Release Notes

SUSE Linux Enterprise Server 12 SP2

This document provides guidance and an overview to high level general features and updates for SUSE Linux Enterprise Server 12 SP2. Besides architecture or product-specific information, it also describes the capabilities and limitations of SUSE Linux Enterprise Server 12 SP2. General documentation may be found at: http://www.suse.com/documentation/sles-12/.

Product to be released: November 2016

Publication Date: 2016-10-19, Version: 12.2.20161019

Contents

1 SUSE Linux Enterprise Server

SUSE Linux Enterprise Server is a highly reliable, scalable, and secure server operating system, built to power mission-critical workloads in both physical and virtual environments. It is an affordable, interoperable, and manageable open source foundation. With it, enterprises can cost-effectively deliver core business services, enable secure networks, and simplify the management of their heterogeneous IT infrastructure, maximizing efficiency and value.

The only enterprise Linux recommended by Microsoft and SAP, SUSE Linux Enterprise Server is optimized to deliver high-performance mission-critical services, as well as edge of network, and web infrastructure workloads.

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

This modular, general purpose operating system runs on three 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, making it the perfect guest operating system for virtual computing.

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 12 has a 13-year life cycle, with 10 years of General Support and 3 years of Extended Support. The current version (SP2) will be fully maintained and supported until 6 months after the release of SUSE Linux Enterprise Server 12 SP2. If you need additional time to design, validate and test your upgrade plans, Long Term Service Pack Support can extend the support you get an additional 12 to 36 months in twelve month increments, giving you a total of 3 to 5 years of support on any given 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.

1.1 What Is New?

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

- Robustness on administrative errors and improved management capabilities with full system rollback based on Btrfs as the default file system for the operating system partition and the Snapper technology of SUSE.
- An overhaul of the installer introduces a new workflow that allows you to register your system and receive all available maintenance updates as part of the installation.
- SUSE Linux Enterprise Server Modules offer a choice of supplemental packages, ranging from tools for Web Development and Scripting, through a Cloud Management module, all the way to a sneak preview of upcoming management tooling called Advanced Systems Management. Modules are part of your SUSE Linux Enterprise Server subscription, are technically delivered as online repositories, and differ from the base of SUSE Linux Enterprise Server only by their life cycle. For more information about modules, see *Section 1.5.1, "Available Modules"*.
- New core technologies like systemd (replacing the time-honored System V-based init process) and Wicked (introducing a modern, dynamic network configuration infrastructure).
- The open-source database system MariaDB is fully supported now.
- Support for open-vm-tools together with VMware for better integration into VMware-based hypervisor environments.
- Linux Containers are integrated into the virtualization management infrastructure (libvirt). Docker is provided as a fully supported technology. For more details, see https://www.suse.com/promo/sle/docker/.
- Support for the AArch64 architecture (64-bit ARMv8) and the 64-bit Little-Endian variant of the IBM POWER architecture. Additionally, we continue to support the Intel 64/AMD64 and IBM z Systems architectures.
- GNOME 3.20 gives users a modern desktop environment with a choice of several different look and feel options, including a special *SUSE Linux Enterprise Classic* mode for easier migration from earlier SUSE Linux Enterprise Desktop environments.

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

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

- *Section 1.4, "Support Statement for SUSE Linux Enterprise Server"*
- *Section 2.3, "Upgrade-Related Notes"*
- *Section 10, "Technical Information"*

1.2 Documentation and Other Information

1.2.1 Available on the Product Media

- Read the READMEs on the media.
- Get the detailed 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 12 SP2. This directory includes PDF versions of the SUSE Linux Enterprise Server 12 SP2 Installation Quick Start and Deployment Guides. Documentation (if installed) is available below the `/usr/share/doc/` directory of an installed system.
- These Release Notes are identical across all architectures, and the most recent version is always available online at http://www.suse.com/releasenotes/. Some entries are listed twice, if they are important and belong to more than one section.

1.2.2 Externally Provided Documentation

- http://www.suse.com/documentation/sles-12/ contains additional or updated documentation for SUSE Linux Enterprise Server 12 SP2.
- Find a collection of White Papers in the SUSE Linux Enterprise Server Resource Library at https://www.suse.com/products/server/resource-library/?ref=b#WhitePapers.

1.3 How to Obtain Source Code

This SUSE product includes materials licensed to SUSE under the GNU General Public License (GPL). The GPL requires SUSE to provide the source code that corresponds to the GPL-licensed material. The source code is available for download at http://www.suse.com/download-linux/source-code.html. Also, for up to three years after distribution of the SUSE product, upon request, 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.

1.4 Support Statement for SUSE Linux Enterprise Server

To receive support, customers need an appropriate subscription with SUSE; for more information, see http://www.suse.com/products/server/services-and-support/.

1.4.1 General Support Statement

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, duplicate customer problems, isolate problem area and provide resolution for problems not resolved by Level 1 or alternatively 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 12 SP2 and its Modules are delivered with L3 support for all packages, except the following:

- Technology Previews
- sound, graphics, fonts and artwork
- packages that require an additional customer contract
- packages provided as part of the Software Development Kit (SDK)

SUSE will only support the usage of original (e.g., unchanged or un-recompiled) packages.

1.4.2 Technology Previews

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

Whether a technical preview will be moved to a fully supported package later, depends on customer and market feedback. A technical preview does not automatically result in support at a later point in time. Technical previews could be dropped at any time and SUSE is not committed to provide a technical preview later in the product cycle.

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

1.4.2.1 Docker Orchestration

Starting with Docker 1.12, the orchestration (swarm) is now an integral part of the engine. It is provided as a Technology Preview within the SLES 12 Containers module.

1.4.2.2 Support for Current AMD Radeon GPUs

As a technical preview, SUSE Linux Enterprise ships the graphics driver `xf86-video-amdgpu` *for current AMD Radeon GPUs.*

Since this driver is still in an experimental state, it is not installed by default. By default, it is only enabled for one GPU on which it was tested successfully.

Important: At this stage, this driver is not supported.

To be able to use the driver, first install the package `xf86-video-amd`. Then, enable it for your GPU by editing `/etc/X11/xorg_pci_ids`.

The required format is: `\<VendorID\>\<DeviceID\>`. It is also described in the configuration file itself.

To find vendor ID and device ID, use the command:

```
lspci -n | grep 0300
```

All supported vendor IDs/device IDs are already in the file but are commented out. For your vendor ID/device ID combination, remove the comment character `#` from the beginning of the line.

1.4.2.3 Support for UEFI in QEMU Virtual Machines

libvirt and KVM/QEMU now support UEFI for virtual machines. UEFI firmware is provided through the `qemu-ovmf-x86_64` package.

1.4.2.4 Converting Physical Machines to KVM Virtual Machines

`libguestfs` *has the tool* `virt-v2v` *to convert virtual machines from Xen to KVM. However, previously, it was not possible to convert physical installations to virtual machine installations.*

As a technology preview, SLES 12 SP2 now ships the tool `virt-p2v` in `libguestfs`. `virt-p2v` allows converting physical machines into KVM guests.

This also means that `libguestfs` has been updated to a more recent version, bringing new features and fixes.

1.4.2.5 Technology Previews: AArch64 (ARMv8)

1.4.2.5.1 GNOME Desktop Environment as a Technology Preview on AArch64

The GNOME desktop environment (including GNOME Shell and GDM) is now available on the AArch64 architecture as an unsupported technology preview.

The only supported graphical environment on the AArch64 architecture is IceWM with XDM as the display manager.

1.4.2.6 Technology Previews: AMD64/Intel 64 64-Bit (x86_64)

1.4.2.6.1 NVDIMM Support

NVDIMM support has been added as Technical Preview. While many of its subsystems are stable, it is recommended to test it for your specific use case and workload before using it in production environments.

1.4.2.6.2 Guest 3D Acceleration With virtio-gpu

In QEMU version 2.5 and before, virtual graphical cards had no 3D support. Therefore, in the past, QEMU guests could not use 3D acceleration.

From the perspective of the host, QEMU 2.5 and later include `virtio-gpu`. `virtio-gpu` allows rendering OpenGL commands from the guest on the GPU of the host. This results in a large improvement of the OpenGL 3D performance of the guest.

From the perspective of the guest, the Linux kernel 4.4 and higher include the `virtio-gpu` driver.

When attaching a `virtio-gpu` to a guest which has the Linux kernel 4.4 or higher and supports OpenGL 3D 3.x acceleration, the guest can use 3D acceleration and will get around 50 percent of native performance.

Unlike VGA pass-through or using an NVIDIA GRID card, `virtio-gpu` does not need a dedicated graphical card or special hardware. Depending on the performance of GPU of the host, it can also provide OpenGL 3D acceleration for multiple guests.

1.5 Modules, Extensions, and Related Products

This section comprises information about modules and extensions for SUSE Linux Enterprise Server 12 SP2. Modules and extensions add parts or functionality to the system.

1.5.1 Available Modules

Modules are fully supported parts of SUSE Linux Enterprise Server with a different life cycle and update timeline. They are a set of packages, have a clearly defined scope and are delivered via online channel only. Release notes for modules are contained in this document.

The following modules are available for SUSE Linux Enterprise Server 12 SP2:

Name	Content	Life Cycle
Advanced Systems Management Module	CFEngine, Puppet and the Machinery tool	Frequent releases
Certifications Module*	FIPS 140-2 certification-specific packages	Certification-dependant
Containers Module	Docker, tools, prepackaged images	Frequent releases
Legacy Module*	Sendmail, old IMAP stack, old Java, …	Until September 2017
Public Cloud Module	Public cloud initialization code and tools	Frequent releases
Toolchain Module	GNU Compiler Collection (GCC)	Yearly delivery
Web and Scripting Module	PHP, Python, Ruby on Rails	3 years, ~18 months overlap

* Module is not available for the AArch64 architecture.

1.5.2 Available Extensions

Extensions add extra functionality to the system and require an own registration key that is usually liable for costs. Extensions are delivered via online channel or physical media. In many cases, extensions have own release notes documents that are available from https://www.suse.com/releasenotes/.

The following extensions are available for SUSE Linux Enterprise Server 12 SP2:

- SUSE Linux Enterprise High Availability Extension: https://www.suse.com/products/high-availability
- Geo Clustering for SUSE Linux Enterprise High Availability Extension: https://www.suse.com/products/highavailability/geo-clustering
- SUSE Linux Enterprise Real Time Extension: https://www.suse.com/products/realtime
- SUSE Linux Enterprise Software Development Kit
- SUSE Linux Enterprise Workstation Extension: https://www.suse.com/products/workstation-extension

1.5.3 Derived and Related Products

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

- SUSE Enterprise Storage: https://www.suse.com/products/suse-enterprise-storage
- SUSE Linux Enterprise Desktop: https://www.suse.com/products/desktop
- SUSE Linux Enterprise Live Patching: https://www.suse.com/products/live-patching
- SUSE Linux Enterprise Server for SAP Applications: https://www.suse.com/products/sles-for-sap
- SUSE Manager: https://www.suse.com/products/suse-manager
- SUSE OpenStack Cloud: https://www.suse.com/products/suse-openstack-cloud

1.6 Security, Standards, and Certification

SUSE Linux Enterprise Server 12 SP2 has been submitted to the certification bodies for:

- Common Criteria Certification (http://www.commoncriteriaportal.org/)
- FIPS 140-2 validation, see: http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140InProcess.pdf

For more information about certification, see https://www.suse.com/security/certificates.html.

2 Installation and Upgrade

SUSE Linux Enterprise Server can be deployed in several ways:

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

2.1 Updating the Installer at the Beginning of the Installation or Upgrade

Until SLES 12 SP1, the only option to update the installer was to apply a driver update disk. This involved manual work such as downloading the driver update and explicitly pointing the installer at it.

Starting with the SLES 12 SP2 installer, at the beginning of the installation or upgrade, the installer contacts the update server to find out whether there are updates for the installer available. If there are, they are automatically applied and YaST is restarted.

The installer is able to download the updates from the regular update server, a local SMT server, or a custom URL. Alternatively, you can disable this functionality completely.

If the automatic update fails for some reason or there is a regression in the installer after installing the updates, disable this feature useing the boot option `self_update=0`.

For more information, see the documentation at https://github.com/yast/yast-installation/blob/SLE-12-SP2/doc/SELF_UPDATE.md.

2.2 Installation

This section includes information related to the initial installation of SUSE Linux Enterprise Server 12 SP2. For information about installing, see *Deployment Guide* at https://www.suse.com/documentation/sles-12/book_sle_deployment/data/book_sle_deployment.html.

2.2.1 Network Interfaces Configured via linuxrc Take Precendence

For some configurations with many network interfaces, it can take several hours until all network interfaces are initialized (see https://bugzilla.suse.com/show_bug.cgi?id=988157). In such cases, the installation is blocked. SLE 12 SP1 and earlier did not offer a workaround for this behavior.

With SLE 12 SP2, you can speed up interactive installations on systems with many network interfaces by configuring them via linuxrc. When a network interface is configured via linuxrc, YaST will not perform automatic DHCP configuration for any interface. Instead, YaST will continue to use the configuration from linuxrc.

To configure a particular interface via linuxrc, add the following to the boot command line before starting the installation:

```
ifcfg=eth0=dhcp
```

In the parameter, replace `eth0` with the name of the appropriate network interface. The `ifcfg` option can be used multiple times.

2.2.2 Media-based Sources Are Disabled After Installation If They Are Not Needed

Previously, when installing from local media, like a CD/DVD or USB drive, these sources remained enabled after the installation.

This could cause problems during software installation, upgrade or migration because an old or obsolete installation source remained there. Additionally, if the source was physically removed (for instance, by ejecting the CD/DVD), Zypper would complain about the source not being available.

After the installation, YaST will now check every local source whether the product they provide is also available through a remote repository. In that case, it will disable them.

2.2.3 Partitioning Proposal: "Flexible Partitioning" Feature Has Been Removed

YaST is a very configurable installer that allows setting very different behaviors for each product using it (SUSE Enterprise Linux, openSUSE, etc.). In previous versions of YaST, it was possible to use a feature called "Flexible Partitioning". This feature has become obsolete, as the more standard proposal mechanism has been used by SLE and openSUSE in all recent releases.

The new version of YaST detects when a (modified) installer tries to use the obsolete "Flexible Partitioning" feature, alerts the user and falls back to the standard proposal mechanism automatically.

2.2.4 YaST Clears New Partitions

Previously, when YaST created a new partition, there could be signatures of previous MD RAIDs on the partition. That caused the MD RAID to be auto-assembled which made the partition busy. Thus, subsequent commands on the new partition failed.

When creating partitions with YaST now, storage signatures are deleted before auto-assembly takes place.

2.2.5 Host Name Setting During Installation

During installation, the hostname is set to `install`, the DHCP-provided value, if any, or the value of the boot option `hostname`. The host name used during installation is not propagated to `/etc/hostname` of the installed system except when set using the boot option `hostname`.

2.2.6 More Explicit and Configurable Importing of SSH Host Keys

During a installation of SUSE Linux Enterprise, existing SSH host keys from a previous installation were imported into the new system. This is convenient in some network scenarios, but as it was done without explicitly informing the user, it could lead to undesired situations.

The installer no longer silently imports the SSH host keys from the most recent Linux installation on the disk. It now allows you to choose whether to import SSH host keys and from which partition they should be imported. It is now also possible to import the rest of the SSH configuration in addition to the keys.

To import previous SSH host keys and configuration during the installation, proceed until the page *Installation Summary*, then choose *Import SSH Host Keys and Configuration*.

2.2.7 Option to Create AutoYaST Profile During Installation Has Been Removed

In earlier versions of SUSE Linux Enterprise, you could clone the system configuration as an AutoYaST profile during installation. But many services and system parameters can only be configured after the installation process has been finished and the system is up and running. This can lead to a situation where parts of the desired configuration are missing in the cloned systems.

The option of creating an AutoYaST profile has been removed. However, you can still create an AutoYaST profile from the running system, after you have made sure that the system configuration fits your needs.

2.2.8 Reading Registration Codes from a USB Drive

During the installation of SUSE products, it can be tedious to remember and type in registration codes.

You can now save the registration codes to a USB drive and have YaST read them automatically.

For more information, see: https://github.com/yast/yast-registration/wiki/Loading-Registration-Codes-From-an-USB-Storage-%28Flash-Drive-HDD%29.

2.3 Upgrade-Related Notes

This section includes upgrade-related information for SUSE Linux Enterprise Server 12 SP2. For information about general preparations and supported upgrade methods and paths, see the documentation at https://www.suse.com/documentation/sles-12/book_sle_deployment/data/cha_update_sle.html.

2.3.1 Online Migration with Live Patching Enabled

The SLES online migration process reports package conflicts when Live Patching is enabled and the kernel is being upgraded. This applies when crossing the SP1/SP2 boundary.

To prevent the conflicts, before starting the migration, execute the following as a super user:

```
zypper rm $(rpm -qa kgraft-patch-*)
```

2.3.2 Support for PIDs cgroup Controller

The version of systemd shipped in SLES 12 SP2 uses the PIDs cgroup controller. This provides some per-service `fork()` *bomb protection, leading to a safer system.*

However, under certain circumstances you may notice regressions. The limits have already been raised above the upstream default values to avoid this but the risk remains.

If you notice regressions, you can change a number of `TasksMax` settings.

To control the default `TasksMax=` setting for services and scopes running on the system, use the `system.conf` setting `DefaultTasksMax=`. This setting defaults to `512`, which means services that are not explicitly configured otherwise will only be able to create `512` processes or threads at maximum.

For thread- or process-heavy services, you may need to set a higher `TasksMax` value. In such cases, set `TasksMax` directly in the specific unit files. Either choose a numeric value or even `infinity`.

Similarly, you can limit the total number of processes or tasks each user can own concurrently. To do so, use the `logind.conf` setting `UserTasksMax` (the default is `12288`).

`nspawn` containers also have a `TasksMax` value set now, the default is `16384`.

2.4 For More Information

For more information, see *Section 3, "Architecture Independent Information"* and the sections relating to your respective hardware architecture.

3 Architecture Independent Information

Information in this section pertains to all architectures supported by SUSE Linux Enterprise Server 12 SP2.

3.1 Kernel

3.1.1 Transparent Huge Page Defragmentation Disabled by Default

Transparent Huge Pages (THP) are an important alternative to `hugetlbfs` *that boosts performance for some applications by reducing the amount of work a CPU must do when translating virtual to physical addresses. It is particularly important for virtual machine performance where there are two translation layers.*

Early in the lifetime of the system there is enough free memory such that these pages can be allocated cheaply. Once the system is running for long enough, memory must be reclaimed and compacted to allocate the THP. This forces applications to stall for potentially long periods of time which many applications cannot tolerate. In many tuning guides, there is simply a recommendation to disable THP in a number of cases.

SLE 12 SP2 disables THP defragmentation by default. THPs will only be used if they are available instead of stalling on defragmentation. Normally, the defragmentation work is deferred and THPs will be created in the future. However, if an application explicitly requests such behavior via `madvise()`, it will stall.

If a system has many applications that are willing to stall to allocate THP, it is possible to restore the previous behavior of SLE via `sysfs`:

```
echo always > /sys/kernel/mm/transparent_hugepage/defrag
```

3.1.2 Enabling Enhanced Information About Physical Memory Page Ownership and Status

Detailed information about physical memory pages can help answer questions such as:

- *Which kernel subsystem or driver has allocated which pages?*
- *What page status flags are set?*

This is useful for L3 support of the kernel and during development and testing of out-of-tree kernel modules, for example, to debug memory leaks. Previously, kernel interfaces could only provide a subset of the page status flags, and only provide a summary about generic memory usage categories.

The Linux kernel shipped with SLE 12 SP2 can provide more detailed information. However, tracking extra information about each page that the kernel allocates creates overhead in terms of code to be executed and memory used. Therefore, this feature is disabled by default.

This feature is shipped with all kernel versions of SLE 12 SP2 and can be enabled during boot using the kernel parameter `page_owner=on`.

To obtain the status of all pages, use:

```
cat /sys/kernel/debug/page_owner > file
```

The file contains the following for each physical page:

- Allocation flags
- Status flags
- Page migration status
- Backtrace leading to the allocation

Additional postprocessing of the output can be used, for example, to count the number of pages for each unique backtrace which can help discover a code path that leaks memory.

3.1.3 Subset of Scheduler Debugging Statistics Disabled by Default

The CPU scheduler maintains a number of statistics for the purposes of debugging, some tracepoints and sleep profiling. They are only useful for detailed analysis but they incur an overhead for all users. They may be disabled at kernel build time but they are enabled as debugging in the field is important and tools like latencytop depend on them.

Some expensive scheduler debugging statistics are disabled by default. Enabling sleep profiling or running `latencytop` will activate them automatically but activating the tracepoints will require user intervention. The affected tracepoints are `sched_stat_wait`, `sched_stat_sleep`, `sched_stat_iowait`, `sched_stat_blocked` and `sched_stat_runtime`.

They can be activated at runtime using:

```
echo 1 > /sys/kernel/debug/tracing/events/sched/enable
```

They can be disabled at runtime using:

```
echo 0 > /sys/kernel/debug/tracing/events/sched/enable
```

The first number of tracepoint activations may contain stale data until the necessary data is collected. If this is undesirable, it is possible to activate them at boot time via the kernel parameter `schedstats=enable`.

3.1.4 Incompatible Changes in the New 4.4 Kernel

The following minor changes have been identified in the 4.4 kernel:

- Support for TCP Limited Slow Start (RFC3742) has been removed. This feature had multiple drawbacks and questionable benefit. Its implementation was inefficient and difficult to configure. The problem that Limited Slow Start was trying to solve is now better covered by the Hybrid Slow Start algorithm which is part of default congestion control algorithm, CUBIC.

- The `kernel.blk_iopoll` sysctl has been removed. This setting allowed toggling some block device drivers between iopoll and non-iopoll mode. This allowed for easier debugging of these drivers during early development. Since using this toggle was dangerous and the toggle is not needed for production setups, it has been removed.

- The `cgroup.event_control` file is only available in cgroups with a memcg attached to it. There was no code using this interface outside of memcg, so this change is considered harmless.

- The `vm.scan_unevictable_pages` sysctl has been removed because the functionality it was backing had been removed in 2011. Any usage of the file has been reported to the kernel log with an explanation that the file has no effect. There were no reports about a use case requiring the functionality.

- The `/sys/devices/system/memory/memory%d/end_phys_index` file has been removed, because the information it exposed is considered internal to the kernel and an implementation detail. This information is not required for the memory hotplug functionality.

3.1.5 Partial Memory Mirroring

Memory mirroring offers increased system reliability. However, full memory mirroring also dramatically decreases available memory size.

Partial memory mirroring addresses this issue by setting up a smaller mirrored memory range and using this range for kernel code and data structures. The remaining memory operates in regular mode which leaves more room for applications. This feature requires support in hardware and EFI firmware and is currently supported on Fujitsu PRIMEQUEST 2000 series systems and its successor models.

3.1.6 Support for CXL Flash Storage Device Driver

The CXL flash storage device provides persistent, flash-based storage using CAPI technology.

3.1.7 Enhanced Accounting and Reporting of shmem Swap Usage

There was a request to provide information about how much of Linux-kernel shared memory (`shmem`) is swapped out, for processes using such memory segments. `shmem` mappings are either System V shared memory segments, mappings created by `mmap()` with `MAP_ANONYMOUS` / `MAP_SHARED` flags, and shared `mmap()` mappings of files residing on the tmpfs RAM disk file system. Prior to the implemented changes, in `/proc/pid/smaps`, swap usage for these segments would have been shown as `0`.

The kernel has been modified to show swap usage of `shmem` segments properly in `/proc/pid/smaps` files. Due to `shmem` implementation limitations, this value will also count swapped-out pages that the process has mapped, but never touched, which differs from anonymous memory accounting. Due to the same limitations and to prevent excessive CPU overhead, the `VmSwap` field in `/proc/pid/status` is unaffected and will not account for swapped-out portions of `shmem` mappings. In addition, the `/proc/pid/status` file has been enhanced to include three new `Rss*` fields as a breakdown of the `VmRSS` field to `anonymous`, `file` and `shmem` mappings. Example excerpt:

```
VmRSS:      5108 kB
RssAnon:              92 kB
RssFile:            1324 kB
RssShmem:           3692 kB
```

3.2 Kernel Modules

An important requirement for every Enterprise operating system is the level of support customers receive for their environment. Kernel modules are the most relevant connector between hardware ("controllers") and the operating system.

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

3.2.1 NVDIMM Kernel Subsystem

Non-volatile DIMMs are byte-addressable memory chips that fit inside a computer's normal memory slot but are, in contrast to DRAM chips persistent and thus can be used as an enhancement or replacement for a computer's hard disk drives. This imposes several challenges, namely:

- *Discovery of hardware*
- *Mapping and addressing of this new memory type*
- *Atomic semantics as with traditional storage media*
- *Page frame addressing like with traditional memory*

The Linux kernel shipped with SLE now includes several drivers to address these challenges:

- Hardware discovery is initiated via the ACPI NFIT (Non-Volatile Memory Firmware Interface Table) mechanism and realized with the device driver `nfit.ko`.
- Mapping and addressing of NVDIMMs is accomplished by the device driver `nd_pmem.ko`.
- The driver `nd_btt.ko` takes care of (optional) atomic read/write semantics to the underlying hardware.
- The pfn portion of `nd_pmem.ko` provides the ability to address NVDIMM memory just like any other DRAM type memory.

3.2.2 Direct Access to Files in Non-Volatile DIMMs

The page cache is usually used to buffer reads and writes to files. It is also used to provide the pages which are mapped into userspace by a call to `mmap`*. For block devices that are memory-like, the page cache pages would be unnecessary copies of the original storage.*

The Direct Access (DAX) kernel code avoids the extra copy by directly reading from and writing to the storage device. For file mappings, the storage device is mapped directly into userspace. This functionality is implemented in the XFS and Ext4 file systems.

3.2.3 ZRAM Block Device

The ZRAM module creates RAM-based block devices. Pages written to these disks are compressed and stored in memory itself. Such disks allow for very fast I/O. Additionally, compression provides memory savings.

ZRAM devices can be managed and configured with the help of the tool `zramctl` (see the man page of `zramctl(8)`). Configuration persistence is ensured by `zramcfg` system service.

3.2.4 Memory Compression with zswap

Usually, when a system's physical memory is exceeded, the system moves some memory onto reserved space on a hard drive, called "swap" space. This frees physical memory space for additional use. However, this process of "swapping" memory onto (and off a hard drive is much slower than direct memory access, so it can slow down the entire system.

The `zswap` driver inserts itself between the system and the swap hard drive, and instead of writing memory to a hard drive, it compresses memory. This speeds up both writing to swap and reading from swap, which results in better overall system performance while using swap.

To enable the `zswap` driver, write `1` or `Y` to the file `/sys/module/zswap/parameters/enabled`.

Storage Back-ends

There are two back-ends available for storing compressed pages, `zbud` (the default), and `zsmalloc`. The two back-ends each have their own advantages and disadvantages:

- The effective compression ratio of `zbud` cannot exceed 50 percent. That is, it can at most store two uncompressed pages in one compressed page. If the workload's compression ratio exceeds 50% for all pages, `zbud` will not be able to save any memory.
- `zsmalloc` can achieve better compression ratios. However, it is more complex and its performance is less predictable.
- `zsmalloc` does not free pages when the limit set in `/sys/module/zswap/parameters/max_pool_percent` is reached. This is reflected by the counter `/sys/kernel/debug/zswap/reject_reclaim_fail`.

It is not possible to give a general recommendation on which storage back-end should be used, as the decision is highly dependent on workload. To change the storage back-end, write either `zbud` or `zsmalloc` to the file `/sys/module/zswap/parameters/zpool`. Pick the back-end before enabling zswap. Changing it later is unsupported.

Setting `zswap` Memory

Compressed memory still uses a certain amount of memory, so `zswap` has a limit to the amount of memory which will be stored compressed, which is controllable through the file `/sys/module/zswap/parameters/max_pool_percent`. By default, this is set to `20`, which indicates `zswap` will use 20 percent of the total system physical memory to store compressed memory.

The `zswap` memory limit has to be carefully configured. Setting the limit too high can lead to premature out-of-memory situations that would not exist without `zswap`, if the memory is filled by non-swappable non-reclaimable pages. This includes mlocked memory and pages locked by drivers and other kernel users.

For the same reason, performance can also be hurt by compression/decompression if the current workload's workset would, for example, fit into 90 percent of the available RAM, but 20 percent of RAM is already occupied by `zswap`. This means that the missing 10 percent of uncompressed RAM would constantly be swapped out of/in to the memory area compressed by `zswap`, while the rest of the memory compressed by `zswap` would hold pages that were swapped out earlier which are currently unused. There is no mechanism that would result in gradual writeback of those unused pages to let the uncompressed memory grow.

Freeing `zswap` Memory

`zswap` will only free its pages in certain situations:

- The processes using the pages free the pages or exit
- When the storage back-end `zbud` is in use, `zswap` will also free memory when its configured memory limit is exceeded. In this case, the oldest `zswap` pages are written back to disk-based swap.

Memory Allocation Issues

In theory, it can happen that `zswap` is not yet exceeding its memory limit, but already fails to allocate memory to store compressed pages. In that case, it will refuse to compress any new pages and they will be swapped to disk immediately. For confirmation whether this issue is occurring, check the value of `/sys/kernel/debug/zswap/reject_alloc_fail`.

3.3 Networking

3.3.1 Better Information About Physical Port IDs Used by Network Interfaces with NPAR/SR-IOV Capabilities

Previously, YaST offered no way to know whether two interfaces with NPAR/SR-IOV capabilities were sharing the same physical port. As a result, users could bond them without realizing that they were not getting the desired effect in terms of redundancy.

Information about the physical port ID has been added to *Interface Overview* and also for each entry of the *Bond Slaves* table, so you can now inspect the physical port ID when selecting an interface.

Additionally, you will be alerted when trying to bond devices sharing the same physical port.

3.4 Systems Management

3.4.1 SASL Integration in sudo

When SUSE Linux Enterprise 12 was first released, the `sudo` *binary did not correctly support SASL authentication for LDAP because the package was built without a build dependency on the package* `cyrus-sasl-devel`.

To be able to use `sudo` with SASL, update to the latest version of the package `sudo`. For information about enabling SASL authentication for `sudo`, see `man 5 sudoers.ldap`.

3.4.2 systemd: Support for System V and LSB Init Scripts Has Been Moved Out of Core Daemon

To ease future maintenance, in SLE 12 SP2, systemd was updated to version 228. This version does not support using System V and LSB init scripts from the `systemd` *daemon itself any more.*

This functionality is now implemented as a generator that creates systemd unit files from System V/LSB init scripts. These unit files are generated at boot or when systemd is reloaded. Therefore, to have changed System V init scripts recognized by systemd, run `systemctl daemon-reload` or reboot the machine.

For more information, see the man page of `systemd-sysv-generator` (`man systemd-sysv-generator`).

If you are packaging software that ships System V init scripts, use the RPM macros documented at https://en.opensuse.org/openSUSE:Systemd_packaging_guidelines (https://en.opensuse.org/openSUSE:Systemd_packaging_guidelines#Register_services_in_install_scripts) (Section "Register Services in Install Scripts").

3.4.3 AutoYaST: Applying the First-Stage Network Configuration to the Installed System

Due to a problem in the AutoYaST version shipped with SLE 12 SP1, the network configuration used during the first stage was always copied to the installed system. This happened regardless of the value of `keep_install_network` *in the AutoYaST profile.*

SLE 12 SP2 behaves as expected and `keep_install_network` will be set to `true` by default.

3.4.4 New YaST VPN module

The new YaST VPN module provides an intuitive and easy to use interface for setting up VPN gateways and clients. It simplifies the setup of typical IPSec VPN gateways and clients.

IPSec is an open and standardized VPN protocol, natively supported by most operating systems and devices, including Linux, Unix, Windows, Android, Blackberry, Apple iOS and MacOS, without the need for third-party software solution.

Using the YaST VPN module, you can create VPN gateways for the following scenarios:

- Provide network access to Linux clients authenticated via a pre-shared key or certificate.
- Provide network access to Windows 7, 8, 10, and Blackberry clients authenticated via a combination of certificate and username/password.
- Provide network access to Android, iOS, and MacOS clients authenticated via a combination of a pre-shared key and username/password.

Additionally, you can set up connections to remote VPN gateways, for the following scenarios:

- Prove client identity with a pre-shared key.

3.4.5 Enrolling in a Microsoft Active Directory Domain via YaST

You can configure a SLES computer to become a member in Microsoft Active Directory to leverage its user account and group management. In previous versions of SLES, enrolling a computer in a Microsoft Active Directory was a lengthy and error-prone procedure.

In SLES 12 SP2, YaST ships with the new configuration tool *User Logon Management* (previously *Authentication Client*) which offers a powerful yet simple user interface for joining an Active Directory domain and allows authenticating users using those domain accounts. In addition to Active Directory, the editor can also set up authentication against a generic Kerberos or LDAP service.

3.4.6 ntp 4.2.8

ntp was updated to version 4.2.8.

- *The ntp server ntpd does not synchronize with its peers anymore and the peers are specified by their host name in* `/etc/ntp.conf`.
- *The output of* `ntpq --peers` *lists IP numbers of the remote servers instead of their host names.*

Name resolution for the affected hosts works otherwise.

Parameter changes

The meaning of some parameters for the sntp command-line tool have changed or have been dropped, for example `sntp -s` is now `sntp -S`. Please review any `sntp` usage in your own scripts for required changes.

After having been deprecated for several years, ntpdc is now disabled by default for security reasons. It can be re-enabled by adding the line `enable mode7` to `/etc/ntp.conf`, but preferably `ntpq` should be used instead.

3.4.7 Installing kGraft Patches with Weak Package Dependency Resolution Disabled

In environments with a clearly defined list of packages to be installed on the system and weak package dependency resolution disabled via `solver.onlyRequires=true` *in* `/etc/zypp/zypp.conf`*, automatic installation of the initial kGraft patch is broken.*

As an aid in this situation, the package `kernel-$FLAVOR-kgraft` is provided. Installing this package pulls the associated kGraft patch into the system.

3.4.8 Sudo Now Respects Groups Added by the pam_group Module

Sudo now respects groups added by the `pam_group` module and adds these groups to the target user.

If there is a user `tux`, you can now use the following to add it to the group `games`:

1. Open `/etc/security/group.conf` and add: `sudo;*;tux;Al0000-2400;games`
2. Open `/etc/pam.d/sudo` and add the following line at the beginning of the file: `auth required pam_group.so`
3. Then run: `sudo -iu tux id`

In SLE 12 SP1 and before, the user `tux` would not have been added to the group `games`:

```
uid=1002(tux) gid=100(users) groups=100(users)
```

In SLE 12 SP2, the user `tux` is added to the group `games`:

```
uid=1002(tux) gid=100(users) groups=100(users),40(games)
```

3.5 Performance Related Information

3.5.1 perf Provides Guest Exit Statistics

This feature enables `perf` to collect guest exit statistics based on the `kvm_exits` made by the threads of a guest-to-host context. The statistics report is grouped by exit reason. This can used as an indicator of the performance of a VM under a certain workload.

Besides `kvm_exits`, hypervisor calls are also reported and grouped by `hcall` reason. The statistics can be shown for an individual guest or all guests running on a system.

3.5.2 Deferred and Parallelized Initialization of Page Structures in Memory Management

Page initialization takes a very long time on large-memory systems. This is one of the reasons why large machines take a long time to boot.

The kernel now provides deferred initialization of page structures on the x86_64 architecture. Only approximately 2 GB per memory node are initialized during boot, the rest is initialized in parallel with the boot process by kernel threads named `pgdatinitX`, where X indicates the node ID.

3.6 Storage

3.6.1 Root File System Conversion to Btrfs Not Supported

If it is not the root file system and if the file system has at least 20 % free space available, in-place conversion of an existing Ext2/Ext3/Ext4 or ReiserFS file system is supported for data mount points.

SUSE does not recommend or support in-place conversion of OS root file systems. In-place conversion to Btrfs of root file systems requires manual subvolume configuration and additional configuration changes that are not automatically applied for all use cases.

To ensure data integrity and the highest level of customer satisfaction, when upgrading, maintain existing root file systems. Alternatively, reinstall the entire operating system.

3.6.2 /var/cache on an Own Subvolume for Snapshots and Rollback

`/var/cache` *contains very volatile data, like the Zypper cache with RPM packages in different versions for each update. As a result of storing data that is mostly redundant but highly volatile, the amount of disk space a snapshot occupies can increase very fast.*

To solve this, move `/var/cache` to a separate subvolume. On fresh installations of SLE 12 SP2 or newer, this is done automatically. To convert an existing root file system, perform the following steps:

1. Find out the device name (`/dev/sda2`, `/dev/sda3` etc.) of the root file system: `df /`
2. Identify the parent subvolume of all the other subvolumes. For SLE 12 installations, this is a subvolume named `@`. To check if you have a `@` subvolume, use: `btrfs subvolume list / | grep '@'`. If the output of this command is empty, you do not have a subvolume named `@`. In that case, you may be able to proceed with subvolume ID 5 which was used in older versions of SLE.
3. Now mount the requisite subvolume.
   - If you have a `@` subvolume, mount that subvolume to a temporary mount point: `mount <root_device> -o subvol=@ /mnt`
   - If you don't have a `@` subvolume, mount subvolume ID 5 instead: `mount <root_device> -o subvolid=5 /mnt`
4. `/mnt/var/cache` can already exist and could be the same directory as `/var/cache`. To avoid data loss, move it: `mv /mnt/var/cache /mnt/var/cache.old`
5. In either case, create a new subvolume: `btrfs subvol create /mnt/var/cache`

6. If there is now a directory `/var/cache.old`, move it to the new location: `mv /var/cache.old/* /mnt/var/cache`. If that is not the case, instead do: `mv /var/cache/* /mnt/var/cache/`

7. Optionally, remove `/mnt/var/cache.old`: `rm -rf /mnt/var/cache.old`

8. Unmount the subvolume from the temporary mount point: `umount /mnt`

9. Add an entry to `/etc/fstab` for the new `/var/cache` subvolume. Use an existing subvolume as a template to copy from. Make sure to leave the UUID untouched (this is the root file system's UUID) and change the subvolume name and its mount point consistently to `/var/cache`.

10. Mount the new subvolume as specified in /etc/fstab: `mount /var/cache`

3.6.3 nvme-cli: A User-Space Tool to Manage NVMe Devices on Linux

The tool `nvme-cli` provides management features to NVMe devices, such as adapter information retrieval, namespace creation/formatting and adapter firmware update.

3.6.4 systemd: The NFS Mount Option bg Is Deprecated

The upstream developers of systemd do not support the NFS mount option `bg` any more. While this mount option is still supported in SLE 12 SP2, it will be removed in the next version of SLE.

It will be replaced by the systemd mount option `nofail`.

3.6.5 Snapper: Cleanup Rules Based on Fill Level

Some programs do not respect the special disk space characteristics of a Btrfs file system containing snapshots. This can result in unexpected situations where no free space is left on a Btrfs filesystem.

Snapper can watch the disk space of snapshots that have automatic cleanup enabled and can try to keep the amount of disk space used below a threshold.

If snapshots are enabled, the feature is enabled for the root file system by default on new installations.

For existing installations, the system administrator must enable quota and set limits for the cleanup algorithm to use this new feature. This can be done using the following commands:

1. `snapper setup-quota`
2. `snapper set-config NUMBER_LIMIT=2-10 NUMBER_LIMIT_IMPORTANT=4-10`

For more information, see the man pages of `snapper` and `snapper-configs`.

3.7 Virtualization

3.7.1 Virtual Machine Driver Pack 2.4 (VMDP 2.4)

SUSE Linux Enterprise Virtual Machine Driver Pack is a set of paravirtualized device drivers for Microsoft Windows operating systems. These drivers improve the performance of unmodified Windows guest operating systems that are run in virtual environments created using Xen or KVM hypervisors with SUSE Linux Enterprise Server 11 SP4 and SUSE Linux Enterprise Server 12 SP2. Paravirtualized device drivers are installed in virtual machine instances of operating systems and represent hardware and functionality similar to the underlying physical hardware used by the system virtualization software layer.

The new features of SUSE Linux Enterprise Virtual Machine Driver Pack 2.4 include:

- Support for SUSE Linux Enterprise Server 12 SP2
- Drivers for Windows Server 2016
- Drivers are no longer dependent on `pvvxbn.sys` being loaded
- Support Windows Multipoint Server

New driver and utility features:

- `pvvxbn.sys`: Issues a Xen shutdown/reboot at the end of the power down sequence unless the PV control flag `dfs` ("disable forced shutdown") is enabled.
- `pvvxblk.sys`: VirtIO: MSI vectors can now be used. Xen: support for indirect descriptors. Queuing, queue depth, and `max_segs` are tunable.
- `pvvxscsi.sys`: VirtIO: MSI vectors can now be used.

- `setup.exe`: Has enhanced support for `virt-v2v`.
- `pvctrl.exe : Can now modify NIC parameters. Enable/disable Xen pvvxblk queuing/queue depth (qdepth). Set Xen pvvxblk maximum number of segments (max_segs). Set debug print mask (dpm). Enable/disable Xen force shutdown after power-down sequence (dfs). Enable/disable virtio_serial MSI usage (vserial_msi).`

3.7.2 KVM

3.7.2.1 KVM Legacy Device Assignment Was Disabled

The legacy device assignment feature of KVM was disabled.

As a replacement, use VFIO. VFIO provides the same functionality and has the following advantages:

- It is actively maintained upstream while the legacy code is not.
- It is more secure.
- It supports new hardware features such as interrupt virtualization.

3.7.2.2 Obtaining Addresses with libvirt-nss

With `libvirt-nss`, you can obtain addresses of `dnsmasq` -backed KVM guests. For more information, see the *Virtualization Guide, Chapter "Obtaining IP Addresses with nsswitch for NAT Networks"*.

3.7.2.3 Post-Copy Live Migration Support in libvirt and QEMU/KVM

Pre-copy live migration can take a lot of time depending on the workload and page dirtying rate of the virtual machine.

libvirt and QEMU/KVM now support post-copy live migration. This means that the virtual machine starts running on the destination host as soon as possible and the RAM from the source host is pagefaulted into the destination over time. This ensures minimal downtime for the virtual machine.

The guest will run on target host immediately, only CPU state and device state are transferred to target host. If the network is down before all missing memory pages are copied from the original guest, the new guest will crash.

3.7.3 Xen

3.7.3.1 qemu-xen Has Been Dropped From the Xen Package

QEMU is a large software project that sees many bug and security fixes. Providing several different `qemu` *binaries is challenging for maintenance, requiring bug and security fixes to be backported to all the different* `qemu` *sources.*

The Xen package now uses `qemu-system-x86_64` from the `qemu` package instead of providing its own `qemu` binary.

3.7.3.2 Support UEFI in Xen HVM Virtual Machines

libvirt and Xen now support UEFI for virtual machines. UEFI firmware is provided through the `qemu-ovmf-x86_64` package.

3.7.3.3 GRUB Does Not Support vfb/vkbd Any More

The version of GRUB shipped with SLES 12 SP1 and SP2 does not support vfb/vkbd any more. This means that in Xen paravirtualized machines, there is no graphical display available while GRUB is active.

To be able to see and interact with GRUB, switch to the text-based xencon protocol: Modify the kernel parameter of the PV guest, add `console=hvc0 xencons=tty`, and connect with the command `console DOMAINNAME` of the `libvirt` toolstack.

3.7.3.4 libvirt XML Now Supports the External Block Scripts of Xen

The external block scripts of Xen, such as `block-drbd`, `block-dmmd` *could formerly only be used with* `xl` / `libxl` *using the disk configuration syntax* `script=`. *libvirt did not support such external scripts and thus could not be used with disks configured with the block scripts.*

External block scripts of Xen can now be used with libvirt by specifying base name of the block script in the `<source>` element of the disk. For example:

```
<source dev='dmmd:md;/dev/md0;lvm;/dev/vgxen/lv-vm01'/>
```

3.7.3.5 Support for the PVUSB Driver in Xen and the libvirt Xen Driver

`libxl` now has a PVUSB API which supports passing a USB device from the host to the guest domain via PVUSB. This functionality is also supported by the command line tool `xl`.

PVUSB support was also added to the libvirt `libxl` driver to use PVUSB functionality from the libvirt toolstack.

3.7.3.6 XEN: PV-OPS Kernel Supersedes kernel-xen

The Xen hypervisor functions have been ported over to the standard PV-OPS mechanism and are now included in the default kernel. As everything necessary is now provided by the default kernel, the package `kernel-xen` were removed.

3.7.4 Others

3.7.4.1 virt-convert: Support for Compressed Files in Within an OVA

According to the OVF 1.1.0 specification, OVA files can contain files compressed using `gzip`, for example, `vmdk` files. This case was previously not handled correctly.

In SLE 12 SP2, `virt-convert` will now correctly decompress `gz` files first and then convert them using `qemu-img`.

3.7.4.2 libiscsi Integration with QEMU

QEMU now integrates with `libiscsi`. This allows QEMU to access iSCSI resources directly and use them as virtual machine block devices. iSCSI-based disk devices can also be specified in the libvirt XML configuration. This feature is only available using the RAW image format, as the iSCSI protocol has some technical limitations.

3.7.4.3 DPDK Support for vhost-user Live Migration

Currently, the common back-end implementation to `vhost-user` *is* `dpdk`*. To support* `vhost-user` *live migration, a feature bit called* `VHOST_USER_PROTOCOL_F_LOG_SHMFD` *is required on both the QEMU side and the* `vhost-user` *back-end side.*

On the QEMU side, upstream version 2.6 already provides the required functionality. But on the DPDK side, the upstream release of DPDK 2.2.0 does not provide it.

The version of DPDK 2.2.0 shipped with SLE 12 SP2 is patched to provide the ability of `vhost-user` live migration.

3.7.4.4 wbemcli Now Allows Configuring the SSL/TLS version

Previously, it could be impossible to monitor certain servers that used very specific versions of the SSL/TLS protocols using `wbemcli`*.*

`wbemcli` can now be configured to use a specific SSL/TLS protocol version. To do so, use the environment variable `WBEMCLI_CURL_SSLVERSION`. Possible values are: `SSLv2`, `SSLv3`, `TLSv1`, `TLSv1_0` (TLSv1.0), `TLSv1_1` (TLSv1.1), `TLSv1_2` (TLSv1.2).

4 AMD64/Intel 64 (x86_64) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 12 SP2 for the AMD64/Intel 64 architectures.

4.1 Kernel NOHZ_FULL Process Scheduler Mode

Under normal operation, the kernel interrupts process execution several hundred times per second for statistics collection and kernel internal maintenance tasks. Despite the interruptions being brief, they add up. This adds an unpredictable amount of time to process run time. Highly timing sensitive applications may be disturbed by this activity.

The SLE kernel now ships with adaptive tick mode (`NOHZ_FULL`) enabled by default to reduce the number of kernel interrupts. With this option enabled and conditions for adaptive tick mode fulfilled, the number of interrupts goes down to ones per second.

4.2 System and Vendor Specific Information

4.2.1 Support for Run-Time Allocation of Huge Pages With 1 GB Size

In previous versions of SLE, huge pages with a size of 1 GB could only be allocated via a kernel parameter at boot. This has the following drawbacks:

- *You cannot specify the NUMA node for allocation.*
- *You cannot free these pages later without a reboot.*

On the x86-64 architecture, SLE can now allocate and free 1 GB huge pages at system run time, using the same methods that are also used for regular huge pages.

However, you should still allocate 1 GB huge pages as early as possible during the run time. Otherwise, physical memory can become fragmented by other uses and the risk of allocation failure grows.

5 POWER (ppc64le) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 12 SP2 for the POWER architecture.

5.1 Cluster Support and High Availability for POWER

Packages to facilitate cluster setup and to enable HA have been added to SUSE Linux High Availability Extension for POWER (LE).

5.2 Device Driver ibmvnic Has Been Added

vNIC (Virtual Network Interface Controller) is a new PowerVM virtual networking technology that delivers enterprise capabilities and simplifies network management. It is a high-performance, efficient technology that when combined with SR-IOV NIC 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.

5.3 Enhanced Support for System Call Filtering on POWER

Mode 2 of `seccomp` is now supported on POWER, allowing for fine-grained filtering of system calls. Support is available in both the kernel and in `libseccomp`.

5.4 Hardware Transactional Memory (HTM) support in glibc for POWER

Lock elision in the GNU C Library is available, but disabled by default. To enable it, set the environment variable `GLIBC_ELISION_ENABLE` to the value "yes".

6 IBM z Systems (s390x) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 12 SP2 for the IBM z Systems architecture. For more information, see http://www.ibm.com/developerworks/linux/linux390/documentation_novell_suse.html

IBM zEnterprise 196 (z196) and IBM zEnterprise 114 (z114) further on referred to as z196 and z114.

6.1 Hardware

6.1.1 Support for IPL Device in Any Subchannel Set

IPL devices are no longer restricted to subchannel set 0. The limitation is removed as of IBM zEnterprise 196 GA2.

6.1.2 Bus Awareness for z Systems in systemd

systemd now provides full and correct support for driver model buses specific to Linux on z Systems, such as `ccw`, `ccwgroup`, and `zfcp`.

6.2 Virtualization

6.2.1 Executing Hypervisor-Specific Actions During Boot

Depending on the hypervisor that a system runs on (such as z/VM, zKVM, or LPAR), during boot, differerent actions can be needed.

The service `virtsetup` is preconfigured to do that. To activate it, execute the following command:

```
systemctl enable virtsetup.service
```

To configure this service in more detail, see the file `/etc/sysconfig/virtsetup`. You can also edit the file through YaST:

```
yast2 sysconfig
```

6.2.2 VMUR Print Spool Options for Linux

Linux guests are now better integrated into the z/VM print solution. It is now possible to specify the spool options `CLASS` and `FORM` together with the `print` command of the VMUR tool.

6.2.3 zKVM: SIE Capability Exposed to User Space

Userspace applications can now query whether the Linux instance can act as a hypervisor by checking for the SIE (Start Interpretive Execution) capability. This is useful, for example, in continuous integration (CI) environments.

6.3 Storage

6.3.1 iSCSI Devices Not Enabled After Installation

When installing SLES 12 SP2, iSCSI devices may not be enabled after installation.

When configuring iSCSI volumes, make sure to set start mode to `automatic`. `onboot` is only valid for iSCSI devices which are supposed to be activated from the initrd, that is, when the system is booted from iSCSI. However, that is currently not supported on z Systems.

6.3.2 Query Host Access to Volume Support

You can now concurrently access DASD volumes from different operating system instances. Applications can now query whether a DASD volume is online within another operating system instance by querying the storage server for the online status of all attached hosts. The command `lsdasd` can display this information, and the commands `zdsfs`, `fdasd`, and `dasdfmt` can evaluate it.

6.4 Network

6.4.1 10GbE RoCE Express Feature for RDMA

SLES 12 SP2 supports the 10GbE RoCE Express feature on zEC12, zBC12 and IBM z13 via the Ethernet device using TCP/IP traffic without restrictions. Before using this feature on an IBM z13, make sure that the minimum required service is applied: z/VM APAR UM34525 and HW ycode N98778.057 (bundle 14). Use the default MTU size (1500).

SLES 12 SP2 now includes support for RDMA enablement and DAPL/OFED for z Systems. With the Mellanox virtualization support (SR-IOV) the limitation for LPAR use only on an IBM zEC12 or zBC12 is removed and RDMA can be used on an IBM z13.

6.4.2 Bridging HiperSockets to Ethernet

A HiperSocket port can now be configured to accept Ethernet frames to unknown MAC addresses. This enables it to be used as a member of a software bridge. Control and report of the bridge port status of the HiperSocket port and the udev events are performed via new sysfs attributes.

6.4.3 IPv6 Priority Queuing Added to qeth Device Driver

Priority queuing is now supported for IPv6, similarly to IPv4. This especially improves Linux Live Guest Migration by using IPv6 to minimize impact on workload traffic and enables priority queuing for all applications that use IPv6 QoS traffic operations.

6.4.4 Layer 2 Offloads Enabled

Classic OSA operation in layer 3 mode provides numerous offload operations, exchanging larger amounts of data between the operating system and the OSA adapter. The `qeth` device driver now also provides large send/receive and checksum offload operations for layer 2 mode.

6.4.5 IPv6 Support in snIPL

The tool for remote systems management for Linux, snIPL, now includes IPv6 support. This broadens the set of environments that snIPL supports and simplifies moving from IPv4 to IPv6.

6.4.6 Enhanced OSA Network to Receive All Frames Through a Network Interface

Enhancements in the OSA device driver enable setting network interfaces into promiscuous mode. The mode can provide outside connectivity for virtual servers by receiving all frames through a network interface.

In OpenStack environments, Open vSwitch is one of the connectivity options that use this feature.

6.5 Security

6.5.1 Support for DBRG in libica

The libica support for the generation of pseudo-random numbers for the "Deterministic Random Bit Generator" (DRBG) was enhanced to comply with updated security specifications (NIST SP 800-90A).

6.5.2 Monitoring CPACF Crypto Activity

This feature enables the monitoring of CPACF crypto activity in the Linux image, in the kernel, and in userspace. A configurable crypto-activity counter allows switching monitoring of CPACF crypto activity on or off for selected areas to verify and monitor specific needs in the crypto stack.

6.5.3 Support for Dynamic Traces in openCryptoki

Dynamic tracing in openCryptoki now allows starting and stopping tracing of all openCryptoki API calls and the related tokens while the application is running. This also allows using cryptography in the Java Security Architecture (JCA/JCE) which transparently falls back to software cryptography. Enhanced tracing can now identify whether cryptographic hardware is actually used.

6.5.4 CPACF MSA 4: Support for the GCM mechanism in openCryptoki

The openCryptoki ICA includes support for a new mechanism supported by CPACF MSA 4. GCM is a highly recommended mechanism for use with TLS 1.2.

6.5.5 Support for CCA Master Key Change for openCryptoki CCA Token

We now provide a tool to change master keys on the CCA co-processor without losing the encrypted data. This helps to stay compliant with enhanced industry regulations and company policies.

6.6 Reliability, Availability, Serviceability (RAS)

6.6.1 CUIR: Enhanced Scope Detection

The Linux support for CUIR (Control Unit Initiated Reconfiguration), which enables concurrent storage service with no or minimized down time, has been extended to include Linux running as a z/VM guest.

6.7 Performance

6.7.1 Extended CPU Performance Metrics in HYPFS for Linux z/VM guests

The HYPFS has been extended to provide the "diag 0C data" also for Linux z/VM guests that distinguish "management time" spent as part of CPU load.

6.7.2 IBM z13 Hardware Instructions in glibc

Support of the IBM z13 hardware instructions in glibc provides improved application performance.

6.7.3 Fake NUMA Support

Splitting the system memory into multiple NUMA nodes and distributing memory without using real topology information about the physical memory can improve performance. This is especially true for large systems. This feature is turned off by default but can be enabled for a system from the command line.

6.8 Miscellaneous

6.8.1 Enable Boot Parameter quiet for Better Visibility of Password Prompts

In the default configuration of SLES 12 SP2 for z Systems, the boot parameter `quiet` *is disabled, so the system console shows more useful log messages. This has the drawback that the increased amount of log messages can hide a password prompt, such as the prompt for decrypting devices at boot.*

To make the password prompt more visible among the system messages, add the boot parameter `quiet` when there are encrypted devices that need to be activated at system boot.

6.8.2 Installing From DVD/USB Drive of the HMC

You can now install from media in the DVD/USB drive of the Hardware Management Console (HMC).

To do so:

- Add `install=hmc:/` to the parm file or kernel options.
- Alternatively, in manual mode, in `linuxrc`, choose *Start Installation* > *Installation* > *Hardware Management Console*. The installation medium must be inserted in the HMC.

Important: Do not forget to configure the network in `linuxrc` before starting the installation. There is no way to pass boot parameters later and it is very likely that you will need network access. In `linuxrc`, go to `Start Installation` > `Network Setup`.

Important: Wait until the Linux system is booting before granting access to the DVD in the HMC. IPLing seems to disrupt the connection between the HMC and the LPAR in some way. If the first attempt to use it fails, you can grant the access and retry the option *HMC*.

Note: The installation medium will not be available in the installed system. If you need an installation repository there, register and use the online repository.

7 ARM 64-Bit (AArch64) Specific Information

Information in this section pertains to the version of SUSE Linux Enterprise Server 12 SP2 for the AArch64 architecture.

7.1 KVM on AArch64

KVM virtualization has been enabled and is supported on some system-on-chip platforms for mutually agreed-upon partner-specific use cases. It is only supported on partner certified hardware and firmware. Not all QEMU options and backends are available on AArch64. The same statement is applicable for other virtualization tools shipped on AArch64.

7.2 Toolchain Module Enabled in Default Installation

On AArch64, the Toolchain Module is now automatically pre-selected after registering SLES during installation. This makes the latest SLE compilers available on all installations.

However, in the AutoYaST installation you have to explicitly add the Toolchain module into the XML installation profile.

7.3 Boot Requirements for AppliedMicro X-Gene 1

The AppliedMicro X-C1 Server Development Platform (Mustang) ships with U-Boot based firmware. To install SUSE Linux Enterprise Server 12 SP2, the firmware needs to be updated to the UEFI based firmware version 3.06.15 or newer.

Other server systems, such as Gigabyte MP30, may also require a firmware update for an optimal experience. For details, contact your vendor.

7.4 ARM AArch64 System-on-Chip Platform Driver Enablement

For ARM based systems to boot SUSE Linux Enterprise Server, some chipset-specific drivers are needed.

The following System-on-Chip (SoC) platforms have been enabled for SP2:

- AMD Opteron A1100
- AppliedMicro X-Gene 1
- AppliedMicro X-Gene 2
- Cavium ThunderX
- NXP QorIQ LS2085A / LS2045A, LS2080A / LS2040A
- Xilinx UltraScale+ MPSoC

8 Driver Updates

8.1 Network Drivers

8.1.1 Support Status of Ethernet Drivers

Ethernet drivers have been added between kernel versions 3.12 (SLES 12 GA) and 4.4 (SLES 12 SP2).

The support status of Ethernet drivers has been updated for SLE 12 SP2 and below is the list of newly supported drivers.

- Agere Systems ET1310 (et131x)
- Qualcomm Atheros AR816x/AR817x PCI-E (alx)
- Broadcom BCM573xx (bnxt_en)
- JMicron JMC2x0 PCI-E (jme)
- QLogic FastLinQ 4xxxx (qede)
- SMC 83c170 EPIC series (epic100)
- SMSC LAN911x/LAN921x (smsc911x)
- SMSC LAN9420 PCI (smsc9420)
- STMMAC 10/100/1000 PCI (stmmac-pci)
- WIZnet W5100 (w5100)
- WIZnet W5300 (w5300)
- FUJITSU Extended Socket Network (fjes)
- SMSC95XX USB (smsc95xx)
- Xilinx LL TEMAC (ll_temac)
- APM X-Gene (xgene-enet)
- Cavium Thunder (nicpf, nicvf, thunder_bgx)

8.2 Other Drivers

8.2.1 Support for New Intel Processors

This Service Pack adds support for the following Intel processors:

- Intel® Xeon® Processor E3-1200/1500 v5 Product Family
- Intel® Xeon Phi™ Product Family x200

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 3.4, "Systems Management"*.

9.1 New Packages

9.1.1 The libcxl Userspace Library for CAPI Has Been Added

SLES now ships with the package `libcxl`. It provides the library of the same name that can be used for userspace CAPI.

The SLE SDK contains the corresponding development package, `libcxl-devel`.

9.1.2 targetcli-fb Has Been Added

In addition to the established tool `targetcli`, there is now also its enhanced version `targetcli-fb` available. New users are encouraged to deploy `targetcli-fb`.

9.1.3 Devilspie 2 Has Been Added

Desktop users often want the size and position of windows to remain the same, even across application restarts. Such functionality usually has to be implemented at the application level but not all applications do so.

In SUSE Linux Enterprise 12 SP2, Devilspie 2 (package `devilspie2`) has been added. Desvilspie 2 is a window matching utility that allow you to script actions on windows as they are created, such as maximizing windows or setting their size and position.

9.1.4 openldap2-ppolicy-check-password Has Been Added: OpenLDAP Password Strength Policy Enforcer

To allow evaluating and enforcing password strength in an OpenLDAP deployment, the package `openldap2-ppolicy-check-password` has been added. It is an OpenLDAP password policy plugin which evaluates and enforces strength in new user passwords, and denies weak passwords in password change operations. Configuration options of the plugin allow system administrators to adjust password strength requirements.

9.2 Updated Packages

9.2.1 Ceph Client Enablement Has Been Upgraded to Ceph Jewel

SUSE Enterprise Storage 3 and later versions expose additional functionality and performance to upgraded clients, such as the use of advanced RBD features and improved CephFS integration. While SUSE Enterprise Storage 3 is backwards-compatible with older clients, the full benefits are only available to newer clients.

As part of SUSE Linux Enterprise Server 12 Service Pack 2, the Ceph client code, as provided by `ceph-common` and the related library packages, has been upgraded to match the latest SUSE Enterprise Storage release.

This update also includes rebuilt versions of the KVM integration to take advantage of these.

9.2.2 Upgrade of libStorageMgmt to Version 1.3.2

libStorageMgmt allows programmatically managing storage hardware in a vendor-neutral way.

In SLES 12 SP2, libStorageMgmt was upgraded to version 1.3.2. This version fixes several bugs and adds the ability to more retrieve disk information, such as information on batteries and the list of local disks.

9.2.3 Glibc Has Been Upgraded to Version 2.22

glibc has been upgraded to meet demands in transactional memory handling and memory protection and to gain performance optimizations for modern platforms.

9.2.4 lsof Has Been Updated to Version 4.89

lsof has been updated from version 4.84 to 4.89. The changelog can be found in the file `/usr/share/doc/packages/lsof/DIST`.

9.2.5 Qt 5 Has Been Updated to 5.6.1

The Qt 5 libraries were updated to 5.6.1, a Qt 5.6 LTS based release. Qt 5.6.1 includes new features and security fixes for known vulnerabilities over Qt 5.5.1 (the version shipped in an upgrade to SP1).

This release includes many bug fixes and changes that improve performance and reduce memory consumption.

For security reasons, the MNG and JPEG2000 image format plugins are not shipped anymore, because the underlying MNG and JPEG2000 libraries have known security issues.

New features include:

- Better support for high-DPI screens
- Update of QtWebEngine which updates the included Chromium snapshot to version 45 and now uses many of the system libraries instead of bundled ones
- New Qt WebEngineCore module for new low-level APIs
- The Qt Location module is not fully supported.
- Improved compatibility with C++11 and the STL
- New QVersionNumber class
- Added support for HTTP redirection in QNetworkAccessManager
- Improved support for OpenGL ES 3
- Qt Multimedia got a new PlayList QML type and an audio role API for the media player
- Qt Canvas 3D now supports Qt Quick Items as textures and can directly render to the QML scenes foreground or background
- Qt 3D has received many improvements and new functionality
- Many other features and bugfixes

As part of this update, Qt Creator has been updated to 4.0.1 (from Qt Creator 3.5.1 shipped as an update to SP1).

New features of Qt Creator include:

- Clang static analyzer integration, extended QML profiler features, path editor of Qt Quick Designer and auto test integration (experimental) are now available
- The Clang code model is now automatically used if the (experimental) plugin is turned on
- Improved workflow for CMake-based projects
- The Analyze mode was merged with Debug mode, so that the new unified Debug mode includes the Debugger, Clang Static Analyzer, Memcheck, Callgrind and QML Profiler tools
- Many other features and bugfixes

9.2.6 RPM Ignores the BuildRoot Directive in Spec Files

In versions of RPM greater than 4.6.0, the behavior of the `BuildRoot` *directive was changed compared to prior versions. RPM now enforces using a build root for all packages and ignores the* `BuildRoot` *directive in spec files. By default,* `rpmbuild` *places the build root inside* `%{_topdir}`*. However, this can be changed through macro configuration.*

In the version of RPM shipped with SUSE Linux Enterprise 12 (and later), the `BuildRoot` directive of spec files is silently ignored. However, it is recommended to keep the `BuildRoot` directive in spec files for backward compatibility with earlier versions of SUSE Linux Enterprise (and RPM).

For more information, see the RPM 4.6.0 release notes at http://rpm.org/wiki/Releases/4.6.0.

9.2.7 OpenSSH Has Been Updated to Version 7.2

OpenSSH received numerous changes and improvements in the last years. To ease further maintenance, OpenSSH was upgraded to a more current release.

Note that the SSHv1 protocol is no longer supported.

9.2.8 Puppet Has Been Updated from 3.6.2 to 3.8.5

Puppet has been updated from 3.6.2 to 3.8.5. All releases between these two versions should only bring Puppet 3 backward-compatible features and bug and security fixes.

For more information, read the following release notes:

- Puppet 3.7 Release Notes: http://docs.puppetlabs.com/puppet/3.7/reference/release_notes.html
- Puppet 3.8 Release Notes: http://docs.puppetlabs.com/puppet/3.8/reference/release_notes.html

In particular, you should pay attention to the following upgrade notes and warnings:

- The new default value of the `environment_timeout` option is `0`: http://docs.puppetlabs.com/puppet/3.7/reference/release_notes.html#new-default-value-environmenttimeout--0.
- You can now set the parser setting per-environment in `environment.conf`: http://docs.puppetlabs.com/puppet/3.7/reference/release_notes.html#new-feature-parser-setting-in-environmentconf.
- Make sure the keepalive timeout is configured to be five or more seconds: http://docs.puppetlabs.com/puppet/3.7/reference/release_notes.html#upgrade-warning-rack-server-config.

9.2.9 Changes in Behavior Between coreutils 8.22 and 8.25

SLE 12 SP1 shipped with coreutils 8.22. SLE 12 SP2 ships with coreutils 8.25. This new release brings a number of changes in behavior:

`chroot`

- `base64`: `base64` no longer supports `--wrap` parameters in hexal or octal format. This improves support for decimals with leading zeros.
- `chroot`: Using `/` as the argument no longer implicitly changes the current directory to `/`. This allows changing user credentials for a single command only.
- `chroot`: `--userspec` will now unset supplemental groups associated with root and instead use the supplemental groups of the specified user.

- `cut`: Using `-d$'\n'` will again output lines identified in the `--fields` list (this behavior had been changed in version 8.21 and 8.22). Note that this functionality is non-portable and will result in the delayed output of lines.

- `date`: The option `--iso-8601` now uses the timezone format `+00:00` rather than `+0000`. This "extended" format is preferred by the ISO 8601 standard.

- `df`: `df` now prefers sources towards the root of a device when eliding duplicate `bind` -mounted entries.

- `df`: `df` no longer suppresses separate exports of the same remote device, as these are generally explicitly mounted. The `--total` option does still suppress duplicate remote file systems.

- `join`, `sort`, `uniq`: When called with `--zero-terminated`, these commands now treat `\n` as a field delimiter.

- `ls`: If neither of the environment variables `LS_COLORS` and `COLORTERM` is set and the environment variable `TERM` is empty or unknown, `ls` will now not output colors even with `--colors=always`.

- `ls`: `ls` now quotes file names unambiguously and appropriate for use in a shell, when outputting to a terminal.

- `mv`: `mv` no longer supports moving a file to a hard link. If you try, it issues an error. The prior implementation was susceptible to races in the presence of multiple `mv` instances which could result in both hard links being deleted. Also, on case-insensitive file systems like HFS, `mv` would remove a hardlinked file if called like `mv file File`.

- `numfmt`: The options `--from-unit` and `--to-unit` now interpret suffixes as SI units, and IEC (power of 2) units are now specified by appending `i`.

- `tee`: If there are no more writable outputs, `tee` will exit early.

- `tee`: `tee` does not treat the file operand `-` as meaning standard output any longer. This allows for better POSIX conformance.

- `timeout`: The option `--foreground` no longer sends `SIGCONT` to the monitored process, as this was seen to cause intermittent issues with GDB for example.

9.2.10 openSSL Has Been Updated to Version 1.0.2

openSSL has been updated from version 1.0.1 to 1.0.2 which is a compatible minor version update. This will help future maintenance, and also brings many bug fixes.

The update to openSSL 1.0.2 should be transparent to existing programs.

However, there were some functional changes were done: SSL 2 support is now fully disabled and certain weak ciphers are no longer built in.

9.3 Removed and Deprecated Functionality

9.3.1 Perl Bindings for Cyrus Have Been Removed

With SLE 12 SP2, the packages `perl-Cyrus-IMAP` and `perl-Cyrus-SIEVE-managesieve` have been removed from the media.

9.3.2 librpcsecgss3 Has Been Removed

`librpcsecgss` (packages: `librpcsecgss3`, `librpcsecgss-devel`) has been removed. With the release of `libtirpc`, the development of libsecgss stopped and it fell out of use. We recommend using `libtirpc` instead.

9.3.3 libusnic_verbs-rdmav2 and libusnic_verbs-rdmav2-pingpong Are Now Obsolete

Functionality previously shipped in the packages `libusnic_verbs-rdmav2` and `libusnic_verbs-rdmav2-pingpong` has been integrated into `libibverbs`.

9.3.4 Packages Removed with SUSE Linux Enterprise Server 12

The packages listed below were removed with the major release of SUSE Linux Enterprise Server 12.

9.3.4.1 Nagios Server Now Part of a SUSE Manager Subscription

Support for Icinga (a successor of Nagios) will not be part of the SUSE Linux Enterprise Server 12 subscription.

Fully supported Icinga packages for SUSE Linux Enterprise Server 12 will be available as part of a SUSE Manager subscription. In the SUSE Manager context we will be able to deliver better integration into the monitoring frameworks.

More frequent updates on the monitoring server parts than in the past are planned.

9.3.5 Packages Removed with SUSE Linux Enterprise Server 12 SP1

The packages listed below were removed with the release of SUSE Linux Enterprise Server 12 SP2.

9.3.5.1 wpa_supplicant Replaces xsupplicant

In SUSE Linux Enterprise 12 SP1 and 12 SP2, `xsupplicant` was removed entirely.

For pre-authentication of systems via network (including RADIUS) and specifically wireless connections, install the `wpa_supplicant` package. `wpa_supplicant` now replaces `xsupplicant`. `wpa_supplicant` provides better stability, security and a broader range of authentication options.

9.4 Changes in Packaging and Delivery

9.4.1 Change of OpenMPI Behavor for Plugin Developers

To be compliant with the upstream version of OpenMPI, the source configuration option `--with-devel-header` has been removed. This only affects developers of OpenMPI plugins outside of the source tree.

Developers of plugins outside of the source tree need to recompile the source with the option `--with-devel-header` added. All other users are not affected.

9.5 SDK

9.5.1 Byebug Has Been Added

Byebug is a simple-to-use, feature-rich Ruby 2 debugger that is also used to debug YaST. It uses the TracePoint API and the Debug Inspector API. For speed, it is implemented as a C extension.

It allows you to see what is going on inside a Ruby program while it executes and offers traditional debugging features such as stepping, breaking, evaluating, and tracking.

10 Technical Information

This section contains information about system limits, a number of technical changes and enhancements for the experienced user.

When talking about CPUs, we use the following terminology:

CPU Socket
: The visible physical entity, as it is typically mounted to a motherboard or an equivalent.

CPU Core
: The (usually not visible) physical entity as reported by the CPU vendor.
On IBM z Systems, 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 Virtualization: Network Devices Supported

SLES 12 supports the following virtualized network drivers:

- Full virtualization: Intel e1000
- Full virtualization: Realtek 8139
- Paravirtualized: QEMU Virtualized NIC Card (virtio, KVM only)

10.2 Virtualization: Devices Supported for Booting

SLE12 support VM guest to boot from:

- Parallel ATA (PATA/IDE)
- Advanced Host Controller Interface (AHCI)
- Floppy Disk Drive (FDD)
- virtio-blk
- virtio-scsi
- Preboot eXecution Environment (PXE) ROMs (for supported Network Interface Cards)

Boot from `USB` and `PCI pass-through` devices are not supported.

10.3 Virtualization: Supported Disks Formats and Protocols

The following disk formats support read-write access (RW):

- `raw`
- `qed` (KVM only)
- `qcow` (Xen only)
- `qcow2`

The following disk formats support read-only access (RO):

- `vmdk`
- `vpc`
- `vhd` / `vhdx`

The following protocols can be used for read-only access (RO) to images:

- `http`, `https`
- `ftp`, `ftps`, `tftp`

When using Xen, the `qed` format will not be displayed as a selectable storage in `virt-manager`.

10.4 Kernel Limits

http://www.suse.com/products/server/technical-information/#Kernel

This table summarizes the various limits which exist in our recent kernels and utilities (if related) for SUSE Linux Enterprise Server 12 SP2.

SLES 12 SP2 (Linux 4.4)	**AMD64/Intel 64 (x86_64)**	**IBM z Systems (s390x)**	**POWER (ppc64le)**	**AArch64 (ARMv8)**
CPU bits	64	64	64	64
Maximum number of logical CPUs	8192	256	2048	128
Maximum amount of RAM (theoretical/certified)	> 1 PiB/64 TiB	4 TiB/256 GiB	1 PiB/64 TiB	256 TiB/n.a.
Maximum amount of user space/kernel space	128 TiB/128 TiB	n.a.	2 TiB/2 EiB	256 TiB/128 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).			

SLES 12 SP2 (Linux 4.4)	AMD64/Intel 64 (x86_64)	IBM z Systems (s390x)	POWER (ppc64le)	AArch64 (ARMv8)
Maximum size per block device	Up to 8 EiB on all 64-bit architectures			
FD_SETSIZE	1024			

10.5 KVM Limits

SLES 12 SP2 Virtual Machine (VM)	Limits
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	240
Maximum Memory per VM	4 TiB
Maximum Virtual Block Devices per VM	20 virtio-blk, 4 IDE
Maximum number of Network Cards per VM	8

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

10.6 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 12 SP2 Virtual Machine (VM)	Limits
Maximum number of VMs per host	64
Maximum number of virtual CPUs per VM	64

SLES 12 SP2 Virtual Machine (VM)	**Limits**
Maximum amount of memory per VM	16 GiB x86_32, 511 GiB x86_64
Maximum virtual block devices per VM	100 PV, 100 FV with PV drivers, 4 FV (emulated IDE)
Maximum virtual network devices per VM	8

SLES 12 SP2 Virtual Host Server (VHS)	**Limits**
Maximum number of physical CPUs	256
Maximum number of virtual CPUs	256
Maximum amount of physical memory	5 TiB
Maximum amount of Dom0 physical memory	500 GiB
Maximum number of block devices	12,000 SCSI logical units
Maximum number of iSCSI devices	128
Maximum number of network cards	8
Maximum number of VMs per CPU core	8
Maximum number of VMs per VHS	64
Maximum number of virtual network cards	64 across all VMs in the system

In Xen 4.4, the hypervisor bundled with SUSE Linux Enterprise Server 12 SP2, Dom0 is able to see and handle a maximum of 512 logical CPUs. However, the hypervisor itself, can access up to logical 256 logical CPUs and schedule those for the VMs.

- **PV:** Paravirtualization
- **FV:** Full virtualization

For more information about acronyms, see the virtualization documentation provided at https://www.suse.com/documentation/sles-12/.

10.7 File Systems

https://www.suse.com/products/server/technical-information/#FileSystem

10.7.1 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	ReiserFS **	OCFS 2 ***
Data/metadata journaling	N/A *	– / +		– / +	– / +
Journal internal/external	N/A *	+ / +	+ / –		
Offline extend/shrink	+ / +	– / –	+ / +	+ / –	
Online extend/shrink	+ / +	+ / –	+ / –	+ / –	+ / –
Inode allocation map	B-tree	B+-tree	table	u. B*-tree	table
Sparse files	+				
Tail packing	+	–	+	–	
Defrag	+	–			
ExtAttr/ACLs	+ / +				

Feature	Btrfs	XFS	Ext4	ReiserFS **	OCFS 2 ***
Quotas	+				
Dump/restore	–	+	–		
Block size default	4 KiB				
Maximum file system size	16 EiB	8 EiB	1 EiB	16 TiB	4 PiB
Maximum file size	16 EiB	8 EiB	1 EiB	1 EiB	4 PiB
Support in products	SLE	SLE	SLE	SLE	SLE HA

- * Btrfs is a copy-on-write file system. Rather than 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 new changes are not “committed”. Due to the nature of the file system, quotas are implemented based on subvolumes (`qgroups`).
  The block size default varies with different host architectures. 64 KiB is used on POWER, 4 KiB on most other systems. The actual size used can be checked with the command **`getconf`** *`PAGE_SIZE`*.
- ** ReiserFS is supported for existing file systems. The creation of new ReiserFS file systems is discouraged.
- *** OCFS2 is fully supported as part of the SUSE Linux Enterprise High Availability Extension.

The maximum file size above can be larger than the file system's actual size due to usage of sparse blocks. Note that unless a file system comes with large file support (LFS), the maximum file size on a 32-bit system is 2 GB (2^31 bytes). Currently all of our standard file systems (including Ext3 and ReiserFS) have LFS, which gives a maximum file size of 2^63 bytes in theory. The numbers in the above tables assume that the file systems are using 4 KiB block size. When using different block sizes, the results are different, but 4 KiB reflects the most common standard.

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 12 SP2 delivers integration with Windows 7 Active Directory domains. In addition, we provide the clustered version of Samba as part of SUSE Linux Enterprise High Availability Extension 12 SP2.

10.7.2 Supported Btrfs Features

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

+ supported
– unsupported

Feature	SLES 11 SP4	SLES 12 GA	SLES 12 SP1	SLES 12 SP2
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 GA	SLES 12 SP1	SLES 12 SP2
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	–	–	–	+

11 Legal Notices

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

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

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classifications to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in U.S. export laws. You

agree to not use deliverables for prohibited nuclear, missile, or chemical/biological weaponry end uses. Refer to http://www.suse.com/company/legal/ for more information on exporting SUSE software. SUSE assumes no responsibility for your failure to obtain any necessary export approvals.

SUSE has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at http://www.suse.com/company/legal/ and one or more additional patents or pending patent applications in the U.S. and other countries.

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

12 Colophon

Thanks for using SUSE Linux Enterprise Server in your business.

The SUSE Linux Enterprise Server Team.