

SUSE Linux Enterprise for High-Performance Computing 15 SP4

Release Notes

SUSE Linux Enterprise for High-Performance Computing is a highly-scalable, high-performance open-source operating system designed to utilize the power of parallel computing. This document provides an overview of high-level general features, capabilities, and limitations of SUSE Linux Enterprise for High-Performance Computing 15 SP4 and important product updates.

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

Publication Date: 2022-11-30, Version: 15.400000000.20221130

Contents

- 1 About the release notes 3
- 2 SUSE Linux Enterprise for High-Performance Computing 3
- 3 Modules, extensions, and related products 7
- 4 Technology previews 8
- 5 Modules 10
- 6 Changes affecting all architectures 11
- 7 Removed and deprecated features and packages 22
- 8 Obtaining source code 23

- 9 Legal notices **23**
- A Changelog for 15 SP4 **25**

1 About the release notes

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

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

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

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

2 SUSE Linux Enterprise for High-Performance Computing

SUSE Linux Enterprise for High-Performance Computing is a highly scalable, high performance open-source operating system designed to utilize the power of parallel computing for modeling, simulation and advanced analytics workloads.

SUSE Linux Enterprise for High-Performance Computing 15 SP4 provides tools and libraries related to High Performance Computing. This includes:

- Workload manager
- Remote and parallel shells
- Performance monitoring and measuring tools
- Serial console monitoring tool
- Cluster power management tool
- A tool for discovering the machine hardware topology
- System monitoring
- A tool for monitoring memory errors
- A tool for determining the CPU model and its capabilities (x86-64 only)

- User-extensible heap manager capable of distinguishing between different kinds of memory (x86-64 only)
- Serial and parallel computational libraries providing the common standards BLAS, LAPACK, ...
- Various MPI implementations
- Serial and parallel libraries for the HDF5 file format

2.1 Hardware Platform Support

SUSE Linux Enterprise for High-Performance Computing 15 SP4 is available for the Intel 64/AMD64 (x86-64) and AArch64 platforms.

2.2 Important Sections of This Document

If you are upgrading from a previous SUSE Linux Enterprise for High-Performance Computing release, you should review at least the following sections:

- *Section 2.4, “Support statement for SUSE Linux Enterprise for High-Performance Computing”*

2.3 Support and life cycle

SUSE Linux Enterprise for High-Performance Computing 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 for High-Performance Computing 15 has a 13-year life cycle, with 10 years of General Support and 3 years of Extended Support. The current version (SP4) will be fully maintained and supported until 6 months after the release of SUSE Linux Enterprise for High-Performance Computing 15 SP5.

Any release package is fully maintained and supported until the availability of the next release. Extended Service Pack Overlay Support (ESPOS) and Long Term Service Pack Support (LTSS) are also available for this product. If you need additional time to design, validate and test your upgrade plans, Long Term Service Pack Support (LTSS) can extend the support you get by an additional 12 to 36 months in 12-month increments, providing a total of 3 to 5 years of support on any given Service Pack.

For more information, see:

- The support policy at <https://www.suse.com/support/policy.html> 
- Long Term Service Pack Support page at <https://www.suse.com/support/programs/long-term-service-pack-support.html> 

2.4 Support statement for SUSE Linux Enterprise for High-Performance Computing

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

The following definitions apply:

L1

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

L2

Problem isolation, which means technical support designed to analyze data, reproduce customer problems, isolate problem area and provide a resolution for problems not resolved by Level 1 or prepare for Level 3.

L3

Problem resolution, which means technical support designed to resolve problems by engaging engineering to resolve product defects which have been identified by Level 2 Support.

For contracted customers and partners, SUSE Linux Enterprise for High-Performance Computing is delivered with L3 support for all packages, except for the following:

- Technology Previews, see [Section 4, "Technology previews"](#)
- Sound, graphics, fonts and artwork
- Packages that require an additional customer contract, see [Section 2.4.1, "Software requiring specific contracts"](#)

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

2.4.1 Software requiring specific contracts

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

2.4.2 Software under GNU AGPL

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

- Ghostscript (including subpackages)

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

- MySpell dictionaries and LightProof
- ArgyllCMS

2.5 Documentation and other information

2.5.1 Available on the product media

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

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

- Check the [ChangeLog](#) file in the top level of the installation medium for a chronological log of all changes made to the updated packages.
- Find more information in the [docu](#) directory of the installation medium of SUSE Linux Enterprise for High-Performance Computing 15 SP4. This directory includes PDF versions of the SUSE Linux Enterprise for High-Performance Computing 15 SP4 Installation Quick Start Guide.

2.5.2 Online documentation

- For the most up-to-date version of the documentation for SUSE Linux Enterprise for High-Performance Computing 15 SP4, see <https://documentation.suse.com/sle-hpc/15-SP4>.
- Find a collection of White Papers in the SUSE Linux Enterprise for High-Performance Computing Resource Library at <https://www.suse.com/products/server#resources>.

3 Modules, extensions, and related products

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

3.1 Modules in the SLE 15 SP4 product line

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

The modules available within the product line based on SUSE Linux Enterprise 15 SP4 at the release of SUSE Linux Enterprise for High-Performance Computing 15 SP4 are listed in the *Modules and Extensions Quick Start* at <https://documentation.suse.com/sles/15-SP3/html/SLES-all/article-modules.html>.

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

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

3.2 Available extensions

The following extension is not covered by SUSE support agreements, available at no additional cost and without an extra registration key: SUSE Package Hub, see <https://packagehub.suse.com/> .

3.3 Related products

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

- SUSE Linux Enterprise Server: <https://www.suse.com/products/server> 
- SUSE Linux Enterprise JeOS: <https://www.suse.com/products/server/jeos> 
- SUSE Linux Enterprise Desktop: <https://www.suse.com/products/desktop> 
- SUSE Linux Enterprise Server for SAP Applications: <https://www.suse.com/products/sles-for-sap> 
- SUSE Linux Enterprise Real Time: <https://www.suse.com/products/realtime> 
- SUSE Manager: <https://www.suse.com/products/suse-manager> 

4 Technology previews

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

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

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

4.1 64K page size kernel flavor has been added

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

As a technology preview, SUSE Linux Enterprise for High-Performance Computing for Arm 15 SP4 adds a kernel flavor `64kb`, offering a page size of 64 KiB and physical/virtual address size of 52 bits. Same as the `default` kernel flavor, it does not use preemption.

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



Important: Swap needs to be re-initialized

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

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



Important: Btrfs file system uses page size as block size

It is currently not possible to use Btrfs file systems across page sizes. Block sizes below page size are not yet supported and block sizes above page size might never be supported. During installation, change the default partitioning proposal and choose another file system, such as Ext4 or XFS, to allow rebooting from the default 4K page size kernel of the Installer into `kernel-64kb` and back.

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



Warning: RAID 5 uses page size as stripe size

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

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

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



Note: Cross-architecture compatibility considerations

The SUSE Linux Enterprise for High-Performance Computing 15 SP4 kernels on x86-64 use 4K page size.

The SUSE Linux Enterprise for High-Performance Computing for POWER 15 SP4 kernel uses 64K page size.

5 Modules

5.1 HPC module

The HPC module contains HPC specific packages. These include the workload manager Slurm, the node deployment tool `clustduct`, `munge` for user authentication, the remote shell `mrsh`, the parallel shell `pdsh`, as well as numerous HPC libraries and frameworks.

This module is available with the SUSE Linux Enterprise for High-Performance Computing only. It is selected by default during the installation. It can be added or removed using the YaST UI or the `SUSEConnect` CLI tool. Refer to the system administration guide for further details.

5.2 NVIDIA Compute Module

The NVIDIA Compute Module provides the NVIDIA CUDA repository for SUSE Linux Enterprise 15. Note that any software within this repository is under a 3rd party EULA. For more information check <https://docs.nvidia.com/cuda/eula/index.html>.

This module is not selected for addition by default when installing SUSE Linux Enterprise for High-Performance Computing. It may be selected manually during installation from the *Extension and Modules* screen. You may also select it on an installed system using YaST. To do so, run from a shell as root `yast registration`, select: `Select Extensions` and search for `NVIDIA Compute Module` and press `Next`.

Important

Do not attempt to add this module with the SUSEConnect CLI tool. This tool is not yet capable of handling 3rd party repositories.

Once you have selected this module you will be asked to confirm the 3rd party license and verify the repository signing key.

6 Changes affecting all architectures

Information in this section applies to all architectures supported by SUSE Linux Enterprise for High-Performance Computing 15 SP4.

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

SUSE is committed to helping provide better insights into the consumption of SUSE subscriptions regardless of where they are running or how they are managed; physical or virtual, on-prem or in the cloud, connected to SCC or Repository Mirroring Tool (RMT), or managed by SUSE Manager. To help you identify or filter out systems in SCC that are no longer running or decommissioned, SUSEConnect now features a daily “ping”, which will update system information automatically. For more details see the documentation at <https://documentation.suse.com/subscription/suseconnect/single-html/SLE-suseconnect-visibility/>.

6.2 Automatically opened ports

Installing the following packages automatically opens the following ports:

- dolly - TCP ports 9997 and 9998
- slurm - TCP ports 6817, 6818, and 6819

Important

These release notes only document changes in SUSE Linux Enterprise for High-Performance Computing compared to the immediate previous service pack of SUSE Linux Enterprise for High-Performance Computing. The full changes and fixes can be found on the respective web site of the packages.

6.3 dolly

dolly has been updated to version 0.63.6. It includes some fixes for hostname resolution, a better documentation and now provides a default configuration for firewall.

6.4 memkind

memkind has been updated to version 1.12.0. The full list of changes is available at <http://memkind.github.io/memkind/> .

6.5 openblas

openblas has been updated to version 0.3.17. It contains performance regression fixes and optimization. For more information see <https://github.com/xianyi/OpenBLAS/releases/tag/v0.3.17> .

6.6 spack

spack has been updated to version 0.17.1. It now includes support to build singularity containers from <https://registry.suse.com/> .

6.7 mpich

mpich has been updated to version 3.4.2. For more information see <https://www.mpich.org/2021/05/28/mpich-3-4-2-released/> .

6.8 Slurm

6.8.1 Important Notes for Upgrading Slurm Releases:

If using the slurmdbd (Slurm DataBase Daemon) you must update this first. If using a backup DBD you must start the primary first to do any database conversion, the backup will not start until this has happened.

6.8.2 Slurm version 22.05

An update to Slurm version 22.05 is available.

6.8.2.1 Important notes for upgrading to version 22.05

Slurmdbd version 22.05 will work Slurm daemons of version 20.11. You will not need to update all clusters at the same time, but it is very important to update slurmdbd first and having it running before updating any other clusters making use of it.

Slurm can be upgraded from version 20.11 to version 22.05 without loss of jobs or other state information. Upgrading directly from an earlier version of Slurm will result in loss of state information.

For more information and a recommended upgrade procedure, see the section "Upgrading Slurm" in the chapter "Slurm — utility for HPC workload management" of the in the SLE HPC 15 "Administration Guide".

All SPANK plugins must be recompiled when upgrading from any Slurm version prior to 22.05. If you are using the Slurm plugin for pdsh you must make sure, pdsh_slurm_22_05 is installed together with slurm_22_05.

6.8.2.2 Highlights of version 20.11

- The template slurmrestd.service unit file now defaults to listen on both the Unix socket and the slurmrestd port.
- The template slurmrestd.service unit file now defaults to enable auth/jwt and the munge unit is no longer a dependency by default.
- Add extra “EnvironmentFile = -/etc/default/\$service” setting to service files.

- Allow jobs to pack onto nodes already rebooting with the desired features.
- Reset job start time after nodes are rebooted, previously only done for cloud/power save boots.
- Node features (if any) are passed to `RebootProgram` if run from `slurmctld`.
- Fail `srun` when using invalid `--cpu-bind` options (e.g. `--cpu-bind=map_cpu:99` when only 10 CPUs are allocated).
- Storing batch scripts and env vars are now in indexed tables using substantially less disk space. Those storing scripts in 21.08 will all be moved and indexed automatically.
- Run `MailProg` through `slurmscriptd` instead of directly `fork+exec()`'ing from `slurmctld`.
- Add `acct_gather_interconnect/sysfs` plugin.
- Future and Cloud nodes are treated as "Planned Down" in usage reports.
- Add new shard plugin for sharing GPUs but not with mps.
- Add support for Lenovo SD650 V2 in `acct_gather_energy/xcc` plugin.
- Remove `cgroup_allowed_devices_file.conf`, since the default policy in modern kernels is to whitelist by default. Denying specific devices must be done through `gres.conf`.
- Node state flags (`DRAIN`, `FAILED`, `POWERING UP`, etc.) will be cleared now if node state is updated to `FUTURE`.
- `srun` will no longer read in `SLURM_CPUS_PER_TASK`. This means you will implicitly have to specify `--cpus-per-task` on your `srun` calls, or set the new `SRUN_CPUS_PER_TASK` environment variable to accomplish the same thing.
- Remove `connect_timeout` and `timeout` options from `JobCompParams` as there's no longer a connectivity check happening in the `jobcomp/elasticsearch` plugin when setting the location off of `JobCompLoc`.
- Add support for hourly reoccurring reservations.
- Allow nodes to be dynamically added and removed from the system. Configure `MaxNodeCount` to accomodate nodes created with dynamic node registrations (`slurmd -Z --conf=""`) and `scontrol`.
- Added support for Cgroup Version 2.

- sacct - allocations made by srun will now always display the allocation and step(s). Previously, the allocation and step were combined when possible.
- cons_tres - change definition of the "least loaded node" (LLN) to the node with the greatest ratio of available CPUs to total CPUs.
- Add support to ship Include configuration files with configless.
- Provide a detailed reason in the job log as to why it has been terminated when hitting a resource limit.
- Pass and use alias_list through credential instead of environment variable.
- Add ability to get host addresses from nss_slurm.
- Enable reverse fanout for cloud+alias_list jobs.
- Add support to delete/update nodes by specifying nodesets or the 'ALL' keyword alongside the delete/update node message nodelist expression (i.e. scontrol delete/update NodeName=ALL or scontrol delete/update NodeName=ns1,nodes[1-3]).
- Expanded the set of environment variables accessible through Prolog/Epilog and PrologSlurmctld/EpilogSlurmctld to include SLURM_JOB_COMMENT, SLURM_JOB_STDERR, SLURM_JOB_STDIN, SLURM_JOB_STDOUT, SLURM_JOB_PARTITION, SLURM_JOB_ACCOUNT, SLURM_JOB_RESERVATION, SLURM_JOB_CONSTRAINTS, SLURM_JOB_NUM_HOSTS, SLURM_JOB_CPUS_PER_NODE, SLURM_JOB_NTASKS, and SLURM_JOB_RESTART_COUNT.
- Attempt to requeue jobs terminated by slurm.conf changes (node vanish, node socket/core change, etc). Processes may still be running on excised nodes. Admin should take precautions when removing nodes that have jobs on running on them.
- Add switch/hpe_slingshot plugin.
- Add new SchedulerParameters option bf_licenses to track licenses as within the backfill scheduler.

6.8.2.3 Configuration File changes (for details, see the appropriate man page)

- AcctGatherEnergyType rsmi is now gpu.
- TaskAffinity parameter was removed from cgroup.conf.

- Fatal if the mutually-exclusive `JobAcctGatherParams` options of `UsePss` and `NoShared` are both defined.
- `KeepAliveTime` has been moved into `CommunicationParameters`. The standalone option will be removed in a future version.
- `preempt/qos` - add support for WITHIN mode to allow for preemption between jobs within the same QOS.
- Fatal error if `CgroupReleaseAgentDir` is configured in `cgroup.conf`. The option has long been obsolete.
- Fatal if more than one burst buffer plugin is configured.
- Added `keepaliveinterval` and `keepaliveprobes` to `CommunicationParameters`.
- Added new `max_token_lifespan=<seconds>` to `AuthAltParameters` to allow sites to restrict the lifespan of any requested ticket by an unprivileged user.
- Disallow `slurm.conf` node configurations with `NodeName=ALL`.

6.8.2.4 Command Changes (for details, see the appropriate man page)

- Remove support for (non-functional) `--cpu-bind=boards`.
- Added `--prefer` option at job submission to allow for 'soft' constraints.
- Add `condflags=open` to `sacctmgr show events` to return open/currently down events.
- `sacct -f` flag implies `-c` flag.
- `srun --overlap` now allows the step to share all resources (CPUs, memory, and GRES), where previously `--overlap` only allowed the step to share CPUs with other steps.

6.8.2.5 API Changes

- `openapi/v0.0.35` - Plugin has been removed.
- `burst_buffer` plugins - `err_msg` added to `bb_p_job_validate()`.
- `openapi` - added flags to `slurm_openapi_p_get_specification()`. Existing plugins only need to update their prototype for the function as manipulating the flags pointer is optional.

- openapi - Added OAS_FLAG_MANGLE_OPID to allow plugins to request that the operationId of path methods be mangled with the full path to ensure uniqueness.
- openapi/[db]v0.0.36 - Plugins have been marked as deprecated and will be removed in the next major release.
- switch plugins - add switch_g_job_complete() function.

6.8.3 Highlights of Slurm version 21.08

6.8.3.1 Highlights

- Removed gres/mic plugin used to support Xeon Phi coprocessors.
- Add LimitFactor to the QOS. A float that is factored into an associations GrpTRES limits. For example, if the LimitFactor is 2, then an association with a GrpTRES of 30 CPUs, would be allowed to allocate 60 CPUs when running under this QOS.
- A job's next_step_id counter now resets to 0 after being requeued. Previously, the step id's would continue from the job's last run.
- API change: Removed slurm_kill_job_msg and modified the function signature for slurm_kill_job2. slurm_kill_job2 should be used instead of slurm_kill_job_msg.
- AccountingStoreFlags=job_script allows you to store the job's batch script.
- AccountingStoreFlags=job_env allows you to store the job's env vars.
- Removed sched/hold plugin.
- cli_filter/lua, jobcomp/lua, job_submit/lua now load their scripts from the same directory as the slurm.conf file (and thus now will respect changes to the SLURM_CONF environment variable).
- SPANK - call slurm_spank_init if defined without slurm_spank_slurmd_exit in slurmd context.
- Add new PLANNED state to a node to represent when the backfill scheduler has it planned to be used in the future instead of showing as IDLE. sreport also has changed it's cluster utilization report column name from 'Reserved' to 'Planned' to match this nomenclature.

- Put node into INVAL state upon registering with an invalid node configuration. Node must register with a valid configuration to continue.
- Remove SLURM_DIST_LLLP environment variable in favor of just SLURM_DISTRIBUTION.
- Make --cpu-bind=threads default for --threads-per-core — can be overridden by the CLI or an environment variable.
- slurmd - allow multiple comma-separated controllers to be specified in configless mode with --conf-server
- Manually powering down of nodes with scontrol now ignores SuspendExc<Nodes|Parts>.
- Distinguish queued reboot requests (REBOOT@) from issued reboots (REBOOT^).
- auth/jwt - add support for RS256 tokens. Also permit the username in the 'username' field in addition to the 'sun' (Slurm UserName) field.
- service files - change dependency to network-online rather than just network to ensure DNS and other services are available.
- Add "Extra" field to node to store extra information other than a comment.
- Add ResumeTimeout, SuspendTimeout and SuspendTime to Partitions.
- The memory.force_empty parameters is no longer set by jobacct_gather/cgroup when deleting the cgroup`. This previously caused a significant delay (~2s) when terminating a job, and is not believed to have provided any perceivable benefit. However, this may lead to slightly higher reported kernel mem page cache usage since the kernel cgroup memory is no longer freed immediately.
- TaskPluginParam=verbose is now treated as a default. Previously it would be applied regardless of the job specifying a --cpu-bind.
- Add node_reg_mem_percent SlurmctldParameter to define percentage of memory nodes are allowed to register with.
- Define and separate node power state transitions. Previously a powering down node was in both states, POWERING_OFF and POWERED_OFF. These are now separated. e.g. IDLE+POWERED_OFF (IDLE~) → IDLE+POWERING_UP (IDLE#) - Manual power up or allocation → IDLE → IDLE+POWER_DOWN (IDLE!) - Node waiting for power down → IDLE+POWERING_DOWN (IDLE%) - Node powering down → IDLE+POWERED_OFF (IDLE~) - Powered off

- Some node state flag names have changed. These would be noticeable for example if using a state flag to filter nodes with `sinfo`. e.g. `POWER_UP` → `POWERING_UP` `POWER_DOWN` → `POWERED_DOWN` `POWER_DOWN` now represents a node pending power down
- Create a new process called `slurmscriptd` which runs `PrologSlurmctld` and `EpilogSlurmctld`. This avoids `fork()` calls from `slurmctld`, and can avoid performance issues if the `slurmctld` has a large memory footprint.
- Pass JSON of job to node mappings to `ResumeProgram`.
- QOS accrue limits only apply to the job QOS, not partition QOS.
- Any return code from SPANK plugin or SPANK function that is not `SLURM_SUCCESS` (zero) will be considered to be an error. Previously, only negative return codes were considered an error.
- Add support for automatically detecting and broadcasting executable shared object dependencies for `sbcst` and `srun --bcst`.
- All SPANK error codes now start at 3000. Where previously SPANK would give a return code of 1, it will now return 3000. This change will break ABI compatibility with SPANK plugins compiled against older version of Slurm.
- SPANK plugins are now required to match the current Slurm release, and must be recompiled for each new Slurm major release. (They do not need to be recompiled when upgrading between maintenance releases.)
- `SLURM_NODE_ALIASES` now has brackets around the node's address to be able to distinguish IPv6 addresses. e.g. `<node_name>:[<node_addr>]:<node_hostname>`
- The `job_container/tmpfs` plugin now requires `PrologFlags=contain` to be set in `slurm.conf`.
- Limit `max_script_size` to 512 MB.

6.8.3.2 Configuration File Changes (for details, see the appropriate man page)

- Errors detected in the parser handlers due to invalid configurations are now propagated and can lead to fatal (and thus exit) the calling process.
- Enforce a valid configuration for `AccountingStorageEnforce` in `slurm.conf`. If the configuration is invalid, then an error message will be printed and the command or daemon (including `slurmctld`) will not run.

- Removed AccountingStoreJobComment option. Please update your config to use AccountingStoreFlags=job_comment instead.
- Removed DefaultStorage{Host,Loc,Pass,Port,Type,User} options.
- Removed CacheGroups, CheckpointType, JobCheckpointDir, MemLimitEnforce, SchedulerPort, SchedulerRootFilter options.
- Added Script to DebugFlags for debugging slurmscriptd (the process that runs slurmctld scripts such as PrologSlurmctld and EpilogSlurmctld).
- Rename SbcastParameters to BcastParameters.
- systemd service files - add new “-s” option to each daemon which will change the working directory even with the -D option. (Ensures any core files are placed in an accessible location, rather than /.)
- Added BcastParameters=send_libs and BcastExclude options.
- Remove the (incomplete) burst_buffer/generic plugin.
- Make SelectTypeParameters=CR_Core_Memory default for cons_tres and cons_res.
- Remove support for TaskAffinity=yes in cgroup.conf. Adding task/affinity to TaskPlugins in slurm.conf is strongly recommended instead.

6.8.3.3 Command Changes (for details, see the appropriate man page)

- Changed the --format handling for negative field widths (left justified) to apply to the column headers as well as the printed fields.
- Invalidate multiple partition requests when using partition based associations.
- scrontab - create the temporary file under the TMPDIR environment variable (if set), otherwise continue to use TmpFS as configured in slurm.conf.
- sbcast / srun --bcast - removed support for zlib compression. lz4 is vastly superior in performance, and (counter-intuitively) zlib could provide worse performance than no compression at all on many systems.
- sacctmgr - changed column headings to ParentID and ParentName instead of Par ID and "Par Name` respectively.

- `SALLOC_THREADS_PER_CORE` and `SBATCH_THREADS_PER_CORE` have been added as input environment variables for `salloc` and `sbatch`, respectively. They do the same thing as `--threads-per-core`.
- Don't display node's comment with `scontrol show nodes` unless set.
- Added `SLURM_GPUS_ON_NODE` environment variable within each job/step.
- `sreport` - change to sorting `TopUsage` by the `--tres` option.
- `slurmrestd` - do not run allow operation as `SlurmUser`/root by default.
- `scontrol show node` now shows `State` as `base_state+flags` instead of shortened state with flags appended. eg. `IDLE#` → `IDLE+POWERING_UP`. Also `POWER` state flag string is `POWERED_DOWN`.
- `scrontab` - add ability to update crontab from a file or standard input.
- `scrontab` - added ability to set and expand variables.
- Make `srun` sensitive to `BcastParameters`.
- Added `sbcast/srun --send-libs`, `sbcast --exclude` and `srun --bcast-exclude`.
- Changed `ReqMem` field in `sacct` to match memory from `ReqTRES`. It now shows the requested memory of the whole job with a letter appended indicating units (M for megabytes, G for gigabytes, etc.). `ReqMem` is only displayed for the job, since the step does not have requested TRES. Previously `ReqMem` was also displayed for the step but was just displaying `ReqMem` for the job.

6.8.3.4 API Changes

- `jobcomp` plugin: change plugin API to `jobcomp_p_*`.
- `sched` plugin: change plugin API to `sched_p_*` and remove `slurm_sched_p_initial_priority()` call.
- `step_ctx` code has been removed from the api.
- `slurm_stepd_get_info()` / `stepd_get_info()` has been removed from the api.
- The v0.0.35 OpenAPI plugin has now been marked as deprecated. Please convert your requests to the v0.0.37 OpenAPI plugin.

6.9 Creating containers from current HPC environment


Usually users use environment modules to adjust their environment (that is, environment variables like `PATH`, `LD_LIBRARY_PATH`, `MANPATH` etc.) to pick exactly the tools and libraries they need for their work. The same can be achieved with containers by including only those components in a container that are part of this environment. This functionality is now provided using the `spack` and `singularity` applications.

7 Removed and deprecated features and packages

This section lists features and packages that were removed from SUSE Linux Enterprise for High-Performance Computing or will be removed in upcoming versions.

7.1 Removed features and packages

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

- Python 2 bindings for `genders` has been removed. These are now provided for Python 3.
- Ganglia is not supported anymore in 15 SP4. It has been replaced with Grafana (<https://grafana.com/> )
- Due to a lack of usage by customers, some library packages have been removed from the HPC module in SLE HPC 15 SP4. On SUSE Linux Enterprise you can build your own library using `spack`. These libraries will continue to be available through SUSE Package Hub. The following libraries have been removed:

- `boost`
- `adios`
- `gsl`
- `fftw3`
- `hypre`
- `metis`
- `mumps`

- [netcdf](#)
- [ocr](#)
- [petsc](#)
- [ptscotch](#)
- [scalapack](#)
- [superlu](#)
- [trilinos](#)

7.2 Deprecated features and packages

The following features and packages are deprecated and will be removed in a future version of SUSE Linux Enterprise for High-Performance Computing.


8 Obtaining source code

This SUSE product includes materials licensed to SUSE under the GNU General Public License (GPL). The GPL requires SUSE to provide the source code that corresponds to the GPL-licensed material. The source code is available for download at <https://www.suse.com/download/sle-hpc/> on Medium 2. For up to three years after distribution of the SUSE product, upon request, SUSE will mail a copy of the source code. Send requests by e-mail to sle_source_request@suse.com (mailto:sle_source_request@suse.com). SUSE may charge a reasonable fee to recover distribution costs.


9 Legal notices


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


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

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classifications to export, re-export, or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical/biological weaponry end uses. Refer to <https://www.suse.com/company/legal/>  for more information on exporting SUSE software. SUSE assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2010-2022 SUSE LLC.

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

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

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

A Changelog for 15 SP4


A.1 2022-11-30

A.1.1 New

- Added *Section 6.1, "Enriched system visibility in the SUSE Customer Center (SCC)"* (Jira (<https://jira.suse.com/browse/SLE-24988>) )

A.2 2022-08-31

A.2.1 New

- Added *Section 6.2, "Automatically opened ports"* (Jira (<https://jira.suse.com/browse/SLE-22743>) )

A.3 2022-05-11

A.3.1 New

- Added this changelog

A.4 2022-03-23

A.4.1 New

- Added *Section 6.8, "Slurm"* (Jira (<https://jira.suse.com/browse/SLE-12352>) )
- Added notes about dolly, memkind, openblas, spack, and mpich in *Section 6, "Changes affecting all architectures"*

- Added note about Ganglia being unsupported in *Section 7, “Removed and deprecated features and packages”* (Jira (<https://jira.suse.com/browse/SLE-17777>) )
- Added note about removal of Python 2 bindings for genders (Jira (<https://jira.suse.com/browse/SLE-23359>) )

A.4.2 Updates

- Added a note about building libraries using spack in *Section 7, “Removed and deprecated features and packages”* (Jira (<https://jira.suse.com/browse/SLE-17776>) )
- Added adios and superlu to the list of removed libraries in *Section 7, “Removed and deprecated features and packages”*

A.5 2021-11-03

- Initial SP4 release