Release Notes

SUSE Linux Enterprise Server 15 SP1

SUSE Linux Enterprise Server is a modern, modular operating system for both multimodal and traditional IT. This document provides an overview of high-level general features, capabilities, and limitations of SUSE Linux Enterprise Server 15 SP1 and important product updates.

These release notes are updated periodically. The latest version is always available at https://www.suse.com/releasenotes ↗. General documentation can be found at: https://www.suse.com/documentation/sles-15 ↗.

Publication Date: 2019-05-13, Version: 15.1.20190513

Contents

1 About the Release Notes

The most recent version of the Release Notes is available online at https://www.suse.com/re-leasenotes ↗.

These Release Notes are identical across all supported architectures.

Entries can be listed multiple times if they are important and belong to multiple sections.

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

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 SP1 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 SP1, 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 What Is New?

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

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 3.1, "Modules in the SLE 15 SP1 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 SLE 15, we support migrating from openSUSE Leap 15 to SUSE Linux Enterprise Server 15. Even if you decide to start out with the free community distribution you can later easily upgrade to a distribution with enterprise-class support.

- **Extended package search:** Use the new Zypper command `zypper search-packages` to search across all SUSE repositories available for your product even if they are not yet enabled. This functionality makes it easier for administrators and system architects to find the software packages needed. To do so, it leverages the SUSE Customer Center.

- **Software Development Kit:** With SLE 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.

Major updates to the software selection:

- **Salt:** SLE 15 SP1 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, SLE 15 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.

2.3 Important Sections of This Document

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

- *Section 2.5, "Support Statement for SUSE Linux Enterprise Server"*

- *Section 4.2, "Upgrade-Related Notes"*

- *Section 11, "Technical Information"*

2.4 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 (SP1) will be fully maintained and supported until 6 months after the release of SUSE Linux Enterprise Server 15 SP2.

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↗.

2.5 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 SP1 and its modules are delivered with L3 support for all packages, except for the following:

- Technology Previews

- 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 packages. That is, packages that are unchanged and not recompiled.

2.6 General Support

To learn about supported kernel, virtualization, and file system features, as well as supported Java versions, see *Section 11, "Technical Information"*.

2.7 Documentation and Other Information

2.7.1 Available on the Product Media

- Read the READMEs on the media.

- Get the detailed change log information about a particular package from the RPM (where `<FILENAME>.rpm` is the name of the RPM):

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

- Check the `ChangeLog` file in the top level of the 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 SP1. This directory includes PDF versions of the SUSE Linux Enterprise Server 15 SP1 Installation Quick Start Guide.

2.7.2 Externally Provided Documentation

- For the most up-to-date version of the documentation for SUSE Linux Enterprise Server 15 SP1, 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 ↗.

3 Modules, Extensions, and Related Products

This section comprises information about modules and extensions for SUSE Linux Enterprise Server 15 SP1. Modules and extensions add parts or functionality to the system.

3.1 Modules in the SLE 15 SP1 Product Line

The SLE 15 SP1 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 SP1 at the release of SUSE Linux Enterprise Server 15 SP1. However, not all modules are available with a subscription for SUSE Linux Enterprise Server 15 SP1 itself (see the column *available with*).

For information the avilability of individual packages within modules, see https://scc.suse.com/packages ↗.

Name and Content	Dependencies on Other Modules	Available with	Support[1]
Base System A basic SLE system, comparable to the installation pattern `minimal` of previous SLE versions	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, packaged container images	Base System	SLES, SLES for SAP, SLE HPC	**life cycle:** 10 years **extended:** no **support level:** L3
Desktop Applications A graphical user interface and essential 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
Development Tools Compilers and libraries for developing and debugging 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
High Availability Clustering support for mission-critical setups, includes tools such as Hawk, crm, Pacemaker, and Corosync	Base System	SLES for SAP (default), SLE HPC	**life cycle:** 10 years **extended:** 3 years of LTSS **support level:** L3

Name and Content	Dependencies on Other Modules	Available with	Support[1]
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 Tools to allow deploying to cloud environments, such as Amazon Web Services (AWS), Microsoft Azure, Google Compute Platform, or SUSE OpenStack Cloud	Base System	SLES, SLES for SAP, SLE HPC	**support level:** L3
Python 2 Python 2 runtime and modules (deprecated in favor of Python 3)	Base System	SLES, SLES for SAP, SLE HPC, SLED	
Real Time Reduce system latencies and boost response time	Base System, Desktop Applications, Development Tools, Server Applications	SLE RT (default)	

Name and Content	Dependencies on Other Modules	Available with	Support[1]
SAP Applications Packages specific to SLES for SAP	Base System	SLES for SAP (default)	**life cycle:** 10 years **extended:** 3.5 years of LTSS, 1.5 years of ESPOS **support level:** L3
Server Applications Basic server functionality (DHCP, DNS, Web), 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
SUSE Cloud Application Platform Tools Tools to interact with SUSE Cloud Application Platform	Base System	SLES, SLES for SAP, SLED	
SUSE Package Hub[2] Community-maintained packages	Base System	SLES, SLES for SAP, SLE HPC, SLE RT, SLED	**life cycle:** individual per package **extended:** none **support level:** none
Transactional Server Tools to *atomically* apply and, if necessary, roll back system updates	Base System	SLES	
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.*

3.2 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/releasenotes ↗ .

The following extensions are available for SUSE Linux Enterprise Server 15 SP1:

- 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/ ↗

3.3 Derived and Related Products

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

- SUSE JeOS: https://www.suse.com/products/server/jeos ↗

- SUSE Enterprise Storage: https://www.suse.com/products/suse-enterprise-storage ↗

- SUSE Linux Enterprise Desktop: https://www.suse.com/products/desktop ↗

- SUSE Linux Enterprise Server for SAP Applications: https://www.suse.com/products/sles-for-sap ↗

- SUSE Linux Enterprise for High-Performance Computing: https://www.suse.com/products/server/hpc ↗

- SUSE Linux Enterprise Real Time: https://www.suse.com/products/realtime ↗

- SUSE Manager: https://www.suse.com/products/suse-manager ↗

- SUSE OpenStack Cloud: https://www.suse.com/products/suse-openstack-cloud ↗

4 Installation and Upgrade

SUSE Linux Enterprise Server can be deployed in several ways:

- Physical machine

- Virtual host

- Virtual machine

- System containers

- Application containers

4.1 Installation

This section includes information related to the initial installation of SUSE Linux Enterprise Server 15 SP1.

🔴 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 ↗ .

4.1.1 JeOS Images for Hyper-V and VMware Are Now Compressed

We are providing different virtual disk images for JeOS, using the `.qcow2`, `.vhdx`, and `.vmdk` file formats respectively for KVM, Xen, OpenStack, Hyper-V, and VMware environments. All JeOS images are setting up the same disk size (24 GB) for the JeOS system but due to the nature of the different file formats, the size of the JeOS images were different.

Starting with SLE 15 SP1, the JeOS images for Hyper-V and VMware using the `.vhdx` and `.vmdk` file formats respectively are now compressed with the LZMA2 compression algorithm by default. Therefore, we are now delivering these images in a `.xz` file format, so you need to uncompress the image before using them in your Hyper-V or VMware environment, for example, using the `unxz` command.

The other JeOS images will remain uncompressed because the `.qcow2` format already optimize the size of the images.

4.1.2 CD/DVD Repositories Will Be Disabled After Installation

In previous versions of SLE, enabled CD/DVD repositories would block upgrades if the media was removed after installation.

CD/DVD repositories are now set to disabled when the installation process is finished.

4.2 Upgrade-Related Notes

This section includes upgrade-related information for SUSE Linux Enterprise Server 15 SP1.

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 .

4.2.1 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`.

4.3 For More Information

For more information, see *Section 5, "Architecture Independent Information"* and the sections relating to your respective hardware architecture.

5 Architecture Independent Information

Information in this section pertains to all architectures supported by SUSE Linux Enterprise Server 15 SP1.

5.1 Kernel

5.1.1 IOMMU Passthrough Is Now Default on All Architectures

Passthrough mode provides improved I/O performance, especially for high-speed devices, because DMA remapping is not needed for the host (bare-metal or hypervisor).

IOMMU passthrough is now enabled by default in SUSE Linux Enterprise products. Therefore, you no longer need to add `iommu=pt` (Intel 64/AMD64) or `iommu.passthrough=on` (AArch64) on the kernel command line. To disable passthrough mode, use `iommu=nopt` (Intel 64/AMD64) or `iommu.passthrough=off` (AArch64), respectively.

5.1.2 The Driver i40evf Has Been Renamed to iavf

Starting with SLE 15 SP1, the module name of the Intel Ethernet Adaptive Virtual Function driver changes from `i40evf` to `iavf`. This new naming is consistent with the mainline Linux kernel and also helps better convey its status as the universal Virtual Function driver for multiple product lines.

5.1.3 New sysctl Option to Configure NUMA Statistics

Generating NUMA page allocator statistics can create considerable overhead.

To allow avoiding this overhead under certain circumstances, the sysctl option `vm.numa_stat` has been added. By default, it is set to `1`, meaning NUMA page allocator statistics will be generated.

For workloads where it is desirable to remove the overhead of these statistics, such as high-speed networking, disable the NUMA page allocator statistics by setting `vm.numa_stat` to `0`. The statistics in `/proc/vmstat`, such as `numa_hit` and `numa_miss` will then be reset to `0` and stop increasing, until the functionality is enabled again.

5.2 Security

5.2.1 LUKS2 Support for pam_mount

The pam_mount tool now supports the handling of LUKS2 encrypted volumes

5.3 Networking

5.3.1 resolv.conf Is Now Located in /run

Starting with SLE 15 SP1, both Wicked and NetworkManager will write the file `resolv.conf` into the `/run` directory instead of in `/etc`. `/etc/resolv.conf` will still exist as a symbolic link.

5.4 Systems Management

5.4.1 Chrony Is Now Installed by Default on JeOS and Raspberry Pi Images

Manual correction of the system time can lead to severe problems because, for example, a backward leap can cause malfunction of critical applications. Within a network, it is usually necessary to synchronize the system time of all machines, but manual time adjustment is a bad approach.

SLE 15 SP1 JeOS and Raspberry Pi images now include Chrony by default. This allow our images to follow the SLES 15 SP1 guidance to use Chrony for time synchronization. For more information, see https://www.suse.com/documentation/sles-15/singlehtml/book_sle_admin/book_sle_admin.html#cha.ntp ↗.

5.4.2 Zypper and the --no-recommends Option

Due to a trend toward minimal systems, systems are increasingly installed with the command-line parameter `--no-recommends` or the configuration option `solver.onlyRequires = true` in `/etc/zypp/zypp.conf`.

Unfortunately this option also prevented the autoselection of appropriate driver or language supporting packages.

This flaw is fixed with libzypp 17.10.2 and Zypper 1.14.18:

- The use of `--no-recommends` should no longer affect the selection of driver and language supporting packages.

- `zypper inr --no-recommends` should add missing driver and language-support packages only but omit all other recommends.

5.4.3 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.

5.5 Performance Related Information

5.5.1 The OProfile Package Has Been Updated

The OProfile package has been updated with the following new features:

- Updated CPU type detection for POWER9 models.

- Fix for a OProfile crash when processing data collected on an exiting process (affects all architectures).

5.6 Storage

5.6.1 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) ↗, Stop Using SMB1 (https://web.archive.org/web/20190227091836/https://blogs.technet.microsoft.com/filecab/2016/09/16/stop-using-smb1/) ↗). For security reasons, the SLE 15 SP1 kernel has been changed in a way that 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`).

5.6.2   Snapper Output Highlights Mount Status of Snapshots

Previously, `snapper list` did not indicate which snapshot was currently mounted and which would be mounted next time.

Starting with SLE 15 SP1, the output of `snapper list` now marks these special snapshots by appending one of the following three characters to the snapshot number:

- `*` (currently mounted and will be mounted on next boot)

- `-` (currently mounted)

- `+` (will be mounted on next boot)

The snapshot number is now also the first column in the output.

5.6.3   Snapper's Space-Aware Snapshot Cleanup Has Been Improved

Previously, the space-aware cleanup of snapshots integrated in Snapper only looked at the disk space used by all snapshots. In certain cases, this narrow focus meant that the file system ran out of space anyway.

Starting with SLE 15 SP1, the space-aware cleanup of Snapper additionally looks at the free space of the file system and keeps the file system at least 20 % free.

5.6.4 Displaying Disk Space Used by Snapper Snapshots

Previously, it was hard to calculate the disk space consumption of an individual Btrfs snapshot when the qgroups (quota groups) feature was enabled.

Starting with SLE 15 SP1, Snapper shows the disk space used by individual snapshots when running `snapper list` even if Btrfs quotas are enabled.

5.7 Virtualization

5.7.1 KVM

5.7.1.1 QED Image Format Is No Longer Supported

The QEMU virtual disk image format is no longer supported.

Existing virtual disks using this format can still be accessed, but should be converted to a RAW or QCOW2 format when possible. Using the QED format for new disks is not supported.

5.7.2 Xen

5.7.2.1 Xen Auto-Ballooning Disabled by Default

In previous versions of SLES, the default memory allocation scheme of a Xen host was to allocate all host physical memory to Domain-0 and enable auto-ballooning. Memory was automatically ballooned from Domain-0 when starting additional domains. This behavior has always been error-prone and disabling it is encouraged in the *Virtualization Best Practices Guide.*

Starting with SLES 15 SP1, Domain-0 auto-ballooning has been disabled by default. Domain-0 is allocated 10 percent of host physical memory + 1 Gb. For example, on a host with 32 Gb of physical memory, Domain-0 is allocated 4.2 Gb of memory. The use of the `dom0_mem` Xen

command-line option is still supported and encouraged. The old behavior can be restored by setting `dom0_mem` to the host physical memory size and enabling the `autoballoon` setting in `/etc/xen/xl.conf`.

5.8 Miscellaneous

5.8.1 HiDPI support in GNOME

Starting with SLE 15 SP1, there are several improvements to HiDPI support. If the DPI of your display is greater than 144, GNOME will scale the Session to a 2:1 ratio automatically and deliver you a crisp and sharp user experience. You can adjust the scaling-factor value manually under GNOME Control Center's display panel.

However, there are limitations to this support:

- Fractional scaling is still considered experimental in GNOME 3.26.2, so you can only set the scaling factor to a whole number.

- X11 apps may appear blury under a HiDPI Wayland session (via XWayland), as per-display scaling is not supported on X11.

- Using multiple monitors with different DPI is not supported, scale-monitor-framebuffer is still an immature feature so you cannot set per-monitor scales in GNOME Control Center.

6 AMD64/Intel 64 (x86_64) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 SP1 for the AMD64/Intel 64 architectures.

6.1   System and Vendor Specific Information

6.1.1   Fake NUMA Emulation in the Linux Kernel Can Now Uniformly Split Physical Nodes

Previously, NUMA emulation capabilities for splitting system RAM by a fixed size or by a set number of nodes could result in some nodes being larger than others. This happened because the implementation prioritized establishing a minimum usable memory size over satisfying the requested number of NUMA nodes.

With SLE 15 SP1, the kernel can now evenly partition each physical NUMA node into N emulated nodes. For example, the boot parameter `numa=fake=3U` creates a total of 6 emulated nodes on a system that has 2 physical nodes. This is useful for debugging and evaluating platform memory-side-cache capabilities as described by the ACPI HMAT.

To use, add the boot parameter `numa=fake=<N>U`. The final `U` means that the kernel will divide each physical node into N emulated nodes.

7   POWER (ppc64le) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 SP1 for the POWER architecture.

7.1   Reduced Memory Usage When Booting FADump Capture Kernel

One of the primary issues with Firmware Assisted Dump (FADump) on IBM POWER systems is that it needs a large amount of memory to be reserved. On large systems with terabytes of memory, this reservation can be quite significant.

Normally, the preserved memory is filtered to extract only relevant areas using the `makedump-file` tool. While the tool allows determining what needs to be part of the dump and what memory to filter out, the default is to capture only kernel data and exclude everything else.

We take advantage of this default and the Linux kernel's Contiguous Memory Allocator (CMA) to fundamentally change the memory reservation model for FADump: Instead of setting aside a significant chunk of memory that cannot otherwise be used, the feature uses CMA instead.

It reserves a significant chunk of memory that the kernel is prevented from using (due to `MI-GRATE_CMA`), but applications are free to use it. With this, FADump will still be able to accurately capture all of the kernel memory and most of the user-space memory except for the user pages that are part of the CMA region reserved for FADump.

To disable this feature, pass the kernel parameter `fadump=nocma` instead of `fadump=on`. This ensures that the memory reserved for FADump is not used by applications. This option may be useful in scenarios where you prefer to also capture application data in the dump file.

7.2   Performance Co-pilot (PCP) Updated, Perfevent Performance Metrics Domain Agent (PMDA) Support Libraries Added

PCP has been updated to v4.3.1 and brings many improvements in its ability to collect performance metrics from various sources. In addition, the Perfevent PMDA is now available and provides access to platform performance counter (PMU) data through the Linux perf_event subsystem.

7.3   Uprobes: Support for SDT events with reference counter (perf)

Userspace Statically Defined Tracepoints (USDT) are dtrace-style markers inside userspace applications. With SLES 15 SP1, Uprobe is enhanced to support SDT events having reference counter (semaphore).

7.4   PAPI Package Update

Update to newer version of PAPI to pick up fixes for POWER8 events, POWER9 events, corrections and clean up for some duplicate event names.

7.5   ibmvnic Device 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 .

7.6 SDT Markers added to libglib

SDT markers for debugging and performance monitoring with tools such as perf and systemtap have been added to libglib.

7.7 Access to Additional POWER Registers in GDB

GDB can now access more POWER architecture registers, including PPR, DSCR, TAR, and Hardware Transactional Memory registers.

8 IBM Z (s390x) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 SP1 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.

8.1 Virtualization

The following new features are supported in SUSE Linux Enterprise Server 15 SP1 under KVM:

8.1.1 Huge Pages

Allow KVM guests to use huge page memory backing for improved memory performance for workloads running with large memory footprints.

8.1.2 zPCI Passthrough Support for KVM

Allow KVM to pass control over any kind of PCI host device (a virtual function) to a KVM guest enabling workloads that require direct access to PCI functions.

8.1.3 Interactive Bootloader

Enable to interactively select boot entries to recover misconfigured KVM guests.

8.1.4 Guest-Dedicated Crypto Adapters

Allow KVM to dedicate crypto adapters (and domains) as passthrough devices to a KVM guest such that the hypervisor cannot observe the communication of the guest with the device.

8.1.5 Expose Detailed Guest Crash Information to the Hypervisor

Provides additional debug data for operating system failures that occur within a KVM guest.

8.1.6 Development-Tools Module: Valgrind IBM z13 Support

Valgrind now include instruction support for IBM z13 instructions. This enables debugging and validation of binaries built and optimized for IBM z13. In particular this covers the vector instruction set extensions introduced with IBM z13.

8.1.7 kvm_stat Package from kernel Tree

kvm_stat allows to display KVM trace events, which can be useful for trouble shooting.

8.2 Network

8.2.1 OSA-Express7S Adapters Are Now Supported

With the OSA 7 network cards a link speed of 25Gb/s is supported.

8.2.2   OSA IPv6 Checksum Offload

Checksum offload now supports IPv6 Configuring checksum offload operations.

8.2.3   Full-blown TCP Segmentation Offload

TCP segmentation offload is now supported on both layer 2 and layer 3 and is extended to IPv6.

8.2.4   Shared Memory Communications - Direct (SMC-Direct)

Internal shared memory devices for fast communication between LPARs can be used via a new socket family and the existing tooling via TCP handshake. A preload library can be used to enable applications to use the new socket family transparently.

8.3   Security

8.3.1   Cryptsetup 2.0.5 for LUKS2 Support

It is possible to use cryptsetup to handle protected keys for dm-crypt disks in plain format and cryptsetup provides LUKS 2 support.

8.3.2   Support Multiple zcrypt Device Nodes

The cryptographic device driver can now provide and maintain multiple zcrypt device nodes. These nodes can be restricted in terms of cryptographic adapters, domains, and available IOCTLs.

8.3.3   SIMD Implementation of Chacha20 in OpenSSL

This enables support for TLS 1.3 via the Chacha20 cipher suite providing good performance using SIMD instructions

8.3.4   dm-crypt with Protected Keys - Change Master Key Tool

Manage LUKS2 encryption keys for protected key crypto if the master key of the associated Crypto Express adapter is changed.

8.3.5   libica: Use TRNG to Seed DRBG (crypto)

Improved generation of high (pseudo) quality random numbers via libica DRBG especially to generate safe random keys by use of the PRNO-TRNG instruction.

8.3.6   Support of CPACF Hashes in ep11 Token in openCryptoki and libica

Provides improved performance for applications computing many digital signatures using EP11 like Blockchain.

8.3.7   In-kernel Crypto: Support Protected Keys Generated by random in paes Module

Enable to generate volatile protected keys to support for example the secure encryption of swap volumes w/o the need of an CryptoExpress adapter.

8.3.8   Partial RELRO Support in binutils

With this feature the global offset table content is rearranged to enable the dynamic linker write-protecting parts of the global offset table after initial program load. That way potential attacks requiring to rewrite such entries are prevented.

8.3.9   OpenSSL: Crucial Enhancements

Improved performance of OpenSSL via extended CPACF for additional ciphers like AES CTR, OFB, CFB, CCM.

8.3.10   SIMD Implementation of Poly1305 in OpenSSL

This enables support for TLS 1.3 via the Poly1305 cipher suite providing good performance using SIMD instructions.

8.3.11 Elliptic Curve Support for Crypto

The strategic elliptic curve asymmetric cryptography that provides strong security with shorter keys is now supported by Crypto Express function offloads with `opencryptoki`, `libica`, `icatoken`, and `openssl-ibmca`.

8.3.12 Support 4K Sectors for Fast Clear Key dm-crypt

Encryption is supported with 4K sectors. Using 4K sector leads to significant performance improvements on IBM Z using CPACF crypto hardware.

8.3.13 Enhanced SIMD Instructions in libica

Faster execution of asymmetric cryptographic algorithms via support of new SIMD instructions available with IBM z13 and later hardware.

8.3.14 Support for the CEX6S Crypto Card

The CEX6S crypto card is fully supported.

8.3.15 Support Architectural Limit of Crypto Adapters in zcrypt Device Driver

The crypto device driver now support the theoretical maximum of 255 adapters.

8.3.16 zcrypt DD: Introduce APQN Tags to Support Deterministic Driver Binding

Provides deterministic hot plugging semantics to enable the virtualization and unique etermination of crypto adapters in KVM environments even if the associated hardware gets intermittently lost and reconnected.

8.3.17 In-kernel Crypto: GCM Enhancements

Kernel services like IPSec now exploit IBM z14 crypto hardware for the AES-GCM cipher.

8.3.18   Protected Key dm-crypt Key Management Tool

Protected key crypto for dm-crypt disks in plain format can be used without a dependency on cryptsetup support for LUKS(2) with protected keys. A key management tool as part of the s390-tools enables to manage a key repository allowing to associate secure keys with disk partitions or logical volumes.

8.4   Reliability, Availability, Serviceability (RAS)

8.4.1   PCI Error Reporting Tool

Defective PCIe devices are now reported via error notification events that include health information of the adapters.

8.4.2   scsi: zfcp: Add Port Speed Capabilities

Provides the possibility to display port speed capabilities for SCSI devices.

8.4.3   Handle Provisioned MAC Addresses

You can now use provisioned MAC addresses for devices supported with IBM z14 and later hardware.

8.4.4   Configurable IFCC Handling

Enables to switch off the actual handling of repeated IFCCs (Interface Control Check), for example, by removing paths, so that only IFCC messages are written to the log when thresholds are exceeded.

8.4.5   Collecting NVMe-related Debug Data

To debug NVMe devices, the debug data gets collected and added to the `dbginfo.sh` script.

8.4.6 Raw Track Access without Prefix CCW

This feature enables seamlessly moving Linux system volumes between zPDT and LPAR, allowing for greater flexibility during deployment of new setups.

8.4.7 I/O Device Pre-Configuration

Linux in LPAR mode can now process device configuration data that is user-defined and obtained during boot.

8.5 Performance

8.5.1 Performance Counters for IBM z14 (CPU-MF)

For optimized performance tuning the CPU-measurement counter facility now supports counters, including the MT-diagnostic counter set, that were introduced with IBM z14.

8.5.2 Network Performance Improvements

Enhanced performance for OSA and Hipersockets via code improvements and exploitation of further kernel infrastructure.

9   ARM 64-Bit (AArch64) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 15 SP1 for the AArch64 architecture.

9.1 System-on-Chip Driver Enablement

SUSE Linux Enterprise Server for Arm 15 SP1 includes driver enablement for the following System-on-Chip chipsets:

- AMD Opteron A1100

- Ampere Computing X-Gene, eMAG

- Broadcom BCM2837

- Huawei Kunpeng 920

- Marvell ThunderX1, ThunderX2, Octeon TX, Armada 7040, Armada 8040

- Mellanox BlueField

- NXP QorIQ LS1043A, LS1046A, LS1088A, LS2088A, LX2160A; i.MX 8M

- Qualcomm Centriq 2400

- Rockchip RK3399

- Socionext SynQuacer SC2A11

- Xilinx Zynq UltraScale + MPSoC

9.2 Driver Enablement for NXP SC16IS7xx UARTs

The Raspberry Pi 3 Model B/B+ has only one serial port available on its 40-pin GPIO connector. SUSE Linux Enterprise Server now includes a device driver for NXP SC16IS7xx series of I^2C or SPI bus connected serial ports. These chipsets are found on multiple third-party expansion boards for the Raspberry Pi. For instructions how to describe such boards in the Device Tree for use with SUSE Linux Enterprise Server for Arm, please refer to the respective vendor's documentation and compare the SUSE Release Notes for the Raspberry Pi (in particular, recommended use of `extraconfig.txt` instead of `config.txt`).

9.3 Boot and Driver Enablement for Raspberry Pi

Bootloaders and a supported microSD card image of SUSE Linux Enterprise Server for Arm 15 SP1 for the Raspberry Pi are available. The selection of preinstalled packages and first-boot assistant in the SUSE image are now aligned with the JeOS images, reducing image size. To aid with installing a minimal graphical desktop as found in previous image versions, a new pattern `x11_raspberrypi` is provided for package installation (`zypper in -t pattern x11_raspberrypi`). The template of the SUSE Linux image is available as profile "RaspberryPi" in the package `kiwi-templates-SLES15-JeOS` to derive custom appliances, including appliances with X11 graphical environment preinstalled.

New Features

The Raspberry Pi 7" Touch Display connected via the MIPI DSI flat ribbon cable is now supported in SUSE Linux Enterprise Server for Arm 15 SP1.

Audio via the 3.5 mm jack or via HDMI connector on Raspberry Pi 3 Model B/B+ is now supported. It requires PulseAudio to be installed and started.

Expansion Boards

The Raspberry Pi 3 Model B/B+ offers a 40-pin General Purpose I/O connector, with multiple software-configurable functions such as UART, I^{2}C and SPI. This pin mux configuration along with any external devices attached to the pins is defined in the Device Tree which is passed by the bootloader to the kernel.

SUSE does not currently provide support for any particular HATs or other expansion boards attached to the 40-pin GPIO connector. However, insofar as drivers for pin functions and for attached chipsets are included in SUSE Linux Enterprise, they can be used. SUSE does not provide support for making changes to the Device Tree, but successful changes will not affect the support status of the operating system itself. Be aware that errors in the Device Tree can stop the system from booting successfully or can even damage the hardware.

The bootloader and firmware in SUSE Linux Enterprise Server 15 SP1 support Device Tree Overlays. The recommended way of configuring GPIO pins is to create a file `extraconfig.txt` on the FAT volume (`/boot/efi/extraconfig.txt` in the SUSE image) with a line `dtoverlay=filename-without-.dtbo` per Overlay. For more information about the syntax, see the documentation by the Raspberry Pi Foundation.

If not already shipped in the `/boot/efi/overlays/` directory (`raspberrypi-firmware-dt` package), `.dtbo` files can be obtained from the manufacturer of the HAT or compiled from self-authored sources.

For More Information

For more information, see the SUSE Best Practices documentation for the Raspberry Pi at https://www.suse.com/documentation/suse-best-practices/#els ↗.

10 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 5.4, "Systems Management"*.

10.1 New Packages

10.1.1 NumaTOP Has Been Added

The NumaTOP tool version 2.1 now ships with SLE 15 SP1 for the architectures x86-64 and ppc64le. NumaTOP is a tool to observe the NUMA locality of processes and threads running on a system. It relies on hardware performance monitoring counters present in a subset of Intel Xeon and IBM POWER 8/POWER 9 processors.

NumaTOP can be used to:

- Characterize the locality of all running processes and threads to identify those with the poorest locality in the system.

- Identify "hot" memory areas, report average memory access latency, and provide the location where accessed memory is allocated.

- Provide the call-chain(s) in the process/thread code that accesses a given hot memory area.

- Provide the call-chain(s) when the process/thread generates certain counter events. The call-chain(s) help(s) to locate the source code that generates the events.

- Provide per-node statistics for memory and CPU utilization.

- Show the list of processes/threads sorted by metrics (by default, by CPU utilization). You can also resort the output by the following metrics: Remote Memory Accesses (RMA), Local Memory Accesses (LMA), RMA/LMA ratio, Cycles Per Instruction (CPI), and CPU utilization.

10.1.2 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.

10.2 Updated Packages

10.2.1 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/post-gresql10/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).

10.3 Deprecated Packages and Features

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

10.3.1 Chelsio T3 Driver (cxgbe3) Is Deprecated

The driver for Chelsio T3 networking equipment (`cxgbe3`) is now deprecated and may become unsupported in a future Service Pack of SLE 15.

10.3.2 TLS 1.0 and 1.1 Will Be Disabled Starting with SUSE Linux Enterprise 15 SP2

The TLS 1.0 and 1.1 standards are superseded by TLS 1.2 and (soon) TLS 1.3. SUSE Linux Enterprise will keep backward compatibility with TLS 1.0 and 1.1 until 2020, but starting with SUSE Linux Enterprise 15 SP2, these old standards will be disabled.

10.4 Modules

This section contains information about important changes to modules. For more information about available modules, see *Section 3.1, "Modules in the SLE 15 SP1 Product Line"*.

10.4.1 Web and Scripting Module: Support for NodeJS 10.x

Older version of NodeJS are approaching their end of life, NodeJS 8.x which is currently shipped is already considered deprecated.

NodeJS 10.x, the current LTS version of NodeJS is now available in the *Web and Scripting* module of SLE.

10.4.2   Python 2 Module: python Executable Is Not Available in Standard Distribution

With SLE 15 SP1, SUSE has started to phase out the support for Python 2 in its enterprise distribution. Within the standard distribution, only Python 3 (executable name `python3`) is available. Python 2 (executable names `python2` and `python`) is now only provided via the *Python 2* module which is disabled by default.

Python scripts usually expect the `python` executable (note the lack of a version number) to refer to the Python 2.x interpreter of the system. If instead the Python 3 interpreter were started, that would likely lead to misbehaving applications. For this reason, SUSE has decided not ship a symbolic link for `/usr/bin/python` to the Python 3 executable by default.

To run Python 2 scripts, make sure to enable the SLE module *Python 2* and install the package `python` from it.

10.4.3   Package supportutils-plugin-salt Has Been Moved to the Base System Module

In SLE 15 GA, the package `supportutils-plugin-salt` was only available from the SUSE Manager module, whereas Salt itself was available from the SLE Base System module.

With SLE 15 SP1, this situation has been corrected: both the packages `salt` and `supportutils-plugin-salt` are now available from the SLE Base System module.

11   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.

11.1 Kernel Limits

This table summarizes the various limits which exist in our recent kernels and utilities (if related) for SUSE Linux Enterprise Server 15 SP1.

SLES 15 SP1 (Linux 4.12)	**AMD64/Intel 64 (x86_64)**	**IBM Z (s390x)**	**POWER (ppc64le)**	**AArch64 (ARMv8)**
CPU bits	64	64	64	64
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)			

SLES 15 SP1 (Linux 4.12)	AMD64/Intel 64 (x86_64)	IBM Z (s390x)	POWER (ppc64le)	AArch64 (ARMv8)
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 mmaps 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".

11.2 Virtualization

11.2.1 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)

11.2.2 KVM Limits

SLES 15 SP1 Virtual Machine (VM)	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.

11.2.3 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.

SLES 15 SP1 Virtual Machine (VM)	Limits
Maximum number of virtual CPUs per VM	128
Maximum amount of memory per VM	16 GiB x86_32, 2 TiB x86_64

SLES 15 SP1 Virtual Host Server (VHS)	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 ↗ .

11.3   File Systems

11.3.1   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	+ / +	+ / –	+ / –	– / –

Feature	Btrfs	XFS	Ext4	OCFS 2 [1]
Inode allocation map	B-tree	B + -tree	table	B-tree
Sparse files	+	+	+	+
Tail packing	−	−	−	−
Small files stored inline	+ (in metada-ta)	−	+ (in inode)	+ (in inode)
Defragmentation	+	+	+	−
Extended file attribut-es/ACLs	+ / +	+ / +	+ / +	+ / +
User/group quotas	− / −	+ / +	+ / +	+ / +
Project quotas	−	+	+	−
Subvolume quotas	+	N/A	N/A	N/A
Data dump/restore	−	+	−	−
Block size default	4 KiB [4]			
Maximum file system size	16 EiB	8 EiB	1 EiB	4 PiB
Maximum file size	16 EiB	8 EiB	1 EiB	4 PiB

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

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

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

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

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 SP1 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 SP1.

Some file system features are available in SUSE Linux Enterprise Server 15 SP1 but are not supported by SUSE. By default, the file system drivers in SUSE Linux Enterprise Server 15 SP1 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.

11.3.2   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	SLES 15 SP1
Copy on Write	+	+	+	+	+
Free Space Tree (Free Space Cache v2)	–	–	–	–	+
Snapshots/Subvolumes	+	+	+	+	+

Feature	SLES 11 SP4	SLES 12 SP3	SLES 12 SP4	SLES 15 GA	SLES 15 SP1
Swap Files	–	–	–	–	+
Metadata Integrity	+	+	+	+	+
Data Integrity	+	+	+	+	+
Online Metadata Scrubbing	+	+	+	+	+
Automatic Defragmentation	–	–	–	–	–
Manual Defragmentation	+	+	+	+	+
In-band Deduplication	–	–	–	–	–
Out-of-band Deduplication	+	+	+	+	+
Quota Groups	+	+	+	+	+
Metadata Duplication	+	+	+	+	+
Changing Metadata UUID	–	–	–	–	+
Multiple Devices	–	+	+	+	+
RAID 0	–	+	+	+	+
RAID 1	–	+	+	+	+
RAID 10	–	+	+	+	+
RAID 5	–	–	–	–	–
RAID 6	–	–	–	–	–
Hot Add/Remove	–	+	+	+	+
Device Replace	–	–	–	–	–
Seeding Devices	–	–	–	–	–
Compression	–	+	+	+	+
Big Metadata Blocks	–	+	+	+	+

Feature	SLES 11 SP4	SLES 12 SP3	SLES 12 SP4	SLES 15 GA	SLES 15 SP1
Skinny Metadata	–	+	+	+	+
Send Without File Data	–	+	+	+	+
Send/Receive	–	+	+	+	+
Inode Cache	–	–	–	–	–
Fallocate with Hole Punch	–	+	+	+	+

11.4 Supported Java Versions

The following table lists Java implementations available in SUSE Linux Enterprise Server 15 SP1.

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

12 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.

13 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.

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.