

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

Publication Date: 2022-09-30, Version: 15.300000000.20220930

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 24
- 8 Legal notices 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 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 .

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://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 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 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.

6.2 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 6.27, "Spack"*](#).

6.3 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 [!\[\]\(a03a7eb2f4046e1d3c76772003e549ea_img.jpg\)](#).

6.4 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>.

6.5 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/>.

6.6 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> .

6.7 `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/> .

6.8 `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/> .

6.9 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> .

6.10 gnu10-compilers-hpc

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

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

6.11 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/> .

6.12 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/>.

6.13 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/>.

6.14 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>.

6.15 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>.

6.16 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>.

6.17 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>.

6.18 mvapich2

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

6.19 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/>.

6.20 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/>.

6.21 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>.

6.22 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>.

6.23 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>.

6.24 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>.

6.25 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>.

6.26 slurm

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

6.26.1 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.

6.26.2 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`.
- `squeue`: 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.
- squeue: 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)

6.26.3 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_sbcast_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.

6.27 Spack

Spack version 0.16.0 has been added.

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>.

6.28 superlu

`superlu` has been updated to version 5.2.2.

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

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 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.

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 © 2010-2022 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.