





Release Notes

SUSE Linux Enterprise Server 15 GA

This document provides guidance and an overview to high level general features and updates for SUSE Linux Enterprise Server 15 GA. Besides architecture or product-specific information, it also describes the capabilities and limitations of SUSE Linux Enterprise Server 15 GA.

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://www.suse.com/documentation/sles-15> .

Publication Date: 2019-03-29, Version: 15.0.20190329

Contents

- 1 About the Release Notes 3
- 2 SUSE Linux Enterprise Server 3
- 3 Installation and Upgrade 13
- 4 Architecture Independent Information 24
- 5 AMD64/Intel 64 (x86_64) Specific Information 48
- 6 POWER (ppc64le) Specific Information 49
- 7 IBM Z (s390x) Specific Information 51
- 8 ARM 64-Bit (AArch64) Specific Information 57
- 9 Packages and Functionality Changes 58

10	Technical Information	73
11	Obtaining Source Code	82
12	Legal Notices	82

1 About the Release Notes

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

Entries can be listed twice, 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 documents 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 GA is a multimodal operating system that paves the way for IT transformation in the software-defined era. The modern and modular OS 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-premise and public cloud environments.

SUSE Linux Enterprise Server 15 GA, 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 performing guest on leading hypervisors and supports an unlimited number of virtual machines per physical system with a single subscription. This makes it the perfect guest operating system for virtual computing.

2.2 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 3 years of Extended Support. The current version (GA) will be fully maintained and supported until 6 months after the release of SUSE Linux Enterprise Server 15 SP1.

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, you receive a total of 3 to 5 years of support per Service Pack.

For more information, check our Support Policy page <https://www.suse.com/support/policy.html> or the Long Term Service Pack Support Page <https://www.suse.com/support/programs/long-term-service-pack-support.html>.

2.3 What Is New?

SUSE Linux Enterprise Server 15 introduces many innovative changes compared to SUSE Linux Enterprise Server 12.

Changes to the installation and the module system:

- **Unified installer:** All SUSE Linux Enterprise 15 products can be installed by the same unified installer media. For information about available modules, see *Section 2.9.1, “Modules in the SLE 15 GA Product Line”*.
- **Installation without network using Packages media:** To install without network connection, all necessary packages are available on the Packages medium. This medium consists of directories with module repositories which need to be added manually as needed. RMT (Repository Mirroring Tool) and SUSE Manager provide additional options for disconnected or managed installation.
- **Migration from openSUSE Leap to SUSE Linux Enterprise Server:** Starting with SUSE Linux Enterprise 15 GA, we support migrating from openSUSE Leap 15 to SUSE Linux Enterprise Server 15. Thus, even if you decide to start out with the free community distribution you can later easily upgrade to a distribution with enterprise-class support.

- **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 SCC.
- **Software Development Kit:** With SUSE Linux Enterprise 15, the Software Development Kit is now integrated into the products. Development packages are packaged alongside regular packages. In addition, the Development Tools module contains the 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. For more information, see [Section 3.2.4, “SMT Has Been Replaced by RMT”](#).

Major updates to the software selection:

- **Salt:** SUSE Linux Enterprise 15 can be managed via salt to help integration into up-to-date management solutions, such as SUSE Manager.
- **Python 3:** As the first enterprise distribution, SUSE Linux Enterprise 15 GA offers full support for Python 3 development in addition to Python 2.
- **Directory Server:** 389 Directory Server replaces OpenLDAP to provide a sustainable directory service.

2.4 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 2.6, “Support Statement for SUSE Linux Enterprise Server”](#)
- [Section 3.2, “Upgrade-Related Notes”](#)
- [Section 10, “Technical Information”](#)

2.5 Documentation and Other Information

2.5.1 No Documentation Installed in SLES JeOS By Default

In SLES JeOS 12 SP3, the Zypper configuration led to the system installing documentation packages such as man pages by default. For SLES JeOS 15 GA and up, this configuration has been changed, so the system is leaner.

To return your SLES JeOS system to the previous behavior of installing documentation packages by default, edit the Zypper configuration file `/etc/zypp/zypp.conf`: Change the configuration line `rpm.install.excludedocs = yes` to `rpm.install.excludedocs = no`.

2.5.2 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 media for a chronological log of all changes made to the updated packages.
- Find more information in the `docu` directory of the media of SUSE Linux Enterprise Server 15 GA.

2.5.3 Externally Provided Documentation

- For the most up-to-date version of the documentation for SUSE Linux Enterprise Server 15 GA, see <https://www.suse.com/documentation/sles-15>.
- Find a collection of White Papers in the SUSE Linux Enterprise Server Resource Library at <https://www.suse.com/products/server/resource-library>.

2.6 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/programs/subscriptions/?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 available documentation.

L2

Problem isolation, which means technical support designed to analyze data, reproduce customer problems, isolate 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 15 GA and its Modules are delivered with L3 support for all packages, except the following:

- Technology Previews, see [Section 2.8.1, “Technology Previews for All Architectures”](#)
- Sound, graphics, fonts and artwork
- Packages that require an additional customer contract
- 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 (that is, unchanged and un-recompiled) packages.

2.7 General Support

To learn about supported kernel, virtualization, and file system features, as well as supported Java versions, see [Section 10, “Technical Information”](#).

2.7.1 Reblink Feature of XFS Is Not Supported

XFS reblink support is currently considered experimental in current upstream Linux and is also not supported in SUSE Linux Enterprise.

2.8 Technology Previews


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

Whether a technology preview becomes a fully supported technology later depends on customer and market feedback. Technology previews can be dropped at any time and SUSE does not commit to providing a supported version of such technologies in the future.

Give your SUSE representative feedback, including your experience and use case.

2.8.1 Technology Previews for All Architectures

2.8.1.1 KVM Nested Virtualization

KVM Nested Virtualization is available in SLE 15 as a technology preview. For more information about nested virtualization, see <https://github.com/torvalds/linux/blob/master/Documentation/virtual/kvm/nested-vmx.txt> .

2.8.2 Technology Previews for AMD64/Intel 64 64-Bit (x86_64)

2.8.2.1 GPU Virtualization under KVM (virtio-gpu)

As a Technology Preview, SLES 15 allows using GPU virtualization in KVM with virtio-gpu. With virtio-gpu, you can use accelerated 2D (user space) and have limited support for 3D acceleration using OpenGL 3.x (user space and kernel space).

2.8.2.2 Support for AMD Secure Encrypted Virtualization

As a technology preview, SLE 15 now supports AMD Secure Encrypted Virtualization (SEV). SEV integrates main memory encryption capabilities (SME) with the existing AMD-V virtualization architecture to support encrypted virtual machines. Encrypting virtual machines helps protect them from physical threats and other virtual machines or even the hypervisor itself. SEV represents a new approach to security that is particularly suited to cloud computing where virtual machines may not fully trust the hypervisor and administrator of their host system. As with SME, no application software modifications are required to support SEV.

See also [Section 4.1.7, “Support for AMD Memory Encryption”](#).

2.8.3 Technology Previews for IBM Z (s390x)

2.8.3.1 Exploitation of Shared Memory Communications

As a technology preview, the kernel of SLE 15 enables Shared Memory Communications over RDMA (SMC-R) enabled. SMC-R allows RDMA network interface controllers (RNICs) to offer RDMA over Converged Ethernet (RoCE).

2.8.3.2 Support for dm-crypt with Protected Keys

As a technology preview in SLE 15, you can now use protected keys to encrypt partitions for effective end-to-end encryption.

2.9 Modules, Extensions, and Related Products

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

2.9.1 Modules in the SLE 15 GA Product Line

The SLE 15 GA product line is made up of modules that contain software packages. Each module has a clearly defined scope. Different modules can have a different life cycles and update timelines.

The following modules are available within the product line based on SUSE Linux Enterprise 15 GA at the release of SUSE Linux Enterprise Server 15 GA. However, not all modules are available with a subscription for SUSE Linux Enterprise Server 15 GA itself (see the column *available with*). For information the availability of individual packages within modules, see <https://sc-c.suse.com/packages>.

Name and Content	Dependencies on Other Modules	Available with	Support ¹
Base System SLE base system	None	SLES, SLES for SAP, SLE HPC, SLE RT, SLED (default on all)	life cycle: 10 years extended: 3 years of LTSS support level: L3
Containers Docker, tools, prepackaged images	Base System	SLES, SLES for SAP, SLE HPC	life cycle: 10 years extended: no support level: L3
Desktop Applications Basic desktop functionality	Base System	SLES, SLES for SAP (default), SLE HPC (default), SLE RT (default), SLED (default)	life cycle: 10 years extended: no support level: L3
Workstation Extension² Office tools and multimedia	Base System, Desktop Applications	SLED (default)	life cycle: 10 years extended: no support level: mixed L2/L3 (depending on package)
Development Tools Helps in developing applications, replaces the SLE SDK	Base System, Desktop Applications	SLES, SLES for SAP, SLE HPC (default), SLE RT (default), SLED	life cycle: 10 years extended: 3 years of LTSS support level: L3

Name and Content	Dependencies on Other Modules	Available with	Support ¹
High Availability HA tools such as Hawk, crm, Pacemaker, Corosync	Base System	SLES for SAP (default)	life cycle: 10 years extended: 3 years of LTSS support level: L3
High Performance Computing Tools and libraries related to High Performance Computing (HPC)	Base System	SLE HPC (default)	life cycle: 10 years extended: 3 years of LTSS, 12 months of ESPOS support level: L3
Legacy Packages for migration purposes with limited support time frame	Base System	SLES, SLES for SAP, SLE HPC	
Public Cloud Public cloud initialization code and tools	Base System	SLES, SLES for SAP, SLE HPC	support level: L3
SAP Applications Packages specific to SLES for SAP	Base System	SLES for SAP (default)	life cycle: 10 years extended: 3 years of LTSS, 12 months of ESPOS support level: L3
Server Applications Basic server functionality, NVDIMM support, OFED	Base System	SLES, SLES for SAP, SLE RT, SLE HPC (default on all)	life cycle: 10 years extended: 3 years of LTSS support level: L3

Name and Content	Dependencies on Other Modules	Available with	Support ¹
SUSE Cloud Application Platform Tools Tools to interact with SUSE Cloud Application Platform	Base System	SLES, SLES for SAP, SLED	
SUSE Package Hub² Community-maintained packages	Base System	SLES, SLES for SAP, SLE HPC, SLE RT, SLED	life cycle: individual per package extended: none support level: none
Web and Scripting Additional Web server functionality	Base System, Server Applications	SLES, SLES for SAP, SLE HPC (default)	

¹ ESPOS: Extended Service Pack Overlay Support, LTSS: Long-Term Service Pack Support

² This is an *Extension*.

2.9.2 Available Extensions

Extensions add extra functionality to the system and require their own registration key, usually at additional cost. Usually, extensions have their own release notes documents that are available from <https://www.suse.com/releasesnotes>.

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









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

Additionally, there is the following extension which 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/> 

2.9.3 Derived and Related Products

This section lists derived and related products. These products have their own release notes documents, available at <https://www.suse.com/releasesnotes> .

- SUSE Linux Enterprise Server JeOS: <https://www.suse.com/products/server/jeos> 
- 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 Manager: <https://www.suse.com/products/suse-manager>  (current version is based on SLE 12 SP3)
- SUSE Linux Enterprise Real Time: <https://www.suse.com/products/realtime>  (current version is based on SLE 12 SP3)
- SUSE Enterprise Storage: <https://www.suse.com/products/suse-enterprise-storage>  (current version is based on SLE 12 SP3)
- SUSE OpenStack Cloud: <https://www.suse.com/products/suse-openstack-cloud>  (current version is based on SLE 12 SP3)

3 Installation and Upgrade

SUSE Linux Enterprise Server can be deployed in several ways:


- Physical machine
- Virtual host
- Virtual machine

- System containers
- Application containers

3.1 Installation

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

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 *Deployment Guide* at https://www.suse.com/documentation/sles-15/singlehtml/book_sle_deployment/book_sle_deployment.html .

3.1.1 System Roles for SUSE Linux Enterprise Server

With SUSE Linux Enterprise Server 15, it is possible to choose specific roles for the system. There are four roles available:

- **Minimal:** Set of packages needed to create a supportable and manageable basic system.
- **Text:** Set of packages commonly used for server environment, including base X server to run graphical applications.
- **KVM host:** Hypervisor and tools to set up a KVM-based virtualization host.
- **Xen host:** Hypervisor and tools to set up a Xen-based virtualization host.

3.1.2 Remote Installation via VNC

The installation can be done remotely using VNC, and there are two options for the client software: A native VNC viewer or a Web browser viewer. For the Web browser viewer we replaced a Java-applet based implementation with an implementation using JavaScript/WebSocket, as Java is no longer supported in mainstream browsers. Unfortunately, that has resulted in the loss of an encryption layer for the Web browser viewer.

The VNC connection on port 5801 is unencrypted. The connection on port 5901 continues to be encrypted.

3.1.3 Parted Supports Linux-Specific GPT GUID for Partitions

When Parted 3.1, the version shipped with earlier versions of SLE, was released, there was no Linux-specific GPT GUID. Therefore, it used the Microsoft Basic Data partition type for all new partitions.

With SLE 15, Parted 3.2 is shipped. This version uses the new Linux GPT GUID by default. If an old Linux GPT partition that uses the Microsoft Basic Data type is found, Parted will set the flag `msftdata` on it.

In partition editors and other GPT-enabled disk tools, such partitions may be mislabeled as *Windows Data Partitions* or similar. This affects the YaST Expert Partitioner, as well as `fdisk`, `gdisk`, etc.

The partition can be converted and the flag be cleared like this:

```
parted [DEVICE] set [PARTITION_NUMBER] msftdata off
```

3.1.4 Handling of Extension Repositories in AutoYaST

Starting with SLE 15, AutoYaST handles extension repositories in a more user-friendly way:

- AutoYaST automatically reorders extension repositories according to their dependent repositories during registration. That means the order of extensions in the AutoYaST profile is not important anymore.
- If dependent extensions are missing from the AutoYaST profile, AutoYaST will automatically register them. However, this only works for extensions that do not require a registration key. Extensions that require a registration key must be listed in the AutoYaST profile, including the registration key.

3.1.5 AutoYaST Configuration Is Done After Reboot

A regular YaST-based installation of SLES 15 is performed in a single stage. However, the AutoYaST installation process is divided into two stages. After the installation of the basic system, the system boots into the second stage during which the system is configured.

Make sure that the packages `autoyast2` and `autoyast2-installation` are installed by the first stage, so the second stage can be executed correctly.

Otherwise, an error will be shown before booting into the installed system.

3.1.6 YaST and AutoYaST Changes Because of Switch to firewalld

SuSEFirewall2 has been removed from SLES 15.

SLE 15 introduces firewalld as the new software firewall. This also incurs changes to installation and system management with YaST and AutoYaST.

YaST

firewalld already provides a command-line interface (`firewall-cmd`) and a new graphical interface (`firewall-config`) for configuration. Therefore, the YaST command-line interface and the GUI for SuSEFirewall2 have largely been removed. However, YaST continues to provide a common interface for opening ports for different services.

AutoYaST

AutoYaST profiles based on SuSEFirewall2 do not fit with the firewalld configuration. This meant that a new AutoYaST schema for configuring firewalld was needed. However, you can still use SuSEFirewall2-based profiles but are limited in terms of supported properties. This configuration will then be translated to firewalld rules. However, we recommend using the new schema and also checking the configuration when the system is installed.

For more information about configuring firewalld in AutoYaST, see *AutoYaST Guide, Firewall Configuration* (a draft version of the guide is provided at https://www.suse.com/documentation/sles-15/singlehtml/book_autoyast/book_autoyast.html#CreateProfile.firewall).

For more information about firewalld, also see *Section 4.3.7, “firewalld Replaces SuSEfirewall2 as Default Software Firewall”*.

3.2 Upgrade-Related Notes


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

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 https://www.suse.com/documentation/sles-15/single-html/book_sle_upgrade/book_sle_upgrade.html .

3.2.1 Product Registration Changes for HPC Customers

 This entry has appeared in a previous release notes document.

For SUSE Linux Enterprise 12, there was a High Performance Computing subscription named "SUSE Linux Enterprise Server for HPC" (SLES for HPC). With SLE 15, this subscription does not exist anymore and has been replaced. The equivalent subscription is named "SUSE Linux Enterprise High Performance Computing" (SLE-HPC) and requires a different license key. Because of this requirement, a SLES for HPC 12 system will by default upgrade to a regular "SUSE Linux Enterprise Server".

To properly upgrade a SLES for HPC system to a SLE-HPC, the system needs to be converted to SLE-HPC first. SUSE provides a tool to simplify this conversion by performing the product conversion and switch to the SLE-HPC subscription. However, the tool does not perform the upgrade itself.

When run without extra parameters, the script assumes that the SLES for HPC subscription is valid and not expired. If the subscription has expired, you need to provide a valid registration key for SLE-HPC.

The script reads the current set of registered modules and extensions and after the system has been converted to SLE-HPC, it tries to add them again.

Important: Providing a Registration Key to the Conversion Script

The script cannot restore the previous registration state if the supplied registration key is incorrect or invalid.

1. To install the script, run `zypper in switch_sles_sle-hpc`.
2. Execute the script from the command line as `root`:

```
switch_sles_sle-hpc -e <REGISTRATION_EMAIL> -r <NEW_REGISTRATION_KEY>
```

The parameters `-e` and `-r` are only required if the previous registration has expired, otherwise they are optional. To run the script in batch mode, add the option `-y`. It answers all questions with `yes`.

For more information, see the man page `switch_sles_sle-hpc(8)` and `README.SUSE`.

3.2.2 The User Space X Drivers `cirrus/mga/ast` Have Been Removed

The packages `xf86-video-cirrus`, `xf86-video-mga`, and `xf86-video-ast` have been removed in SLE 15. Kernel mode setting and mode-setting X drivers for these graphics cards have been available throughout the SLE 12 cycle and were used for all new SLE 12 installations. The user space X driver packages were only retained to ease upgrades from SLE 11.

If you are upgrading a machine from SLE 12 to SLE 15 that has previously been upgraded from SLE 11 to SLE 12, X may no longer start after the upgrade to SLE 15. If that is the case, rename or remove the file `/etc/X11/xorg.conf`, for example using:

```
sudo old /etc/X11/xorg.conf
```

3.2.3 Migrating from SLES 12 to SLES 15 when the HPC Module is Registered

If you are using SLES 12 with the HPC module registered, you cannot immediately upgrade to SLES 15, as only SLE HPC 15 will be offered as a migration target.

SLE HPC 15 will be offered as the only migration target whenever the HPC Module is registered on SLES-12 prior to migration. To migrate to SLES 15 instead, unregister the HPC Module before starting the migration. To do so, open a root shell and execute:

```
SUSEConnect -d -p sle-module-hpc/12/<ARCH>
```

Replace `<arch>` with the appropriate architecture, that is, `x86_64` or `aarch64`.

3.2.4 SMT Has Been Replaced by RMT

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

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

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

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

For more feature comparison between RMT and SMT, see <https://github.com/SUSE/rmt#rmt-and-smt>.


For more information about RMT, also see the new RMT Guide at <https://www.suse.com/documentation/sles-15>.

3.2.5 `/etc/SuSE-release` Has Been Removed

With SLE 15, the file `/etc/SuSE-release` has been removed. Previously, this file contained information on the version of SUSE Linux Enterprise that you were using.

Version information can now be found in /etc/os-release. The advantages of /etc/os-release are:

- The file exists across all major Linux distributions
- its format is well-specified and easily parseable
- it can be sourced by a shell script

For more information, see the man page of os-release: man 5 os-release in the installed system or online at <https://www.freedesktop.org/software/systemd/man/os-release.html> .

3.2.6 NIS Supports IPv6

With SLE 15, NIS is now fully IPv6 enabled. When migrating from SLE 12, existing configuration becomes outdated, for example, access control lists.

Check your configuration for IPv4-only configuration. Especially access control lists may need to be amended by IPv6 addresses. For example, check /var/yp/securenets.

3.2.7 qemu-kvm Wrapper Not Installed by Default

By default, the qemu-kvm wrapper binary is no longer installed on SLE 15. This change is transparent in new installations. However, in pre-SLE 15 environments, there may be VM configurations which use the legacy qemu-kvm wrapper. Migrating such a VM to a SLE 15 host will fail because the legacy wrapper qemu-kvm is not available.

Instead of using qemu-kvm, use QEMU by directly starting the qemu-system-ARCH binary.


To resolve the issue during migration:

- Change the VM configuration on the original host to use the qemu-system-ARCH emulator directly (preferred).
- Manually install the package qemu-kvm on the destination SLE 15 host.

3.2.8 Registration Rollback When Migrating From SLE 11 to SLE 15

When the migration is aborted after the registration step, a re-registration of the SLE 11 system is necessary.

Since the migration from SLE 11 to SLE 15 also involves switching the registration server from NCC to SCC, a rollback is not possible. Try to avoid aborting the upgrade or going back in the workflow.

If the migration is aborted, the system needs to be re-registered against NCC for further upgrade attempts. After re-registration, the synchronization of data between NCC and SCC can take some time. Make sure the re-registered system shows up on [before running the migration. \(https://scc.suse.com\)](https://scc.suse.com) 

3.2.9 `cryptconfig` Has Been Removed

Previous versions of SLE supported encrypting home directories individually via `cryptconfig`. This feature and the `cryptconfig` package have been removed in SLE 15.

To encrypt user data on SLE 15, encrypt the whole partition or volume which contains the home directories.

Important: Decrypt Before Upgrading

Before performing an upgrade from SLE 12, encrypted home directory images need to be decrypted. Otherwise, users will not have access to them after the upgrade.

3.2.10 Migrating Systems with BIOS RAID from SLES 11 to SLES 15

Systems with BIOS RAID handled via device mapper (DM-RAID) cannot be upgraded from SLES 11 to SLES 15 directly.

BIOS RAID, as provided by some chipsets or additional cards, is managed by the Linux kernel with either Device Mapper or via MD-RAID arrays. In SLES11, DM-RAID was used for some systems which is not supported in SLES 15 anymore. We recommend reinstalling these systems from scratch. Alternatively upgrade to SLES 12 first and then to SLES 15.

3.2.11 ReiserFS Support Removed

ReiserFS support for new installations was removed from YaST in SUSE Linux Enterprise 12 but upgrades were still supported.

With SUSE Linux Enterprise 15, support for ReiserFS will be completely removed from YaST and the installer will block the upgrade when it detects a ReiserFS file system.

For existing data partitions formatted with ReiserFS, we suggest converting them to Btrfs before migrating your system to SUSE Linux Enterprise 15.

3.2.12 Manually Selecting Repositories When Upgrading from SLES 12 to SLES 15

When upgrading a system which is registered against SCC, the registration server drives the selection of modules and repositories to be used during upgrade. This works well in most cases. However, there are scenarios in which modules or extensions are not selected as desired. For example, this can be the case when third-party software is installed and needs to be upgraded.

To allow manual selection of repositories, when booting the upgrade ISO, add `media_upgrade=1` to the kernel command line. This will make YaST skip the communication with SCC and you will have full control over the selection of repositories.

After the upgrade the system needs to be registered again, to update the registration data in SCC and to get access to the update channels for SUSE Linux Enterprise 15.

Notes regarding the standard migration (via SCC)

At the beginning of the upgrade, YaST registers the system to SLE 15. If the upgrade is aborted, YaST automatically rolls back the registration to the SLE 12 state. However, if the upgrade is aborted unexpectedly (for example, because of a power failure or hard reset), the registration state is not rolled back. In that case, you might need to run the rollback manually after booting the original system using the command `SUSEConnect --rollback`.

3.2.13 System-wide Locale/Keymap/Font Settings are not read from /etc/sysconfig/ anymore

Previously, there were different places for configuring a given setting.

For example, to set the system-wide locale, you could either:

- *write the settings in /etc/locale.conf*
- *use localectl*
- *write ROOT_USES_LANG in /etc/sysconfig/language if LANG was not already configured in /etc/locale.conf.*

This could be confusing, especially since settings in /etc/sysconfig/language usually override the locale settings used by users's shells only and therefore should not influence the system-wide locale.

Similar situations and similar problems could also be seen for the keymap/font settings:

- *The keyboard layout could be configured in both /etc/vconsole.conf and /etc/sysconfig/keyboard, the former having a higher priority.*
- *the font used by virtual consoles could be read from both /etc/vconsole.conf and /etc/sysconfig/console.conf, the former having a higher priority.*

With SLE 15, systemd does not read certain settings from the following files anymore:

- /etc/sysconfig/language for the system-wide locale settings (ROOT_USES_LANG)
- /etc/sysconfig/keyboard for the keyboard layout used by the virtual consoles (CONSOLE_FONT, CONSOLE_SCREENMAP, CONSOLE_UNICODEMAP)
- /etc/sysconfig/console for the font used by the virtual consoles (KEYTABLE)

All variables defined in /etc/sysconfig/language will still be used to override the system-wide locale and to define a different locale settings for users's shells as it is currently described in the official documentation.

To keep backward compatibility with the old systems, during the update of the systemd package, all variables mentioned will be migrated from sysconfig to their final destinations if they are not already defined there.

Replacement settings:

Locale:

- The system-wide locale can be changed via localectl(1) or YaST.
- The settings are stored in /etc/locale.conf, see man 5 locale.conf.

Virtual Consoles: The settings can instead be written directly in /etc/vconsole.conf. Also see man 5 vconsole.conf.

Keyboard:

- The system-wide locale can be changed via localect(1).
- The settings are stored in /etc/vconsole.conf, see man vconsole.conf(5).

3.3 For More Information

For more information, see *Section 4, “Architecture Independent Information”* and the sections relating to your respective hardware architecture.

4 Architecture Independent Information

Information in this section pertains to all architectures supported by SUSE Linux Enterprise Server 15 GA.

4.1 Kernel

4.1.1 Support for SAP HANA Workloads on Intel Optane DC Memory

SUSE Linux Enterprise Server 15 and SUSE Linux Enterprise Server for SAP Applications 15 add support for Intel Optane DC memory. This enables SAP workloads, such as SAP HANA to benefit from persistent memory in the future to shorten start times of the system and provide better overall system stability. Currently, configurations up to 12 TB of NVDIMMs plus 3 TB of regular DIMMS of supported memory and 4 socket machines have been tested. Additional configurations will be tested over time.

From a file system perspective, the XFS file system is supported for the NVDIMMs, with SAP HANA running in DAX mode. SUSE intends to keep the leading position as technology provider, working closely with SAP on future developments.

If there are `pmem` namespaces, these need to be destroyed before the installation. To mount persistent memory directly on boot, we recommend adding the `nofail` mount option in `/etc/fstab` as it can take a long time for the `/dev/pmem` devices to become usable.

For example:

<code>/dev/pmem0</code>	<code>/mnt/pmem0</code>	<code>xfs</code>	<code>dax,nofail</code>	<code>0</code>	<code>0</code>
<code>/dev/pmem1</code>	<code>/mnt/pmem1</code>	<code>xfs</code>	<code>dax,nofail</code>	<code>0</code>	<code>0</code>

Namespaces need to be created individually. That means, you need to execute the following command for each namespace you want to create:

```
ndctl create-namespace --mode=fsdax --map=dev
```


4.1.2 Device Error Prevention Enabled (CONFIG_IO_STRICT_DEVMEM)

The kernel build option `CONFIG_IO_STRICT_DEVMEM` has been enabled in the SLE kernel to prevent device errors. This option disables tampering with device state while a kernel driver is using the device.

Unfortunately, some vendor tools currently use such functionality. If you depend on such a tool, make sure to set the kernel boot parameter `iomem=relaxed`. Among others, this affects several firmware flash tools for POWER9 machines.

4.1.3 Intel Resource Director Technology Interface Update and Skylake Errata

Due to CPU defects identified in the Intel Skylake platform, most of the Resource Director Technology features are switched off by default on Skylake. Additionally, the mainline kernel is adopting a new interface for the resource management functions.

The features can be re-enabled using the kernel parameter `rdt`. For information on its usage, see `/usr/src/linux/Documentation/admin-guide/kernel-parameters.txt`. The old perf-based interface has been deprecated in favor of the new `resctrl` file system.

4.1.4 Page Cache Limit Is Now Opt-in cgroup Isolation

The kernel swaps out rarely accessed memory pages to use freed memory pages as cache to speed up file system operations, for example during backup operations. Certain applications use large amounts of memory for accelerated access to business data. Rarely accessed parts of this memory are subject of this swap out. Later access to swapped out memory regions results in poor application response times.

In previous SUSE Linux Enterprise versions there was a tunable known as page cache limit to mitigate this problem. This has now been replaced with a more mature mainline mechanism known as opt-in memory cgroup isolation.

A memory cgroup can define its so-called low limit (`memory.low_limit_in_bytes`) which works as a protection against memory pressure. Work loads that need to be isolated from outside memory management activity should set the value to the expected Resident Set Size (RSS) plus some head room. If a memory pressure condition triggers on the system and the particular group is still under its low limit, its memory is protected against being reclaimed. As a result, work loads outside of the cgroup do not need the aforementioned capping.

4.1.5 Kernel Address Space Randomization (KASLR) Enabled by Default

Kernel Address Space Randomization is one of several kernel hardening techniques that raise a practical hurdle for exploiting memory corruption vulnerabilities. Starting with SLE 15, this feature is enabled in the kernel by default. The feature can be switched off by specifying `nokaslr` option on kernel command line.

4.1.6 Support for Scalable Machine Check Architecture (Scalable MCA)

Scalable MCA improves hardware error reporting to better diagnose issues in AMD Zen processors. It provides a clearer, easier to use rules for the kinds of information supplied by the hardware when reporting errors.

This clearer separation of architectural and implementation-specific functions allows operating systems to better take advantage of architectural features.

In addition, it expands information logged in MCA banks to allow for improved error handling, better diagnosability, and future scalability.

4.1.7 Support for AMD Memory Encryption

To provide protection against physical attacks on a system, AMD SME can provide full or partial memory encryption depending on the use case, on AMD family 17h CPU processors. Full memory encryption means all DRAM contents are encrypted using random keys. This provides strong protection against cold boot, DRAM interface snooping and similar types of attacks. This technology is especially prominent for systems equipped with NVDIMMs whose contents remain intact after powering down the system.

Memory encryption support is present in SLE 15 kernels but not enabled by default. To enable it on compatible hardware (AMD family 17h CPU, with proper BIOS/UEFI support), supply the boot option `mem_encrypt=on`.

4.2 Kernel Modules

An important requirement for every enterprise operating system is the level of support available for specific environments. Kernel modules are the most relevant connector between hardware (“controllers”) and the operating system.

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

4.2.1 Removed Kernel Modules

In SLE 15, kernel drivers for the following device types and devices have been removed:

- analog TV/radio
- DMAPI
- IrDA
- ISDN
- miscellaneous obsolete USB devices
- obsolete Ethernet devices (AMD PCnet32, Atheros L2, DEC Tulip, PCI NE2000, AT-LAN-TEC/Realtek pocket adapter, SiS 900, SMC EtherPower II, Sun Microsystems "Happy Meal" Ethernet)
- obsolete sound devices
- PCMCIA
- Wireless USB (WUSB)

In addition, support for the ReiserFS file system has been removed, though there is a KMP that allows converting ReiserFS file system. (For more information, see [Section 3.2.11, "ReiserFS Support Removed"](#) .

4.2.2 Using Thunderbolt Devices

SLE 15 supports the Thunderbolt hardware interface. However, there are limitations to this support:

- Hot-unplugging is supported on a bus level but we cannot guarantee that every PCI driver also used for Thunderbolt devices supports it.
- GPUs must not be hotplugged/hot-unplugged, as doing so can crash the system.
- Thunderbolt monitors are not supported.

4.2.3 IPVS Has Been Moved From the HA Extension to the Base OS

IPVS (IP Virtual Server) implements transport-layer load balancing (Layer 4 LAN switching) in the Linux kernel. In SLES 12 and prior versions, IPVS was shipped only with the SUSE Linux Enterprise High Availability extension. However, IPVS is increasingly used outside the HA context, for example by Docker.

With SLES 15, IPVS has been moved into the base system. Other HA-related functionality that relies on IPVS remains part of the HA extension.

4.3 Security

4.3.1 GPG Does Not Support GPG V3 Keys Anymore, Resulting in Zypper/rpm Warnings

SLE 12 shipped with GPG 2.0, while SLE 15 includes GPG 2.2. In between these GPG versions, support for GPG V3 keys was removed. If your system's key database still contains GPG V3 keys, you may receive warnings about this when executing Zypper or `rpm` commands, as these commands are checking the integrity of the package database. These warnings take the form `warning: Unsupported version of key: V3`.

Usually, these warnings are benign, as these keys may have been used for repositories that are no longer enabled on the system or that have since had key updates. However, if these keys are still in active use by the upstream repository, they must be replaced as soon as possible:

- Package management tools in SLE 15 can no longer use them to verify package integrity.
- The keys in themselves are insecure. Hence, even though older package management tools will use them to verify integrity of packages, the result of this check cannot be trusted anymore.

To delete such keys, perform the following:

1. Run an `rpm` `tt` with high verbosity and check its output:

```
rpm -vv -qf /etc

ufdio: 1 reads, 18883 total bytes in 0.000006 secs
[...]
D: read h# 168 Header sanity check: OK
warning: Unsupported version of key: V3
```

```
[...]
```

In the example, header 168 is associated with an outdated key - the warning appears directly after the message that this specific header is being checked.

2. Find out the key number associated with the header: `rpm -q --querybynumber HEADER`
Replace `HEADER` with the required header number. In the example, that would be `168`.
This command returns a key identifier starting with `gpg-pubkey-`.
3. (Optional) Use the key identifier (`KEY_ID`) to learn more about the key: `rpm -qi KEY_ID`
4. Remove the key from the system: `sudo rpm -e KEY_ID`
5. If you continue to see warnings on subsequent uses of package management tools, repeat the procedure.

4.3.2 `su` Does Not Preserve the Value of `PATH`

For security reasons, the `su` command does not preserve the value of the environment variable `PATH` any more by default. To return to the behavior from SLE 12, open the file `/etc/default/su` and change the option `ALWAYS_SET_PATH` to `no`.

4.3.3 `systemctl stop apparmor` Does Not Work

In the past, there could be confusion over the difference between how the very similarly named `systemctl subcommands reload` and `restart` worked for AppArmor:

- `systemctl reload apparmor` properly reloaded all AppArmor profiles. (It was and continues to be the recommended way of reloading AppArmor profiles.)
- `systemctl restart apparmor` meant that AppArmor would stop, thereby unloading all AppArmor profiles and then restart which left all existing processes unconfined. Only newly started processes would then be confined again.

Unfortunately, `systemd` does not provide a solution within its unit file format for the issue posed by the `restart` scenario.

Starting with AppArmor 2.12, the command `systemctl stop apparmor` will not work. As a consequence, `systemctl restart apparmor` will now correctly reload AppArmor profiles.

To unload all AppArmor profiles, use the new command `aa-teardown` instead which matches the previous behavior of `systemctl stop apparmor`.

For more information, see https://bugzilla.opensuse.org/show_bug.cgi?id=996520 and https://bugzilla.opensuse.org/show_bug.cgi?id=853019.

4.3.4 New AppArmor Features to Restrict Processes

To properly protect processes, they must be safeguarded not only from files and network connections, but also from other process. For example, processes can be arbitrarily terminated by signals from other processes.

The version of AppArmor shipped with SLE 15 includes new features to further safeguard and restrict your processes. These features include:

- mount
- pivot_root
- ptrace
- signal

4.3.5 GnuPG Uses SHA-2 Family of Digests by Default

Research was published that showed weaknesses in the SHA-1 family of hashes for some applications. The use of stronger digests is advised for most applications.

The default behavior for GnuPG (`gpg2`) has been changed to use SHA-2 family digests for key certificates, default preferences stored in keys, and signature generation in the absense of a configuration file. GnuPG no longer generates a new configuration when called in an empty home. Existing GnuPG configurations are not altered. GnuPG continues to support SHA-1 digest generation and verification as mandated by OpenPGP standards.

4.3.6 All SLE 15 Packages Are Enabled for Address Space Layout Randomization

Security consists of layers of defense. One of those layers of defense is randomizing address for programs, so offsets and functions and similar are at randomized addresses on every start.

All SUSE Linux Enterprise 15 binaries are built with support for PIE (Position-Independent Executables) which will randomize all code layout in memory on every startup of the binary.

4.3.7 firewalld Replaces SuSEfirewall2 as Default Software Firewall

SuSEfirewall2 was originally tailored towards running a router with forwarding and/or NAT rules. This use case is rarely required anymore. Furthermore, the static nature of SuSEfirewall2 made it difficult to react to today's dynamic networking events like hotplugged network interfaces or virtual networking.

To allow greater flexibility in SLE 15, the default firewall has been switched to the firewalld upstream solution. It provides a resident daemon process which can dynamically adjust firewall rules on behalf of the user or other programs. SuSEfirewall2 is no longer available.

There is no automatic migration from SuSEfirewall2 to firewalld. To migrate an existing SuSEfirewall2 configuration to firewalld, you can use the script from the package `susefirewall2-to-firewalld`. However, after running the script, you still need to manually adjust and verify the resulting firewalld rules.

More technical information about firewalld could be found in the Security Guide at https://www.suse.com/documentation/sles-15/singlehtml/book_security/book_security.html#sec.security.firewall.firewalld.

4.4 Networking

4.4.1 Wicked: Using RFC 4361 DHCPv4 client-id on Ethernet

RFC 4361 updates the `client-id` defined in RFC 2132, section 9.14 to be compatible with DHCPv6 `client-id` (`duid`).

The use of an RFC 4361 is mandatory on Infiniband (RFC 4390) and is also required to perform DNS record updates in the same zone for DHCPv4 and DHCPv6 addresses also on Ethernet.

In SLE 15:

- ISC DHCP 4.3.x server supports the new RFC 4361 (required for DNS update)
- Wicked provides an option to send such a `client-id` and to automatically use a DHCPv6-based `client-id` in DHCPv4 (used on Infiniband).

To send the `client-id` during the installation, use `linuxrc` (also see <https://en.opensuse.org/SDB:Linuxrc>) with the following `ifcfg`:

```
ifcfg=eth0=dhcp,DHCLIENT_CLIENT_ID=01:03:52:54:00:02:c2:67,DHCLIENT6_CLIENT_ID=00:03:52:54:00:02:c2:67
```

For more information, see the documentation for the options `dhcp4 "create-cid"`, `dhcp6 "default-duid"` in `man 5 wicked-config`, `wicked duid --help`, and `wicked iaid --help`. The traditionally used RFC 2132 DHCPv4 `client-id` on Ethernet is constructed from the hardware type (`01` for Ethernet) and followed by the hardware address (the MAC address), for example:

```
01:52:54:00:02:c2:67
```

The RFC 4361 `client-id` starts with `0xff` (instead of the hardware type), followed by the DHCPv6 IAID (the interface-address association ID that describes the interface on the machine), followed by the DHCPv6 DUID (`client-id` which identifies the machine).

Using the above hardware type-based and hardware address-based DUID (LLT type used by default), the new RFC 4361 DHCPv4 `client-id` would be:


- Using the last bytes of the MAC address as the IAID:
`ff:00:02:c2:67:00:01:xx:xx:xx:xx:52:54:00:02:c2:67`
- When the IAID is a simple incremented number:
`ff:00:00:00:01:00:01:xx:xx:xx:xx:52:54:00:02:c2:67`

The `xx:xx:xx:xx` in the DUID-LLT is a creation timestamp. A DUID-LL (`00:03:00:01:$MAC`) does not have a timestamp.

4.4.2 Open vSwitch Has Been Updated to Version 2.8

Open vSwitch has been updated to version 2.8. Major changes are:

- `ovs-ofctl` can now accept and display port names in place of numbers. By default it always accepts names and in interactive use it displays them; use `--names` or `--no-names` to override. For more information, see `ovs-ofctl(8)` for details.
- `ovs-ofctl dump-flows` now accepts `--no-stats` to omit flow statistics.
- New `ovs-dpctl` command `ct-stats-show` to show connection tracking stats.
- DPDK log messages are redirected to the OVS logging subsystem. The log level can be changed by using `ovs-appctl vlog` commands for the `>dpdk` module (as for other modules). The lower bound can still be configured via extra arguments for DPDK EAL.
- `dpdkvhostuser` ports are marked as deprecated. They will be removed in an upcoming release.

- Support for DPDK v17.05.1.
- New support for multiple VLANs (802.1ad or "QinQ"), including a new dot1q-tunnel port VLAN mode.
- Added NAT support for user-space data path.
- Added FTP and TFTP support with NAT for user-space data path.
- Experimental NSH (Network Service Header) support in user-space data path.
- Tracing with ofproto/trace now traces through recirculation.
- New support for role-based access control
- New commands stp/show and rstp/show
- All features required by OpenFlow 1.4 are now implemented, so ovs-vswitchd now enables OpenFlow 1.4 by default (in addition to OpenFlow 1.0 to 1.3).
- Better support for OpenFlow 1.6 (draft).
- The learn action now supports a limit option
- OpenFlow 1.5 packet-out is now supported.
- Support for OpenFlow 1.5 field packet_type and packet-type-aware pipeline (PTAP).
- Added generic encapsulation and decapsulation actions (EXT-382). First supported use case is encapsulation/decapsulation for Ethernet.
- Added NSH (Network Service Header) support in userspace Used generic encap and de-cap actions to implement encapsulation and decapsulation of NSH header. For more info-provider, see the IETF NSH draft at <https://datatracker.ietf.org/doc/draft-ietf-sfc-nsh/> .
- ovs-vswitchd and ovsdb-server run as non-root users by default.
- Add --cleanup option to command ovs-appctl exit
- Use new tunnel port option packet_type to configure L2 vs. L3.
- In conjunction with PTAP tunnel ports can handle a mix of L2 and L3 payload.
- New vxlan tunnel extension gpe to support VXLAN-GPE tunnels.
- New support for non-Ethernet (L3) payloads in GRE and VXLAN-GPE.
- Added experimental support for hardware offloading.

- Hardware offloading is disabled by default.
- Hardware offloading is done through the TC interface.
- The next major version of OVS will introduce a change in the Conntrack API. Conntrack state is only available to the processing path that follows the `recirc_table` argument of the `ct()` action. Starting in OVS 2.9, this state will be cleared for the current processing path whenever `ct()` is called.

4.4.3 Intel® Omni-Path Architecture (OPA) Host Software

Intel Omni-Path Architecture (OPA) host software is fully supported in SUSE Linux Enterprise Server 15. Intel OPA provides Host Fabric Interface (HFI) hardware with initialization and setup for high performance data transfers (high bandwidth, high message rate, low latency) between compute and I/O nodes in a clustered environment.

For more information, see the Intel Omni-Path Architecture documentation at https://www.intel.com/content/dam/support/us/en/documents/network-and-i-o/fabric-products/Intel_OP_Software_SLES_15_RN_J98644.pdf.

4.4.4 Support for the IDNA2008 Standard for Internationalized Domain Names

The original method for implementing Internationalized Domain Names was IDNA2003. This has been replaced by the IDNA2008 standard, the use of which is mandatory for some top-level domains.

The network utilities `wget` and `curl` have been updated to support IDNA2008 through the use of `libidn2`. This update also affects consumers of the `libcurl` library.

4.5 Systems Management

4.5.1 Searching Packages Across All SLE Modules

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

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

This command contacts the SCC and searches all modules for matching packages.

4.5.2 SDT Markers in Select Applications and Libraries

SDT markers are static tracepoints included in the source code that expose certain information deemed useful by the application/library developers for various purposes including debugging and performance monitoring. Tools such as `perf`, `systemtap` and `bcc` can be used to record data provided at these tracepoints, and for subsequent processing.

In SLES 15, certain userspace applications and libraries (for example, `glibc`) are built with SDT markers enabled. This enhances the serviceability tooling.

4.5.3 Better AutoYaST Device Resize Handling

In SLE 12 and earlier, AutoYaST only supported resizing partitions but not logical volumes. Moreover, specifying the new size was quite limited, as values like `max` or `auto` were not allowed. Additionally, when using a percentage as the value of a resizing operation, the percentage was assumed to refer to the previous size of the partition. This was different from creating a partition where percentages refer to device size.

AutoYaST now supports resizing both partitions and logical volumes. Additionally, the `size` element will behave the same way, no matter whether a partition or logical volume is being created or resized: Percentages always refer to size of the whole device.

4.5.4 Zypper Return Code for Non-Fatal Failures

Similarly to other commands, Zypper signifies success exits with a return code of `0` and signifies failures with an error-specific non-zero return code. Prior to SLE 15 GA, Zypper would return `0` for some non-fatal failures. For example, this happened when a package was installed but there were issues with a post-installation script.

In cases of such non-fatal failures, Zypper now returns exit code 100 or higher. The list of exit codes is available in the man page (`man zypper`).

4.5.5 New SAP Applications Module

The SAP Applications module contains specialized tools for SAP Applications administration. The module is maintained and supported through the SUSE Linux Enterprise Server for SAP Applications product subscription. It can be installed using the online repository or the installer media.

Important

The default installation workflow of the SUSE Linux Enterprise Server for SAP base product depends on a graphical environment. If you decide to proceed with the "SLES for SAP" system role without installing the Desktop Applications Module, the message "Failed to select default product pattern gnome_basic. Pattern has not been found." will appear.

The package `kiwi-sap-template` provides a SLES for SAP package list template for customers that want to build custom images for a specialized SAP Applications use case. This package is provided as part of the SAP Applications module but is used for development purposes. To use it, the "Development Tools" module needs to be installed.

4.5.6 389 Directory Server Is Now the Primary LDAP Server

389 Directory Server is a full-featured LDAPv3-compliant server. It is the fit for modern environments and supports very large LDAP deployments.

The 389 Directory Server was included and set as the primary LDAP server on SLE 15. The YaST modules provided by the package `yast2-auth-server` were updated to ease the deployment and Kerberos integration of the new `389-ds`.

The OpenLDAP server is still available on the Legacy Module for migration purposes, but it will not be maintained for the entire SLE 15 lifecycle.

4.5.7 Support for UEFI HTTPS Boot

SLE 15 now supports the use of UEFI HTTP protocol for network booting in IPv4 or IPv6 environment. This uses a new extension to the DHCP options for URI-based identification for Network Boot Program (NBP), you can specify it with the `https://` scheme to boot remotely from an HTTPS server.

The authentication is done one-way. That means that the server authenticates an unauthenticated client. The server CA certificate needs to be enrolled in the client-side firmware to enable the HTTPS Boot feature. Securely encrypted connections can only be established with authenticated server using its enrolled certificate.

4.5.8 Support for Socket-Based Services Activation

Systemd allows for new ways of starting services, such as the so-called socket-based activation. Services which are configured to be started on demand will not run until it is needed, for example, when a new request comes in.

The YaST Services Manager has been extended to allow setting services to be started on-demand. Currently, only a subset of services supports this configuration. The current start mode is displayed in the column *Start* of the YaST Services Manager. In the drop-down box *Start Mode* of the YaST Services Manager, the mode *On-demand* will only be shown when it is available for the selected service.

Additionally, the table column *Active* has been adapted to show the correct value provided by Systemd.

4.5.9 Refactored YaST iSCSI LIO server

lio-utils, the former back-end of the YaST module iSCSI LIO Server was incompatible with current kernel modules and configFS interfaces. It was also based on Python 2. In addition, SUSEfirewall2 has been replaced by firewalld in SLE 15.

The SLE 15 version of the YaST module iSCSI LIO Server has been completely refactored. The following changes were made:

- *lio-utils* has been replaced with *targetcli-fb* now.
- SUSEfirewall2-related settings have been replaced by firewalld settings.
- The *Edit* button for LUNs has been removed in both *Add iSCSI target* and *Edit iSCSI target* page, because users should not edit a LUN configuration: If, for example, a LUN path is changed, that will lead a mismatch situation and errors on the initiator side
- In *Modify Initiator's ACLs* page, the *Copy* button has been removed, to avoid accidental leaking ACLs.

4.5.10 Support for Floppy Disks Has Been Removed from YaST

Starting with the version shipped in SLE 15, YaST does not have support for floppy disks anymore. For example, this means, that you can no longer install the boot loader to floppy disk or use AutoYaST files from floppy disk.

4.5.11 YaST Partitioner: Redesigned Back-end and UI Changes

Back-end Changes

The partitioning back-end previously used by YaST, `libstorage`, has been replaced by `lib-storage-ng` that is architected to allow new capabilities that were not possible before. For example, it is now possible to install a fully encrypted system without LVM using the automatic proposal and to correctly handle file systems placed directly on a disk without any partitioning.

UI Changes

Along with the library replacement, the `yast2-storage` module will be replaced by `yast2-storage-ng` which reimplements the storage code of YaST. Several outdated and less useful system views were removed from the partitioner in that rewrite:

- *Crypt Files* - Use LUKS-based encryption instead
- *Device Mapper*
- *Unused Devices*
- *Mount Graph*
- *Tmpfs* - Largely managed by systemd now

In addition, the *Hard Disks* system view does not display devices that cannot be manipulated using the partitioner. That includes:

- Unformatted DASDs
- Individual devices (that is, wires) of a multipath device
- Disks that are part of a BIOS RAID

Moreover, the *Configure* menu in the initial partitioner screen no longer includes the options *Provide Crypt Password* and *Configure Multipath*. During installation, any operation that necessitates a system rescan (such as using the *Rescan Devices* button in the same screen) will always ask the user what to do with inactive multipath systems and closed encrypted devices.

AutoYaST Changes

While the back-end that handles the AutoYaST section `<partitioning>` has changed, we kept compatibility in mind and there should be no changes to the XML layout.

4.6 Performance Related Information

4.6.1 `sapconf` SAP Tuning Tool Sets All Specified Tuning Values on OS Irrespective of Current Value

The previous solution only allowed `sapconf` to increase values, but in some cases a lower value may be the correct path to take. Therefore, `sapconf` needed to set all values irrespective of whether the current value is greater than or less than what `sapconf` wants to set.

`sapconf` provides a default set of values for SAP workloads which should apply to the majority of use cases. If a default `sapconf` value is not appropriate for any reason (for example, special workloads, support cases), then `sapconf` offers the possibility to enter own values.

4.6.2 NFS Tuning

On systems with a high NFS load, connections may block.

To work around such performance regressions with NFSv4, you can open more than one TCP connection to the same physical host.

Earlier versions of SLE supported the option `sharetransport` to remedy this. That option is no longer supported on SLE 15.

On SLE 15, this can be accomplished with the following mount options that will request multiple transport connections:

```
mount -o nconnect=N server:/path /mountpoint
```

In this case, `N` can be any number between 1 and 16 and defines the number of connections.

4.7 Storage

4.7.1 SMB Shares Used via mount or /etc/fstab Are Now Expected to use SMB 2.1 or Higher

The first version of the SMB network protocol, SMB1 is an old and insecure protocol and has been deprecated by its originator Microsoft (also see [SMBv1 is not installed by default](https://aka.ms/smb1rs3) (<https://aka.ms/smb1rs3>), [Stop Using SMB1](https://web.archive.org/web/20190227091836/https://blogs.technet.microsoft.com/filecab/2016/09/16/stop-using-smb1/) (<https://web.archive.org/web/20190227091836/https://blogs.technet.microsoft.com/filecab/2016/09/16/stop-using-smb1/>)). For security reasons, a kernel update for SLE 15 GA has changed the SMB kernel module (`cifs.ko` / `mount.cifs`) in a way that will break some existing setups: By default, the `mount` command on will now only mount SMB shares using newer protocol versions by default, namely SMB 2.1, SMB 3.0, or SMB 3.02.

Note that this change does not affect your installed Samba server or `smbclient` programs.

If possible, use an SMB 2.1 server. Depending on your SMB server, you may have to enable SMB 2.1 specifically:

- Windows has offered SMB 2.1 support since Windows 7 and Windows Server 2008 and it is enabled by default.
- If you are using a Samba server, make sure SMB 2.1 is enabled on it. To do so, set the global parameter `server max protocol` in `/etc/samba/smb.conf` to `SMB2_10` (for more possible values, see `man smb.conf`).

If your SMB server does not support any of the modern SMB versions and cannot be upgraded or you rely on SMB1's/CIFS's Unix extensions, you can mount SMB1 shares even with the current kernel. To do so, explicitly enable them using the option `vers=1.0` in your `mount` command line (or in `/etc/fstab`).

4.7.2 Updated Btrfs Subvolume Layout

SLE 15 introduces a new default Btrfs subvolume layout that aims for the following:

- Simplified snapshots and rollbacks
- Prevention of accidental data loss
- Better performance of databases and VM images stored in `/var`

Instead of using multiple Btrfs subvolumes for different subdirectories of `/var`, SLE 15 ships with a single subvolume for all of `/var`. This new subvolume has copy-on-write functionality disabled.

There is no defined way of upgrading to this new Btrfs subvolume layout. Therefore, if you want to take advantage of it, make sure to freshly install SLE 15 instead of upgrading.

For more information on the default Btrfs subvolume layout before and after this change, see <https://en.opensuse.org/SDB:BTRFS>.

4.7.3 `lvmlockd` Has Been Added As an Alternative to `clvm`

`clvmd` is a daemon that makes basic LVM2 functionality cluster-aware. However, due to its design, `clvmd` could make it hard to troubleshoot cluster-wide errors, such as cluster deadlock issues.

`lvmlockd` is a redesign that addresses design issues of `clvmd` and to add more functionality. For example, it supports both `sanlock` and `dlm` as the cluster locking managers, and it supports `lvmetad` in the cluster. It also allows scaling LVM2 for use in large virtualization/storage clusters.

For more information about `lvmlockd`, see the man page `lvmlockd(8)`.

4.7.4 XFS Realtime Volumes Feature Has Been Disabled

The XFS feature realtime volumes was abandoned upstream.

Starting with SLE 15, the kernel module for XFS is compiled without support for realtime volumes.

4.7.5 Alignment of Partition Size in the Storage Proposal and Expert Partitioner

To avoid the performance drop caused by excessive read-modify-write cycles, the partitions in a system must be properly aligned. Even beyond performance considerations, some types of partition tables require alignment to work. For example, DASD partition tables need to be aligned to tracks (usually 12 sectors).

In SLE 15, the expert partitioner of YaST takes alignment into consideration when creating and resizing partitions. It ensures alignment where it is required (such as DASD tracks) and encourages alignment where it can help performance, avoiding gaps between partitions in the process.

If you specify a size in the expert partitioner, the start and end of the partition will be aligned to ensure optimal performance and to minimize gaps. This may result in a slightly smaller partition (the difference is usually less than 1 MiB).

If you specify a custom region, the start and end will be honored as closely as possible, without attempting to optimize performance. However, mandatory alignment, such as DASD tracks, will take place. This option is best suited to creating very small partitions.

The same considerations for optimal alignment will also be taken into account while resizing an existing partition and calculating the minimal and maximal sizes suggested by the partitioner during that process.

4.8 Virtualization

4.8.1 Support for Nested Virtualization Performance Features in Newer AMD Processors

In nested virtualization, the hypervisor has to intercept and emulate most virtualization instructions in KVM guests in software. This slows down nested virtualization.

Newer AMD processors have support for hardware virtualization of common virtualization instructions, making software emulation unnecessary. These features in newer AMD processors are now supported, making nested virtualization faster.

4.8.2 KVM

4.8.2.1 KVM/libvirt: Support for Setting OEM Strings Table in SMBIOS

The SMBIOS defines an "OEM Strings" table whose purpose is to allow exposing arbitrary vendor specific strings to the operating system.

Management and orchestration services often need to convey information to virtual machines in the environment. For example, to set up new machines, OpenStack conveys sysprep information to the setup command cloud-init running in the instances it creates. SUSE Manager could also use SMBIOS information for inventory purposes.

With SLE 15, libvirt and QEMU now support setting the OEM strings table in the SMBIOS presented to virtual machines. Applications running inside KVM virtual machines can now access the information using dmidecode.

4.8.3 Xen

4.8.3.1 QEMU Guest Agent Is Now Supported on Xen

The QEMU guest agent has been available for KVM guests for some time. Custom tooling using the guest agent could not be used for Xen guests, posing a challenge in environments with both KVM and Xen.

With SLE 15, the QEMU guest agent is now supported in Xen guests. The package xen-tools-domU now contains a udev rule and a script that configures the guest agent device.

To use the guest agent in a Xen guest:

- Install the packages xen-tools-domU and qemu-guest-agent.
- Add a channel device for the agent to the libvirt XML.
- Restart the guest. After the restart, the guest agent will automatically use the channel device.

4.8.4 libvirt

4.8.4.1 virt-bootstrap: Creating libvirt LXC root file systems

In the past, creating a libvirt-lxc system container required using the virt-create-rootfs tool which has a few limitations.

In SLE 15, the new virt-bootstrap tool provides an easier way to create root file systems for system containers to enhance the user experience. Integrated with virt-manager, it uses Docker images or virt-builder templates to create the root file system either as a plain folder or as a qcow2 image with a backing chain for layers.

4.8.5 Others

4.8.5.1 Supported Offline Migration Scenarios

The following host operating system combinations will be fully supported (L3) for migrating guests from one host to another for SLES 15:

- SLES 12 SP2 to SLES 15
- SLES 12 SP3 to SLES 15
- SLES 12 SP4 to SLES 15 (once released)

4.8.5.2 Supported JeOS Image Formats

With SLE 15, there are now supported JeOS images for the following virtualization platforms:

- KVM and Xen PV
- Xen HVM
- VMware
- Microsoft Hyper-V
- OpenStack

4.9 Miscellaneous

4.9.1 MariaDB: Dates After 2038

In MariaDB, the data type TIMESTAMP is limited to dates until 2038.

For dates beyond 2038, use the data type DATETIME.

4.9.2 MariaDB: Default Encoding and Collation Changed to utf8mb4

The utf8 encoding of MariaDB only support Unicode codepoints up to three bytes. The string would be truncated at the first encountered codepoint that would be encoded with four bytes.

The default encoding and collation of MariaDB was changed to utf8mb4 which supports all codepoints.

4.9.3 MariaDB: Support for TokuDB Storage Engine

The version of MariaDB shipped with SLE 15 on Intel 64/AMD64 now supports TokuDB. TokuDB is a high-performance storage engine focused on scalability and operational efficiency.

Note that the TokuDB storage engine cannot be used when the transparent hugepages feature of the Linux kernel is enabled. To disable transparent hugepages, follow the instructions at <https://mariadb.com/kb/en/library/enabling-tokudb/#check-for-transparent-hugepage-support-on-linux>.

4.9.4 Plymouth/GDM May Hang If No Display Is Connected

When you are using the graphical boot target (with GDM) but there is no display connected, Plymouth may be unable to quit. This affects the start of systemd services that are normally started subsequent to Plymouth.

To diagnose whether a system is in the problematic status, remotely log in to it and run the command `systemctl list-jobs`. The system is affected if the `plymouth-quit-wait.service` is shown as running.

Any of the following methods can be used as a workaround:

- Connect the machine to a monitor.
- Add `plymouth.enable=0` in kernel boot options.
- Run command `plymouth quit` when the system is running to the status.

4.9.5 Graphics Chipset Compatibility under Wayland

The drivers for the following graphics chipsets do not yet support Wayland sessions:

- Nvidia GPUs running under the proprietary driver from Nvidia
- Cirrus Logic chipsets in QEMU virtual machines
- Matrox mgag20 chipset
- Aspeed graphics chipsets

In all of these cases, even if the Wayland stack is fully installed, GNOME will automatically fall back to starting an X session.

4.9.6 CUPS No Longer Supports System V-Style Interface Scripts

In SLE 15 GA, the CUPS print server has been updated to version 2.2.x.

Among other things, this version also includes a backward-incompatible change: For security reasons, it does no longer support System V-style interface scripts. Hence, the directory that stored them, /etc/cups/interfaces has also been removed.

To replicate the functionality of System V-style Interface Scripts in with CUPS 2.2.0 and later, create an interface script that is called as a normal CUPS filter.

To do so, create a PPD file that specifies the interface script in a cupsFilter line. You can copy an existing PPD file and strip out most of the options but you need to leave at least one paper size for your printer (with PageSize, PageRegion, ImageableArea, and PaperDimension) in the PPD file so that it passes the cupstestppd program's checks.

For example, a minimal PPD file with a cupsFilter line pointing to an existing interface script for text-only printing (for example, called /usr/lib/cups/filter/TextToPrinter) could look similar to this:

```
*PPD-Adobe: "4.3"
*FormatVersion: "4.3"
*FileVersion: "1.0"
*LanguageVersion: English
*LanguageEncoding: ISOLatin1
*PCFileName: "txtprntr.ppd"
*Product: "(Text Printer)"
*Manufacturer: "Text"
*ModelName: "Printer"
*NickName: "Text Printer"
*ShortNickName: "Text Printer"
*PSVersion: "(none) 0"
*cupsFilter: "text/plain 0 /usr/lib/cups/filter/TextToPrinter"
*OpenUI *PageSize/Media Size: PickOne
*OrderDependency: 10 AnySetup *PageSize
*DefaultPageSize: A4
*PageSize A4: ""
*CloseUI: *PageSize
*OpenUI *PageRegion/Media Size: PickOne
*OrderDependency: 10 AnySetup *PageRegion
*DefaultPageRegion: A4
```

```
*PageRegion A4: ""
*CloseUI: *PageRegion
*DefaultImageableArea: A4
*ImageableArea A4: "0 0 595 842"
*DefaultPaperDimension: A4
*PaperDimension A4: "595 842"
```

For more general information about Print Filters, see https://en.opensuse.org/SDB:Using_Your_Own_Filters_to_Print_with_CUPS.

4.9.7 No Default Compose Key Combination

In previous versions of SLE, the compose key combination allowed typing characters that were not part of the regular keyboard layout. For example, to produce "å", you could press and release Shift-Right Ctrl and then press a twice.

Starting with SLE 15, there is no longer a predefined compose key combination because Shift-Right Ctrl does not work as expected anymore.

- To define a system-wide custom compose key combination, use the file /etc/X11/Xmodmap and look for the following lines:

```
[...]
!! Third example: Change right Control key to Compose key.
!! To do Compose Character, press this key and afterwards two
!! characters (e.g. `a' and `^' to get 342).
!remove Control = Control_R
!keysym Control_R = Multi_key
!add Control = Control_R
[...]
```

To uncomment the example code, remove the ! characters at the beginning of lines. However, note that the setup from Xmodmap will be overwritten if you are using setxkbmap.

- To define a user-specific compose key combination, use your desktop's keyboard configuration tool or the command-line tool setxkbmap:

```
setxkbmap [...] -option compose:COMPOSE_KEY
```

For the variable COMPOSE_KEY, use your preferred character, for example ralt, lwin, rwin, menu, rctl, or caps.

- Alternatively, use an IBus input method that allows typing the characters you need without a Compose key.

4.9.8 MariaDB Has Been Upgraded to 10.2

The MariaDB package has been upgraded to the 10.2 series that brings many new features and bug fixes. The list of major changes can be found at <https://mariadb.com/kb/en/library/changes-improvements-in-mariadb-102/>.

The update to the new MariaDB package generally does not cause issues. However, there are certain incompatible changes that need to be considered during this process. For example, InnoDB is the default storage engine now, some options have updated default values, some options have been removed/renamed etc. For more information about upgrading, see the upgrade notes at <https://mariadb.com/kb/en/library/upgrading-from-mariadb-100-to-mariadb-101/> and <https://mariadb.com/kb/en/library/upgrading-from-mariadb-101-to-mariadb-102/>.

MariaDB now uses the client library `libmariadb3` instead of the library `libmysqlclient`. The library `libmariadb3` is provided by the package `mariadb-connector-c`.

5 AMD64/Intel 64 (x86_64) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 GA for the AMD64/Intel 64 architectures.

5.1 System and Vendor Specific Information

5.1.1 Support for 32-bit Runtimes on Intel 64/AMD64 (x86-64)

Some tools to set up software or hardware are still compiled as 32-bit binaries.

SUSE Linux Enterprise 15 contains a 32-bit environment to run such applications on the architectures Intel 64/AMD64 (x86-64). The support targets tools to set up software or hardware. Other 32-bit applications may work with the given environment, but the environment is not intended to be a full replacement for a 32-bit installation.

Building 32-bit applications is not supported on SUSE Linux Enterprise 15.

5.1.2 TPM 2.0 Software Stack Has Been Updated

The upstream projects for Intel's TPM 2.0 Software Stack have introduced major changes to the project structure. Notably, the resource manager daemon has been replaced by a new implementation that fixes stability and security issues.

The packaging has been adjusted to the upstream changes. The previous resource manager daemon, `resourcemgr`, which was previously part of the `tpm2-0-tss` package has been dropped. The new package `tpm2.0-abrmd` provides the new resource manager implementation (`tpm2-abrmd` / `tabrmd`).

6 POWER (ppc64le) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 GA for the POWER architecture.

6.1 Support for ibmvnic Networking Driver

The kernel device driver `ibmvnic` provides support for vNIC (virtual Network Interface Controller) which is a PowerVM virtual networking technology that delivers enterprise capabilities and simplifies network management on IBM POWER systems. It is an efficient high-performance technology.

When combined with SR-IOV NIC, it provides bandwidth control Quality of Service (QoS) capabilities at the virtual NIC level. vNIC significantly reduces virtualization overhead resulting in lower latencies and fewer server resources (CPU, memory) required for network virtualization. For a detailed support statement of `ibmvnic` in SLES, see <https://www.suse.com/support/kb/doc/?id=7023703>.

6.2 Support for POWER9 24x7 Counters Has Been Added

SLES 15 adds support for the new version of the hypervisor API which is used to access 24x7 performance counters on POWER9 systems.

6.3 Support for POWER9 Radix Page Tables When Running as KVM Guest

Using hash memory page tables is not efficient when running as a KVM guest.

With SLES 15, radix page tables are supported on POWER9.

6.4 GLIBC Support for POWER9

GLIBC provides full support for the POWER9 architecture and exposes its features via AT_PLATFORM (power9) and AT_HWCAP2 (darn, ieee128 and arch_3_00).

6.5 Optimized GLIBC for POWER9

SLES 15 provides POWER9-optimized versions of `strcmp` and `strncmp`, via GNU IFUNC, which are automatically used without requiring an application rebuild.

6.6 Support for POWER9 Has Been Added to GDB

GDB now supports disassembling, recording and replaying new Power ISA 3.0 instructions. Additionally, GDB now supports basic encoding and decoding of 128-bit IEEE floating-point types.

6.7 Support for POWER9 PMU Events Has Been Added to the perf Tool

When executed on a POWER9 system, the perf tool now supports listing or specifying POWER9 PMU events by name.

6.8 GLIBC/GCC Support of Standard Math Functions for _Float128 (POWER9)

The math library now implements 128-bit floating point operations as defined by ISO/IEC/IEEE 60559:2011.


6.9 Support for POWER9 Has Been Added to PAPI

PAPI package updated to include support for POWER9 processors.

6.10 Support for POWER9 Has Been Added to LibPFM

The LibPFM package has been updated to include support for POWER9 processors.

7 IBM Z (s390x) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 GA for the IBM Z architecture. For more information, see https://www.ibm.com/developerworks/linux/linux390/documentation_novell_suse.html 

IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114), subsequently called z196 and z114.

7.1 Hardware

Support for IBM z13 and IBM z14 Machines

SUSE Linux Enterprise Server 15 GA has improved support for IBM z13 and IBM z14 machines. This following new features are supported:

- With z14, the hardware provides an indication of the configuration level of SIE, for example LPAR or KVM. IBM z14 sample configurations help to analyze and optimize KVM performance.
- CPU-MF Hardware Counters are added for IBM z13 and z13s. You can now access counters from the MT-diagnostic counter set that is available with IBM z13. You can also specify z13 specific counters using their symbolic event names and obtain counter descriptions with the `lscpumf` utility.
- For IBM z14 machines, access to the host's INVALIDATE PAGE TABLE is provided via the Guest Address-Space-Control Element (ASCE). The KVM host kernel can use the host DAT-enhancement 1 facility to avoid unnecessary purging of guest TLB entries.

- Store Hypervisor Information (STHYI) from LPAR is available in KVM. Non-privileged user-space applications running on KVM can retrieve hypervisor capacity data through the LPAR if not provided by the Linux kernel.
- KVM guests can now use CPU features, including CPACF functions, that were introduced with IBM z14.
- The LLVM compiler supports IBM z14 instructions for improved performance.
- OpenSSL, `ibmca`, and `libica` support IBM z14 instructions for AES-GCM-based encryption of data in flight.
- Support for the True Random Number Generator (TRNG, CPACF MSA7) in IBM z14 machines via CPACF. This improves the availability of random data in the kernel entropy pool.
- `libica` supports hardware acceleration for the SHA3 algorithm (CPACF MSA6) using CPACF hardware in IBM z14 machines.
- Improved performance on IBM z14 machines through enhanced instruction set support in the toolchain.
- The SIMD instructions for IBM z14 can be used in user space for improved performance of analytic workloads and math libraries.

7.2 Virtualization

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

- Standard network boot setups can be used to deploy KVM guests.
- `LOADPARM` and `BOOTPROG` are fully supported. A boot menu selection is available during IPL, for example, to recover from a defective KVM guest operating system.
- Keyless Guests are supported for performance gains through improved memory handling for workloads running on Linux.
- Guarded Storage Facility is supported for improved performance of all Java workloads on KVM virtual servers.
- Store Hypervisor Information (STHYI) from LPAR is available in KVM. Non-privileged user-space applications running on KVM can retrieve hypervisor capacity data through the LPAR if not provided by the Linux kernel.

- Machine checks caused by failing KVM guests are now targeted at the KVM virtual server instead of the KVM hypervisor, thus making the hypervisor more resilient.
- TLB Purge Enhancements are supported under KVM. This improves performance for KVM guests, in particular when subject to memory pressure.
- Transparent Facility Bit Handling is supported. Hardware functions that do not need a specific configuration in the KVM hypervisor are enabled for KVM guests.
- The IBM Call Home feature is enabled for KVM.

7.2.1 SIMD Extensions for IBM z14

The SIMD instructions for IBM z14 can be used in user space for improved performance of analytic workloads and math libraries.

7.3 Storage

7.3.1 No dasd_reload During Installation

In previous versions of SLES, when a DASD was activated during the installation, other DASDs may be renumbered. This renumbering could be confusing to users.

Starting with SLES 15, the call to `dasd_reload` has been removed from the installation. For addressing disks on z Systems, YaST now uses udev device names primarily. This prevents issues with disk name changes and is similar to the behavior on other hardware architectures.

7.4 Network

7.4.1 Support for SET VNIC_CHARS in qeth

qeth now supports `SET VNIC_CHARS`. You can configure MAC address flooding, learning, forwarding, and takeover behavior for HiperSockets devices.

7.5 Security

7.5.1 libica Supports FIPS 140-2 Mode

The FIPS PUB 140-2 Security Requirements for Cryptographic Modules specify that cryptographic modules in FIPS mode must only use NIST-approved algorithms and perform integrity checks and a self-test upon activation.

In SLES 15, `libica` is enabled for FIPS 140-2 certification and supports a FIPS mode. To enable this mode, add the boot parameter `fips=1` which will set the flag `/proc/sys/crypto/fips_enabled` to `1`

7.5.2 OpenSSH Supports Using Hardware Security Modules (HSMs) for Connecting via OpenSSL

In SLE 15, OpenSSH can use hardware security modules (HSM) that are available through the use of the OpenSSL engines `openssl-ibmca` and `openssl-ibmpkcs11`. Previously, the OpenSSH `seccomp` filter denied some system calls from being made, preventing the use of `openssl-ibmca` and `openssl-ibmpkcs11`. The update on `seccomp` filter enabled such system calls on IBM Z.

7.5.3 Support for libseccomp in systemd

With SLES 15, you can restrict the runtime environment, system calls and network use per application, in particular for containers.

7.5.4 Concurrent Support of Multiple Crypto Domains

With SLES 15, the support for generic cryptographic device drivers supports multiple cryptographic domains simultaneously.

7.5.5 Support for Transforming Secure Keys Into Protected Keys

Using a new device driver, cryptographic applications can generate protected keys from secure keys or from clear keys. Protected keys can be used by CPACF for accelerated encryption and decryption.

7.6 Reliability, Availability, Serviceability (RAS)

7.6.1 Zone Awareness for lsmem / chmem

When setting hotplug memory online or offline using chmem, you can now specify a memory zone. The output of lsmem shows the available memory zones.

7.6.2 lsmem and chmem Have Been Moved to util-linux

In SUSE Linux Enterprise Server 15, the IBM Z-specific commands lsmem and chmem are available in the common Linux tool package util-linux.

7.6.3 Parted Has Been Upgraded to Use fdasd/vtoc Code Base from s390-tools

The partitioning tool GNU Parted now uses the fdasd/vtoc code base from a current version of the IBM Z-specific package s390-tools.

7.6.4 Support for Uprobes

SLES 15 eases debugging user-space programs by enabling the uprobes architecture backend for IBM Z architecture.

7.6.5 Persistent Device Configuration

The following SUSE-supplied commands are now deprecated:

- ctc_configure
- dasd_configure
- qeth_configure
- zfcplib_configure
- zfcplib_host_configure

With SLES 15, as an intermediate step, these scripts have been modified to use the IBM-supplied commands chzdev and lszdev. These commands will be removed in a future release.

If you are using the SUSE-supplied scripts, discontinue their use and directly use the commands chzdev and lszdev provided by IBM in the package s390-tools.

7.6.6 CPU Speed Information Provided on procfs Interface

Statistics and dynamic CPU speed can now be obtained through the cpuinfo interface for improved debugging.

7.6.7 Support for DASD Block Layer Discard

SLES 15 includes support for the Linux discard function that releases unused space on z/VM VDISKS.

7.7 Performance

7.7.1 Single Increment Assignment of Memory

A new option for the “Attach Storage Element” SCLP command to speed up memory hotplug is available.

7.7.2 Backend Support for eBPF

In SLES 15, the Extended Berkeley Packet Filter (eBPF) is supported on IBM Z. It provides enhanced performance over other packet filters.

7.7.3 Improved CRC-32 Algorithm in the Kernel

SLES 15 ships with an optimized CRC-32 implementation that uses the Vector Extension Facility on the IBM Z architecture.

7.7.4 IBM z13 Lock Optimizations in glibc

Improved performance for applications via enhanced instruction support in glibc for IBM z13 machines.

7.7.5 Guest Kernel Support to Avoid Unnecessary TLB Purges

The Linux kernel now tags pages that are not used as part of a page table, so that the hypervisor can avoid unnecessary purging of guest TLB (translation lookaside buffer) entries.

7.7.6 Kernel Interface for the Guarded Storage Facility Added to Improve Java Performance

Optimized Java processes improve performance for many Java applications.

7.8 Miscellaneous

7.8.1 31-Bit Libraries Have Been Removed

In SLES 15, 31-bit libraries are not available anymore.

8 ARM 64-Bit (AArch64) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 GA for the AArch64 architecture.

8.1 Raspberry Pi Using Device Tree From Firmware

The U-Boot bootloader shipped with SUSE Linux Enterprise Server for ARM 12 SP3 for the Raspberry Pi passed its own Device Tree to Linux. That way, all user-configured Device Tree modifications by its firmware were discarded. The only way to modify the Device Tree then was to override it with a user-provided bcm2837-rpi-3-b.dtb file.

The updated u-boot-rpi3 package reuses the Device Tree passed on by raspberrypi-firmware. The new raspberrypi-firmware-dt package must be present for this to work as expected.

If you have installed custom Device Tree files (bcm2837-rpi-3-b.dtb), you should remove and replace them with Device Tree Overlay files (.dtbo) that can be enabled via the new user-supplied /boot/efi/extraconfig.txt configuration file. For more information, see the Raspberry Pi config.txt documentation.

8.2 Cavium ThunderX2 CRB Firmware Requirements

Unlike SUSE Linux Enterprise Server 12 SP3, SUSE Linux Enterprise Server 15 relies on the firmware to indicate which SMMU version is available in hardware. Early versions of the Cavium ThunderX2 "Sabre" reference boards did not indicate the SMMU model in the IORT ACPI table.

Update the firmware of boards that do not correctly report the SMMU model:

- For the ThunderX2 "Sabre" reference board, make sure the AMI firmware version is `0ACKL006` or newer.
- For ThunderX CRB systems, make sure the AMI firmware version is `0ACGA018` or newer.

9 Packages and Functionality Changes

This section comprises changes to packages, such as additions, updates, removals and changes to the package layout of software. It also contains information about modules available for SUSE Linux Enterprise Server. For information about changes to package management tools, such as Zypper or RPM, see [Section 4.5, "Systems Management"](#).

9.1 New Packages

9.1.1 Package `insserv-compat` Has Been Added to SAP Application Server Base Pattern

SAP Applications depends on a `sapinit` System V script that is widely used in all SAP Applications.

The package `insserv-compat` was included to be installed on the SAP Applications Server Base pattern by default, as this pattern is used and recommended to configure a SAP Application Server. That way, it will provide the System V compatibility until SAP fully adopts systemd unit scripts.

9.1.2 OpenJDK 10 Has Been Added

SUSE Linux Enterprise 15 ships with OpenJDK 10. Note that due to the accelerated development cycle of Java, we expect to move OpenJDK 10 to the Legacy Module and replace it by a then-current OpenJDK version in SUSE Linux Enterprise 15 SP1.

For more information about supported Java versions, also see [Section 10.5, “Supported Java Versions”](#).

9.1.3 PostgreSQL: psqlODBC Driver Has Been Added to Replace unixODBC

The unixODBC driver for PostgreSQL is not maintained anymore.

SLE 15 does not include the unixODBC driver anymore. However, we have added the psqlODBC driver from the upstream PostgreSQL project which is much better supported.

9.1.4 wodim Has Been Replaced by cdrtools

wodim was created as fork of cdrtools. Unfortunately, the wodim project stagnated over the years.

SLE 15 migrates back to using cdrtools. This means that some tools have been renamed. The following package names have changed:

- genisoimage has been renamed to mkisofs
- wodim has been renamed to cdrecord
- icedax has been renamed to cdda2wav

cdrkit-cdrtools-compat is no longer supplied. It only provided symbolic links for compatibility between cdrtools and wodim. If you were using it, no changes are necessary. If you were using the replaced packages above, executable binaries were renamed accordingly.

9.1.5 nginx Has Been Added

Starting with SUSE Linux Enterprise 15, the Web server software nginx is supported.

9.1.6 UnRAR Has Been Replaced by unar

UnRAR is freeware command-line application for extracting RAR archives. Unfortunately, it is non-free.

In SLE 15, The Unarchiver command-line tool, which is LGPL-licensed (package unar, binaries unar and lsar), has replaced UnRAR.

Unarchiver supports the same archive formats (including RAR5), except for UUE, JAR, and limited support for ARJ (no multi-part) and ACE (no support for Ace 2.0).

UnRAR and Unarchiver are not completely CLI-compatible, as they have a different set of options. Because of this, a simple wrapper script was added within the `unrar_wrapper` package (with a symbolic link to `/usr/bin/unrar`). This script transforms a subset of `unrar` commands to `unar` and `lsar` to provide a backwards compatibility:

- Supported commands: `l[t[a],b]`, `t`, `v[t[a],b]`, `x`.
- Supported options: `-o+`, `-o-`, `-or`, `-p`
- Other: `files`, `@listfiles` and `path_to_extract/` (only for extracting)
- Return codes: `0` (success), `1` (error), `2` (invalid argument)

For more information about functionality supported by the wrapper, see https://github.com/openSUSE/unrar_wrapper.

9.1.7 ntpd Has Been Replaced With Chrony

The time server synchronization daemon `ntpd` has been replaced with the more modern daemon Chrony.

This change means that AutoYaST files with an `ntp_client` section need to be updated to a new format for this section. For more information about the new AutoYaST `ntp_client` format, see *AutoYaST Guide*, section *NTP Client* (https://www.suse.com/documentation/sles-15/singlehtml/book_autoyast/book_autoyast.html#Configuration.Network.Ntp).

To sync time in intervals, YaST sets up a cron configuration file. From SLE 15 on, the configuration file used for this is owned by the package `yast2-ntp-client` (previously no package owned it). The configuration file has been renamed from `novell.ntp-synchronization` to `suse-ntp_synchronization` to be consistent with other cron configuration files. The upgrade from `ntpd` to Chrony is performed automatically during the SLE upgrade: If a file with the old name is found, it will be renamed and references to `ntpd` in it will be replaced by `chrony`.

`ntpd` has been moved to the Legacy module. For more information, see *Section 9.6.4, "Legacy Module: ntpd is now part of the Legacy Module"*.

9.1.8 Open MPI Has Been Replaced by Open MPI 2

Open MPI 1.x has reached its end-of-life upstream, as Open MPI 3.0 has been released.

With SLE 15, Open MPI 2 (package `openmpi2`) has been added and should be used.

Open MPI 1 (package `openmpi`) has been moved to the Legacy module.

9.1.9 `lshw` Has Been Added

`lshw` provides detailed information on the hardware configuration of IBM Z and POWER machines (it is not available on Intel 64/AMD64). It can report exact memory configuration, firmware version, mainboard configuration, CPU version and speed, cache configuration, bus speed, etc.

9.1.10 BPF Compiler Collection (BCC) Has Been Added

BCC is a toolkit for creating efficient kernel tracing and manipulation programs. It uses Extended Berkeley Packet Filters (eBPF) and includes tools to capture and analyze system or application data such as disk I/O stats, database query latencies, filesystem stats, network stats, method calls in high-level languages (often used to develop enterprise applications), and many more.

9.2 Updated Packages

9.2.1 KIWI Has Been Updated to 9.15.3

In SLE 15, KIWI has been updated to version 9.15.3. Version 9 of KIWI is a complete rewrite of the software that, while keeping general compatibility, also includes many new features. This update also fixes several bugs related to building JeOS. Given the newer codestream, this update also simplifies support of the tool.

For a comparison between KIWI 7 and KIWI 9, see https://opensource.suse.com/kiwi/overview/legacy_kiwi.html.

9.2.2 LibreOffice Has Been Updated to Version 6.1

LibreOffice has been updated to the new major version 6.1. For information about major changes, see the LibreOffice 6.1 release notes at <https://wiki.documentfoundation.org/ReleaseNotes/6.1>.

9.2.3 PostgreSQL Has Been Upgraded to Version 10



This entry has appeared in a previous release notes document.

SLES 12 SP4 and SLES 15 ship with PostgreSQL 10 by default. To enable an upgrade path for customers, SLE 12 SP3 now includes PostgreSQL 10 in addition to PostgreSQL 9.6 (the version that was originally shipped).

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



Important: PostgreSQL Upgrade Needs to Be Performed Before Upgrade to New SLES Version

Neither SLES 12 SP4 nor SLES 15 include PostgreSQL 9.6. However, availability of PostgreSQL 9.6 is a requirement for performing the database upgrade to the PostgreSQL 10 format. Therefore, you must upgrade the database to the PostgreSQL 10 format before upgrading to the desired new SLES version.

Major New Features

The following major new features are included in PostgreSQL 10:

- Logical replication: a publish/subscribe framework for distributing data
- Declarative table partitioning: convenience in dividing your data
- Improved query parallelism: speed up analyses
- Quorum commit for synchronous replication: distribute data with confidence
- SCRAM-SHA-256 authentication: more secure data access

PostgreSQL 10 also brings an important change to the versioning scheme that is used for PostgreSQL: It now follows the format *major.minor*. This means that minor releases of PostgreSQL 10 are for example 10.1, 10.2, ... and the next major release will be 11. Previously, both the parts of the version number were significant for the major version. For example, PostgreSQL 9.3 and PostgreSQL 9.4 were different major versions.

For the full PostgreSQL 10 release notes, see <https://www.postgresql.org/docs/10/release-10.html>.

Upgrading

Before starting the migration, make sure the following preconditions are fulfilled:

1. The packages of your current PostgreSQL version must have been upgraded to their latest maintenance update.
2. The packages of the new PostgreSQL major version need to be installed. For SLE 12, this means installing `postgresql10-server` and all the packages it depends on. Because `pg_upgrade` is contained in the package `postgresql10-contrib`, this package must be installed as well, at least until the migration is done.
3. Unless `pg_upgrade` is used in link mode, the server must have enough free disk space to temporarily hold a copy of the database files. If the database instance was installed in the default location, the needed space in megabytes can be determined by running the following command as `root`: `du -hs /var/lib/pgsql/data`. If there is little disk space available, run the command `VACUUM FULL` SQL command on each database in the PostgreSQL instance that you want to migrate. This command can take very long.

Upstream documentation about `pg_upgrade` including step-by-step instructions for performing a database migration can be found locally at `file:///usr/share/doc/packages/postgresql10/html/pgupgrade.html` (if the `postgresql10-docs` package is installed), or online at <https://www.postgresql.org/docs/10/pgupgrade.html>. The online documentation explains how you can install PostgreSQL from the upstream sources (which is not necessary on SLE) and also uses other directory names (`/usr/local` instead of the `update-alternatives` based path as described above).

9.2.4 OpenSSH Has Been Updated to Version 7.6p1

OpenSSH was updated to 7.6p1 for updated and improved security features. The package follows the upstream recommendation of disabling less secure legacy cryptography options. This includes the `ssh-dss` (DSA) public key algorithm.

The `ssh-dss` (DSA) public key algorithm can be re-enabled using the configuration option `HostKeyAlgorithms`.

- On the command line:

```
ssh -oHostKeyAlgorithms=+ssh-dss user@legacyhost
```

- In the configuration file `~/.ssh/config`:

```
Host somehost.example.org
    HostKeyAlgorithms +ssh-dss
```

9.2.5 Squid Has Been Updated to Version 4

Squid has been updated to the 4.x branch. Current users updating from squid 3.5.x (and earlier) should update their configuration files and get acquainted with the new features.

For more information, see the Squid 4 Release Notes at <http://www.squid-cache.org/Versions/v4/RELEASENOTES.html>.

9.2.6 Ceph Client Packages Have Been Updated to Upstream Release "Mimic" (13.0)

Keep Ceph client packages aligned with latest version of SUSE Enterprise Storage.

In SLE 15, the Ceph client packages (`ceph-common`, `librados`, `librbd`, etc.) have been updated to be based on the latest upstream Ceph release "Mimic" (13.0) to ensure they work optimally with the upcoming SUSE Enterprise Storage 6.

9.2.7 LIO target utilities and libraries Have Been Updated

In SLE 15, LIO-related packages in userspace have been updated to current versions. This affects the packages `python-configshell-fb` (now: version 1.1), `python-rtslib-fb` (now: version 2.1), `targetcli-fb` (now: version 2.1), and `tcmu-runner` (now: version 1.3). For more information about the updates, see the package change log.

9.2.8 Wireshark Qt UI Replaces Deprecated GTK+ UI

The GTK+ user interface of the Wireshark network protocol analyzer has been deprecated by the upstream project.

The Wireshark Qt interface is now shipped in the package `wireshark-ui-qt`.

9.2.9 DPDK Has Been Updated, libdpdk Package Has Been Added

SLE 15 ships with DPDK (Data Plane Development Kit) 17.11.

Because certain applications may need specific versions of the DPDK library, SLE 15 ships with a versioned package of that library (currently, `libdpdk-17_11`). For the future, this enables installing more than one version of the DPDK library at the same time.

9.3 Removed Packages and Features

The following packages have been removed in this release.

The following packages and groups of packages have been removed or replaced in SUSE Linux Enterprise 15:

- 31-bit libraries for the IBM Z architecture. For more information, see [Section 7.8.1, “31-Bit Libraries Have Been Removed”](#)
- `atftp`. For more information, see [Section 9.3.4, “atftp Has Been Removed”](#).
- `cfengine`
- CIM providers (but not the CIM infrastructure)
- Support for dial-up networks via `rszs` and `wvdial`. Use `modemmanager` instead.
- `finger`
- `gkrellm`. Use `conky` instead
- `libcgroup1`. Use the cgroup functionality of `systemd` instead.
- `libibcm`. For more information, see [Section 9.3.3, “libibcm Has Been Removed”](#).
- `puppet`. Use Salt instead.
- `reiserfs`. For more information, see [Section 3.2.11, “ReiserFS Support Removed”](#).
- `rsh`. Use `ssh` instead. In addition, SLE HPC includes `mrsh`, so that you can pick depending on your use case.
- `sapvnc`. Connect via RDP instead.
- `smt` and associated packages. Use `rmt-server` instead. For more information, see [Section 3.2.4, “SMT Has Been Replaced by RMT”](#).

- squidguard
- unrar. For more information, see *Section 9.1.6, “UnRAR Has Been Replaced by unrar”*.
- wodim, genisoimage, and icedax. For more information, see *Section 9.1.4, “wodim Has Been Replaced by cdrtools”*.
- x11vnc. For more information, see *Section 9.3.2, “x11vnc Has Been Removed”*.
- xinetd and yast2-inetd. For more information, see *Section 9.3.5, “xinetd and yast2-inetd Have Been Removed”*.
- xf86-video-cirrus, xf86-video-mga, and xf86-video-ast. For more information, see *Section 3.2.2, “The User Space X Drivers cirrus/mga/ast Have Been Removed”*
- yast2-add-on-creator and yast2-product-creator. To build add-ons or base products, use the Open Build Service and/or Kiwi instead.
- yast2-ca-management

9.3.1 The User Space X Drivers cirrus/mga/ast Have Been Removed

The packages xf86-video-cirrus, xf86-video-mga, and xf86-video-ast have been removed in SLE 15. Kernel mode setting and mode-setting X drivers for these graphics cards have been available throughout the SLE 12 cycle and were used for all new SLE 12 installations. The user space X driver packages were only retained to ease upgrades from SLE 11.

If you are upgrading a machine from SLE 12 to SLE 15 that has previously been upgraded from SLE 11 to SLE 12, X may no longer start after the upgrade to SLE 15. If that is the case, rename or remove the file /etc/X11/xorg.conf, for example using:

```
sudo old /etc/X11/xorg.conf
```

9.3.2 x11vnc Has Been Removed

In SLE 15, the package x11vnc is not available anymore. Instead, use x0vncserver. The command x11vnc is now a compatibility wrapper that internally starts x0vncserver. It does not have all features that x11vnc had, but it is faster, more secure, and built from better tested and maintained code.

9.3.3 libibcm Has Been Removed

libibcm is deprecated and unmaintained in the upstream community. Starting with rdma-core 17, it will also not be distributed anymore by upstream and it is planned to remove kernel support of ucm.

With SLE 15 and the update to rdma-core 16, the tool libibcm has been removed, along with the package that contained it, libibcm1.

9.3.4 atftp Has Been Removed

SLE 12 included two implementations of server for the TFTP protocol, tftp and atftp. Of these two implementations, only tftp is actively developed.

With SLE 15, atftp has been removed. Therefore, tftp is the recommended solution for most use cases.

However, there are important feature differences between the two implementations:

- tftp does not support multicast.
- tftp does not support symbolic links within a directory tree that point to locations outside of that directory tree. This can lead to unusable systems.

To work around both issues, you can use the TFTP server built into dnsmasq instead. If you do not need the DHCP and DNS features that dnsmasq provides, set it up in a TFTP-only mode. To do so, configure the following settings in /etc/dnsmasq.conf:

```
# disable DNS and DHCP
port=0
# Enable dnsmasq's built-in TFTP server
enable-tftp
# Set the root directory for files available via FTP.
tftp-root=/srv/tftpboot
# Do not abort if the tftp-root is unavailable
tftp-no-fail
```

If you need to work with symbolic links pointing outside the directory tree:

- If the listed TFTP_DIR values do not contain the path referred to by a symbolic link below the tftp-root directory, make sure to adapt its value.
- If you are additionally using AppArmor, update the configuration in /etc/apparmor.d/usr.sbin.dnsmasq:

```
@{TFTP_DIR}=/var/tftp /srv/tftpboot [YOUR_SYMBOLIC_LINK_DIRECTORY]
```

9.3.5 `xinetd` and `yast2-inetd` Have Been Removed

In SLE 15, `xinetd` and `yast2-inetd` have been removed, in favor of `systemd` sockets. All software provided in SLE is already adapted to use `systemd` sockets and YaST modules activate socket instead of `xinetd`. If you are working with third-party software, it might have to be updated.



Note: `only_from` Feature Not Supported

`systemd` does not support the `only_from` feature of `xinetd`. If you were previously using this setting in your `xinetd` configuration, you now need to configure your firewall accordingly instead.

9.4 Deprecated Packages and Features

The following packages are deprecated and will be removed with a future service pack of SUSE Linux Enterprise Server.

9.4.1 `dmraid` Is Deprecated

With SLE 15, `dmraid` is considered deprecated. It will be removed in a future service pack. Instead of `dmraid`, use `mdadm`.

9.4.2 `SunRPC` Is Deprecated

The `SunRPC` code in `glibc` has been deprecated upstream since 2012 and is no longer active developed. `SunRPC` provides an outdated RPC API which only works with IPv4.

SLE 15 ships with a TI-RPC based library that provides a flexible RPC API which understands both IPv4 and IPv6.

All RPC based applications should migrate to `libtirpc` instead of using `sunrpc`. SUSE will follow upstream and disable compiling new code using `SunRPC` with one of the next releases. However, at runtime, `SunRPC` code will continue to be available for old binaries.

9.4.3 net-tools Has Been Split into net-tools and net-tools-deprecated

The tools arp, route, netstat, iptunnel, ipmaddr, and ifconfig from the package net-tools are now deprecated.

With SLE 15, to separate the deprecated tools from the others in the package net-tools, the package was split into net-tools and net-tools-deprecated.

The package net-tools retains the following tools:

- ether-wake
- nameif
- plipconfig
- slattach

The package net-tools-deprecated contains the obsolete tools that can be replaced with ip subcommands as below:

- arp -> ip [r] neigh
- route -> ip route
- netstat -> ss [-r]
- iptunnel -> ip tunnel
- ipmaddr -> ip maddress
- ifconfig -> ip address

The tools hostname, domainname, dnsdomainname have been moved to the package host-name which is required by net-tools and net-tools-deprecated.

9.4.4 YaPI Is Deprecated and Should Not Be Used Anymore

YaPI, a Perl API for YaST, is now considered deprecated and must not be called by applications outside of YaST/Installer anymore. It will be removed completely from YaST in a close future.

9.5 Changes in Packaging and Delivery

The following packages and groups of packages have been renamed or have had major packaging changes in SLE 15:

- The package `libpsm_infinipath1` now provides the default implementation of the library `libpsm_infinipath1`. For more information, see [Section 9.5.2, “The libpsm_infinipath1 Implementation from the Package libpsm_infinipath1 Will Be Picked By Default”](#).
- The package `sap_suse_cluster_connector` has been renamed to `sap-suse-cluster-connector`.
- In the package `xen-tools`, the executable `qemu-dm` has been removed.

9.5.1 MPI Implementations Identify in mpi-selector's list Command By Name Only

Previously, the MPI implementations `openmpi`, `mvapich2`, and `mpich` (and their variants) were configured to identify themselves with their name and exact version in `mpi-selector --list`. However, this behavior created a package upgrade issue where newly updated MPI packages would not be registered automatically.

As part of a maintenance update to SLE, the registration issue was fixed by making the packages identify only by their name but not their exact version number in `mpi-selector --list`. As this functionality was never meant as a way to support multiple versions of the same MPI implementation side by side, this should not cause practical issues.

Note that as an exception from the rule, it continues to be possible to install all available major versions of openMPI side by side (those can be, depending on the operating system version and installed extensions, `openmpi`, `openmpi2`, and `openmpi3`).

9.5.2 The libpsm_infinipath1 Implementation from the Package libpsm_infinipath1 Will Be Picked By Default

The library `libpsm_infinipath1` is provided by two packages: `libpsm_infinipath1` and `libpsm2-compat`. Both provide a library with the same file name, API and ABI, but `libpsm_infinipath1` is targeted for TrueScale hardware, while `libpsm2-compat` provides a compatibility layer to run on OmniScale/PSM2 hardware.

Because both RPM packages provide the same library, in the past, Zypper would pick either of the packages when trying to install `libpsm_infinipath1` to satisfy the dependency of another package. However, `libpsm2-compat` is only targeted at PSM1 users testing software on PSM2 hardware and should therefore not be picked by default.

The SLE 15 version of the package `libpsm_infinipath1` is marked as obsoleting `libpsm2-compat`. This ensures that it will be picked by Zypper by default. To try PSM1 applications on PSM2 hardware, make sure to downgrade to `libpsm2-compat`.

9.6 Modules

This section contains information about important changes to modules. For more information about available modules, see [Section 2.9.1, “Modules in the SLE 15 GA Product Line”](#).

9.6.1 Advanced Systems Management Module Has Been Removed

The Advanced Systems Management module that was available in SLE 12 has been removed. Salt has been incorporated into the Base System module, other packages have been removed from the distribution.

9.6.2 Development Module: OProfile Support for POWER9

The OProfile package has been updated to include support for POWER9 processors.

9.6.3 Development Module: Valgrind Support for POWER9

Valgrind has been updated to include support for POWER9 processors.

9.6.4 Legacy Module: `ntpd` is now part of the Legacy Module

With SLE 15, the network time daemon `ntpd` has been replaced by `chrony`. `ntpd` has been moved to the Legacy module instead.

9.6.5 Legacy Module: OpenSSL 0.9.8 Has Been Removed

In December 2015, OpenSSL 0.9.8 reached its end of life. The code was maintained via source back-ports in the SUSE Linux Enterprise Server 12 Legacy Module.

With SLE 15, OpenSSL 0.9.8 has been removed from the product. Users of OpenSSL must upgrade code to version 1.0 or 1.1, both of which bring many improvements such as updated protocol support.

9.6.6 Legacy Module: OpenSSL 1.0.x Has Been Moved to the Legacy Module

The lifetime of OpenSSL versions 1.0.x does not cover the full lifetime of the product. Additionally, OpenSSL will not support TLS 1.3. However, some applications may require this older version for a transitional period.

OpenSSL libraries version 1.0.x were moved to the Legacy Module. The module has a different lifecycle from SUSE Linux Enterprise Server itself. This version is not expected to receive feature updates or security certifications. For new development, make sure to use the default OpenSSL version 1.1.x.

9.6.7 Legacy Module: OpenLDAP Client Libraries Have Been Moved to the Legacy Module

Even though 389 Directory Server (package `389-ds`) is the new default LDAP server on SLE, the OpenLDAP client libraries are widely used for LDAP integrations.

As `389-ds` is compatible with the OpenLDAP client libraries, they will be still provided and supported on SLE 15 to provide an easier transition for customers that currently use the OpenLDAP Server. For more information on 389 Directory Server, see [Section 4.5.6, “389 Directory Server Is Now the Primary LDAP Server”](#).

In the past, the `libldap` package provided two main libraries: `libldap.so` and `libldap_r.so`, the latter being thread-safe and therefore recommended to be used. For that reason, in SLE 15, the package `libldap` only provides the `libldap_r.so` library. For backward compatibility, there is the additional package `libldap-legacy`, that provides a `libldap.so` library that redirects to the `libldap_r.so` library.

The packages `pam_ldap` and `nss_ldap` are considered legacy and were therefore moved to the Legacy Module on SLE 15. Consider using the System Security Services Daemon (SSSD) instead.

10 Technical Information

This section contains information about system limits, technical changes and enhancements for experienced users.

When talking about CPUs, we use the following terminology:

CPU Socket

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

CPU Core

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

On IBM Z, this is equivalent to an IFL.

Logical CPU

This is what the Linux Kernel recognizes as a “CPU”.

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

Virtual CPU

A logical CPU as seen from within a virtual machine.

10.1 Supported Host Operating Systems and Hypervisors

SUSE Hosts

The following SUSE host operating systems and hypervisors are supported to run SLE 15 GA guests:

- SLES 11 SP4 with KVM or Xen
- SLES 12 SP1 with KVM or Xen
- SLES 12 SP2 with KVM or Xen
- SLES 12 SP3 with KVM or Xen
- SLES 12 SP4 with KVM or Xen (after SLES 12 SP4 has been released)
- SLES 15 with KVM or Xen
- SLES 15 SP1 with KVM or Xen (after SLES 15 SP1 has been released)

When running on one of the SUSE host operating system listed above, your installation will enjoy full L3 support from SUSE, both for the guest and host.

Third-Party Hosts

The following third-party host operating systems and hypervisors are supported to run SLE 15 GA guests:

- VMware ESXi 6.0
- VMware ESXi 6.5
- Microsoft Windows 2008 SP2 and later
- Microsoft Windows 2008 R2 SP1 and later
- Microsoft Windows 2012 and later
- Microsoft Windows 2012 R2 and later
- Microsoft Windows 2016
- Citrix XenServer 6.5
- Oracle VM 3.3

When running on one of the third-party host operating system listed above, your installation will enjoy full L3 support from SUSE for the guest. To learn about support options for the third-party hosts, contact the third-party vendor in question.

10.2 Kernel Limits

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

<i>SLES 15 GA (Linux 4.12)</i>	AMD64/Intel 64 (x86_64)	IBM Z (s390x)	POWER (ppc64le)	AArch64 (AR- Mv8)
CPU bits	64	64	64	64

SLES 15 GA (Linux 4.12)	AMD64/Intel 64 (x86_64)	IBM Z (s390x)	POWER (ppc64le)	AArch64 (AR- Mv8)
Maximum number of logical CPUs	8192	256	2048	256
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 (x86_64) or 30 * 64 GB (other architectures)			
Maximum number of processes	1048576			
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”.

10.3 Virtualization

10.3.1 Supported Live Migration Scenarios

You can migrate a virtual machine from one physical machine to another. The following live migration scenarios are supported under both KVM and Xen:

- SLE 12 SP3 to SLE 15
- SLE 12 SP4 to SLE 15 (after SLE 12 SP4 has been released)
- SLE 15 to SLE 15
- SLE 15 to SLE 15 SP1 (after SLE 15 SP1 has been released)

10.3.2 KVM Limits

<i>SLES 15 GA Virtual Machine (VM)</i>	Limits
Maximum Physical Memory per Host	64 TiB
Maximum Physical CPUs per Host	8192
Maximum VMs per Host	Unlimited (total number of virtual CPUs in all guests being no greater than 8 times the number of CPU cores in the host)
Maximum Virtual CPUs per VM	288
Maximum Memory per VM	4 TiB

Virtual Host Server (VHS) limits are identical to those of SUSE Linux Enterprise Server.

10.3.3 Xen Limits

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

<i>SLES 15 GA Virtual Machine (VM)</i>	Limits
Maximum number of virtual CPUs per VM	128
Maximum amount of memory per VM	16 GiB x86_32, 2 TiB x86_64

<i>SLES 15 GA Virtual Host Server (VHS)</i>	Limits
Maximum number of physical CPUs	1024
Maximum number of virtual CPUs	Unlimited (total number of virtual CPUs in all guests being no greater than 8 times the number of CPU cores in the host)
Maximum amount of physical memory	16 TiB
Maximum amount of Dom0 physical memory	500 GiB

- PV: Paravirtualization
- FV: Full virtualization

For more information about acronyms, see the virtualization documentation provided at <https://www.suse.com/documentation/sles-15>.

10.4 File Systems

10.4.1 Comparison of Supported File Systems

SUSE Linux Enterprise was the first enterprise Linux distribution to support journaling file systems and logical volume managers back in 2000. Later, we introduced XFS to Linux, which today is seen as the primary work horse for large-scale file systems, systems with heavy load and multiple parallel reading and writing operations. With SUSE Linux Enterprise 12, we went the next step of innovation and started using the copy-on-write file system Btrfs as the default for the operating system, to support system snapshots and rollback.

+ supported

– unsupported

Feature	Btrfs	XFS	Ext4	OCFS 2 ¹
Support in products	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
Online extend/shrink	+ / +	+ / –	+ / –	– / –
Inode allocation map	B-tree	B+ -tree	table	B-tree
Sparse files	+	+	+	+
Tail packing	–	–	–	–
Small files stored inline	+ (in metadata)	–	+ (in inode)	+ (in inode)
Defragmentation	+	+	+	–
Extended file attributes/ACLs	+ / +	+ / +	+ / +	+ / +
User/group quotas	– / –	+ / +	+ / +	+ / +
Project quotas	–	+	+	–
Subvolume quotas	+	N/A	N/A	N/A
Data dump/restore	–	+	–	–

Feature	Btrfs	XFS	Ext4	OCFS 2 ¹
Block size default	4 KiB ⁴			
Maximum file system size	16 EiB	8 EiB	1 EiB	4 PiB
Maximum file size	16 EiB	8 EiB	1 EiB	4 PiB

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

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

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

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

Additional Notes

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

The numbers in the above table 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>.

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

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

Some file system features are available in SUSE Linux Enterprise Server 15 GA but are not supported by SUSE. By default, the file system drivers in SUSE Linux Enterprise Server 15 GA 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.

10.4.2 Supported Btrfs Features

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

+ supported

– unsupported

Feature	SLES 11 SP4	SLES 12 SP3	SLES 12 SP4	SLES 15 GA
Copy on Write	+	+	+	+
Snapshots/Subvolumes	+	+	+	+
Metadata Integrity	+	+	+	+
Data Integrity	+	+	+	+
Online Metadata Scrubbing	+	+	+	+
Automatic Defragmentation	–	–	–	–
Manual Defragmentation	+	+	+	+
In-band Deduplication	–	–	–	–
Out-of-band Deduplication	+	+	+	+
Quota Groups	+	+	+	+
Metadata Duplication	+	+	+	+
Multiple Devices	–	+	+	+
RAID 0	–	+	+	+
RAID 1	–	+	+	+
RAID 10	–	+	+	+
RAID 5	–	–	–	–

Feature	SLES 11 SP4	SLES 12 SP3	SLES 12 SP4	SLES 15 GA
RAID 6	–	–	–	–
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	–	+	+	+

10.5 Supported Java Versions

The following table lists Java implementations available in SUSE Linux Enterprise Server 15 GA.

Name (Package Name)	Version	SUSE Linux Enterprise Server Module	Support
OpenJDK (java-10-openjdk)	10	Base System	SUSE, L3
OpenJDK (java-1_8_0-openjdk)	1.8.0	Legacy	SUSE, L3
IBM Java (java-1_8_0-ibm)	1.8.0	Legacy	External

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 <http://www.suse.com/download-linux/source-code.html>. Also, for up to three years after distribution of the SUSE product, upon request, SUSE will mail a copy of the source code. Requests should be sent by e-mail to mailto:sle_source_request@suse.com or as otherwise instructed at <http://www.suse.com/download-linux/source-code.html>. SUSE may charge a reasonable fee to recover distribution costs.



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- 2019 SUSE LLC. This release notes document is licensed under a Creative Commons Attribution-NoDerivs 3.0 United States License (CC-BY-ND-3.0 US, <https://creativecommons.org/licenses/by-nd/3.0/us/>).

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