

## SUSE Linux Enterprise for High-Performance Computing 15 SP3

# 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 SP3 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/releasesnotes>. General documentation can be found at <https://documentation.suse.com/sle-hpc/15-SP3>.

Publication Date: 2025-10-31, Version: 15.300000000.20251031

## Contents

- 1 About the release notes 2
- 2 SUSE Linux Enterprise for High-Performance Computing 2
- 3 Modules, extensions, and related products 6
- 4 Technology previews 7
- 5 Modules 9
- 6 Changes affecting all architectures 10
- 7 Obtaining source code 53
- 8 Legal notices 54

# 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 SP3 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 SP3 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 (SP3) will be fully maintained and supported until 6 months after the release of SUSE Linux Enterprise for High-Performance Computing 15 SP4.

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](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 SP3 (and the SUSE Linux Enterprise modules) includes the following software that is shipped *only* under a GNU AGPL software license:

- Ghostscript (including subpackages)

SUSE Linux Enterprise for High-Performance Computing 15 SP3 (and the SUSE Linux Enterprise modules) includes the following software that is shipped under multiple licenses that include a GNU AGPL software license:

- MySpell dictionaries and LightProof
- ArgyllCMS

## 2.5 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 SP3. This directory includes PDF versions of the SUSE Linux Enterprise for High-Performance Computing 15 SP3 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 SP3, see <https://documentation.suse.com/sle-hpc/15-SP3>.
- 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 SP3. Modules and extensions add functionality to the system.

### 3.1 Modules in the SLE 15 SP3 product line


The SLE 15 SP3 product line is made up of modules that contain software packages. Each module has a clearly defined scope. Modules differ in their life cycles and update timelines.

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

Not all SLE modules are available with a subscription for SUSE Linux Enterprise for High-Performance Computing 15 SP3 itself (see the column *Available for*).







For information about the availability of individual packages within modules, see <https://sc-c.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://package-hub.suse.com/> .

## 3.3 Related products

This section 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 SP3 adds a kernel flavor `64kb`, offering a page size of 64 KiB and physical/virtual address size of 52 bits. Same as the `default` kernel flavor, it does not use preemption.

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



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

The SUSE Linux Enterprise for High-Performance Computing for POWER 15 SP3 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 SP3.

### 6.1 SLE HPC no longer a separate product

As of 15 SP3, SUSE Linux Enterprise for High-Performance Computing is no longer a separate product. As a result:

- the HPC Module can now be enabled in SUSE Linux Enterprise Server
- when migrating from SUSE Linux Enterprise for High-Performance Computing 15 SP3, SP4, and SP5, only SUSE Linux Enterprise Server 15 SP6 will be available as migration target. The result of such a migration will be an installation of SUSE Linux Enterprise Server with all the previously enabled modules.

### 6.2 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/>.



## 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 gnu11-compiler-hpc

This package adds an environment module for the GNU toolchain version 11.

It is needed to build and run applications against this version of the GNU toolchain with environment module support.

To build software, install the package `gnu11-compiler-hpc-devel`. This will install all required packages (including the GNU C, C++ and Fortran compilers) and set up the system accordingly. To build packages with these compilers, run `module load gnu`. This sets the standard binaries for the GNU compiler toolchain: `gcc`, `g++`, and `gfortran` to the appropriate version and sets environment variables `CC` and `CPP` to the desired binaries.



## Note

The environment module system always takes the latest version of the toolchain installed as the default. To load the environment for an older toolchain version, the version number needs to be specified when loading the module. For example: `module load gnu/7`.

To only run binaries, it is sufficient to install `gnu11-compiler-hpc` by running `zypper install gnu11-compiler-hpc`. However, building in this environment will not be possible.

## 6.4 Deprecation of packages

Due to a lack of usage by customers, some library packages have been deprecated and will not be included in the HPC module starting with SLE HPC 15 SP4. These libraries will continue to be available through SUSE Package Hub. The following libraries will be deprecated and will be dropped from SLE HPC 15 SP4:

- [boost](#)
- [gsl](#)
- [fftw3](#)
- [hypre](#)
- [metis](#)
- [mumps](#)
- [netcdf](#)
- [ocr](#)
- [petsc](#)
- [ptscotch](#)
- [scalapack](#)
- [trilinos](#)

In SLE HPC 15 SP3, Spack is now available which allows the user to easily download and install the entire HPC solution stack. For more information, see [Section 6.29, “Spack”](#).

## 6.5 boost

[boost](#) has been updated to version 1.75.0. For the full change log, see [https://www.boost.org/users/history/version\\_1\\_75\\_0.html](https://www.boost.org/users/history/version_1_75_0.html) [↗](#).

## 6.6 clustduct

Support for dolly has been added to clustduct. dolly is used to clone the installation of one machine to other machines. It can distribute image files, partitions or whole hard disk drives. For more information, see <http://www.cs.inf.ethz.ch/stricker/CoPs/patagonia/dolly.html> and <https://github.com/mslacken/clustduct>.

## 6.7 conman

conman has been updated to version 0.3.0.

The main changes include:

- fixed slow connection to Unix socket consoles triggered from inotif
- IPMI defaults can now be overridden via libipmiconsole.conf

For more information, see <http://dun.github.io/conman/>.

## 6.8 cpuid

cpuid has been updated to version 20201006. Support was added for:

- 14nm Zen
- Alder Lake
- Cato
- Comet Lake
- Cyrix MediaGX
- Elkhart Lake B0
- Golden Cove
- Hygon
- Jasper Lake A0 stepping (from Coreboot\*)
- Matisse B0 stepping
- Picasso A1

- Renoir A1
- Rocket Lake
- Sapphire Rapids
- Tiger Lake-U B0
- Zhaoxin KaiXian KX-6000

Other changes include:

- Update 4/eax CPU & core count terminology in the same way
- Zhaoxin decoding
- Added SEV cpuid bit

For more information, see <http://etallen.com/cpuid.html> .

## 6.9 `fftw3`

`fftw3` has been updated to version 3.3.9. The changes include:


- New API `fftw_planner_nthreads()` returns the number of threads currently being used by the planner.
- New `fftw_threads_set_callback` function to change the threading backend at runtime.
- Tweak usage of FMA instructions in a way that favors newer processors (Skylake and Ryzen) over older processors (Haswell)

For more information, see <http://www.fftw.org/> .

## 6.10 `ganglia`


`ganglia` has been updated to version 3.7.5. The changes include:

- added support for a global user config under `/etc/ganglia/conf.d/`
- added `download_js.sh` which can download the external JavaScript libraries
- added the available memory to the graph

For more information, see <http://ganglia.info/> .



## Important

ganglia relies on Python 2 which will no longer be available in SLE HPC 15 SP4. Thus in SLE HPC 15 SP4, ganglia will be replaced with graphana (<https://grafana.com/> .

## 6.11 genders

genders has been updated to version 1.27.3. The changes include:

- fixed parsing of a corner case
- newlines are not required at end of lines

For more information, see <https://github.com/chaos/genders> .

## 6.12 gnu10-compilers-hpc

Support for gcc 10 has been added to SLE for HPC.

<https://github.com/openhpc/ohpc> .

## 6.13 gsl

gsl has been updated to version 2.6. The changes include:

- added some statistics functions
- updated algorithms and implementations
- removed multiple previously deprecated functions
- removed -u flag to gsl-histogram

For more information, see <http://git.savannah.gnu.org/cgit/gsl.git/tree/NEWS>  and <https://www.gnu.org/software/gsl/> .

## 6.14 hdf5

hdf5 has been updated to version 1.10.7.

For more information, see <https://support.hdfgroup.org/ftp/HDF5/releases/hdf5-1.10/hdf5-1.10.7/src/hdf5-1.10.7-RELEASE.txt> and <http://www.hdfgroup.org/HDF5/>.

## 6.15 hypre

hypre has been updated to version 2.20.0. The changes include:

- Added matrix-based interpolation routines for AMG (CPU and GPU)
- Added GPU support for aggressive coarsening in AMG
- Added AMG-DD solver
- Added GPU support for AMG setup and several interpolation approaches.
- Added parallel ILU solvers and smoothers
- Added MGR features
- Moved reference manual API to online documentation
- IJMatrix/Vector assembly on GPUs (with pointers to GPU memory)
- Separated C and C++ headers

For more information, see <https://www.llnl.gov/casc/hypre/>.

## 6.16 imb

imb has been updated to version 2019.6. The changes include:

- Added IMB-P2P Stencil2D and Stencil3D benchmarks
- Added Visual Studio projects for IMB-P2P

For more information, see <https://software.intel.com/en-us/articles/intel-mpi-benchmarks>.

## 6.17 memkind

memkind has been updated to version 1.10.0. The changes include:

- Provided another way to use Persistent Memory in memkind (MEMKIND\_DAX\_KMEM\_\* kinds)
- Added C++ allocator for static kinds (including MEMKIND\_DAX\_KMEM\_\* kinds)
- Added support for background thread
- Extended hbw interface with hbw\_malloc\_usable\_size()

For more information, see <http://memkind.github.io/memkind>.

## 6.18 mpiP

mpiP has been updated to version 3.5. The changes include:

- Update deprecated functions
- Additional I/O routines
- Add Multi-Threading (MT) test suite
- Introduce additional statistics layer
- Change default stack frame unwinding count
- Adding additional RMA functions to cover all MPI3.1 RMA functions. Expanded test coverage
- Added MPI3 One-sided operation support, from changes by Jeff Hammond
- Include mpiP wrapper scripts for mpirun and Slurm srun
- Package API

For more information, see <http://mpip.sourceforge.net>.

## 6.19 munge

munge has been updated to version 0.5.14. The changes include:

- Added mungekey command for key generation via HKDF
- Added negative caching of user lookups for processing supplementary groups
- Added munged --origin and --stop cmdline opt
- Added unmunge --numeric cmdline opt
- Changed default name of munged seedfile

For more information, see <http://dun.github.io/munge>.

## 6.20 mvapich2

mvapich2 updated to version 2.3.4. For more information, see <http://mvapich.cse.ohio-state.edu/overview/#mv2>.

## 6.21 netcdf

netcdf has been updated to version 4.7.4. The changes include:

- Increased default size of cache buffer to 16 MB, from 4 MB. Increased number of slots to 4133.
- Allow zlib compression to be used with parallel I/O writes, if HDF5 version is 1.10.3 or greater.
- Restore use of szip compression when writing data (including writing in parallel if HDF5 version is 1.10.3 or greater).

For more information, see <https://github.com/Unidata/netcdf-c/releases/tag/v4.7.4> and <http://www.unidata.ucar.edu/software/netcdf/>.

## 6.22 netcdf-cxx4

netcdf-cxx4 has been updated to version 4.3.1. For more information, see <https://www.unidata.ucar.edu/blogs/news/entry/netcdf-c-4-3-1> and <http://www.unidata.ucar.edu/software/netcdf/>.

## 6.23 netcdf-fortran

netcdf-fortran has been updated to version 4.5.3. The changes include:

- Add a new libnetcdff.settings file
- Add build support for gcc10 to HPC build (bsc#1174439).
- Add build support for openmpi4.

For more information, see <https://github.com/Unidata/netcdf-fortran/releases/tag/v4.5.3> and <http://www.unidata.ucar.edu/software/netcdf/>.

## 6.24 openblas

openblas has been updated to version 0.3.13. Since version 0.3.8, the changes include:

- Added an optimized bfloat16 SBGEMV kernel for SkylakeX and Cooperlake
- Improved the performance in several areas
- Fixed missing BLAS/LAPACK functions (inadvertently dropped during the build system restructuring to support selective compilation)
- Update \_constraints to use 12GB RAM on x86\_64
- Improved performance of TRMM and TRSM for certain problem sizes
- Added support for Intel Cooperlake
- Add build support for gcc10 to HPC build (bsc#1174439).
- Improved thread locking behaviour in blas\_server and parallel getrf
- Added API (openblas\_setaffinity) to set thread affinity programmatically on Linux

- Improved speed of the AVX512 GEMM3M code, added an AVX512 kernel for
- STRMM and improved performance of the AVX2 GEMM kernels
- LAPACK has been updated to 3.9.0 (plus patches up to January 2nd, 2020)
- The "generic" (plain C) gemm beta kernel used by many targets has been sped up
- A new AVX512 DGEMM kernel was added and the AVX512 SGEMM kernel was significantly improved
- Optimized AVX512 kernels for CGEMM and ZGEMM have been added
- AVX2 kernels for STRMM, SGEMM, and CGEMM have been significantly sped up and optimized CGEMM3M and ZGEMM3M kernels have been added
- Added support for QEMU virtual cpus

For more information, see <http://www.openblas.net>.

## 6.25 `petsc`

`petsc` has been updated to version 3.14.5. The changes include:

- Add build support for gcc10 to HPC build (bsc#1174439).
- Add openmpi4 flavors.
- Binaries have changed path from bin/ to lib/petsc/bin/

For more information, see <https://www.mcs.anl.gov/petsc/documentation/changes/314.html> and <http://www.mcs.anl.gov/petsc>.

## 6.26 `ptscotch`

`ptscotch` has been updated to version 6.1.0. The change in release number reflects a change in behavior of the software, regarding two features:

- Algorithms for sparse matrix reordering now fully take into account vertex weights. While the previous versions only accounted for them in the nested dissection method to compute and optimize separator size/weight, they are now also accounted for in the minimum

degree and minimum fill algorithms. Consequently, vertex amalgamation can be tuned so that the size of (column-)blocks is tailored according to the real amount of computation to be performed on supervariables. The esmumps interface with MUMPS has been updated so as to take into account this new feature.

- The routines of the vertex graph partitioning (with overlap) module have been redesigned, leading to much smaller run time and, most often, higher quality

For more information, see <http://www.labri.fr/perso/pelegrin/scotch>.

## 6.27 python-numpy

python-numpy has been updated to version 1.17.3. This version added a new extensible random module.

For more information, see <http://www.numpy.org>.

## 6.28 Slurm

### 6.28.1 Deprecation of old Versions of Slurm

SLE receives a new Slurm version for roughly every second upstream release. To facilitate this, old version of Slurm will go out of maintenance successively. Users of these versions are encouraged to migrate to a later version before the expiry date. Note, that Slurm only allows migration two versions upwards without loss of data. To migrate to the latest version, migrations to intermediate versions may be required.

Once Slurm versions are out of maintenance, updates to packages depending on this Slurm version will no longer be provided. In particular, this will be the case for pdsh-slurm - the Slurm plugin to Pdsh. Note that pdsh-slurm is compatible to the Slurm initially shipped with a product/service pack) while packages pdsh-slurm\_<slurm\_version> is compatible with an upgrade version of Slurm.

TABLE 1: TABLE SUNSET SCHEDULE FOR SLURM

Slurm Version	Released for Service Pack	Support End Date
17.02	SLE-12-SP2	May 2024

18.08	SLE-15-SP1 and older	October 2024
20.02	SLE-15-SP2 and older	Januar 2025
20.11	SLE-15-SP3/4 and older	January 2026
22.05	SLE-15-SP4 and older	December 2026
23.02	SLE-15-SP5/6 and older	January 2028

### 6.28.2 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.

### 6.28.3 Slurm version 22.05

An update to Slurm version 22.05 is available.

#### 6.28.3.1 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`.

### 6.28.3.2 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 `MaxNode-Count` 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.

### 6.28.3.3 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`.

### 6.28.3.4 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.

### 6.28.3.5 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.

## 6.28.4 Highlights of Slurm version 21.08

### 6.28.4.1 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 `slurmshd` which runs `PrologSlurmctld` and `Epi-logSlurmctld`. 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 -bcast`.
- 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.

#### 6.28.4.2 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.

#### 6.28.4.3 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.

#### 6.28.4.4 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.

## 6.28.5 Version 20.11.05

Slurm has been updated to version 20.11.05.

The following changes have been made in 20.11:

- slurmctld is now set to fatal in case of computing node configured with CPUs == #Sockets. CPUs has to be either total number of cores or threads
- the FastSchedule option has been removed. The FastSchedule=2 functionality (used for testing and development) is available as the new SlurmdParameters=config\_overrides option
- slurmdbd is now set to fatal if the slurmdbd.conf file is not owned by SlurmUser or its mode is not set to 0600

### 6.28.5.1 Highlights of version 20.11

- Log messages enabled by the various DebugFlags have been overhauled, and will all print at the verbose() level, and prepend the flag name that is associated with a given log message
- accounting\_storage/filetxt has been removed as an option. Consider using accounting\_storage/slurmdbd as an alternative.
- Setting of number of Sockets per node was standardized for configuration line with and without Boards=. Specifically in case of Boards=1 and #CPUs given the default value of Sockets will be set to #CPUs/#Cores/#Threads
- Dynamic Future Nodes: slurmds started with -F[<feature>] will be associated with a node name in Slurm that matches the same hardware configuration
- SlurmctldParameters=cloud\_reg\_addrsa: Cloud nodes automatically get NodeAddr and NodeHostname set from slurmd registration

- SlurmctldParameters=power\_save[\_min]\_interval: Configure how often the power save module looks to do work`
- By default, a step started with `srun` will be granted exclusive (or non-overlapping) access to the resources assigned to that step. No other parallel step will be allowed to run on the same resources at the same time. This replaces one facet of the `--exclusive` option's behavior, but does not imply the `--exact` option described below. To get the previous default behavior, which allowed parallel steps to share all resources, use the new `srun --overlap` option.
- In conjunction to this non-overlapping step allocation behavior being the new default, there is an additional new option for step management `--exact`, which will allow a step access to only those resources requested by the step. This is the second half of the `--exclusive` behavior. Otherwise, by default all non-gres resources on each node in the allocation will be used by the step, making it so no other parallel step will have access to those resources unless `--overlap` is specified for both steps.
- The option `--threads-per-core` now influences task layout/binding, not just allocation.
- `AutoDetect` in `gres.conf` can now be specified for some nodes while not for others via the `NodeName` option.
- `gres.conf`: Add new `MultipleFiles` configuration entry to allow a single GRES to manage multiple device files simultaneously.
- The option `SallocDefaultCommand` has been removed.
- Support for an "Interactive Step" has been added, designed to be used with `salloc` to launch a terminal on an allocated compute node automatically. Enable this by setting `use_interactive_step` as part of `LaunchParameters`.
- IPv6 support has been added. This must be explicitly enabled with `EnableIPv6` in `CommunicationParameters`. IPv4 support can be disabled with `DisableIPv4`.
- Allow use of a target directory with `srun --bcast`, and change the default file name to include the node name as well.
- The new option `--mail-type=INVALID_DEPEND` has been added to `salloc`, `sbatch`, and `srun`.

- Differences between hardware (memory size, number of CPUs) discovered on node vs configured in `slurm.conf` will now throw an error only when the node state is set to drain. Previously it was done on every node registration, those messages were demoted to debug level.
- The command `scrontab` has been added. It permits `crontab`-compatible job scripts to be defined. These scripts will recur automatically (at most) on the intervals described.
- Enable the `-lnodes=:gpus=` in `#PBS/qsub -l` nodes syntax.
- Any user with `AdminLevel` equal or higher than `Operator` can see any hidden partition by default, as `SlurmUser` or root already did.
- `select/linear` will now allocate up to nodes `RealMemory` as specified in `slurm.conf` when configured with `SelectTypeParameters=CR_Memory` and `--mem=0` specified. Previous behavior was no memory accounted and no memory limits implied to job.
- `slurmrestd`, an API to interface with `slurmdbd`.
- The option `--ntasks-per-gpu` has been added to `sbatch` and `srun`.
- The `--gpu-bind=single` option has been added to `sbatch` and `srun`.
- Fix: `scontrol takeover [backup]` hangs when specifying a backup > 1 have been fixed. All `slurmctlds` below the "backup" will be shutdown.

#### 6.28.5.2 Version 20.11 Command Changes

- `sacct`: get the UID from database instead of from the user name and a system call. Add `--use-local-uid` option to `sacct` to use the previous behavior.
- `sbatch`: the `%s` format in `-e/-i/-o` options will expand to `batch` rather than `4294967294`.
- `queue`: added `pendingtime` as a option for `--Format`.
- `sacct`: `AllocGres` and `ReqGres` were removed. `Alloc/ReqTres` should be used instead.
- `scontrol`: added the "Reserved" license count to `scontrol show licenses`.
- `queue`: put sorted start times of `N/A` or `0` at the end of the list.

- scontrol: Change scontrol reboot ASAP to use next\_state=resume logic.
- scontrol: added an admin-settable "Comment" field to each Node.
- squeue and sinfo: -0 no longer repeat the last suffix specified.
- salloc: wait for PrologSlurmctld to finish before entering the shell.
- Add time specification: now-X (that is, subtract X from the present)

### 6.28.5.3 Version 20.11 API Changes

- slurm\_ctl\_conf\_t has been renamed to slurm\_conf\_t.
- slurm\_free\_kvs\_comm\_set() has been renamed to slurm\_pmi\_free\_kvs\_comm\_set(), slurm\_get\_kvs\_comm\_set() has been renamed to slurm\_pmi\_get\_kvs\_comm\_set().
- slurm\_job\_step\_layout\_get() parameters has changed to use slurm\_step\_id\_t see slurm.h for new implementation. If not running hetsteps just put NO\_VAL as the value for step\_het\_comp.
- slurm\_job\_step\_stat() parameters has changed to use slurm\_step\_id\_t see slurm.h for new implementation. If not running hetsteps just put NO\_VAL as the value for step\_het\_comp.
- slurm\_job\_step\_get\_pids() parameters has changed to use slurm\_step\_id\_t see slurm.h for new implementation. If not running hetsteps just put NO\_VAL as the value for step\_het\_comp.
- slurm\_job\_step\_get\_pids() parameters has changed to use slurm\_step\_id\_t see slurm.h for new implementation. If you are not running hetsteps, use NO\_VAL as the value for step\_het\_comp.
- slurmdb\_selected\_step\_t has been renamed slurm\_selected\_step\_t.
- slurm\_sbcst\_lookup() arguments have changed. It now takes a populated slurm\_selected\_step\_t instead of job\_id, het\_job\_offset, step\_id.
- Due to internal restructuring ahead of the 20.11 release, applications calling libslurm must call slurm\_init(NULL) before any API calls. Otherwise the API call is likely to fail due to libslurm's internal configuration not being available.

## 6.29 Spack

This is a package manager for supercomputers. It allows user to build pre-created recipes and deploy the software in a cluster. Spack creates environment modules and thus fits well with the rest of SLE for HPC.

For more information, see <https://spack.io>.

### 6.29.1 v0.19.1

#### 6.29.1.1 Spack Bugfixes

- buildcache create: make file exists less verbose
- spack mirror create: don't change paths to urls
- Improve error message for requirements
- uninstall: fix accidental cubic complexity
- scons: fix signature for install\_args
- Fix combine\_phase\_logs text encoding issues
- Use a module-like object to propagate changes in the MRO, when setting build env
- PackageBase should not define builder legacy attributes
- Forward lookup of the run\_tests attribute
- Bugfix for timers
- Fix path handling in prefix inspections
- Fix libtool filter for Fujitsu compilers
- FileCache: delete the new cache file on exception
- Propagate exceptions from Spack python console
- Tests: Fix a bug/typo in a config\_values.py fixture
- Various CI fixes
- Docs: remove monitors and analyzers, typos
- bump release version for tutorial command

## 6.29.2 v0.19.0

v0.19.0 is a major feature release.

### 6.29.2.1 Major features in this release

#### 1. Package requirements

Spack's traditional `package preferences` ([https://spack.readthedocs.io/en/latest/build\\_settings.html#package-preferences](https://spack.readthedocs.io/en/latest/build_settings.html#package-preferences)) are soft, but we've added hard requirements to `packages.yaml` and `spack.yaml`. Package requirements use the same syntax as specs:

```
packages:
  libfabric:
    require: "@1.13.2"
  mpich:
    require:
      - one_of: ["+cuda", "+rocm"]
```

More details in [the docs](https://spack.readthedocs.io/en/latest/build_settings.html#package-requirements) ([https://spack.readthedocs.io/en/latest/build\\_settings.html#package-requirements](https://spack.readthedocs.io/en/latest/build_settings.html#package-requirements)).

#### 2. Environment UI Improvements

- Fewer surprising modifications to `spack.yaml` :
  - `spack install` in an environment will no longer add to the `specs:` list; you'll need to either use `spack add <spec>` or `spack install --add <spec>`.
  - Similarly, `spack uninstall` will not remove from your environment's `specs:` list; you'll need to use `spack remove` or `spack uninstall --remove`. This will make it easier to manage an environment, as there is clear separation between the stack to be installed (`spack.yaml` / `spack.lock`) and which parts of it should be installed (`spack install` / `spack uninstall`).
- `concretizer:unify:true` is now the default mode for new environments  
We see more users creating `unify:true` environments now. Users who need `unify:false` can add it to their environment to get the old behavior. This will concretize every spec in the environment independently.
- Include environment configuration from URLs ([docs](https://spack.readthedocs.io/en/latest/environments.html#included-configurations) (<https://spack.readthedocs.io/en/latest/environments.html#included-configurations>))

You can now include configuration in your environment directly from a URL:

```
spack:
  include:
    - https://github.com/path/to/raw/config/compilers.yaml
```

### 3. Compiler and variant propagation

Currently, compiler flags and variants are inconsistent: compiler flags set for a package are inherited by its dependencies, while variants are not. We should have these be consistent by allowing for inheritance to be enabled or disabled for both variants and compiler flags. Example syntax: `* package ++variant`: enabled variant that will be propagated to dependencies `* package +variant`: enabled variant that will NOT be propagated to dependencies `* package ~~variant`: disabled variant that will be propagated to dependencies `* package ~variant`: disabled variant that will NOT be propagated to dependencies `* package cflags==g`: `cflags` will be propagated to dependencies `* package cflags=-g`: `cflags` will NOT be propagated to dependencies

+ Syntax for non-boolean variants is similar to compiler flags. More in the docs for [variants](https://spack.readthedocs.io/en/latest/basic_usage.html#variants) ([https://spack.readthedocs.io/en/latest/basic\\_usage.html#variants](https://spack.readthedocs.io/en/latest/basic_usage.html#variants)) and [compiler flags](https://spack.readthedocs.io/en/latest/basic_usage.html#compiler-flags) ([https://spack.readthedocs.io/en/latest/basic\\_usage.html#compiler-flags](https://spack.readthedocs.io/en/latest/basic_usage.html#compiler-flags)).

### 4. Enhancements to git version specifiers

- `v0.18.0` added the ability to use git commits as versions. You can now use the `git.` prefix to specify git tags or branches as versions. All of these are valid git versions in `v0.19` :

```
foo@abcdef1234abcdef1234abcdef1234abcdef1234      # raw commit
foo@git.abcdef1234abcdef1234abcdef1234abcdef1234  # commit with git prefix
foo@git.develop                                    # the develop branch
foo@git.0.19                                       # use the 0.19 tag
```

- `v0.19` also gives you more control over how Spack interprets git versions, in case Spack cannot detect the version from the git repository. You can suffix a git version with `=<version>` to force Spack to concretize it as a particular version :

```
# use mybranch, but treat it as version 3.2 for version comparison
foo@git.mybranch=3.2

# use the given commit, but treat it as develop for version comparison
foo@git.abcdef1234abcdef1234abcdef1234abcdef1234=develop
```

More in the docs ([https://spack.readthedocs.io/en/latest/basic\\_usage.html#version-specifier](https://spack.readthedocs.io/en/latest/basic_usage.html#version-specifier)) ↗

## 5. Changes to Cray EX Support

Cray machines have historically had their own platform within Spack, because we needed to go through the module system to leverage compilers and MPI installations on these machines. The Cray EX programming environment now provides standalone `craycc` executables and proper `mpicc` wrappers, so Spack can treat EX machines like Linux with extra packages .

We expect this to greatly reduce bugs, as external packages and compilers can now be used by prefix instead of through modules. We will also no longer be subject to reproducibility issues when modules change from Cray PE release to release and from site to site. This also simplifies dealing with the underlying Linux OS on cray systems, as Spack will properly model the machine's OS as either SuSE or RHEL.

## 6. Improvements to tests and testing in CI

- `spack ci generate --tests` will generate a `.gitlab-ci.yml` file that not only does builds but also runs tests for built packages . Public GitHub pipelines now also run tests in CI.
- `spack test run --explicit` will only run tests for packages that are explicitly installed, instead of all packages.

## 7. Experimental binding link model

You can add a new option to `config.yaml` to make Spack embed absolute paths to needed shared libraries in ELF executables and shared libraries on Linux ( [docs \(https://spack.readthedocs.io/en/latest/config\\_yaml.html#shared-linking-bind\)](https://spack.readthedocs.io/en/latest/config_yaml.html#shared-linking-bind) ↗):

```
config:
  shared_linking:
    type: rpath
    bind: true
```

This can improve launch time at scale for parallel applications, and it can make installations less susceptible to environment variables like `LD_LIBRARY_PATH`, even especially when dealing with external libraries that use `RUNPATH`. You can think of this as a faster, even higher-precedence version of `RPATH`.

### 6.29.2.2 Other new features of note

- `spack spec` prints dependencies more legibly. Dependencies in the output now appear at the *earliest* level of indentation possible
- You can override `package.py` attributes like `url`, directly in `packages.yaml` ( [docs \(https://spack.readthedocs.io/en/latest/build\\_settings.html#assigning-package-attributes\)](https://spack.readthedocs.io/en/latest/build_settings.html#assigning-package-attributes) )
- There are a number of new architecture-related format strings you can use in Spack configuration files to specify paths ( [docs \(https://spack.readthedocs.io/en/latest/configuration.html#config-file-variables\)](https://spack.readthedocs.io/en/latest/configuration.html#config-file-variables) )

### 6.29.2.3 Performance Improvements

- Major performance improvements for installation from binary caches
- Test suite can now be parallelized using `xdist` (used in GitHub Actions)
- Reduce lock contention for parallel builds in environments

### 6.29.2.4 New binary caches and stacks

- We now build nearly all of E4S with `oneapi` in our buildcache
- Added 3 new machine learning-centric stacks to binary cache: `x86_64_v3`, `CUDA`, `ROCm`

### 6.29.2.5 Removals and Deprecations

- Support for Python 3.5 is dropped . Only Python 2.7 and 3.6+ are officially supported.
- This is the last Spack release that will support Python 2 . Spack `v0.19` will emit a deprecation warning if you run it with Python 2, and Python 2 support will soon be removed from the `develop` branch.
- `LD_LIBRARY_PATH` is no longer set by default by `spack load` or module loads. Setting `LD_LIBRARY_PATH` in Spack environments/modules can cause binaries from outside of Spack to crash, and Spack's own builds use `RPATH` and do not need `LD_LIBRARY_PATH` set in order to run. If you still want the old behavior, you can run these commands to configure Spack to set `LD_LIBRARY_PATH`:

```
spack config add modules:prefix_inspections:lib64:[LD_LIBRARY_PATH]
```

```
spack config add modules:prefix_inspections:lib:[LD_LIBRARY_PATH]
```

- The `spack:concretization:[together|separately]` has been removed after being deprecated in `v0.18`. Use `concretizer:unify:[true|false]`.
- `config:module_roots` is no longer supported after being deprecated in `v0.18`. Use configuration in module sets instead ( docs ([https://spack.readthedocs.io/en/latest/module\\_file\\_support.html](https://spack.readthedocs.io/en/latest/module_file_support.html)) ↗).
- `spack activate` and `spack deactivate` are no longer supported, having been deprecated in `v0.18`. Use an environment with a view instead of activating/deactivating (docs (<https://spack.readthedocs.io/en/latest/environments.html#configuration-in-spack-yaml>) ↗).
- The old YAML format for buildcaches is now deprecated . If you are using an old buildcache with YAML metadata you will need to regenerate it with JSON metadata.
- `spack bootstrap trust` and `spack bootstrap untrust` are deprecated in favor of `spack bootstrap enable` and `spack bootstrap disable` and will be removed in `v0.20`.
- The `graviton2` architecture has been renamed to `neoverse_n1`, and `graviton3` is now `neoverse_v1`. Buildcaches using the old architecture names will need to be rebuilt.
- The terms `blacklist` and `whitelist` have been replaced with `include` and `exclude` in all configuration files . You can use `spack config update` to automatically fix your configuration files.

#### 6.29.2.6 Notable Bugfixes

- Permission setting on installation now handles effective uid properly
- `buildable:true` for an MPI implementation now overrides `buildable:false` for `mpi`
- Improved error messages when attempting to use an unconfigured compiler
- Do not punish explicitly requested compiler mismatches in the solver
- `spack stage: add missing` `-fresh` and `-reuse`
- Fixes for adding build system executables like `cmake` to package scope
- Bugfix for binary relocation with aliased strings produced by newer `binutils`

### 6.29.3 v0.18.1

#### 6.29.3.1 Spack Bugfixes

- Fix several bugs related to bootstrapping
- Fix a regression that was causing spec hashes to differ between Python 2 and Python 3
- Fixed compiler flags for oneAPI and DPC++
- Fixed several issues related to concretization
- Improved support for Cray manifest file and `spack external find`
- Assign a version to openSUSE Tumbleweed according to the GLIBC version in the system
- Improved Dockerfile generation for `spack containerize`
- Fixed a few bugs related to concurrent execution of commands

#### 6.29.3.2 Package updates

- WarpX: add v22.06, fixed libs property
- openPMD: add v0.14.5, update recipe for @develop

### 6.29.4 v0.18.0

v0.18.0 is a major feature release.

#### 6.29.4.1 Major features in this release

##### 1. Concretizer now reuses by default

`spack install --reuse` was introduced in v0.17.0, and `--reuse` is now the default concretization mode. Spack will try hard to resolve dependencies using installed packages or binaries.

To avoid reuse and to use the latest package configurations, (the old default), you can use `spack install --fresh`, or add configuration like this to your environment or `concretizer.yaml`:

```
concretizer:
```

```
reuse: false
```

## 2. Finer-grained hashes

Spack hashes now include link, run, *and* build dependencies, as well as a canonical hash of package recipes. Previously, hashes only included link and run dependencies (though build dependencies were stored by environments). We coarsened the hash to reduce churn in user installations, but the new default concretizer behavior mitigates this concern and gets us reuse *and* provenance. You will be able to see the build dependencies of new installations with spack find. Old installations will not change and their hashes will not be affected.

## 3. Improved error messages

Error handling with the new concretizer is now done with optimization criteria rather than with unsatisfiable cores, and Spack reports many more details about conflicting constraints.

## 4. Unify environments when possible

Environments have thus far supported concretization: together or concretization: separately. These have been replaced by a new preference in concretizer.yaml:

```
concretizer:
  unify: [true|false|when_possible]
```

concretizer:unify:when\_possible will *try* to resolve a fully unified environment, but if it cannot, it will create multiple configurations of some packages where it has to. For large environments that previously had to be concretized separately, this can result in a huge speedup (40-50x).

## 5. Automatically find externals on Cray machines

Spack can now automatically discover installed packages in the Cray Programming Environment by running spack external find (or spack external read-cray-manifest to *only* query the PE). Packages from the PE (e.g., cray-mpich are added to the database with full dependency information, and compilers from the PE are added to compiler-s.yaml. Available with the June 2022 release of the Cray Programming Environment.

## 6. New binary format and hardened signing

Spack now has an updated binary format, with improvements for security. The new format has a detached signature file, and Spack verifies the signature before untarring or decompressing the binary package. The previous format embedded the signature in a tar file,

which required the client to run `tar` before verifying. Spack can still install from build caches using the old format, but we encourage users to switch to the new format going forward.

Production GitLab pipelines have been hardened to securely sign binaries. There is now a separate signing stage so that signing keys are never exposed to build system code, and signing keys are ephemeral and only live as long as the signing pipeline stage.

## 7. Bootstrap mirror generation

The `spack bootstrap mirror` command can automatically create a mirror for bootstrapping the concretizer and other needed dependencies in an air-gapped environment.

## 8. Makefile generation

`spack env depfile` can be used to generate a `Makefile` from an environment, which can be used to build packages the environment in parallel on a single node. e.g.:

```
spack -e myenv env depfile > Makefile
make
```

Spack propagates `gmake` jobserver information to builds so that their jobs can share cores.

## 9. New variant features

In addition to being conditional themselves, variants can now have [conditional values](https://spack.readthedocs.io/en/latest/packaging_guide.html#conditional-possible-values) (https://spack.readthedocs.io/en/latest/packaging\_guide.html#conditional-possible-values) that are only possible for certain configurations of a package.

Variants can be [declared sticky](https://spack.readthedocs.io/en/latest/packaging_guide.html#sticky-variants) (https://spack.readthedocs.io/en/latest/packaging\_guide.html#sticky-variants), which prevents them from being enabled or disabled by the concretizer. Sticky variants must be set explicitly by users on the command line or in `packages.yaml`.

- Allow conditional possible values in variants
- Add a sticky property to variants

### 6.29.4.2 Other new features of note

- Environment views can optionally link only `run` dependencies with `link:run`
- `spack external find --all` finds library-only packages in addition to build dependencies
- Customizable `config:license_dir` option

- `spack external find --path PATH` takes a custom search path
- `spack spec` has a new `--format` argument like `spack find`
- `spack concretize --quiet` skips printing concretized specs
- `spack info` now has cleaner output and displays test info
- Package-level submodule option for git commit versions
- Using `/hash` syntax to refer to concrete specs in an environment now works even if `/hash` is not installed.

#### 6.29.4.3 Major internal refactors

- full hash (see above)
- new develop versioning scheme `0.19.0-dev0`
- Allow for multiple dependencies/dependents from the same package
- Splice differing virtual packages

#### 6.29.4.4 Performance Improvements

- Concretization of large environments with `unify: when_possible` is much faster than concretizing separately (see above)
- Single-pass view generation algorithm is 2.6x faster

#### 6.29.4.5 Archspec improvements

- `oneapi` and `dpcpp` flag support
- better support for `M1` and `a64fx`

#### 6.29.4.6 Removals and Deprecations

- Spack no longer supports Python `2.6`
- Removed deprecated `--run-tests` option of `spack install`; use `spack test`

- Removed deprecated `spack flake8`; use `spack style`
- Deprecate `spack:concretization` config option; use `concretizer:unify`
- Deprecate top-level module configuration; use module sets
- `spack activate` and `spack deactivate` are deprecated in favor of environments; will be removed in `0.19.0`

#### 6.29.4.7 Notable Bugfixes

- Fix bug that broke locks with many parallel builds
- Many bugfixes and consistency improvements for the new concretizer and `--reuse`

#### 6.29.4.8 Packages

- `CMakePackage` uses `CMAKE_INSTALL_RPATH_USE_LINK_PATH`
- Refactored `lua` support: `lua-lang` virtual supports both `lua` and `luajit` via new `LuaPackage` build system
- `PythonPackage`: now installs packages with `pip`
- `Python`: improve `site_packages_dir` handling
- `Extends`: support spec, not just package name
- Use stable URLs and `?full_index=1` for all github patches

### 6.30 superlu

`superlu` has been updated to version 5.2.2.

For more information, see <http://crd.lbl.gov/~xiaoye/SuperLU>.

### 6.31 Slurm 23.02

#### 6.31.1 Important Notes on Upgrading Slurm from a Previous Version

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.

The 23.02 `slurmdbd` will work with Slurm daemons of version 21.08 and above. 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 22.05 to version 23.02 without loss of jobs or other state information. Upgrading directly from an earlier version of Slurm will result in loss of state information.

All `SPANK` plugins must be recompiled when upgrading from any Slurm version prior to 23.02.



## Note

PMIx v1.x is no longer supported.

### 6.31.2 Highlights

- `slurmctld` - Add new RPC rate limiting feature. This is enabled through `SlurmctldParameters=rl_enable`, otherwise disabled by default.
- Make `scontrol` reconfigure and sending a `SIGHUP` to the `slurmctld` behave the same. If you were using `SIGHUP` as a 'lighter' `scontrol` reconfigure to rotate logs please update your scripts to use `SIGUSR2` instead.
- Change cloud nodes to show by default. `PrivateData=cloud` is no longer needed.
- `sreport` - Count planned (FKA reserved) time for jobs running in `IGNORE_JOBS` reservations. Previously was lumped into `IDLE` time.
- `job_container/tmpfs` - Support running with an arbitrary list of private mount points (`/tmp` and `/dev/shm` are the default, but not required).
- `job_container/tmpfs` - Set more environment variables in `InitScript`.
- Make all cgroup directories created by Slurm owned by root. This was the behavior in `cgroup/v2` but not in `cgroup/v1` where by default the step directories ownership were set to the user and group of the job.
- `accounting_storage/mysql` - change purge/archive to calculate record ages based on end time, rather than start or submission times.

- job\_submit/lua - add support for log\_user() from slurm\_job\_modify().
- Run the following scripts in slurmshd instead of slurmctld: ResumeProgram, ResumeFailProgram, SuspendProgram, ResvProlog, ResvEpilog, and RebootProgram (only with SlurmctldParameters=reboot\_from\_controller).
- Only permit changing log levels with srun --slurmd-debug by root or SlurmUser.
- slurmctld will fatal() when reconfiguring the job\_submit plugin fails.
- Add PowerDownOnIdle partition option to power down nodes after nodes become idle.
- Add “[jobid.stepid]” prefix from slurmstepd and “slurmshd” prefix from slurmshd to Syslog logging. Previously was only happening when logging to a file.
- Add purge and archive functionality for job environment and job batch script records.
- Extend support for Include files to all "configless" client commands.
- Make node weight usable for powered down and rebooting nodes.
- Removed “launch” plugin.
- Add “Extra” field to job to store extra information other than a comment.
- Add usage gathering for AMD (requires ROCM 5.5+) and NVIDIA gpus.
- Add job’s allocated nodes, features, oversubscribe, partition, and reservation to SLURM\_RESUME\_FILE output for power saving.
- Automatically create directories for stdout/stderr output files. Paths may use %j and related substitution characters as well.
- Add --tres-per-task to salloc/sbatch/srun.
- Allow nodefeatures plugin features to work with cloud nodes. e.g. - Powered down nodes have no active changeable features.
  - Nodes can’t be changed to other active features until powered down.
  - Active changeable features are reset/cleared on power down.
- Make slurmstepd cgroups constrained by total configured memory from slurm.conf (NodeName=<> RealMemory=#) instead of total physical memory.

- node\_features/helpers - add support for the OR and parentheses operators in a --constraint expression.
- slurmctld will fatal() when [Prolog|Epilog]Slurmctld are defined but are not executable.
- Validate node registered active features are a super set of node's currently active changeable features.
- On clusters without any PrologFlags options, batch jobs with failed prologs no longer generate an output file.
- Add SLURM\_JOB\_START\_TIME and SLURM\_JOB\_END\_TIME environment variables.
- Add SuspendExcStates option to slurm.conf to avoid suspending/powering down specific node states.
- Add support for DCMI power readings in IPMI plugin.
- slurmrestd served /slurm/v0.0.39 and /slurmdb/v0.0.39 endpoints had major changes from prior versions. Almost all schemas have been renamed and modified. Sites using OpenAPI Generator clients are highly suggested to upgrade to using atleast version 6.x due to limitations with prior versions.
- Allow for --nodelist to contain more nodes than required by --nodes.
- Rename “nodes” to “nodes\_resume” in SLURM\_RESUME\_FILE job output.
- Rename “all\_nodes” to “all\_nodes\_resume” in SLURM\_RESUME\_FILE output.
- Add jobcomp/kafka plugin.
- Add new PreemptParameters=reclaim\_licenses option which will allow higher priority jobs to preempt jobs to free up used licenses. (This is only enabled for with Preempt-Modes of CANCEL and REQUEUE, as Slurm cannot guarantee suspended jobs will release licenses correctly.)
- hpe/slingshot - add support for the instant-on feature.
- Add ability to update SuspendExc\* parameters with scontrol.
- Add ability to restore SuspendExc\* parameters on restart with slurmctld -R option.
- Add ability to clear a GRES specification by setting it to "0" via “scontrol update job”.

- Add `SLURM_JOB_OVERSUBSCRIBE` environment variable for `Epilog`, `Prolog`, `EpilogSlurmctld`, `PrologSlurmctld`, and mail output.
- System node down reasons are appended to existing reasons, separated by '!'.
- New command `scrunch` has been added. `scrunch` acts as an Open Container Initiative (OCI) runtime proxy to run containers seamlessly via Slurm.
- Fixed `GpuFreqDef` option. When set in `slurm.conf`, it will be used if `--gpu-freq` was not explicitly set by the job step.

### 6.31.3 Configuration File Changes (see appropriate man page for details)

- `job_container.conf` - Added “Dirs” option to list desired private mount points.
- `node_features` plugins - invalid users specified for `AllowUserBoot` will now result in `fatal()` rather than just an error.
- Deprecate `AllowedKmemSpace`, `ConstrainKmemSpace`, `MaxKmemPercent`, and `MinKmemSpace`.
- Allow jobs to queue even if the user is not in `AllowGroups` when `EnforcePartLimits=no` is set. This ensures consistency for all the Partition access controls, and matches the documented behavior for `EnforcePartLimits`.
- Add `InfluxDBTimeout` parameter to `acct_gather.conf`.
- `job_container/tmpfs` - add support for expanding `%h` and `%n` in `BasePath`.
- `slurm.conf` - Removed `SlurmctldPlugstack` option.
- Add new `SlurmctldParameters=validate_nodeaddr_threads=<number>` option to allow concurrent hostname resolution at `slurmctld` startup.
- Add new `AccountingStoreFlags=job_extra` option to store a job’s extra field in the database.
- Add new “`defer_batch`” option to `SchedulerParameters` to only defer scheduling for batch jobs.
- Add new `DebugFlags` option “`JobComp`” to replace “`Elasticsearch`”.

- Add configurable job requeue limit parameter - MaxBatchRequeue - in slurm.conf to permit changes from the old hard-coded value of 5.
- helpers.conf - Allow specification of node specific features.
- helpers.conf - Allow many features to one helper script.
- job\_container/tmpfs - Add “Shared” option to support shared namespaces. This allows autofs to work with the job\_container/tmpfs plugin when enabled.
- acct\_gather.conf - Added EnergyIPMIPowerSensors=Node=DCMI and Node=D-CMI\_ENHANCED.
- Add new “getnameinfo\_cache\_timeout= <number>” option to CommunicationParameters to adjust or disable caching the results of getnameinfo().
- Add new PrologFlags=ForceRequeueOnFail option to automatically requeue batch jobs on Prolog failures regardless of the job --requeue setting.
- Add HealthCheckNodeState=NONDRAINED\_IDLE option.
- Add “explicit” to Flags in gres.conf. This makes it so the gres is not automatically added to a job’s allocation when --exclusive is used. Note that this is a per-node flag.
- Moved the “preempt\_” options from SchedulerParameters to PreemptParameters, and dropped the prefix from the option names. (The old options will still be parsed for backwards compatibility, but are now undocumented.)
- Add LaunchParameters=ulimit\_pam\_adapt, which enables setting RLIMIT\_RSS in adopted processes.
- Update SwitchParameters=job\_vni to enable/disable creating job VNIs for all jobs, or when a user requests them.
- Update SwitchParameters=single\_node\_vni to enable/disable creating single node VNIs for all jobs, or when a user requests them.
- Add ability to preserve SuspendExc\* parameters on reconfig with ReconfigFlags=Keep-PowerSaveSettings.
- slurmdbd.conf - Add new AllResourcesAbsolute to force all new resources to be created with the Absolute flag.

- `topology/tree` - Add new `TopologyParam=SwitchAsNodeRank` option to reorder nodes based on switch layout. This can be useful if the naming convention for the nodes does not naturally map to the network topology.
- Removed the default setting for `GpuFreqDef`. If unset, no attempt to change the GPU frequency will be made if `--gpu-freq` is not set for the step.

#### 6.31.4 Command Changes (see man pages for details)

- `sacctmgr` - no longer force updates to the `AdminComment`, `Comment`, or `SystemComment` to lower-case.
- `sinfo` - Add `-F/--future` option to `sinfo` to display future nodes.
- `sacct` - Rename “Reserved” field to “Planned” to match `sreport` and the nomenclature of the 'Planned' node.
- `scontrol` - advanced reservation flag `MAINT` will no longer replace nodes, similar to `STATIC_ALLOC`
- `sbatch` - add parsing for `#PBS -d` and `#PBS -w`.
- `scontrol show assoc_mgr` will show `username(uid)` instead of `uid` in QoS section.
- Add `strigger --draining` and `-R/--resume` options.
- Change `--oversubscribe` and `--exclusive` to be mutually exclusive for job submission. Job submission commands will now fatal if both are set. Previously, these options would override each other, with the last one in the job submission command taking effect.
- `scontrol` - Requested `TRES` and allocated `TRES` will now always be printed when showing jobs, instead of one `TRES` output that was either the requested or allocated.
- `srun --ntasks-per-core` now applies to job and step allocations. Now, use of `--ntasks-per-core=1` implies `--cpu-bind=cores` and `--ntasks-per-core>1` implies `--cpu-bind=threads`.
- `salloc/sbatch/srun` - Check and abort if `ntasks-per-core > threads-per-core`.
- `scontrol` - Add `ResumeAfter=<secs>` option to “`scontrol update nodename=`”.
- Add a new “`nodes=`” argument to `scontrol setdebug` to allow the debug level on the `slurmd` processes to be temporarily altered.

- Add a new “nodes=” argument to “scontrol setdebugflags” as well.
- Make it so `scrntab` prints client-side the `job_submit()` error message (which can be set i.e. by using the `log_user()` function for the lua plugin).
- `scontrol` - Reservations will not be allowed to have `STATIC_ALLOC` or `MAINT` flags and `REPLACE[_DOWN]` flags simultaneously.
- `scontrol` - Reservations will only accept one reoccurring flag when being created or updated.
- `scontrol` - A reservation cannot be updated to be reoccurring if it is already a floating reservation.
- `squeue` - removed unused “%s” and “SelectJobInfo” formats.
- `squeue` - align print format for exit and derived codes with that of other components (`<exit_status>:<signal_number>`).
- `sacct` - Add `--array` option to expand job arrays and display array tasks on separate lines.
- Partial support for “--json” and “--yaml” formatted outputs have been implemented for `sacctmgr`, `sdiag`, `sinfo`, `squeue`, and `scontrol`. The resultant data output will be filtered by normal command arguments. Formatting arguments will continue to be ignored.
- `salloc`/`sbatch`/`srun` - extended the `--nodes` syntax to allow for a list of valid node counts to be allocated to the job. This also supports a "step count" value (e.g., `--nodes=20-100:20` is equivalent to `--nodes=20,40,60,80,100`) which can simplify the syntax when the job needs to scale by a certain "chunk" size.
- `srun` - add user requestible vnis with “--network=job\_vni” option.
- `srun` - add user requestible single node VNIs with the “--network=single\_node\_vni” option.


### 6.31.5 API Changes

- `job_container` plugins - `container_p_stepd_create()` function signature replaced `uint32_t uid` with `stepd_step_rec_t* step`.
- `gres` plugins - `gres_g_get_devices()` function signature replaced `pid_t pid` with `stepd_step_rec_t* step`.

- `cgroup` plugins - `task_cgroup_devices_constrain()` function signature removed `pid_t` `pid`.
- `task` plugins - replace `task_p_pre_set_affinity()`, `task_p_set_affinity()`, and `task_p_post_set_affinity()` with `task_p_pre_launch_priv()` like it was back in slurm 20.11.
- Allow for concurrent processing of `job_submit_g_submit()` and `job_submit_g_modify()` calls. If your plugin is not capable of concurrent operation you must add additional locking within your plugin.
- Removed return value from `slurm_list_append()`.
- The `List` and `ListIterator` types have been removed in favor of `list_t` and `list_itr_t` respectively.
- `burst_buffer` plugins - add `bb_g_build_het_job_script()`. `bb_g_get_status()` - added authenticated UID and GID. `bb_g_run_script()` - added `job_info` argument.
- `burst_buffer.lua` - Pass UID and GID to most hooks. Pass `job_info` (detailed job information) to many hooks. See `etc/burst_buffer.lua.example` for a complete list of changes. WARNING: Backwards compatibility is broken for `slurm_bb_get_status`: UID and GID are passed before the variadic arguments. If UID and GID are not explicitly listed as arguments to `slurm_bb_get_status()`, then they will be included in the variadic arguments. Backwards compatibility is maintained for all other hooks because the new arguments are passed after the existing arguments.
- `node_features` plugins - `node_features_p_reboot_weight()` function removed. `node_features_p_job_valid()` - added parameter `feature_list`. `node_features_p_job_xlate()` - added parameters `feature_list` and `job_node_bitmap`.
- New `data_parser` interface with v0.0.39 plugin.

## 7 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 re-

quest, 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) ([mailto:sle\\_source\\_request@suse.com](mailto:sle_source_request@suse.com)) . SUSE may charge a reasonable fee to recover distribution costs.


## 8 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 © 2010-2025 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.