue and no-write-workqueue options. The options can be supplied through the /etc/crypttab file. See the crypttab(5) man page for more information.

5.10.2 ClamAV has been updated to version 0.103

ClamAV 0.103 provides better on-access scanning and improvements that reduce the attack surface.

5.10.3 tpm2-tss has been updated to version 2.3.3

The `tpm2-tss` package has been updated to version 2.3.3.

5.10.4 Information about Workstation Extension security policies has been added

SLES and SLED have different security policies but installing the Workstation Extension on SLES does not change this. This is not mentioned anywhere.

Now, when installing the Workstation Extension in SUSE Linux Enterprise Server 15 SP3, you will be informed that the SLES security policies still apply.

5.10.5 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 Storage and file systems

Also see the following additional note:

- *Section 5.12.17, “Snapper cleanup has new algorithms”*

5.11.1 Btrfs RAID 1 not fully supported

In case of a disk failure, the remaining disk needs to be mounted with the degraded option manually:

```
mount -odegraded,ro
```

5.11.2 bcache-tools has been added

The package bcache-tools has been added. It provides tools for analyzing bcache devices.

5.11.3 exFAT tools have been added

The package exfatprogs has been added to SUSE Linux Enterprise Server 15 SP3. It provides the utilities for working with exFAT file systems.

5.11.4 Per-inode DAX flag

In previous SUSE Linux Enterprise Server releases, the DAX mode (direct access mode for Ext4 and XFS) was either enabled or disabled for the whole storage volume with the dax mount option.

SUSE Linux Enterprise Server 15 SP3 adds the possibility to enable DAX on individual files. The corresponding file system mount options are dax={always, never, inode}. The old dax option corresponds to the new dax=always option. This option reflects in the content of the /proc/mounts file.

For SUSE Linux Enterprise Server 15 SP3, there is a transitional change to show dax,dax=always in /proc/mounts for compatibility with applications that detect DAX by the presence of the standalone dax option. Future SUSE Linux Enterprise Server releases will remove this transitional behavior, and the option will be shown as dax=<option> in /proc/mounts.

5.11.5 Serialization of Btrfs operations

Certain operations cannot be performed concurrently on a Btrfs file system, namely: balancing, device removal, device addition, and file-system resizing. In previous releases, when attempting to perform these operations concurrently, they conflicted, one operation failed, and a message was added to the kernel log.

The Btrfs utilities (package `btrfsprogs`) now provide conflict reporting and allow serializing these exclusive operations using the `--enqueue` option. For more information, see the man pages from the `btrfsprogs` package.

5.11.6 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	–	–	–	–

Feature	Btrfs	XFS	Ext4	OCFS 2 ¹
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 2⁶³ 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 SP3 but are not supported by SUSE. By default, the file system drivers in SUSE Linux Enterprise Server 15 SP3 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.

5.11.7 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)	–	–	–	+	+	+

Feature	SLES 11 SP4	SLES 12 SP5	SLES 15 GA	SLES 15 SP1	SLES 15 SP2	SLES 15 SP3
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	–	+	+	+	+	+ ^a

Feature	SLES 11 SP4	SLES 12 SP5	SLES 15 GA	SLES 15 SP1	SLES 15 SP2	SLES 15 SP3
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	–	+	+	+	+	+

a See [Section 5.11.1, “Btrfs RAID 1 not fully supported”](#).

5.11.8 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

To find out whether XFS filesystem has V4, run the following command:

```
xfs_info <mounted-filesystem>
```

Then check for the presence of the `crc=1` string.

5.12 System management

5.12.1 `systemd KillMode=none` is deprecated

Support for `KillMode=none` is deprecated and will be eventually removed in a future SLE version. However, the obsolete feature will be kept long enough to give customers time to migrate to another `KillMode` value. Please see the SUSE TID at <https://www.suse.com/support/kb/doc/?id=000020394> for more information.

5.12.2 Salt has been updated to version 3002

The `salt` package has been updated to version 3002. This update also includes patches, backports, and enhancements by SUSE for the SUSE Manager Server, Proxy and Client Tools. This applies to client operating systems with Python 3.5+. Otherwise Salt 3000 or 2016.11 is used.

We intend to regularly upgrade Salt to more recent versions.

For more details about changes in your manually-created Salt states, see <https://docs.saltproject.io/en/latest/topics/releases/3002.html>.

5.12.3 Additional settings now easily available during installation

During installation, there are some settings that were only accessible from certain screens.

With this change, these settings are now available at any point during the installation. The dialog provides access to these options: network devices, network proxy, software repositories, and expert console. Currently, they are only accessible using these keyboard shortcuts:

- Ctrl + Alt + Shift + C (in graphical mode)
- Ctrl + D Shift + C (in text mode)

5.12.4 Disable automatic updating of NVRAM in YaST and AutoYaST

Before this change, NVRAM was updated every time GRUB was installed or updated. This set the running SUSE OS as the new primary boot entry. Among other issues, this caused custom boot order to be lost every time that happened.

After this change, you can set the `UPDATE_NVRAM` parameter to `no` in `/etc/sysconfig/bootloader`. This will prevent NVRAM from being updated automatically.

For AutoYaST, you can use this configuration snippet:

```
<bootloader>
  <global>
    <update_nvram>false</update_nvram>
  </global>
</bootloader>
```



Note: Affected architectures

This only applies to UEFI on x86-64, AArch64, and PowerPC. SLES cannot modify the boot order on other architectures and set the BIOS to directly boot the newly installed OS.

5.12.5 SELinux support has been added to YaST

During installation, AutoYaST now allows you enable Security Enhanced Linux (SELinux). You can choose between `enforcing` and `permissive` mode.



Warning

SLES 15 SP3 does not include a default SELinux policy. You need your own policy for SELinux to work. Enabling SELinux without a policy will probably result in an unbootable system.

For more information, see <https://github.com/SELinuxProject/selinux> .

5.12.6 `xca` has been added

`xca` (X Certificate and Key Management) has been added as the new Certificate Authority (CA) management tool. `xca` replaces the old YaST CA management tool. It allows to:

- create CA and keys
- create, sign, and revoke certificates
- import and export keys and certificates in PEM, DER, and PKCS8 formats
- sign and revoke certificates in PEM, DER, and PKCS12 formats with select x509v3 extensions

It also provides a graphical interface and a tree-like view of certificates.

5.12.7 Shorter and more effective AutoYaST profiles

Previously, when AutoYaST generated a profile from an existing system, it included a lot of information to reproduce the installation. As a consequence, profiles were usually long, which made working with them more difficult. However, much of that information was not needed as it corresponded to default values or disabled features.

Now AutoYaST tries to skip irrelevant information, producing shorter and more manageable profiles. You can ask AutoYaST to additionally reduce the size of the profile by applying simple heuristics with the new `compact` mode. Bear in mind that in that case, some relevant information could be missing (for example, manually-created system users).

The new `compact` mode is disabled by default but it can be enabled by passing the `target` option to the `clone_system` command:

```
yast2 clone_system modules target=compact
```

Additionally, it is now possible to use `t` instead of `config:type` to add type annotations, reducing the size of the profile and making it easier to modify it manually.

5.12.8 Export registration information is included in the AutoYaST profile

Previously, although AutoYaST profiles used to contain a lot of information, the registration settings were not included. Additionally, the list of registered add-ons was wrongly exported as a regular repository.

AutoYaST now includes the `<suse_register>` section, containing the registration keys and the list of registered add-ons.

5.12.9 Improved scripting support in AutoYaST

Scripting support provides a powerful way to extend AutoYaST with custom behavior. Previously, Shell, Perl, and Python were the only supported scripting languages.

This limitation has been removed and it is now possible to use any interpreter which is available during the installation.

In addition to that, scripting has seen other improvements such as:

- ensuring that all artifacts are copied to the installed system
- reporting an error when the script returns a non-zero value.

5.12.10 Dynamic AutoYaST profiles using ERB

AutoYaST offers different ways of modifying a profile at runtime: asking the user for values during installation, running pre-installation scripts, or using rules and classes to merge different profiles. However, dealing with XML with basic tools might be hard.

In order to make it easier to modify the profile, AutoYaST now has support for ERB, which stands for **Embedded Ruby**. This allows to use the Ruby programming language to alter the profile at installation time. Additionally, AutoYaST offers a set of helpers to inspect the system (disks, network cards, etc.) and modify the profile accordingly.

5.12.11 AutoYaST profile validation at runtime

The AutoYaST documentation recommends using `xmllint` or `jing` to perform an XML-based validation of the profile. Although it is not mandatory, having to perform this step outside of the AutoYaST workflow can be annoying.

To make this easier, AutoYaST now validates the profile at runtime, reporting issues to the user. However, you can disable this behavior by setting the `YAST_SKIP_XML_VALIDATION` boot parameter to `1`.

5.12.12 Reducing the need for the AutoYaST second stage

AutoYaST uses two stages to perform the installation. Most of the work is done during the **first stage**: partitioning, system registration, software installation, network configuration, etc. After the first reboot, the **second stage** comes into play to configure additional services (for example, the firewall).

To reduce the need for a **second stage**, we have been moving the processing of several AutoYaST sections to the **first stage**. At this point, these sections are processed during the first stage:

- bootloader
- configuration_management
- files
- firewall
- host
- kdump
- keyboard
- language
- networking
- partitioning
- runlevel
- scripts (except post-scripts and init-scripts which are processed during the **second stage**)
- security
- services-manager

- [software](#)
- [ssh_import](#)
- [suse_register](#)
- [timezone](#) and
- [users](#)

If your profile does not contain any section not mentioned above, the **second stage** can be disabled.

5.12.13 [Extended support for customizing the AutoYaST partitioning schema from the UI](#)

Previously, the support for defining the partitioning schema in the AutoYaST user interface was limited. The tool only supported a subset of devices (disks, partitions, and LVM volume groups) and properties. In addition, the interface was somewhat confusing.

This interface has been greatly improved and extended to support software RAID devices, non-partitioned drives, and Bcache and multi-device Btrfs file systems.

5.12.14 [Disabling the automatic creation of bridges for virtual networks in AutoYaST](#)

When a virtualization package is selected for installation, for example, Xen, QEMU or KVM, AutoYaST sets up a bridge as part of the network configuration.

Now it is possible to disable this behavior by setting the `virt_bridge_proposal` element to `false`. This causes AutoYaST to delegate the creation of the bridge to the selected virtualization package.

5.12.15 [DOCUMENTATION_URL has been added to /etc/os-release](#)

`/etc/os-release` now contains the tag `DOCUMENTATION_URL`, which points to the online documentation of SUSE Linux Enterprise Server. The `DOCUMENTATION_URL` tag is used by certain tools, such as Cockpit.

5.12.16 `fwupd` has been updated

`fwupd` is simple daemon which allows session software to update firmware. In SUSE Linux Enterprise Server 15 SP3, we have updated `fwupd` from version 1.2 to version 1.5, which includes many new features and bug fixes.

5.12.17 `Snapper cleanup` has new algorithms

The `Snapper cleanup` command now has a new cleanup algorithm, `--free-space` that tries to free the requested amount of space. To clean up `/`, you can use for example:

```
snapper cleanup --path / --free-space "20 GiB" all
```

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

`systemd` in SUSE Linux Enterprise Server 15 SP3 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.

5.12.19 SUSE-specific RPM macros have been split from rpm package

The package `rpm-config-SUSE` is available on SUSE Linux Enterprise Server 15 SP3. This package allows adding or updating macros used at build-time without having to touch the core `rpm` package. This simplifies backporting packages that rely on newer macros.

5.13 Virtualization

For more information about acronyms used below, see <https://documentation.suse.com/sles/15-SP3/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-SP3/html/SLES-all/cha-virt-support.html>.

5.13.1 KVM

5.13.1.1 6 TiB memory support

KVM now supports 6 TiB of maximum memory per virtual machine.

5.13.1.2 swtpm has been added

The `swtpm` package has been added. It provides a software TPM (Trusted Platform Module) emulator.

QEMU can use `swtpm` as an external provider of a virtual TPM device. For more information, see <https://qemu-project.gitlab.io/qemu/specs/tpm.html>.

5.13.1.3 2nd generation AMD EPYC processor support has been added

Support for 2nd generation AMD EPYC processors has been added to QEMU/KVM. The model display name is `EPYC-Rome`.

5.13.1.4 haltpoll driver and governor for latency-sensitive virtual guests have been added

On bare-metal, a task waiting for a spinlock can use the `mwait` instruction to detect a change. This avoids an expensive Inter Processor Interrupt (IPI) when a waiting task must be woken. On virtual guests, `mwait` is difficult to emulate and IPIs are generally required (though this cost can be reduced with `halt_poll_ns`).

The SUSE Linux Enterprise Server 15 SP3 kernel for x86_64 includes `haltpoll`, a guest driver that polls a virtual CPU within the guest for an auto-tuned duration.

`haltpoll` improves the performance of some latency-sensitive, virtualized applications. `haltpoll` can only be used on physical hosts with a recent x86_64 CPU.

To use it:


- On the physical host, the QEMU commands that starts the virtual machine has to contain the parameter `-cpu host,kvm-hint-dedicated=on`. `virsh` allows specifying this parameter using `<hint-dedicated state='on'/>` and `<cpu mode='host-passthrough' check='none'/>`. For more information, see the [libvirt Documentation \(https://libvirt.org/formatdomain.html#elementsFeatures\)](https://libvirt.org/formatdomain.html#elementsFeatures).
- Load the driver in the virtual host: `modprobe cpuidle-haltpoll`. If it cannot be loaded, check `journalctl -k`. If something went wrong, you may see an `-ENODEV` error.

If you are using `libvirt/virsh`, verify that the `kvm-hint-dedicated` parameter is actually passed to QEMU. There are two complimentary ways of checking whether the parameter is successfully applied:

- On the host: Check the `qemu` command in the process list.
- On the guest: Check whether the QEMU KVM parameter above is active with `cpuid` (from the package `cpuid`): If it is active, `cpuid -1 -l 0x40000001` will show that the first bit of `edx` is set: `edx=0x00000001`.

5.13.1.5 QEMU has been updated to version 5.2

QEMU has been updated to [version 5.2](https://wiki.qemu.org/ChangeLog/5.2) (<https://wiki.qemu.org/ChangeLog/5.2>) .

In an effort to bridge the gap between openSUSE Leap and SLE, we have removed uses of the `is_opensuse` macro from the RPM spec file. This means that the packages built for SLE can be reused for openSUSE Leap. Some subpackages which are included for openSUSE Leap will not be included with SLE. Such packages will be provided in SUSE Package Hub for SLE users as unsupported packages (see also <https://packagehub.suse.com/> ).

Also review upstream [feature removals](https://wiki.qemu.org/Features/RemovedFeatures) (<https://wiki.qemu.org/Features/RemovedFeatures>) .

5.13.1.6 Fixed UIDs and GIDs for the `kvm`, `qemu`, and `libvirt` groups

With previous versions of SLES, if disks for KVM guests had been stored on NFS and the UID and GID were the same on both hosts, the guest disks became read-only after migration.

Starting with SUSE Linux Enterprise Server 15 SP3, we rely on `system-user-qemu` and `system-group-kvm` to provide these users and groups. These packages provide fixed UID and GID are now set for the `kvm`, `qemu`, and `libvirt` groups which avoids the migration problem.

5.13.1.7 Virtual machines support more than 256 CPUs

Virtual environments without virtualized IOMMU now support more than 256 CPUs. This, for example, helps support large AWS instances of SAP HANA.

5.13.2 Xen

- Xen: NetWare Support has been removed
- Update to Xen 4.14.0 FCS release
- Linux stub domain improvements
- Control-flow Enforcement Technology (CET) Shadow Stack support
- Support for running Xen as a Hyper-V Guest
- Domain ID randomization, persistence across save/restore
- Automatic generation of Go language bindings
- The debugging tool for Windows guests, KDD, now supports Windows 7, 8.x, and 10

For more information, see the [upstream Xen release notes \(https://wiki.xenproject.org/wiki/Xen_Project_4.14_Feature_List\)](https://wiki.xenproject.org/wiki/Xen_Project_4.14_Feature_List).

5.13.3 libvirt

libvirt has been updated to version 7.0.0. Major new features are:

- QEMU: Tolerate non-existent files such as `/dev/kvm` when populating domain private namespace
- Add all new APIs and constants in libvirt 7.0.0

For more information, see the [upstream libvirt release notes \(https://libvirt.org/news.html\)](https://libvirt.org/news.html).

5.13.3.1 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/>

5.13.3.2 kubevirt-virt-* packages have been moved

All the `kubevirt-virt-*` packages have been moved to the Containers module. As such, these packages are not maintained anymore. Everything else is shipped only as containers.

5.13.4 spice

5.13.4.1 spice-gtk

The new version 0.38 provides fixes and new features:

- Added CD/DVD redirection, to allow mounting ISO images from client
- Improved clipboard functionality, related to host/guest races and clipboard managers

5.13.4.2 spice-protocol

The version has been updated to 0.14.3:

- Added support for mouse side-buttons
- Added a `MonitorsMM` field to `VDAgentMonitorsConfig` to allow passing physical monitor dimensions
- Updated `VD_AGENT_*` capabilities
- Deprecated CELT support

For more information, see the [upstream change log \(https://gitlab.freedesktop.org/spice/spice-protocol/-/blob/master/CHANGELOG.md\)](https://gitlab.freedesktop.org/spice/spice-protocol/-/blob/master/CHANGELOG.md).

5.13.4.3 spice-gtk PulseAudio back-end has been removed

The PulseAudio back-end of spice-gtk has been removed in SUSE Linux Enterprise Server 15 SP3.

5.13.5 virt-manager has been updated to version 3.2.0

`virt-manager` has been updated to `virt-manager 3.2.0`. Major changes since the version included with the previous service pack of SUSE Linux Enterprise Server include:

- Display information about the NVRAM file used instead of only displaying the path
- Support for `virt-install --cloud-init`.
- The `virt-convert` tool has been removed. Use `virt-v2v` instead.
- A handful of UI XML configuration options have been removed. The XML editor can be used instead. For a larger discussion, see <https://www.redhat.com/archives/virt-tools-list/2019-June/msg00117.html>.
- The *New VM* UI now has a *Manual Install* option which creates a VM without any required install media.
- In the *New VM* UI, the network/PXE install option has been removed. If you need network boot, choose *Manual Install* and set the boot device after initial VM creation.
- *Migrate VM* UI now has an XML editor for the destination VM.

- Global and per-VM option to disable graphical console autoconnect. This makes it easier to use virt-manager alongside another client like `virt-viewer`.
- `virt-install`: Added `--reinstall=DOMAIN` option
- `virt-install`: Added `--autoconsole text|graphical|none` option
- `virt-install`: Added `--os-variant detect=on,require=on` suboptions
- CLI: Added `--xml XPATH=VAL` option for making direct XML changes
- CLI: Added `--clock`, `--keywrap`, `--blkio tune`, `--cputune` options
- CLI: Added `--features kvm.hint-dedicated.state=` feature.
- CLI: Added `--iommu` option.
- CLI: Added `--graphics websocket=` support.
- CLI: Added `--disk type=nvme source.*` suboptions.
- CLI: Fill in all `--filesystem` suboptions.
- New VMs are created by default with audio enabled

5.13.5.1 Vagrant boxes for SUSE Linux Enterprise Server

We are providing official Vagrant Boxes for SUSE Linux Enterprise Server x86-64 and AArch64 using the VirtualBox and `libvirt` providers. These boxes come with the bare minimum of packages to reduce their size and are not registered. Thus, you need to register the boxes prior to further provisioning.

These boxes are only available for direct download from <https://download.suse.com>⁷. Therefore, downloaded boxes must be registered manually with Vagrant as follows:

```
vagrant box add --name SLES-15-SP3 SLES15-SP3-Vagrant.x86_64-15.3-libvirt-
*.vagrant.libvirt.box
```

The box is then available under the name *SLES-15-SP3* and can be used like other Vagrant boxes:

```
vagrant init SLES-15-SP3
vagrant up
vagrant ssh
```

5.13.5.2 AArch64 support

The SUSE Linux Enterprise Server box is also available for the AArch64 architecture using the `libvirt` provider. It has been pre-configured for usage on SUSE Linux Enterprise Server on AArch64 and might not launch on other operating systems without additional settings. Running it on architectures other than AArch64 is not supported.

In case the box fails to start with a `libvirt` error message, add the following to your `Vagrantfile` and adjust the variables according to the guest operating system:

```
config.vm.provider :libvirt do |libvirt|
  libvirt.driver = "kvm"
  libvirt.host = 'localhost'
  libvirt.uri = 'qemu:///system'
  libvirt.host = "main"
  libvirt.features = ["apic"]
  # path to the UEFI loader for aarch64
  libvirt.loader = "/usr/share/qemu/aavmf-aarch64-code.bin"
  libvirt.video_type = "vga"
  libvirt.cpu_mode = "host-passthrough"
  libvirt.machine_type = "virt-3.1"
  # path to the qemu aarch64 emulator
  libvirt.emulator_path = "/usr/bin/qemu-system-aarch64"
end
```

5.13.6 VMware

5.13.6.1 High video resolutions in VMware ESXi need more VRAM

Virtual machines with less than 32 MB video memory can fail on resolutions higher than 1024x768.

If you are using VMs with resolutions higher than 1024x768, reserve 32 MB or more video memory.

5.13.7 Others

- support for NVIDIA Virtual GPU v12 has been added. This support uses the SR-IOV framework for the Ampere (A100/A10) architecture and the mediated device (`mdev`) framework for Volta and earlier architectures. The support does NOT include NVIDIA vGPU live migration support.
- Microsoft Azure: Support for hibernation of Linux VMs on Microsoft Azure has been added.
- The `os-dbinfo` database has been updated to version 20201218.
- `open-vm-tools` has been updated to version 11.2.5. For more information, see the [upstream change log \(https://github.com/vmware/open-vm-tools/blob/stable-11.2.5/open-vm-tools/ChangeLog\)](https://github.com/vmware/open-vm-tools/blob/stable-11.2.5/open-vm-tools/ChangeLog).
- `vm-install`: Modified the PV PXE booting feature to only allow a PXE server address to be passed on command line. The use of `udhcp` to look up PXE servers has been removed.

5.13.8 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 will be 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.

5.14 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.



Note: Package dependencies on additional SLE modules

When installing packages from SUSE Package Hub, you may need to activate additional SLE modules to solve dependency issues.

5.14.1 NVIDIA Compute module

The repositories for NVIDIA* CUDA* are available as the *NVIDIA Compute* module for x86-64 and AArch64. These repositories are provided by NVIDIA and the software in them is not supported by SUSE. All software in these repositories is licensed under the third-party [NVIDIA CUDA EULA \(https://docs.nvidia.com/cuda/eula/index.html\)](https://docs.nvidia.com/cuda/eula/index.html) .

The *NVIDIA Compute* module is not enabled by default when installing SUSE Linux Enterprise Server. During installation, the module can be selected from the *Extension and Module Selection* screen in YaST. Within an installed system, you can add it as follows: Run `yast registration` from a shell as root, select *Select Extensions*, search for *NVIDIA Compute Module* and continue with *Next*. Verify and accept the NVIDIA repository GPG key.



Important: Do not use the SUSEConnect tool to add this repository

Do not try to add this module with the `SUSEConnect` CLI tool. `SUSEConnect` is not yet capable of handling third-party repositories.



Important: Combining *Workstation Extension* and *NVIDIA Compute* module is unsupported

The *Workstation Extension* module includes some of the same drivers for NVIDIA graphics cards as the *NVIDIA Compute* module. However, their package versions may differ. As SUSE package management installs the latest package versions by default, enabling both modules at the same time can lead to a system with a mixture of packages from both modules.

Such a setup can result in drivers not working as expected and is not supported by SUSE.

5.14.2 Important package additions to SUSE Package Hub

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

python-anymarkup

Parse or serialize different markup formats. Currently supports INI, JSON, JSON5, TOML, XML and YAML.

pgaudit

An auditing module for PostgreSQL that collects audit events from various sources and logs them in CSV format. The generated logs include a timestamp, user information, details of objects affected (if any), and the fully-qualified command text (whenever available).

rshim

Access the serial console of an NVIDIA* BlueField* or BlueField-2 Data Processing Unit (DPU) over PCIe or USB.

5.15 Miscellaneous

5.15.1 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/>.

5.15.2 Access to logs via the audit group

The audit group has been added.

Its purpose is a better separation of permissions for access to audit logs. With this change, users can be given access to logs without the need to change sudo rules.

5.15.3 Mounting multipath devices via `by-label`

In 15 SP3, mounting multipath devices using `by-label` mounts might fail during boot.

To resolve this, the `multipath` module needs to be manually added to the initial RAM disk:

1. Create a new file called `999-multipath.conf` in `/etc/dracut.conf.d/` with the following content: `add_dracutmodules+=multipath`
2. Re-generate the initial RAM disk with this command: `dracut /boot/initrd-$(uname -r)`.

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

Information in this section applies to SUSE Linux Enterprise Server 15 SP3 for the AMD64/Intel 64 architectures.

6.1 Intel platforms and technologies

SUSE Linux Enterprise Server 15 SP3 introduces support for the following Intel platforms and technologies:

- Initial enabling for platforms based on the Intel 4th generation Scalable XEON Processors (known as Eagle Stream / Sapphire Rapids)
- Prepare support for next generation Intel Optane Persistent Memory (known as Crow Pass)
- Platforms based on next generation Xeon-D Processors (known as Idaville)
- Platforms based on latest Intel XEON E3 Processors (known as Tatlow)
- Platforms using 11th Gen Intel Core i Processors (known as Tiger Lake-UP3/-UP4/-H)
- Platforms using 11th Gen Intel Core S-series desktop processors (known as Rocket Lake-S)

7 POWER-specific changes (ppc64le)

Information in this section applies to SUSE Linux Enterprise Server for POWER 15 SP3.

7.1 IBM POWER10 support

On SLES 15 SP3, 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 ServiceReport has been added

A new tool named ServiceReport has been added. The tool allows you to quickly validate the FFDC (First Failure Data Capture) configuration and optionally fix the incorrect configurations automatically. This automation drastically reduces the time required to set up the FFDC and improves serviceability.

7.3 Rebuild capture kernel initrd after migration and/or hardware changes

The initrd for the `kdump` kernel is generated against the system it will run on to save memory usage and disk space. It contains the minimum set of kernel modules and utilities to boot the machine to a stage where the dump target could be mounted.

With the `kdump` service enabled, `kdump` will try to detect system changes and rebuild the `kdump` initrd if needed. But it can not guarantee to cover every possible case. If there was a hardware change, disk migration, storage setup update, or any similar system level change, it is highly recommended to rebuild the initrd manually with following command:

```
mkdumprd -f; systemctl restart kdump
```

7.4 Increased memory when running fadump

Firmware-assisted dump (`fadump`) in PowerVM was crashing due to low memory.

To resolve this, in SLES 15 SP3 the memory has been increased to 4 GB when running `fadump`.

7.5 Speed of `ibmveth` interface not reported accurately

The `ibmveth` interface is a paravirtualized interface. When communicating between LPARs within the same system, the interface's speed is limited only by the system's CPU and memory bandwidth. When the virtual Ethernet is bridged to a physical network, the interface's speed is limited by the speed of that physical network.

Unfortunately, the `ibmveth` driver has no way of determining automatically whether it is bridged to a physical network and what the speed of that link is. `ibmveth` therefore reports its speed as a fixed value of 1 Gb/s which in many cases will be inaccurate. To determine the actual speed of the interface, use a benchmark. Using `ethtool`, you can then set a more accurate displayed speed.

7.6 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 SP3. For more information, see https://www.ibm.com/support/knowledgecenter/en/linuxonibm/liaaf/lnz_r_suse.html ↗

8.1 Hardware

There were the following hardware-related changes:

- support has been added for IPL and re-IPL from local PCI NVMe storage
 - currently a workaround is required for installation, see <https://www.suse.com/support/kb/doc/?id=000020258> ↗
- support has been added for IBM z14 instructions in Valgrind

- the following new commands have been added to the `qclib` package:
 - `zhypinfo` - displays the virtualization stack
 - `zname` - displays information on the hardware platform
- s390x CPU topology masks have been made consistent with all other architectures
- improved performance of re-IPL by not clearing memory
- improved performance of the GNU C Library's `libm` math library by using IBM Z instructions
- the OpenBLAS library has been optimized with IBM z13 and IBM z14 instructions

8.2 Networking

8.2.1 Degraded performance on RoCE ConnectX-4 hardware

Using default settings of SUSE Linux Enterprise Server 15 SP1, 15 SP2, and 15 SP3, the performance of RoCE ConnectX-4 hardware on IBM z14 and IBM z15 systems is degraded compared to when used under SUSE Linux Enterprise Server 15 GA.

To improve performance to the same level as with SUSE Linux Enterprise Server 15 GA, set the following flag for all RoCE ethernet interfaces: `ethtool --set-priv-flags DEVNAME rx_striding_rq`. This needs to be done for each RoCE interface and at each boot.

8.2.2 geth: Converged HiperSockets/Ethernet Interface

Support for HiperSockets Converged Interface functionality has been added. This provides a converged interface that forms a single LAN based on HiperSockets and OSA/RoCE. This feature only supports a single registered MAC address for now.

8.2.3 SMC-R: Link failover support

Provides Link Group failover support which enables HA setups and makes the zLinux implementation reach full protocol compliance.

8.2.4 SMC-Dv2 support

SMC-Dv2 lifts the limitation to traffic within a single IP subnet only that SMC-D had, allowing traffic to peers in any IP subnet. It also simplifies ISM device configuration.

8.2.5 `smc-tools`: Integrate SMC-R Link Group (LG) support

SMC-R LG support has been fully integrated into the `smc-tools` package through proper userspace tooling.

8.3 Performance

There were the following performance-related changes:

- use z15 instructions for the kernel's zlib implementation which is used, for example, for Btrfs compression
- when placed at the beginning of a function, kprobes will use the ftrace infrastructure, which increases performance

8.4 Security

There were the following miscellaneous security-related changes:

- the `zkey` tool from `s390-tools` has been extended to import keys and recreate a repository based on keys generated by the EKMF web enterprise key management system
- self-test has been added to the `paes_s390` module to allow loading and using the PAES cipher if the kernel FIPS flag is switched on
- The `cpacfstats` tool from `s390-tools` has been enhanced to display Elliptic Curve Cryptography (ECC) CPU-MF counters

8.4.1 openCryptoki

There were the following openCryptoki-related changes:

- the `pkcstok_migrate` tool has been added

- the tool is able to convert all token data including PINs from using PINs encrypted with the method of v3.11 and earlier to being encrypted with a FIPS 140-2 compliant method
- it allows to migrate old key repositories to use data structures that support FIPS 140-2 compliant methods
- enhancements introduced with IBM z15 have been added, including Dilithium signing (quantum-safe support), and the Reencrypt function to the openCryptoki EP11 token
- support has been added for new identifiers and the PKCS #11 Baseline Provider Profile
- the `p11sak` tool has been added for generating, listing and deleting token keys in an openCryptoki token repository

8.4.2 Support for EP11 secure keys

The `pkey` module and the `zkey` tool have been extended to support EP11 secure keys. This allows the use of protected keys derived from EP11 secure keys with `dm-crypt`.

8.4.3 Enhanced error handling for zcrypt device driver

The error handling for the `zcrypt` device driver has been enhanced, for example, by adding a device offline state. This allows to distinguish between devices being offline due to external events and devices configured to be offline.

8.5 Storage

8.5.1 zdsfs: Coordinated read access

The `zdsfs` tool from `s390-tools` can now read from z/OS data sets while the containing DASD volume is online in z/OS.

8.5.2 Support for querying FICON link IBM Fibre Channel Endpoint Security

You can now find out if a FICON DASD is accessed using authenticated or encrypted links via the new sysfs `fc_security` attribute.

8.5.3 Support for querying FCP link IBM Fibre Channel Endpoint Security

You can now find out if an FCP remote port is accessed using authenticated or encrypted links, both in the running system and through kernel logs.

8.6 Virtualization

8.6.1 Added IBM Z LPAR fence agent `fence_ibmz` for Pacemaker

An IBM Z LPAR fence agent has been added for KVM setups with high-availability requirements which are often based on Corosync/Pacemaker.

8.6.2 Enhanced hardware diagnosis data of guest kernel

KVM now makes available additional data to improve hardware diagnoses for guest kernels.

8.6.3 `kvm_stat`: Improvements to sampling and logging

The sampling and logging capabilities of `kvm_stat` have been refined to provide improved RAS capabilities for both test/development and production environments.

8.6.4 Enablement of channel path handling for `vfio-ccw`

Improved handling of channel paths in `vfio-ccw` has been added. For example, this includes passing through channel-path operations and notifying of channel path changes.

8.6.5 Transparent CCW IPL from DASD (`vfio-ccw`) has been enabled

The existing support for native CCW IPL required the setting of a per-device property to enforce unlimited prefetch. This feature removes the necessity to specify the additional property and thus enables Linux IPL from `vfio-ccw` attached DASDs transparently.

8.6.6 Enable host key document verification

The tool `genprotimg` from the package `s390-tools` can now be used for host-key document verification. This removes the extra manual verification step that was needed before.

8.6.7 Support for `virtio-fs` on IBM Z

`virtio-fs` can now share a host file system with a guest.

8.6.8 Support for libvirt node device for `vfio-ap` matrix device

Enable and simplify the passthrough of crypto devices through use of libvirt mediated device management.

8.6.9 Support for DASD in libvirt node device driver

Enable and simplify the passthrough of DASD devices through use of libvirt mediated device management.

8.6.10 Implementation of full set of zPCI function properties

All properties of host PCI devices are now passed down to the guest, except for properties that are overridden by the user. This improves the support for all PCI devices except network adapters.

8.7 Miscellaneous

8.7.1 Server Time Protocol (STP) leap second handling

When using STP, leap seconds will now be handled correctly.

9 Arm 64-bit-specific changes (AArch64)

Information in this section applies to SUSE Linux Enterprise Server for Arm 15 SP3.

9.1 System-on-Chip driver enablement

SUSE Linux Enterprise Server for Arm 15 SP3 includes driver enablement for the following System-on-Chip (SoC) chipsets:

- AMD* Opteron* A1100
- Ampere* X-Gene*, eMAG*, Altra*
- AWS* Graviton, Graviton2
- 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 Driver enablement for Arm GIC v4.1

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel updates the Arm* Generic Interrupt Controller (GIC) driver `irq-gic-v4` to prepare for upcoming chips with GIC version 4.1.

KVM support for GIC v4.1 is still missing, see [Section 9.3.1, “No KVM support for Arm GIC v4.1”](#).

9.2.2 Driver enablement for NVIDIA Xavier

SUSE Linux Enterprise Server for Arm 15 SP2 added initial enablement for the NVIDIA* Tegra* X1 (T210) and Tegra X2 (T186) System-on-Chip (SoC) chipsets.

SUSE Linux Enterprise Server for Arm 15 SP3 adds enablement for the NVIDIA Xavier* SoC (T194), which is found on Jetson AGX Xavier* and Jetson Xavier NX System-on-Modules (SoM).

Drivers for the integrated, NVIDIA *Volta* microarchitecture-based Graphics Processor Unit (GPU) are not included ([Section 9.3.3, “No graphics drivers on NVIDIA Jetson”](#)).



Note: UEFI firmware may need to be flashed for NVIDIA Jetson

The NVIDIA Jetson AGX Xavier and Jetson Xavier NX SoMs by default ship with a *CBoot* bootloader. CBoot does not implement the Unified Extensible Firmware Interface (UEFI) and will thereby not boot the SUSE Linux Enterprise Server for Arm 15 SP3 installation media (compare [Section 9.1, “System-on-Chip driver enablement”](#)).

For more information, see the *NVIDIA Jetson Linux Developer Guide*, section "Jetson Xavier NX and Jetson AGX Xavier Series Boot Flow" (https://docs.nvidia.com/jetson/l4t/#page/Tegra%20Linux%20Driver%20Package%20Development%20Guide/bootflow_jetson_xavier.html%23wwpID0E0FB0HA) ↗.

NVIDIA offers an alternative bootloader firmware for the NVIDIA Jetson AGX Xavier and Jetson Xavier NX Developer Kits: <https://developer.nvidia.com/embedded/downloads#?search=uefi> ↗ (at the time of writing: *NVIDIA UEFI/ACPI Experimental Firmware for Jetson AGX Xavier and Jetson Xavier NX*, version 1.1.0)

For other devices based on NVIDIA Xavier SoCs, check with the respective hardware vendor whether a UEFI firmware is available.



Note: No UEFI support on NVIDIA DRIVE AGX platforms

The NVIDIA DRIVE* AGX Xavier and NVIDIA DRIVE AGX Pegasus* Developer Kits use a NVIDIA DRIVE OS hypervisor. Its virtual guest bootloader *OSLoader*, as of NVIDIA DRIVE OS version 5.2, does not implement UEFI but a custom guest partition image format.

For more information, see the *NVIDIA DRIVE OS Linux SDK Developer Guide* chapter *Bootloader Programming*, sections [Understanding the Boot Flow: OSLoader](https://docs.nvidia.com/drive/drive-os-5.2.3.0L/drive-os/index.html#page/DRIVE_OS_Linux_SDK_Development_Guide/Bootloader/bootloader_bootflow.html%23wwpID0ETHA) (https://docs.nvidia.com/drive/drive-os-5.2.3.0L/drive-os/index.html#page/DRIVE_OS_Linux_SDK_Development_Guide/Bootloader/bootloader_bootflow.html%23wwpID0ETHA) ↗ and [Flashing with Bootburn: Virtualization Behavior](https://docs.nvidia.com/drive/drive-os-5.2.3.0L/drive-os/index.html#page/DRIVE_OS_Linux_SDK_Development_Guide/Bootloader/bootloader_bootburn.html%23wwpID0EVHA) (https://docs.nvidia.com/drive/drive-os-5.2.3.0L/drive-os/index.html#page/DRIVE_OS_Linux_SDK_Development_Guide/Bootloader/bootloader_bootburn.html%23wwpID0EVHA) ↗.

Contact NVIDIA to discuss how to use SUSE Linux Enterprise Server for Arm 15 SP3 on NVIDIA DRIVE AGX platforms.

9.2.3 Driver enablement for NXP i.MX 8M Mini

SUSE Linux Enterprise Server for Arm 15 SP1 added initial enablement for the NXP* i.MX 8M System-on-Chip (SoC), also referred to as 8MQ (quad-core).

SUSE Linux Enterprise Server for Arm 15 SP3 adds enablement for the i.MX 8M Mini (8MM) and further prepares 8M Nano (8MN) and 8M Plus (8MP).

9.2.4 Driver enablement for NXP Layerscape LS1012A

SUSE Linux Enterprise Server for Arm 15 SP3 adds initial enablement for the NXP* Layerscape* LS1012A System-on-Chip (SoC).

Known limitations for the built-in network interfaces are detailed in [Section 9.3.5, “No PFE network drivers on NXP Layerscape LS1012A”](#).

9.3 Known limitations

9.3.1 No KVM support for Arm GIC v4.1

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel does not support KVM on the Arm* Generic Interrupt Controller (GIC) version 4.1.

Contact your SUSE representative if you have a System-on-Chip with GICv4.1 and need KVM virtualization support.

9.3.2 No ACPI support on NXP Layerscape LX2160A

For the NXP* Layerscape* LX2160A System-on-Chip NXP provides an alternative bootloader firmware based on TianoCore EDK II. This firmware can be configured to use both Device Tree and ACPI.

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel drivers for NXP LX2160A do not yet support ACPI. Continue to use the Device Tree booting method for now, or contact your SUSE representative if that is not possible.

9.3.3 No graphics drivers on NVIDIA Jetson

The NVIDIA* Tegra* System-on-Chip chipsets include an integrated Graphics Processor Unit (GPU).

SUSE Linux Enterprise Server for Arm 15 SP3 does not include graphics drivers for any of the NVIDIA Jetson* or NVIDIA DRIVE* platforms.

Contact the chip vendor NVIDIA for whether third-party graphics drivers are available for SUSE Linux Enterprise Server for Arm 15 SP3.

9.3.4 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 2.8.1.6, “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 SP3.

Alternatively, contact your hardware vendor for whether a bootloader update is available that implements graphics output, allowing to instead use `efi fb` framebuffer graphics in SUSE Linux Enterprise Server for Arm 15 SP3.



Note

The Vivante GC7000UL GPU driver (`etnaviv`) is available as a technology preview ([Section 2.8.1.4, “etnaviv drivers for Vivante GPUs are available”](#)).

9.3.5 No PFE network drivers on NXP Layerscape LS1012A

The NXP* Layerscape* LS1012A System-on-Chip contains a Packet Forwarding Engine (PFE) for up to two Ethernet ports.

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel does not include drivers for PFE.

The bootloader firmware provided by your hardware vendor should allow you to load and use the *GRUB* bootloader from SUSE Linux Enterprise Server for Arm 15 SP3 over the PFE Ethernet ports. Check with your hardware vendor for any firmware updates.

But the Installer and installed system will not be able to access built-in PFE-connected Ethernet ports.

Contact the chip vendor NXP for whether third-party PFE network drivers are available for SUSE Linux Enterprise Server for Arm 15 SP3.

Alternatively, your bootloader may be configured to support PCI-based Ethernet adapters based on mutually supported chipsets, such as e1000e.



Note

The use of PCI-based Ethernet adapters on LS1012A may require to run pci_enum from the *U-Boot* bootloader prompt before continuing to boot.

9.3.6 Some drivers not ready for Raspberry Pi

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel does not include a driver for VideoCore* Host Interface Queue (VCHIQ), which was still in *staging*. The tool vcgencmd depends on VCHIQ and is therefore not included. Any drivers depending on vchiq driver are not included either, in particular snd-bcm2835 for 3.5 mm TRRS **audio jack** and bcm2835-camera (kernel module bcm2835-v4l2) for MIPI* CSI-2* **camera connector** are unavailable. Also dependent on VCHIQ is the Multi-Media Abstraction Layer (MMAL) driver vchiq-mmal (kernel module bcm2835-mmal-vchiq), whose absence precludes you from using OpenMAX* (OMX) API based tools using MMAL, such as raspivid and raspistill.

A performance monitoring driver for the Advanced eXtensible Interface (AXI) bus on the Raspberry Pi (raspberrypi_axi_monitor) is not available.

9.3.6.1 Raspberry Pi 3 missing drivers

On Raspberry Pi 3, **video codec** hardware acceleration (bcm2835_codec) depends on VCHIQ and is unavailable. Applications will need to use software decoding for playback.

9.3.6.2 Raspberry Pi 4 missing drivers

The vc4 Display Rendering Manager (DRM) driver and the v3d Display Rendering Infrastructure (DRI) driver for the Broadcom* VideoCore VI Graphics Processor Unit (GPU) are available in the SUSE Linux Enterprise Server for Arm 15 SP3 kernel, but the Mesa graphics library code for it was not stable. Software-based rendering should be used instead of 3D hardware acceleration.

The Direct Memory Access (DMA) engine driver `bcm2835-dma` does not implement 40-bit transfers and is limited to 30 bits, that is, the lower 1 GiB of RAM. Transfers to higher areas of RAM on applicable models (2/4/8 GiB) will transparently use bounce buffers in low memory, so that functionality is not impaired but performance will be impacted.

Video codec hardware acceleration support (H.264, HEVC, VP9) is missing. Applications will need to use software decoding for playback.

9.4 Deprecation 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 (Mobicore based vs. Synopsys DesignWare* based respectively).

The SUSE Linux Enterprise Server for Arm 15 SP3 kernel supports 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
```

SUSE Linux Enterprise Server for Arm 15 SP4 will remove the support for rev. 1.0 silicon by dropping patches from the kernel. This may then 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.

9.5 Removal of early Marvell ThunderX2 silicon support

Marvell* ThunderX2* System-on-Chip silicon revisions Ax had errata for the SATA controller. Silicon revisions B0 and later are not affected.

SUSE Linux Enterprise Server for Arm 12 SP3 up to 15 SP2 included kernel patches with a recommended workaround. This allowed evaluation of early server systems with the affected silicon revisions.

As announced with SUSE Linux Enterprise Server for Arm 15 SP2, the SUSE Linux Enterprise Server for Arm 15 SP3 kernel no longer includes the patches with those workarounds. Production servers should not be affected by that change. For early systems with pre-production silicon check with the hardware vendor whether CPU upgrade kits are available.

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 SP3

For more information about all package and module changes since the last version, see [Section 2.2.3, “Package and module changes in 15 SP3”](#).

10.1 Removed features and packages

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

- NodeJS 8 and NodeJS 10 have been removed. For more information, see [Section 5.5.5, “Web and Scripting Module: NodeJS 14 has been added, NodeJS 8 and NodeJS 10 has been removed”](#).
- The PulseAudio back-end of spice-gtk has been removed. For more information, see [Section 5.13.4.3, “spice-gtk PulseAudio back-end has been removed”](#).
- The `rxecfg` binary has been removed from the package `libibverbs` (part of `rdma-core`).
- Kernel support for early Marvell* ThunderX2* silicon has been removed. For more information, see [Section 9.5, “Removal of early Marvell ThunderX2 silicon support”](#).

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.

Also see the following release notes:

- [Section 5.11.8, “XFS V4 format file systems have been deprecated”](#)
- The `imgen` package, containing Mellanox firmware generator, is deprecated and will be removed with SLES 15 SP4.
- The OpenLDAP server is deprecated and will be removed with SLES 15 SP4. It will no longer be available from the *Legacy SLE* module. For more information, see [Section 5.1.1, “389 Directory Server is the primary LDAP server, the OpenLDAP server is deprecated”](#).
- Python 2 will be removed entirely from SLE with SLE 15 SP4 and will no longer be available via the *Python 2* SLE module. For more information, see [Section 5.5.11, “Python 2 is deprecated”](#).
- 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 5.10.5, “TLS 1.1 and 1.0 are no longer recommended for use”](#).
- NXP LX2160A revision 1 silicon quirks will be removed with SUSE Linux Enterprise Server for Arm 15 SP4. For more information, see [Section 9.4, “Deprecation of NXP Layerscape LX2160A rev. 1 silicon support”](#).

- 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 5.12.18, “Support for System V init.d scripts is deprecated”](#).
- Support for libvirt LXC containers is deprecated and will be removed with SUSE Linux Enterprise Server 15 SP4. For more information, see [Section 5.13.8, “VM installer of YaST can no longer install LXC containers”](#).
- `systemd KillMode=none` is deprecated. For more information see [Section 5.12.1, “systemd KillMode=none is deprecated”](#).
- `lftp_wrapper` is deprecated. Use `lftp` directly instead.
- `pam_ldap` and `nss_ldap` are deprecated. Use SSSD instead.
- PostgreSQL 10 is deprecated and has been moved to the Legacy module. For more information about PostgreSQL, see [Section 5.4.3, “PostgreSQL 13 has been added”](#).
- On the POWER architecture, transactional memory is deprecated. For more information, see [Section 7.6, “Transactional memory is deprecated and disabled”](#).
- System containers using LXC have been deprecated and will be removed in SUSE Linux Enterprise Server 15 SP4. For more information, see [Section 5.3.4, “LXC containers have been deprecated”](#).

10.2.1 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 (mailto: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 © 2010-2022 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.

A Kernel parameter changes



Warning

This list of changes may not be exhaustive.

A.1 During 15 SP3

These Linux kernel parameters have been changed during SLES 15 SP3.

Setting	Previous value	New value	Package	More details
<code>sysctl_net_ipv6_conf_lo_accept_redirects</code>	0	1	wicked-0.6.68-1.50300.1.1.5.3	https://bugzilla.suse.com/show_bug.cgi?id=1181163

Setting	Previous value	New value	Package	More details
<u>sysctl_net_ipv6_conf_lo_use_tempaddr</u>	1	-1	wicked-0.6.68-150300.4.1.5.3	Bug 1181163 (https://bugzilla.suse.com/show_bug.cgi?id=1181163) ↗

A.2 Since 15 SP2

These Linux kernel parameters have been changed since SLES 15 SP2.

Setting	Previous value	New value	Package	More details
<u>sysctl_fs_file-max</u>	490215	9223372036854775807		Bug 1180823 (https://bugzilla.suse.com/show_bug.cgi?id=1180823) ↗
<u>sysctl_fs_suid_dumpable</u>	0	2	systemd 246	Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339) ↗
<u>sysctl_kernel_cap_last_cap</u>	37	39		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c18) ↗

Setting	Previous value	New value	Package	More details
<u>sysctl_kernel_core_pattern</u>	/usr/ lib/ systemd/ systemd- coredump %P %u %g %s %t %c %e	/usr/ lib/ systemd/ systemd- coredump %P %u %g %s %t %c %h		Bug 1183450 (https:// bugzilla.suse.com/ show_bug.cgi? id=1183450) ↗
<u>sysctl_kernel_core_pipe_limit</u>	0	16		Bug 1183450 (https:// bugzilla.suse.com/ show_bug.cgi? id=1183450) ↗
<u>sysctl_kernel_printk_devkmsg</u>	ratelimit	on	systemd 248	Bug 1183339 (https:// bugzilla.suse.com/ show_bug.cgi? id=1183339#c21) ↗
<u>sysctl_kernel_suid_dumpable</u>	0	2	systemd 246	Bug 1183450 (https:// bugzilla.suse.com/ show_bug.cgi? id=1183450#c1) ↗
<u>sysctl_kernel_usermodehelper_bsu</u>	63	255		Bug 1183339 (https:// bugzilla.suse.com/ show_bug.cgi? id=1183339#c18) ↗

Setting	Previous value	New value	Package	More details
<u>sysctl_kernel_usermodehelper_inheritable</u>	0	1		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c18) ↗
<u>sysctl_net_core_bpf_jit_kallsyms</u>	0	1		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c9) ↗
<u>sysctl_net_ipv4_ping_group_range</u>	0	0		
		2147483647		
		a		
<u>sysctl_net_ipv4_tcp_available_uhp</u>		espintcp		Bug 1181745 (https://bugzilla.suse.com/show_bug.cgi?id=1181745) ↗
<u>sysctl_net_ipv4_tcp_mem</u>	22689 30255 45378	22683 30246 45366		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
		b		
<u>sysctl_net_ipv4_udp_mem</u>	45381 60510 90762	45369 60492 90738		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗

Setting	Previous value	New value	Package	More details
<u>sysctl_fs_epoll_max_user_watches</u>	410398	410275 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_kernel_threads-max</u>	15650	15655 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_cgroup_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_ipc_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_mnt_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗

Setting	Previous value	New value	Package	More details
<u>sysctl_user_max_net_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_pid_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_time_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_user_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗
<u>sysctl_user_max_uts_namespaces</u>	7827	7825 [b]		Bug 1183339 (https://bugzilla.suse.com/show_bug.cgi?id=1183339#c11) ↗

a Also see [Section 5.8.1, “All users can now create ICMP Echo sockets”](#).

b All `sysctl` parameters based on the total amount of memory or available memory can vary between service packs or even maintenance updates. Small variations are to be expected. Either the accounting might have changed for the former, or more early allocations are necessary for the latter.