



SingularityCE Admin Guide

Release 3.11

SingularityCE Project Contributors

Aug 25, 2023

CONTENTS

1	Admin Quick Start	2
1.1	Architecture of SingularityCE	2
1.2	SingularityCE Security	2
1.3	OCI Compatibility	3
1.4	Installation from Source	3
1.5	Installation from RPM/Deb Packages	6
1.6	Configuration	6
1.7	Test SingularityCE	6
2	Installing SingularityCE	7
2.1	Installation on Linux	7
2.2	Installation on Windows or Mac	18
3	SingularityCE Configuration Files	20
3.1	singularity.conf	20
3.2	cgroups.toml	27
3.3	ecl.toml	30
3.4	GPU Library Configuration	31
3.5	capability.json	32
3.6	seccomp-profiles	34
3.7	remote.yaml	34
4	User Namespaces & Fakeroot	37
4.1	User Namespace Requirements	37
4.2	Unprivileged Installations	38
4.3	–userns option	38
4.4	–sif-fuse option	38
4.5	Fakeroot feature	39
4.6	Unprivileged Builds Without User Namespaces	42
5	Security in SingularityCE	44
5.1	Security Policy	44
5.2	Background	44
5.3	Setuid & User Namespaces	45
5.4	Runtime & User Privilege Model	46
5.5	Singularity Image Format (SIF)	46
5.6	Plugins	47
5.7	Configuration & Runtime Options	47
6	Installed Files	48

Welcome to the SingularityCE Admin Guide!

This guide aims to cover installation instructions, configuration detail, and other topics important to system administrators working with SingularityCE.

See the [user guide](#) for more information about how to use SingularityCE.

ADMIN QUICK START

This quick start gives an overview of installation of SingularityCE from source, a description of the architecture of SingularityCE, and pointers to configuration files. More information, including alternate installation options and detailed configuration options can be found later in this guide.

1.1 Architecture of SingularityCE

SingularityCE is designed to allow containers to be executed as if they were native programs or scripts on a host system. No daemon is required to build or run containers, and the security model is compatible with shared systems.

As a result, integration with clusters and schedulers such as Univa Grid Engine, Torque, SLURM, SGE, and many others is as simple as running any other command. All standard input, output, errors, pipes, IPC, and other communication pathways used by locally running programs are synchronized with the applications running locally within the container.

SingularityCE favors an ‘integration over isolation’ approach to containers. By default only the mount namespace is isolated for containers, so that they have their own filesystem view. Access to hardware such as GPUs, high speed networks, and shared filesystems is easy and does not require special configuration. Default access to user home directories, `/tmp` space, and installation specific mounts makes it simple for users to benefit from the reproducibility of containerized applications without major changes to their existing workflows. Where more complete isolation is important, SingularityCE can use additional Linux namespaces and other security and resource limits to accomplish this.

1.2 SingularityCE Security

Note: See also the *security section* (page 44) of this guide, for more detail.

SingularityCE uses a number of strategies to provide safety and ease-of-use on both single-user and shared systems. Notable security features include:

- The effective user inside a container is the same as the user who ran the container. This means access to files and devices from the container is easily controlled with standard POSIX permissions.
- Container filesystems are mounted `nosuid` and container applications run with the `prctl NO_NEW_PRIVS` flag set. This means that applications in a container cannot gain additional privileges. A regular user cannot `sudo` or otherwise gain root privilege on the host via a container.
- The Singularity Image Format (SIF) supports encryption of containers, as well as cryptographic signing and verification of their content.

- SIF containers are immutable and their payload is run directly, without extraction to disk. This means that the container can always be verified, even at runtime, and encrypted content is not exposed on disk.
- Restrictions can be configured to limit the ownership, location, and cryptographic signatures of containers that are permitted to be run.

To support the SIF image format, automated networking setup etc., and older Linux distributions without user namespace support, Singularity runs small amounts of privileged container setup code via a `starter-setuid` binary. This is a 'setuid root' binary, so that SingularityCE can perform filesystem loop mounts and other operations that need privilege. The setuid flow is the default mode of operation, but *can be disabled* (page 14) on build, or in the `singularity.conf` configuration file if required.

Note: Running SingularityCE in non-setuid mode requires unprivileged user namespace support in the operating system kernel and does not support all features. This impacts integrity/security guarantees of containers at runtime.

See the *non-setuid installation section* (page 14) for further detail on how to install SingularityCE to run in non-setuid mode.

1.3 OCI Compatibility

SingularityCE allows users to run, and build from, the majority of OCI containers created with tools such as Docker. Beginning with SingularityCE 3.11, there are two modes of operation that support OCI containers in different ways.

SingularityCE's *native runtime*, used by default, supports all features that are exposed via the `singularity` command. It builds and runs containers in SingularityCE's own on-disk formats. When an OCI container is pulled or built into a SingularityCE image, a translation step occurs. While most OCI images are supported as-is, there are some limitations and compatibility options may be required.

SingularityCE 3.11's *experimental OCI runtime*, enabled with the `--oci` flag, runs containers using a low-level OCI runtime - either `crun` or `runc`. The container is executed from a native OCI format on-disk. Not all CLI features are currently implemented, but OCI containers using the `USER` directive or which are otherwise incompatible with SingularityCE's native runtime are better supported. This mode is considered experimental, ahead of full OCI support in version 4.0. Functionality may change across 3.11 patch releases of SingularityCE.

1.4 Installation from Source

SingularityCE can be installed from source directly, or by building an RPM package from the source. Linux distributions may also package SingularityCE, but their packages may not be up-to-date with the upstream version on GitHub.

To install SingularityCE directly from source, follow the procedure below. Other methods are discussed in the *Installation* (page 7) section.

Note: This quick-start that you will install as `root` using `sudo`, so that SingularityCE uses the default `setuid` workflow, and all features are available. See the *non-setuid installation* (page 14) section of this guide for detail of how to install as a non-root user, and how this affects the functionality of SingularityCE.

1.4.1 Install Dependencies

On Red Hat Enterprise Linux or CentOS install the following dependencies:

```
# Install basic tools for compiling
sudo yum groupinstall -y 'Development Tools'
# Install RPM packages for dependencies
sudo yum install -y \
  wget \
  libseccomp-devel \
  glib2-devel \
  squashfs-tools \
  cryptsetup \
  runc
```

On Ubuntu or Debian install the following dependencies:

```
# Ensure repositories are up-to-date
sudo apt-get update
# Install debian packages for dependencies
sudo apt-get install -y \
  wget \
  build-essential \
  libseccomp-dev \
  libglib2.0-dev \
  pkg-config \
  squashfs-tools \
  cryptsetup \
  runc
```

Note - *runc* can be omitted if you will not use the *singularity oci* commands.

1.4.2 Install Go

SingularityCE v3 is written primarily in Go, and you will need Go installed to compile it from source. Versions of Go packaged by your distribution may not be new enough to build SingularityCE.

{SingularityCE} aims to maintain support for the two most recent stable versions of Go. This corresponds to the [Go Release Maintenance Policy](#) and [Security Policy](#), ensuring critical bug fixes and security patches are available for all supported language versions.

The method below is one of several ways to [install and configure Go](#).

Note: If you have previously installed Go from a download, rather than an operating system package, you should remove your `go` directory, e.g. `rm -r /usr/local/go` before installing a newer version. Extracting a new version of Go over an existing installation can lead to errors when building Go programs, as it may leave old files, which have been removed or replaced in newer versions.

Visit the [Go download page](#) and pick a package archive to download. Copy the link address and download with `wget`. Then extract the archive to `/usr/local` (or use other instructions on [go installation page](#)).

```
$ export VERSION=1.20.4 OS=linux ARCH=amd64 && \
  wget https://dl.google.com/go/go$VERSION.$OS-$ARCH.tar.gz && \
```

(continues on next page)

(continued from previous page)

```
sudo tar -C /usr/local -xzvf go$VERSION.$OS-$ARCH.tar.gz && \
rm go$VERSION.$OS-$ARCH.tar.gz
```

Finally, add `/usr/local/go/bin` to the `PATH` environment variable:

```
echo 'export PATH=$PATH:/usr/local/go/bin' >> ~/.bashrc
source ~/.bashrc
```

1.4.3 Download SingularityCE from a GitHub release

You can download SingularityCE from one of the releases. To see a full list, visit [the GitHub release page](#). After deciding on a release to install, you can run the following commands to proceed with the installation.

```
$ export VERSION=3.11.0 && # adjust this as necessary \
  wget https://github.com/sylabs/singularity/releases/download/v${VERSION}/singularity-
↪ce-${VERSION}.tar.gz && \
  tar -xzf singularity-ce-${VERSION}.tar.gz && \
  cd singularity-ce-${VERSION}
```

1.4.4 Compile & Install SingularityCE

SingularityCE uses a custom build system called `makeit`. `mconfig` is called to generate a `Makefile` and then `make` is used to compile and install.

```
$ ./mconfig && \
  make -C ./builddir && \
  sudo make -C ./builddir install
```

By default SingularityCE will be installed in the `/usr/local` directory hierarchy. You can specify a custom directory with the `--prefix` option, to `mconfig`:

```
$ ./mconfig --prefix=/opt/singularity
```

This option can be useful if you want to install multiple versions of Singularity, install a personal version of SingularityCE on a shared system, or if you want to remove SingularityCE easily after installing it.

For a full list of `mconfig` options, run `mconfig --help`. Here are some of the most common options that you may need to use when building SingularityCE from source.

- `--sysconfdir`: Install read-only config files in `sysconfdir`. This option is important if you need the `singularity.conf` file or other configuration files in a custom location.
- `--localstatedir`: Set the state directory where containers are mounted. This is a particularly important option for administrators installing SingularityCE on a shared file system. The `--localstatedir` should be set to a directory that is present on each individual node.
- `-b`: Build SingularityCE in a given directory. By default this is `./builddir`.
- `--without-common`: Do not build `common`, a container monitor that is used by the `singularity oci` commands. `common` is bundled with SingularityCE and will be built and installed by default. Use `--without-common` if you wish to use a version of `common` `>=2.0.24` that is provided by your distribution rather than the bundled version. You can also specify `--without-common` if you know you will not use the `singularity oci` commands.

1.5 Installation from RPM/Deb Packages

Sylabs provides `.rpm` packages of SingularityCE, for mainstream-supported versions of RHEL and derivatives (e.g. Alma Linux / Rocky Linux). We also provide `.deb` packages for current Ubuntu LTS releases.

These packages can be downloaded from the [GitHub release page](#) and installed using your distribution's package manager.

The packages are provided as a convenience for users of the open source project, and are built in our public CircleCI workflow. They are not signed, but SHA256 sums are provided on the release page.

1.6 Configuration

SingularityCE is configured using files under `etc/singularity` in your `--prefix`, or `--syconfdir` if you used that option with `mconfig`. In a default installation from source without a `--prefix` set you will find them under `/usr/local/etc/singularity`. In a default installation from RPM or Deb packages you will find them under `/etc/singularity`.

You can edit these files directly, or using the SingularityCE `config global` command as the root user to manage them.

`singularity.conf` contains the majority of options controlling the runtime behavior of SingularityCE. Additional files control security, network, and resource configuration. Head over to the [Configuration files](#) (page 20) section where the files and configuration options are discussed.

1.7 Test SingularityCE

You can run a quick test of SingularityCE using a container in the Sylabs Container Library:

```
$ singularity exec library://alpine cat /etc/alpine-release
3.9.2
```

See the [user guide](#) for more information about how to use SingularityCE.

INSTALLING SINGULARITYCE

This section will guide you through the process of installing SingularityCE 3.11.0 via several different methods. (For instructions on installing earlier versions of SingularityCE please see [earlier versions of the docs.](#))

2.1 Installation on Linux

SingularityCE can be installed on any modern Linux distribution, on bare-metal or inside a Virtual Machine. Nested installations inside containers are not recommended, and require the outer container to be run with full privilege.

2.1.1 System Requirements

SingularityCE requires ~163MiB disk space once compiled and installed.

There are no specific CPU or memory requirements at runtime, though 2GB of RAM is recommended when building from source.

Full functionality of SingularityCE requires that the kernel supports:

- **OverlayFS mounts** - (minimum kernel ≥ 3.18) Required for full flexibility in bind mounts to containers, and to support persistent overlays for writable containers.
- **Unprivileged user namespaces** - (minimum kernel ≥ 3.8 , ≥ 3.18 recommended) Required to run containers without root or setuid privilege. Required to build containers unprivileged in `--fakeroot` mode. Required to run containers using the experimental `--oci` mode.
- **Unprivileged overlay** - (minimum kernel ≥ 5.11 , ≥ 5.13 recommended) Required to use `--overlay`, to mount a persistent overlay directory onto the container, when running without root or setuid.

External Binaries

Singularity depends on a number of external binaries for full functionality. From SingularityCE 3.9, the methods that are used to find these binaries have been standardized as below.

Configurable Paths

The following binaries are found on `$PATH` during build time when `./mconfig` is run, and their location is added to the `singularity.conf` configuration file. At runtime this configured location is used. To specify an alternate executable, change the relevant path entry in `singularity.conf`.

- `cryptsetup` version 2 with kernel LUKS2 support is required for building or executing encrypted containers.
- `ldconfig` is used to resolve library locations / symlinks when using the `-nv` or `--rocm` GPU support.
- `nvidia-container-cli` is used to configure a container for Nvidia GPU / CUDA support when running with the experimental `--nvcccli` option.

For the following additional binaries, if the `singularity.conf` entry is left blank, then `$PATH` will be searched at runtime.

- `go` is required to compile plugins, and must be an identical version as that used to build SingularityCE.
- `mksquashfs` from `squashfs-tools 4.3+` is used to create the squashfs container filesystem that is embedded into SIF container images. The `mksquashfs procs` and `mksquashfs mem` directives in `singularity.conf` can be used to control its resource usage.
- `unsquashfs` from `squashfs-tools 4.3+` is used to extract the squashfs container filesystem from a SIF file when necessary.

Searching \$PATH

The following utilities are always found by searching `$PATH` at runtime:

- `true`
- `mkfs.ext3` is used to create overlay images.
- `cp`
- `dd`
- `newuidmap` and `newgidmap` are distribution provided `setuid` binaries used to configure `subuid/gid` mappings for `--fakeroot` in non-`setuid` installs.
- `crun` or `runc` are OCI runtimes used for the `singularity oci` commands and experimental `--oci` mode for `run / shell / exec`. `crun` is preferred over `runc` if it is available. `runc` is provided by a package in all common Linux distributions. `crun` is packaged in more recent releases of common Linux distributions.
- `proot` is an optional dependency that can be used to permit limited unprivileged builds without user namespace / `subuid` support. It is packaged in the community repositories for common Linux distributions, and is available as a static binary from [proot-me.github.io](https://github.com/proot-me/proot).

Bootstrap Utilities

The following utilities are required to bootstrap containerized distributions using their native tooling:

- `mount`, `umount`, `pacstrap` for Arch Linux.
- `mount`, `umount`, `mknod`, `debootstrap` for Debian based distributions.
- `dnf` or `yum`, `rpm`, `curl` for EL derived RPM based distributions.
- `uname`, `zypper`, `SUSEConnect` for SLES derived RPM based distributions.

Non-standard ldconfig / Nix & Guix Environments

If SingularityCE is installed under a package manager such as Nix or Guix, but on top of a standard Linux distribution (e.g. CentOS or Debian), it may be unable to correctly find the libraries for `--nv` and `--rocm` GPU support. This issue occurs as the package manager supplies an alternative `ldconfig`, which does not identify GPU libraries installed from host packages.

To allow SingularityCE to locate the host (i.e. CentOS / Debian) GPU libraries correctly, set `ldconfig path` in `singularity.conf` to point to the host `ldconfig`. I.E. it should be set to `/sbin/ldconfig` or `/sbin/ldconfig.real` rather than a Nix or Guix related path.

Filesystem support / limitations

SingularityCE supports most filesystems, but there are some limitations when installing SingularityCE on, or running containers from, common parallel / network filesystems. In general:

- We strongly recommend installing SingularityCE on local disk on each compute node.
- If SingularityCE is installed to a network location, a `--localstatedir` should be provided on each node, and Singularity configured to use it.
- The `--localstatedir` filesystem should support overlay mounts.
- `TMPDIR` / `SINGULARITY_TMPDIR` should be on a local filesystem wherever possible.

Note: Set the `--localstatedir` location by providing `--localstatedir my/dir` as an option when you configure your SingularityCE build with `./mconfig`.

Disk usage at the `--localstatedir` location is negligible (<1MiB). The directory is used as a location to mount the container root filesystem, overlays, bind mounts etc. that construct the runtime view of a container. You will not see these mounts from a host shell, as they are made in a separate mount namespace.

Overlay support

Various features of SingularityCE, such as the `--writable-tmpfs` and `--overlay`, options use the Linux overlay filesystem driver to construct a container root filesystem that combines files from different locations. Not all filesystems can be used with the overlay driver, so when containers are run from these filesystems some SingularityCE features may not be available.

Overlay support has two aspects:

- `lowerdir` support for a filesystem allows a directory on that filesystem to act as the 'base' of a container. A filesystem must support overlay `lowerdir` for you be able to run a Singularity sandbox container on it, while using functionality such as `--writable-tmpfs` / `--overlay`.
- `upperdir` support for a filesystem allows a directory on that filesystem to be merged on top of a `lowerdir` to construct a container. If you use the `--overlay` option to overlay a directory onto a container, then the filesystem holding the overlay directory must support `upperdir`.

Note that any overlay limitations mainly apply to sandbox (directory) containers only. A SIF container is mounted into the `--localstatedir` location, which should generally be on a local filesystem that supports overlay.

Fakeroot / (sub)uid/gid mapping

When SingularityCE is run using the *fakeroot* (page 39) option it creates a user namespace for the container, and UIDs / GIDs in that user namespace are mapped to different host UID / GIDs.

Most local filesystems (ext4/xfs etc.) support this uid/gid mapping in a user namespace.

Most network filesystems (NFS/Lustre/GPFS etc.) *do not* support this uid/gid mapping in a user namespace. Because the fileservers are not aware of the mappings it will deny many operations, with 'permission denied' errors. This is currently a generic problem for rootless container runtimes.

SingularityCE cache / atomic rename

SingularityCE will cache SIF container images generated from remote sources, and any OCI/docker layers used to create them. The cache is created at `$HOME/.singularity/cache` by default. The location of the cache can be changed by setting the `SINGULARITY_CACHEDIR` environment variable.

The directory used for `SINGULARITY_CACHEDIR` should be:

- A unique location for each user. Permissions are set on the cache so that private images cached for one user are not exposed to another. This means that `SINGULARITY_CACHEDIR` cannot be shared.
- Located on a filesystem with sufficient space for the number and size of container images anticipated.
- Located on a filesystem that supports atomic rename, if possible.

In SingularityCE version 3.6 and above the cache is concurrency safe. Parallel runs of SingularityCE that would create overlapping cache entries will not conflict, as long as the filesystem used by `SINGULARITY_CACHEDIR` supports atomic rename operations.

Support for atomic rename operations is expected on local POSIX filesystems, but varies for network / parallel filesystems and may be affected by topology and configuration. For example, Lustre supports atomic rename of files only on a single MDT. Rename on NFS is only atomic to a single client, not across systems accessing the same NFS share.

If you are not certain that your `$HOME` or `SINGULARITY_CACHEDIR` filesystems support atomic rename, do not run `singularity` in parallel using remote container URLs. Instead use `singularity pull` to create a local SIF image, and then run this SIF image in a parallel step. An alternative is to use the `--disable-cache` option, but this will result in each SingularityCE instance independently fetching the container from the remote source, into a temporary location.

NFS

NFS filesystems support overlay mounts as a `lowerdir` only, and do not support user-namespace (sub)uid/gid mapping.

- Containers run from SIF files located on an NFS filesystem do not have restrictions.
- You cannot use `--overlay mynfsdir/` to overlay a directory onto a container when the overlay (`upperdir`) directory is on an NFS filesystem.
- When using `--fakeroot` to build or run a container, your `TMPDIR` / `SINGULARITY_TMPDIR` should not be set to an NFS location.
- You should not run a sandbox container with `--fakeroot` from an NFS location.

Lustre / GPFS / PanFS

Lustre, GPFS, and PanFS do not have sufficient `upperdir` or `lowerdir` overlay support for certain SingularityCE features, and do not support user-namespace (sub)uid/gid mapping.

- You cannot use `--overlay` or `--writable-tmpfs` with a sandbox container that is located on a Lustre, GPFS, or PanFS filesystem. SIF containers on Lustre, GPFS, and PanFS will work correctly with these options.
- You cannot use `--overlay` to overlay a directory onto a container, when the overlay (`upperdir`) directory is on a Lustre, GPFS, or PanFS filesystem.
- When using `--fakeroot` to build or run a container, your `TMPDIR/SINGULARITY_TMPDIR` should not be a Lustre, GPFS, or PanFS location.
- You should not run a sandbox container with `--fakeroot` from a Lustre, GPFS, or PanFS location.

2.1.2 Install from Provided RPM / Deb Packages

Sylabs provides `.rpm` packages of SingularityCE, for mainstream-supported versions of RHEL and derivatives (e.g. Alma Linux / Rocky Linux). We also provide `.deb` packages for current Ubuntu LTS releases.

These packages can be downloaded from the [GitHub release page](#) and installed using your distribution's package manager.

The packages are provided as a convenience for users of the open source project, and are built in our public CircleCI workflow. They are not signed, but SHA256 sums are provided on the release page.

2.1.3 Install from Source

To use the latest version of SingularityCE from GitHub you will need to build and install it from source. This may sound daunting, but the process is straightforward, and detailed below.

If you have an earlier version of SingularityCE installed, you should *remove it* (page 16) before executing the installation commands. You will also need to install some dependencies and install [Go](#).

Install Dependencies

On Red Hat Enterprise Linux or CentOS install the following dependencies:

```
# Install basic tools for compiling
sudo yum groupinstall -y 'Development Tools'
# Install RPM packages for dependencies
sudo yum install -y \
  libseccomp-devel \
  glib2-devel \
  squashfs-tools \
  cryptsetup \
  runc
```

On Ubuntu or Debian install the following dependencies:

```
# Ensure repositories are up-to-date
sudo apt-get update
# Install debian packages for dependencies
```

(continues on next page)

(continued from previous page)

```
sudo apt-get install -y \  
  build-essential \  
  libseccomp-dev \  
  libglib2.0-dev \  
  pkg-config \  
  squashfs-tools \  
  cryptsetup \  
  runc
```

Note: You can build SingularityCE without `cryptsetup` available, but will not be able to use encrypted containers without it installed on your system.

If you will not use the `singularity oci` commands, `runc` is not required.

Install Go

SingularityCE is written in Go, and aims to maintain support for the two most recent stable versions of Go. This corresponds to the Go Release Maintenance Policy and Security Policy, ensuring critical bug fixes and security patches are available for all supported language versions.

Building SingularityCE may require a newer version of Go than is available in the repositories of your distribution. We recommend installing the latest version of Go from the [official binaries](<https://golang.org/dl/>).

This is one of several ways to [install and configure Go](#).

Note: If you have previously installed Go from a download, rather than an operating system package, you should remove your `go` directory, e.g. `rm -r /usr/local/go` before installing a newer version. Extracting a new version of Go over an existing installation can lead to errors when building Go programs, as it may leave old files, which have been removed or replaced in newer versions.

Visit the [Go download page](#) and pick a package archive to download. Copy the link address and download with `wget`. Then extract the archive to `/usr/local` (or use other instructions on [go installation page](#)).

```
$ export VERSION=1.20.4 OS=linux ARCH=amd64 && \  
  wget https://dl.google.com/go/go$VERSION.$OS-$ARCH.tar.gz && \  
  sudo tar -C /usr/local -xzf go$VERSION.$OS-$ARCH.tar.gz && \  
  rm go$VERSION.$OS-$ARCH.tar.gz
```

Then, set up your environment for Go.

```
$ echo 'export GOPATH=${HOME}/go' >> ~/.bashrc && \  
  echo 'export PATH=/usr/local/go/bin:${PATH}:${GOPATH}/bin' >> ~/.bashrc && \  
  source ~/.bashrc
```

Download SingularityCE from a release

You can download SingularityCE from one of the releases. To see a full list, visit the [GitHub release page](#). After deciding on a release to install, you can run the following commands to proceed with the installation.

```
$ export VERSION=3.11.0 && # adjust this as necessary \
  wget https://github.com/sylabs/singularity/releases/download/v${VERSION}/singularity-
ce-${VERSION}.tar.gz && \
  tar -xzf singularity-ce-${VERSION}.tar.gz && \
  cd singularity-ce-${VERSION}
```

Checkout Code from Git

The following commands will install SingularityCE from the [GitHub repo](#) to `/usr/local`. This method will work for `>=v3.11.0`. To install an older tagged release see [older versions of the docs](#).

When installing from source, you can decide to install from either a **tag**, a **release branch**, or from the **main branch**.

- **tag**: GitHub tags form the basis for releases, so installing from a tag is the same as downloading and installing a [specific release](#). Tags are expected to be relatively stable and well-tested.
- **release branch**: A release branch represents the latest version of a minor release with all the newest bug fixes and enhancements (even those that have not yet made it into a point release). For instance, to install v3.10 with the latest bug fixes and enhancements checkout `release-3.10`. Release branches may be less stable than code in a tagged point release.
- **main branch**: The main branch contains the latest, bleeding edge version of SingularityCE. This is the default branch when you clone the source code, so you don't have to check out any new branches to install it. The main branch changes quickly and may be unstable.

To ensure that the SingularityCE source code is downloaded to the appropriate directory use these commands.

```
$ git clone --recurse-submodules https://github.com/sylabs/singularity.git && \
  cd singularity && \
  git checkout --recurse-submodules v3.11.0
```

Compile Singularity

SingularityCE uses a custom build system called `makeit`. `mconfig` is called to generate a Makefile and then `make` is used to compile and install.

To support the SIF image format, automated networking setup etc., and older Linux distributions without user namespace support, Singularity must be `make install`ed as root or with ``sudo`, so it can install the `libexec/singularity/bin/starter-setuid` binary with root ownership and `setuid` permissions for privileged operations. If you need to install as a normal user, or do not want to use `setuid` functionality [see below](#) (page 14).

```
$ ./mconfig && \
  make -C ./builddir && \
  sudo make -C ./builddir install
```

By default SingularityCE will be installed in the `/usr/local` directory hierarchy. You can specify a custom directory with the `--prefix` option, to `mconfig` like so:

```
$ ./mconfig --prefix=/opt/singularity
```


This option can be useful if you want to install multiple versions of SingularityCE, install a personal version of SingularityCE on a shared system, or if you want to remove SingularityCE easily after installing it.

For a full list of `mconfig` options, run `mconfig --help`. Here are some of the most common options that you may need to use when building SingularityCE from source.

- `--sysconfdir`: Install read-only config files in `sysconfdir`. This option is important if you need the `singularity.conf` file or other configuration files in a custom location.
- `--localstatedir`: Set the state directory where containers are mounted. This is a particularly important option for administrators installing SingularityCE on a shared file system. The `--localstatedir` should be set to a directory that is present on each individual node.
- `-b`: Build SingularityCE in a given directory. By default this is `./builddir`.
- `--without-common`: Do not build the common OCI container monitor. Use this option if you are certain you will not use the `singularity oci` commands, or wish to use `common >=2.0.24` provided by your distribution, and available on `$PATH`.
- **`--reproducible`: Enable support for reproducible builds. Ensures** that the compiled binaries do not include any temporary paths, the source directory path, etc. This disables support for building plugins.

Unprivileged (non-setuid) Installation

If you need to install SingularityCE as a non-root user, or do not wish to allow the use of a setuid root binary, you can configure SingularityCE with the `--without-suid` option to `mconfig`:

```
$ ./mconfig --without-suid --prefix=/home/dave/singularity-ce && \
  make -C ./builddir && \
  make -C ./builddir install
```

If you have already installed SingularityCE you can disable the setuid flow by setting the option `allow setuid = no` in `etc/singularity/singularity.conf` within your installation directory.

When SingularityCE does not use setuid all container execution will use a user namespace. This requires support from your operating system kernel, and imposes some limitations on functionality. You should review the [requirements](#) (page 37) and [limitations](#) (page 38) in the [user namespace](#) (page 37) section of this guide.

Relocatable Installation

Since SingularityCE 3.8, an unprivileged (non-setuid) installation is relocatable. As long as the structure inside the installation directory (`--prefix`) is maintained, it can be moved to a different location and SingularityCE will continue to run normally.

Relocation of a default setuid installation is not supported, as restricted location / ownership of configuration files is important to security.

Source bash completion file

To enjoy bash shell completion with SingularityCE commands and options, source the bash completion file:

```
$ . /usr/local/etc/bash_completion.d/singularity
```

Add this command to your `~/.bashrc` file so that bash completion continues to work in new shells. (Adjust the path if you installed SingularityCE to a different location.)

2.1.4 Build and install an RPM

If you use RHEL, CentOS or SUSE, building and installing a Singularity RPM allows your SingularityCE installation be more easily managed, upgraded and removed. In SingularityCE $\geq v3.0.1$ you can build an RPM directly from the [release tarball](#).

Note: Be sure to download the correct asset from the [GitHub releases page](#). It should be named `singularity-ce-<version>.tar.gz`.

After installing the [dependencies](#) (page 11) and installing [Go](#) (page 12) as detailed above, you are ready to download the tarball and build and install the RPM.

```
$ export VERSION=3.11.0 && # adjust this as necessary \
  wget https://github.com/sylabs/singularity/releases/download/v${VERSION}/singularity-
  ce-${VERSION}.tar.gz && \
  rpmbuild -tb singularity-ce-${VERSION}.tar.gz && \
  sudo rpm -ivh ~/rpmbuild/RPMS/x86_64/singularity-ce-${VERSION}-1.el7.x86_64.rpm && \
  rm -rf ~/rpmbuild singularity-ce-${VERSION}*.tar.gz
```

If you encounter a failed dependency error for go lang but installed it from source, build with this command:

```
rpmbuild -tb --nodeps singularity-ce-${VERSION}.tar.gz
```

Options to `mconfig` can be passed using the familiar syntax to `rpmbuild`. For example, if you want to force the local state directory to `/mnt` (instead of the default `/var`) you can do the following:

```
rpmbuild -tb --define='_localstatedir /mnt' singularity-ce-${VERSION}.tar.gz
```

Note: It is very important to set the local state directory to a directory that physically exists on nodes within a cluster when installing SingularityCE in an HPC environment with a shared file system.

Build an RPM from Git source

Alternatively, to build an RPM from a branch of the Git repository you can clone the repository, directly make an rpm, and use it to install Singularity:

```
$ ./mconfig && \
make -C builddir rpm && \
sudo rpm -ivh ~/rpmbuild/RPMS/x86_64/singularity-ce-3.11.0.el7.x86_64.rpm # or whatever_
version you built
```

To build an rpm with an alternative install prefix set `RMPREFIX` on the make step, for example:

```
$ make -C builddir rpm RMPREFIX=/usr/local
```

For finer control of the `rpmbuild` process you may wish to use `make dist` to create a tarball that you can then build into an rpm with `rpmbuild -tb` as above.

2.1.5 Remove an old version

In a standard installation of SingularityCE 3.0.1 and beyond (when building from source), the command `sudo make install` lists all the files as they are installed. You must remove all of these files and directories to completely remove SingularityCE.

```
$ sudo rm -rf \
  /usr/local/libexec/singularity \
  /usr/local/var/singularity \
  /usr/local/etc/singularity \
  /usr/local/bin/singularity \
  /usr/local/bin/run-singularity \
  /usr/local/etc/bash_completion.d/singularity
```

If you anticipate needing to remove SingularityCE, it might be easier to install it in a custom directory using the `--prefix` option to `mconfig`. In that case SingularityCE can be uninstalled simply by deleting the parent directory. Or it may be useful to install SingularityCE *using a package manager* (page 15) so that it can be updated and/or uninstalled with ease in the future.

2.1.6 Testing & Checking the Build Configuration

After installation you can perform a basic test of Singularity functionality by executing a simple container from the Syllabs Cloud library:

```
$ singularity exec library://alpine cat /etc/alpine-release
3.10.0
```

See the [user guide](#) for more information about how to use SingularityCE.

singularity buildcfg

Running `singularity buildcfg` will show the build configuration of an installed version of SingularityCE, and lists the paths used by SingularityCE. Use `singularity buildcfg` to confirm paths are set correctly for your installation, and troubleshoot any ‘not-found’ errors at runtime.

```
$ singularity buildcfg
PACKAGE_NAME=singularity
PACKAGE_VERSION=3.11.0
BUILDDIR=/home/dtrudg/Syllabs/Git/singularity/builddir
PREFIX=/usr/local
EXECPREFIX=/usr/local
BINDIR=/usr/local/bin
SBINDIR=/usr/local/sbin
LIBEXECDIR=/usr/local/libexec
```

(continues on next page)

(continued from previous page)

```
DATAROOTDIR=/usr/local/share
DATADIR=/usr/local/share
SYSCONFDIR=/usr/local/etc
SHAREDSTATEDIR=/usr/local/com
LOCALSTATEDIR=/usr/local/var
RUNSTATEDIR=/usr/local/var/run
INCLUDEDIR=/usr/local/include
DOCDIR=/usr/local/share/doc/singularity
INFODIR=/usr/local/share/info
LIBDIR=/usr/local/lib
LOCALEDIR=/usr/local/share/locale
MANDIR=/usr/local/share/man
SINGULARITY_CONFDIR=/usr/local/etc/singularity
SESSIONDIR=/usr/local/var/singularity/mnt/session
```

Note that the LOCALSTATEDIR and SESSIONDIR should be on local, non-shared storage.

The list of files installed by a successful `setuid` installation of SingularityCE can be found in the [appendix, installed files section](#) (page 48).

Test Suite

The SingularityCE codebase includes a test suite that is run during development using CI services.

If you would like to run the test suite locally you can run the test targets from the `builddir` directory in the source tree:

- `make check` runs source code linting and dependency checks
- `make unit-test` runs basic unit tests
- `make integration-test` runs integration tests
- `make e2e-test` runs end-to-end tests, which exercise a large number of operations by calling the SingularityCE CLI with different execution profiles.

Note: Running the full test suite requires a `docker` installation, and `nc` in order to test `docker` and `instance/networking` functionality.

SingularityCE must be installed in order to run the full test suite, as it must run the CLI with `setuid` privilege for the `starter-suid` binary.

Warning: `sudo` privilege is required to run the full tests, and you should not run the tests on a production system. We recommend running the tests in an isolated development or build environment.

2.2 Installation on Windows or Mac

Linux container runtimes like SingularityCE cannot run natively on Windows or Mac because of basic incompatibilities with the host kernel. (Contrary to a popular misconception, MacOS does not run on a Linux kernel. It runs on a kernel called Darwin originally forked from BSD.)

For this reason, the SingularityCE community maintains a set of Vagrant Boxes via [Vagrant Cloud](#), one of Hashicorp's open source tools. The current versions can be found under the [sylabs](#) organization.

2.2.1 Windows

Install the following programs:

- [Git for Windows](#)
- [VirtualBox for Windows](#)
- [Vagrant for Windows](#)
- [Vagrant Manager for Windows](#)

2.2.2 Mac

SingularityCE is available via Vagrant (installable with [Homebrew](#) or manually)

To use Vagrant via Homebrew:

```
$ /usr/bin/ruby -e "$(curl -fsSL https://raw.githubusercontent.com/Homebrew/install/master/install)"
$ brew install --cask virtualbox vagrant vagrant-manager
```

2.2.3 SingularityCE Vagrant Box

Run Git Bash (Windows) or open a terminal (Mac) and create and enter a directory to be used with your Vagrant VM.

```
$ mkdir vm-singularity-ce && \
  cd vm-singularity-ce
```

If you have already created and used this folder for another VM, you will need to destroy the VM and delete the Vagrantfile.

```
$ vagrant destroy && \
  rm Vagrantfile
```

Then issue the following commands to bring up the Virtual Machine. (Substitute a different value for the \$VM variable if you like.)

```
$ export VM=sylabs/singularity-ce-3.8-ubuntu-bionic64 && \
  vagrant init $VM && \
  vagrant up && \
  vagrant ssh
```

You can check the installed version of SingularityCE with the following:

```
vagrant@vagrant:~$ singularity version
3.11.0
```

Of course, you can also start with a plain OS Vagrant box as a base and then install SingularityCE using one of the above methods for Linux.

2.2.4 SingularityCE Docker Image

It is possible to use a Dockerized Singularity, here is a sample `compose.yaml` (Singularity version 3.7.4) for use with Docker Compose:

```
services:
  singularity:
    image: quay.io/singularity/singularity:v3.7.4-slim
    stdin_open: true
    tty: true
    privileged: true
    volumes:
      - .:/root
    entrypoint: ["/bin/sh"]
```

Singularity in Docker can have various disadvantages, but basic container operations will work. Currently, the intended use case is continuous integration, meaning that you should be able to build a Singularity container using this Docker Compose file. For more information see [issue#5](#) and the image's source [repo](#)

SINGULARITYCE CONFIGURATION FILES

As a SingularityCE Administrator, you will have access to various configuration files, that will let you manage container resources, set security restrictions and configure network options etc, when installing SingularityCE across the system. All of these files can be found in `/usr/local/etc/singularity` by default for installations from source (though the location may differ based on options passed during the installation). For installations from RPM or Deb packages you will find the configuration files in `/etc/singularity`. This section will describe the configuration files and the various parameters contained by them.

3.1 singularity.conf

Most of the configuration options are set using the file `singularity.conf` that defines the global configuration for SingularityCE across the entire system. Using this file, system administrators can influence the behavior of SingularityCE and restrict the functionality that users can access. As a security measure, for `setuid` installations of SingularityCE, `singularity.conf` must be owned by root and must not be writable by users or SingularityCE will refuse to run. This is not the case for non-`setuid` installations that will only ever execute with user privilege and thus do not require such limitations.

The options set via `singularity.conf` are listed below. Options are grouped together based on relevance. The actual order of options within `singularity.conf` may differ.

3.1.1 Setuid and Capabilities

`allow setuid`: To use all features of SingularityCE containers, SingularityCE will need to have access to some privileged system calls. SingularityCE achieves this by using a helper binary with the `setuid` bit enabled. The `allow-setuid` option lets you enable/disable users ability to utilize these binaries within SingularityCE. By default, it is set to “yes”, but when disabled, various SingularityCE features will not function. Please see [Unprivileged Installations](#) (page 38) for more information about running SingularityCE without `setuid` enabled.

`root default capabilities`: SingularityCE allows the specification of capabilities kept by the root user when running a container by default. Options include:

- `full`: all capabilities are maintained, this gives the same behavior as the `--keep-privs` option.
- `file`: only capabilities granted for root in `etc/singularity/capability.json` are maintained.
- `no`: no capabilities are maintained, this gives the same behavior as the `--no-privs` option.

Note: The root user can manage the capabilities granted to individual containers when they are launched through the `--add-caps` and `drop-caps` flags. Please see [Linux Capabilities](#) in the user guide for more information.

3.1.2 Loop Devices

SingularityCE uses loop devices to facilitate the mounting of container file systems from SIF and other images.

`max loop devices`: This option allows an admin to limit the total number of loop devices SingularityCE will consume at a given time.

`shared loop devices`: This allows containers running the same image to share a single loop device. This minimizes loop device usage and helps optimize kernel cache usage. Enabling this feature can be particularly useful for MPI jobs.

3.1.3 Namespace Options

`allow pid ns`: This option determines if users can leverage the PID namespace when running their containers through the `--pid` flag.

Note: Using the PID namespace can confuse the process tracking of some resource managers, as well as some MPI implementations.

3.1.4 Configuration Files

SingularityCE can automatically create or modify several system files within containers to ease usage.

Note: These options will have no effect if the file does not exist within the container, or overlay or underlay support are enabled.

`config passwd`: This option determines if SingularityCE should automatically append an entry to `/etc/passwd` for the user running the container.

`config group`: This option determines if SingularityCE should automatically append the calling user's group entries to the containers `/etc/group`.

`config resolv_conf`: This option determines if SingularityCE should automatically bind the host's `/etc/resolv.conf` within the container.

3.1.5 Session Directory and System Mounts

`sessiondir max size`: In order for the SingularityCE runtime to run a container it needs to create a temporary in-memory `sessiondir` as a location to assemble various components of the container, including mounting filesystems over the base image. In addition, this `sessiondir` will hold files that are written to the container when `--writable-tmpfs` is used, plus isolated temporary filesystems in `--contain` mode. The default value from SingularityCE 3.11 is 64MiB. If users commonly run containers with `--writable-tmpfs`, `--contain`, or in `--oci` mode, this value should be increased to accommodate their workflows. The `tmpfs` will grow to the specified maximum size, as required. Memory is not allocated ahead of usage.

`mount proc`: This option determines if SingularityCE should automatically bind mount `/proc` within the container.

`mount sys`: This option determines if SingularityCE should automatically bind mount `/sys` within the container.

`mount dev`: Should be set to "YES", if you want SingularityCE to automatically bind mount a complete `/dev` tree within the container. If set to `minimal`, then only `/dev/null`, `/dev/zero`, `/dev/random`, `/dev/urandom`, and `/dev/shm` will be included.

mount devpts: This option determines if SingularityCE will mount a new instance of `devpts` when there is a `minimal /dev` directory as explained above, or when the `--contain` option is passed.

Note: This requires either a kernel configured with `CONFIG_DEVPTS_MULTIPLE_INSTANCES=y`, or a kernel version at or newer than 4.7.

mount home: When this option is enabled, SingularityCE will automatically determine the calling user's home directory and attempt to mount it into the container.

mount tmp: When this option is enabled, SingularityCE will automatically bind mount `/tmp` and `/var/tmp` into the container from the host. If the `--contain` option is passed, SingularityCE will create both locations within the `sessiondir` or within the directory specified by the `--workdir` option if that is passed as well.

mount hostfs: This option will cause SingularityCE to probe the host for all mounted filesystems and bind those into containers at runtime.

mount slave: SingularityCE automatically mounts a handful host system directories to the container by default. This option determines if filesystem changes on the host should automatically be propagated to those directories in the container.

Note: This should be set to `yes` when `autofs` mounts occurring on the host system should be reflected up in the container.

memory fs type: This option allows admins to choose the temporary filesystem used by SingularityCE. Temporary filesystems are primarily used for system directories like `/dev` when the host system directory is not mounted within the container.

Note: For Cray CLE 5 and 6, up to CLE 6.0.UP05, there is an issue (kernel panic) when Singularity uses `tmpfs`, so on affected systems it's recommended to set this value to `ramfs` to avoid a kernel panic.

3.1.6 Bind Mount Management

bind path: This option is used to define a list of files or directories to automatically be made available when SingularityCE runs a container. In order to successfully mount listed paths the file or directory must exist within the container, or SingularityCE must be configured with either `overlay` or `underlay` support enabled.

Note: This option is ignored when containers are invoked with the `--contain` option.

You can define the a bind point where the source and destination are identical:

```
bind path = /etc/localtime
```

Or you can specify different source and destination locations using:

```
bind path = /etc/singularity/default-nsswitch.conf:/etc/nsswitch.conf
```

user bind control: This allows admins to decide if users can define bind points at runtime. By Default, this option is set to `YES`, which means users can specify bind points, `scratch` and `tmp` locations.

3.1.7 Limiting Container Execution

Warning: If unprivileged user namespace creation is allowed on a system, a user can supply and use their own unprivileged installation of Singularity or another container runtime to bypass container limits. They may also be able to use standard system tools such as `unshare`, `nsenter`, and FUSE mounts to access / execute arbitrary containers without installing any runtime.

The ‘limit container’ and ‘allow container’ directives are not effective if unprivileged user namespaces are enabled. They are only effectively applied when Singularity is running using the native runtime in `setuid` mode, and unprivileged container execution is not possible on the host.

You must disable unprivileged user namespace creation on the host if you rely on these directives to limit container execution. This will disable OCI mode, which is unprivileged and cannot enforce these limits.

There are several ways to limit container execution as an admin listed below. If stricter controls are required, check out the [Execution Control List](#) (page 30).

`limit container owners`: This restricts container execution to only allow containers that are owned by the specified user.

Note: This feature will only apply when SingularityCE is running in SUID mode and the user is non-root. By default this is set to NULL.

`limit container groups`: This restricts container execution to only allow containers that are owned by the specified group.

Note: This feature will only apply when SingularityCE is running in SUID mode and the user is non-root. By default this is set to NULL.

`limit container paths`: This restricts container execution to only allow containers that are located within the specified path prefix.

Note: This feature will only apply when SingularityCE is running in SUID mode and the user is non-root. By default this is set to NULL.

`allow container {type}`: This option allows admins to limit the types of image formats that can be leveraged by users with SingularityCE.

- `allow container sif` permits / denies execution of unencrypted SIF containers.
- `allow container encrypted` permits / denies execution of SIF containers with an encrypted root filesystem.
- `allow container squashfs` permits / denies execution of bare SquashFS image files. E.g. Singularity 2.x images.
- `allow container extfs` permits / denies execution of bare EXT image files.
- `allow container dir` permits / denies execution of sandbox directory containers.

Note: These limitations do not apply to the root user.

This behavior differs from SingularityCE versions before 3.9, where the `allow container squashfs/extfs` directives also applied to the filesystem embedded in a SIF image.

3.1.8 Disabling Kernel Filesystem Mounts

When running in `setuid` mode, SingularityCE will mount `extfs` and `squashfs` filesystems using the kernel's filesystem drivers. These mounts are performed for standalone or SIF container images, overlay images or partitions, that use `extfs` or `squashfs` formats.

Options in `singularity.conf` allow these mounts to be disabled, to e.g. work around a kernel vulnerability that cannot be patched in a timely manner. Note that disabling kernel mounts will result in a significant loss of functionality in `setuid` mode.

`allow kernel squashfs`: Defaults to yes. When set to no, SingularityCE will not mount `squashfs` filesystems using the kernel `squashfs` driver.

`allow kernel extfs`: Defaults to yes. When set to no, SingularityCE will not mount `extfs` filesystems using the kernel `extfs` driver.

3.1.9 Networking Options

The `--network` option can be used to specify a CNI networking configuration that will be used when running a container with `network virtualization`. Unrestricted use of CNI network configurations requires root privilege, as certain configurations may disrupt the host networking environment.

SingularityCE 3.8 allows specific users or groups to be granted the ability to run containers with administrator specified CNI configurations.

`allow net users`: Allow specified root administered CNI network configurations to be used by the specified list of users. By default only root may use CNI configuration, except in the case of a `fakeroot` execution where only `40_fakeroot.conflist` is used. This feature only applies when SingularityCE is running in `SUID` mode and the user is non-root.

`allow net groups`: Allow specified root administered CNI network configurations to be used by the specified list of users. By default only root may use CNI configuration, except in the case of a `fakeroot` execution where only `40_fakeroot.conflist` is used. This feature only applies when SingularityCE is running in `SUID` mode and the user is non-root.

`allow net networks`: Specify the names of CNI network configurations that may be used by users and groups listed in the `allow net users` / `allow net groups` directives. Thus feature only applies when SingularityCE is running in `SUID` mode and the user is non-root.

3.1.10 GPU Options

SingularityCE provides integration with GPUs in order to facilitate GPU based workloads seamlessly. Both options listed below are particularly useful in GPU only environments. For more information on using GPUs with SingularityCE checkout *GPU Library Configuration* (page 31).

`always use nv`: Enabling this option will cause every action command (`exec/shell/run/instance`) to be executed with the `--nv` option implicitly added.

`always use rocm`: Enabling this option will cause every action command (`exec/shell/run/instance`) to be executed with the `--rocm` option implicitly added.

3.1.11 Supplemental Filesystems

enable fusemount: This will allow users to mount fuse filesystems inside containers using the `--fusemount` flag.

enable overlay: This option will allow SingularityCE to create bind mounts at paths that do not exist within the container image. This option can be set to `try`, which will try to use an overlaysfs. If it fails to create an overlaysfs in this case the bind path will be silently ignored.

enable underlay: This option will allow SingularityCE to create bind mounts at paths that do not exist within the container image, just like `enable overlay`, but instead using an underlay. This is suitable for systems where overlay is not possible or not working. If the overlay option is available and working, it will be used instead.

3.1.12 CNI Configuration and Plugins

cni configuration path: This option allows admins to specify a custom path for the CNI configuration that SingularityCE will use for [Network Virtualization](#).

cni plugin path: This option allows admins to specify a custom path for SingularityCE to access CNI plugin executables. Check out the [Network Virtualization](#) section of the user guide for more information.

3.1.13 External Binaries

SingularityCE calls a number of external binaries for full functionality. The paths for certain critical binaries can be set in `singularity.conf`. At build time, `mconfig` will set initial values for these, by searching on the `$PATH` environment variable. You can override which external binaries are called by changing the value in `singularity.conf`.

cryptsetup path: Path to the `cryptsetup` executable, used to work with encrypted containers. Must be owned by root for security reasons.

ldconfig path: Path to the `ldconfig` executable, used to find GPU libraries. Must be owned by root for security reasons.

nvidia-container-cli path: Path to the `nvidia-container-cli` executable, used to find GPU libraries and configure the container when running with the `--nvcccli` option. Must be owned by root for security reasons.

For the following additional binaries, if the `singularity.conf` entry is left blank, then `$PATH` will be searched at runtime.

go path: Path to the `go` executable, used to compile plugins.

mksquashfs path: Path to the `mksquashfs` executable, used to create SIF and SquashFS containers.

mksquashfs procs: Allows the administrator to specify the number of CPUs that `mksquashfs` may use when building an image. The fewer processors the longer it takes. To use all available CPU's set this to 0.

mksquashfs mem: Allows the administrator to set the maximum amount of memory that `mksquashfs` may use when building an image. e.g. 1G for 1gb or 500M for 500mb. Restricting memory can have a major impact on the time it takes `mksquashfs` to create the image. NOTE: This functionality did not exist in `squashfs-tools` prior to version 4.3. If using an earlier version you should not set this.

unsquashfs path: Path to the `unsquashfs` executable, used to extract SIF and SquashFS containers.

3.1.14 Concurrent Downloads

SingularityCE 3.9 and above will pull `library://` container images using multiple concurrent downloads of parts of the image. This speeds up downloads vs using a single stream. The defaults are generally appropriate for the Sylabs Cloud, but may be tuned for your network conditions, or if you are pulling from a different library server.

`download concurrency`: specifies how many concurrent streams when downloading (pulling) an image from cloud library.

`download part size`: specifies the size of each part (bytes) when concurrent downloads are enabled.

`download buffer size`: specifies the transfer buffer size (bytes) when concurrent downloads are enabled.

3.1.15 Cgroups Options

`systemd cgroups`: specifies whether to use `systemd` to manage container cgroups. Required (with cgroups v2) for unprivileged users to apply resource limits on containers. If set to `no`, SingularityCE will directly manage cgroups via the `cgroupfs`.

3.1.16 Experimental Options

`sif fuse`: If set to `yes`, always attempt to mount a SIF image using `squashfuse` when running in unprivileged / user namespace flows. Requires `squashfuse` and `fusermount` on `$PATH`. Will fall back to extracting the SIF file on failure.

3.1.17 Updating Configuration Options

In order to manage this configuration file, SingularityCE has a `config global` command group that allows you to get, set, reset, and unset values through the CLI. It's important to note that these commands must be run with elevated privileges because the `singularity.conf` can only be modified by an administrator.

Example

In this example we will changing the `bind path` option described above. First we can see the current list of bind paths set within our system configuration:

```
$ sudo singularity config global --get "bind path"
/etc/localtime,/etc/hosts
```

Now we can add a new path and verify it was successfully added:

```
$ sudo singularity config global --set "bind path" /etc/resolv.conf
$ sudo singularity config global --get "bind path"
/etc/resolv.conf,/etc/localtime,/etc/hosts
```

From here we can remove a path with:

```
$ sudo singularity config global --unset "bind path" /etc/localtime
$ sudo singularity config global --get "bind path"
/etc/resolv.conf,/etc/hosts
```

If we want to reset the option to the default at installation, then we can reset it with:

```
$ sudo singularity config global --reset "bind path"
$ sudo singularity config global --get "bind path"
/etc/localtime,/etc/hosts
```

And now we are back to our original option settings. You can also test what a change would look like by using the `--dry-run` option in conjunction with the above commands. Instead of writing to the configuration file, it will output what would have been written to the configuration file if the command had been run without the `--dry-run` option:

```
$ sudo singularity config global --dry-run --set "bind path" /etc/resolv.conf
# SINGULARITY.CONF
# This is the global configuration file for Singularity. This file controls
[...]
# BIND PATH: [STRING]
# DEFAULT: Undefined
# Define a list of files/directories that should be made available from within
# the container. The file or directory must exist within the container on
# which to attach to. you can specify a different source and destination
# path (respectively) with a colon; otherwise source and dest are the same.
# NOTE: these are ignored if singularity is invoked with --contain.
bind path = /etc/resolv.conf
bind path = /etc/localtime
bind path = /etc/hosts
[...]
$ sudo singularity config global --get "bind path"
/etc/localtime,/etc/hosts
```

Above we can see that `/etc/resolv.conf` is listed as a bind path in the output of the `--dry-run` command, but did not affect the actual bind paths of the system.

3.2 cgroups.toml

The cgroups (control groups) functionality of the Linux kernel allows you to limit and meter the resources used by a process, or group of processes. Using cgroups you can limit memory and CPU usage. You can also rate limit block IO, network IO, and control access to device nodes.

There are two versions of cgroups in common use. Cgroups v1 sets resource limits for a process within separate hierarchies per resource class. Cgroups v2, the default in newer Linux distributions, implements a unified hierarchy, simplifying the structure of resource limits on processes.

- v1 documentation: <https://www.kernel.org/doc/Documentation/cgroup-v1/cgroups.txt>
- v2 documentation: <https://www.kernel.org/doc/Documentation/cgroup-v2.txt>

SingularityCE 3.9 and above can apply resource limitations to systems configured for both cgroups v1 and the v2 unified hierarchy. Resource limits are specified using a TOML file that represents the `resources` section of the OCI runtime-spec: <https://github.com/opencontainers/runtime-spec/blob/master/config-linux.md#control-groups>

On a cgroups v1 system the resources configuration is applied directly. On a cgroups v2 system the configuration is translated and applied to the unified hierarchy.

Under cgroups v1, access restrictions for device nodes are managed directly. Under cgroups v2, the restrictions are applied by attaching eBPF programs that implement the requested access controls.

3.2.1 Examples

To apply resource limits to a container, use the `--apply-cgroups` flag, which takes a path to a TOML file specifying the cgroups configuration to be applied:

```
$ singularity shell --apply-cgroups /path/to/cgroups.toml my_container.sif
```

Note: The `--apply-cgroups` option requires cgroups v2 to be used without root privileges.

Limiting memory

To limit the amount of memory that your container uses to 500MB (524288000 bytes), set a `limit` value inside the `[memory]` section of your cgroups TOML file:

```
[memory]
  limit = 524288000
```

Start your container, applying the toml file, e.g.:

```
$ singularity run --apply-cgroups path/to/cgroups.toml library://alpine
```

Limiting CPU

CPU usage can be limited using different strategies, with limits specified in the `[cpu]` section of the TOML file.

shares

This corresponds to a ratio versus other cgroups with cpu shares. Usually the default value is `1024`. That means if you want to allow to use 50% of a single CPU, you will set `512` as value.

```
[cpu]
  shares = 512
```

A cgroup can get more than its share of CPU if there are enough idle CPU cycles available in the system, due to the work conserving nature of the scheduler, so a contained process can consume all CPU cycles even with a ratio of 50%. The ratio is only applied when two or more processes conflicts with their needs of CPU cycles.

quota/period

You can enforce hard limits on the CPU cycles a cgroup can consume, so contained processes can't use more than the amount of CPU time set for the cgroup. `quota` allows you to configure the amount of CPU time that a cgroup can use per period. The default is 100ms (100000us). So if you want to limit amount of CPU time to 20ms during period of 100ms:

```
[cpu]
  period = 100000
  quota = 20000
```

cpus/mems

You can also restrict access to specific CPUs (cores) and associated memory nodes by using `cpus/mems` fields:

```
[cpu]
cpus = "0-1"
mems = "0-1"
```

Where container has limited access to CPU 0 and CPU 1.

Note: It's important to set identical values for both `cpus` and `mems`.

Limiting IO

To control block device I/O, applying limits to competing container, use the `[blockIO]` section of the TOML file:

```
[blockIO]
weight = 1000
leafWeight = 1000
```

`weight` and `leafWeight` accept values between 10 and 1000.

`weight` is the default weight of the group on all the devices until and unless overridden by a per device rule.

`leafWeight` relates to weight for the purpose of deciding how heavily to weigh tasks in the given cgroup while competing with the cgroup's child cgroups.

To apply limits to specific block devices, you must set configuration for specific device major/minor numbers. For example, to override `weight/leafWeight` for `/dev/loop0` and `/dev/loop1` block devices, set limits for device major 7, minor 0 and 1:

```
[blockIO]
  [[blockIO.weightDevice]]
    major = 7
    minor = 0
    weight = 100
    leafWeight = 50
  [[blockIO.weightDevice]]
    major = 7
    minor = 1
    weight = 100
    leafWeight = 50
```

You can also limit the IO read/write rate to a specific absolute value, e.g. 16MB per second for the `/dev/loop0` block device. The rate is specified in bytes per second.

```
[blockIO]
  [[blockIO.throttleReadBpsDevice]]
    major = 7
    minor = 0
    rate = 16777216
  [[blockIO.throttleWriteBpsDevice]]
    major = 7
    minor = 0
    rate = 16777216
```


Other limits

SingularityCE can apply all resource limits that are valid in the OCI runtime-spec `resources` section, including unified cgroups v2 constraints. It is most compatible, however, to use the cgroups v1 limits, which will be translated to v2 format when applied on a cgroups v2 system.

See <https://github.com/opencontainers/runtime-spec/blob/master/config-linux.md#control-groups> for information about the available limits. Note that SingularityCE uses TOML format for the configuration file, rather than JSON.

3.3 ecl.toml

The execution control list that can be used to restrict the execution of SIF files by signing key is defined here. You can authorize the containers by validating both the location of the SIF file in the filesystem and by checking against a list of signing entities.

Warning: If unprivileged user namespace creation is allowed on a system, a user can supply and use their own unprivileged installation of Singularity or another container runtime to bypass container limits. They may also be able to use standard system tools such as `unshare`, `nsenter`, and FUSE mounts to access / execute arbitrary containers without installing any runtime.

The ECL is not effective if unprivileged user namespaces are enabled. It is only effectively applied when Singularity is running using the native runtime in `setuid` mode, and unprivileged container execution is not possible on the host.

You must disable unprivileged user namespace creation on the host if you rely on the ECL limit container execution. This will disable OCI mode, which is unprivileged and cannot enforce these limits.

Warning: The ECL configuration applies to SIF container images only. To lock down execution fully you should disable execution of other container types (`squashfs/extfs/dir`) via the `singularity.conf` file `allow container settings`.

```
[[execgroup]]
  tagname = "group2"
  mode = "whitelist"
  dirpath = "/tmp/containers"
  keyfp = ["7064B1D6EFF01B1262FED3F03581D99FE87EAFD1"]
```

Only the containers running from and signed with above-mentioned path and keys will be authorized to run.

Three possible list modes you can choose from:

Whitestrict: The SIF must be signed by all of the keys mentioned.

Whitelist: As long as the SIF is signed by one or more of the keys, the container is allowed to run.

Blacklist: Only the containers whose keys are not mentioned in the group are allowed to run.

Note: The ECL checks will use the new signature format introduced in SingularityCE 3.6.0. Containers signed with older versions of Singularity SingularityCE will not pass ECL checks.

To temporarily permit the use of legacy insecure signatures, set `legacyinsecure = true` in `ecl.toml`.

3.3.1 Managing ECL public keys

Since SingularityCE 3.7.0 a global keyring is used for ECL signature verification. This keyring can be administered using the `--global` flag for the following commands:

- `singularity key import` (root user only)
- `singularity key pull` (root user only)
- `singularity key remove` (root user only)
- `singularity key export`
- `singularity key list`

Note: For security reasons, it is not possible to import private keys into this global keyring because it must be accessible by users and is stored in the file `SYSCONFDIR/singularity/global-pgp-public`.

3.4 GPU Library Configuration

When a container includes a GPU enabled application, SingularityCE (with the `--nv` or `--rocm` options) can properly inject the required Nvidia or AMD GPU driver libraries into the container, to match the host's kernel. The GPU / dev entries are provided in containers run with `--nv` or `--rocm` even if the `--contain` option is used to restrict the in-container device tree.

Compatibility between containerized CUDA/ROCm/OpenCL applications and host drivers/libraries is dependent on the versions of the GPU compute frameworks that were used to build the applications. Compatibility and usage information is discussed in the 'GPU Support' section of the [user guide](#)

3.4.1 NVIDIA GPUs / CUDA

The `nvliblist.conf` configuration file is used to specify libraries and executables that need to be injected into the container when running SingularityCE with the `--nv` Nvidia GPU support option. The provided `nvliblist.conf` is suitable for CUDA 11, but may need to be modified if you need to include additional libraries, or further libraries are added to newer versions of the Nvidia driver/CUDA distribution.

When adding new entries to `nvliblist.conf` use the bare filename of executables, and the `xxxx.so` form of libraries. Libraries are resolved via `ldconfig -p`, and executables are found by searching `$PATH`.

Experimental nvidia-container-cli Support

The `nvidia-container-cli` tool is Nvidia's officially support method for configuring containers to use a GPU. It is targeted at OCI container runtimes.

SingularityCE 3.9 introduces an experimental `--nvcccli` option, which will call out to `nvidia-container-cli` for container GPU setup, rather than use the `nvliblist.conf` approach.

To use `--nvcccli` a root-owned `nvidia-container-cli` binary must be present on the host. The binary that is run is controlled by the `nvidia-container-cli` directive in `singularity.conf`. During installation of SingularityCE, the `./mconfig` step will set the correct value in `singularity.conf` if `nvidia-container-cli` is found on the `$PATH`. If the value of `nvidia-container-cli path` is empty, SingularityCE will look for the binary on `$PATH` at runtime.

Note: To prevent use of `nvidia-container-cli` via the `--nvccli` flag, you may set `nvidia-container-cli` path to `/bin/false` in `singularity.conf`.

`nvidia-container-cli` is run as the `root` user during `setuid` operation of SingularityCE. The container starter process grants a number of Linux capabilities to `nvidia-container-cli`, which are required for it to configure the container for GPU operation. The operations performed by `nvidia-container-cli` are broadly similar to those which SingularityCE carries out when setting up a GPU container from `nvliblist.conf`.

3.4.2 AMD Radeon GPUs / ROCm

The `rocmliblist.conf` file is used to specify libraries and executables that need to be injected into the container when running SingularityCE with the `--rocm` Radeon GPU support option. The provided `rocmliblist.conf` is suitable for ROCm 4.0, but may need to be modified if you need to include additional libraries, or further libraries are added to newer versions of the ROCm distribution.

When adding new entries to `rocmlist.conf` use the bare filename of executables, and the `xxxx.so` form of libraries. Libraries are resolved via `ldconfig -p`, and executables are found by searching `$PATH`.

3.4.3 GPU liblist format

The `nvliblist.conf` and `rocmliblist` files list the basename of executables and libraries to be bound into the container, without path information.

Binaries are found by searching `$PATH`:

```
# put binaries here
# In shared environments you should ensure that permissions on these files
# exclude writing by non-privileged users.
rocm-smi
rocminfo
```

Libraries should be specified without version information, i.e. `libname.so`, and are resolved using `ldconfig`.

```
# put libs here (must end in .so)
libamd_comgr.so
libcomgr.so
libCXLActivityLogger.so
```

If you receive warnings that binaries or libraries are not found, ensure that they are in a system path (binaries), or available in paths configured in `/etc/ld.so.conf` (libraries).

3.5 capability.json

Warning: It is extremely important to recognize that **granting users Linux capabilities with the `capability` command group is usually identical to granting those users root level access on the host system.** Most if not all capabilities will allow users to “break out” of the container and become root on the host. This feature is targeted toward special use cases (like cloud-native architectures) where an admin/developer might want to limit the attack surface within a container that normally runs as root. This is not a good option in multi-tenant HPC environments

where an admin wants to grant a user special privileges within a container. For that and similar use cases, the *fakeroot feature* (page 39) is a better option.

SingularityCE provides full support for admins to grant and revoke Linux capabilities on a user or group basis. The `capability.json` file is maintained by SingularityCE in order to manage these capabilities. The `capability` command group allows you to add, drop, and list capabilities for users and groups.

For example, let us suppose that we have decided to grant a user (named `pingr`) capabilities to open raw sockets so that they can use `ping` in a container where the binary is controlled via capabilities.

To do so, we would issue a command such as this:

```
$ sudo singularity capability add --user pingr CAP_NET_RAW
```

This means the user `pingr` has just been granted permissions (through Linux capabilities) to open raw sockets within SingularityCE containers.

We can check that this change is in effect with the `capability list` command.

```
$ sudo singularity capability list --user pingr
CAP_NET_RAW
```

To take advantage of this new capability, the user `pingr` must also request the capability when executing a container with the `--add-caps` flag. `pingr` would need to run a command like this:

```
$ singularity exec --add-caps CAP_NET_RAW \
  library://sylabs/tests/ubuntu_ping:v1.0 ping -c 1 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=52 time=73.1 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 73.178/73.178/73.178/0.000 ms
```

If we decide that it is no longer necessary to allow the user `pingr` to open raw sockets within SingularityCE containers, we can revoke the appropriate Linux capability like so:

```
$ sudo singularity capability drop --user pingr CAP_NET_RAW
```

Now if `pingr` tries to use `CAP_NET_RAW`, SingularityCE will not give the capability to the container and `ping` will fail to create a socket:

```
$ singularity exec --add-caps CAP_NET_RAW \
  library://sylabs/tests/ubuntu_ping:v1.0 ping -c 1 8.8.8.8
WARNING: not authorized to add capability: CAP_NET_RAW
ping: socket: Operation not permitted
```

The `capability add` and `drop` subcommands will also accept the case insensitive keyword `all` to grant or revoke all Linux capabilities to a user or group.

For more information about individual Linux capabilities check out the [man pages](#) or use the `capability avail` command to output available capabilities with a description of their behaviors.

3.6 seccomp-profiles

Secure Computing (seccomp) Mode is a feature of the Linux kernel that allows an administrator to filter system calls being made from a container. Profiles made up of allowed and restricted calls can be passed to different containers. *Seccomp* provides more control than *capabilities* alone, giving a smaller attack surface for an attacker to work from within a container.

You can set the default action with `defaultAction` for a non-listed system call. Example: `SCMP_ACT_ALLOW` filter will allow all the system calls if it matches the filter rule and you can set it to `SCMP_ACT_ERRNO` which will have the thread receive a return value of *errno* if it calls a system call that matches the filter rule. The file is formatted in a way that it can take a list of additional system calls for different architecture and SingularityCE will automatically take syscalls related to the current architecture where it's been executed. The `include/exclude-> caps` section will include/exclude the listed system calls if the user has the associated capability.

Use the `--security` option to invoke the container like:

```
$ sudo singularity shell --security seccomp:/home/david/my.json my_container.sif
```

For more insight into security options, network options, cgroups, capabilities, etc, please check the [Userdocs](#) and it's [Appendix](#).

3.7 remote.yaml

System-wide remote endpoints are defined in a configuration file typically located at `/usr/local/etc/singularity/remote.yaml` (this location may vary depending on installation parameters) and can be managed by administrators with the `remote` command group.

3.7.1 Remote Endpoints

Sylabs introduced the online [Sylabs Cloud](#) to enable users to [Create](#), [Secure](#), and [Share](#) their container images with others.

SingularityCE allows users to login to an account on the Sylabs Cloud, or configure SingularityCE to use an API compatible container service such as a local installation of SingularityCE Enterprise, which provides an on-premise private Container Library, Remote Builder and Key Store.

Note: A fresh installation of SingularityCE is automatically configured to connect to the public [Sylabs Cloud](#) services.

Examples

Use the `remote` command group with the `--global` flag to create a system-wide remote endpoint:

```
$ sudo singularity remote add --global company-remote https://enterprise.example.com
INFO: Remote "company-remote" added.
INFO: Global option detected. Will not automatically log into remote.
```

Conversely, to remove a system-wide endpoint:

```
$ sudo singularity remote remove --global company-remote
INFO: Remote "company-remote" removed.
```

Note: Once users log in to a system wide endpoint, a copy of the endpoint will be listed in a their `~/ .singularity/remote.yaml` file. This means modifications or removal of the system-wide endpoint will not be reflected in the users configuration unless they remove the endpoint themselves.

Exclusive Endpoint

SingularityCE 3.7 introduces the ability for an administrator to make a remote the only usable remote for the system by using the `--exclusive` flag:

```
$ sudo singularity remote use --exclusive company-remote
INFO: Remote "company-remote" now in use.
$ singularity remote list
Cloud Services Endpoints
=====
```

NAME	URI	ACTIVE	GLOBAL	EXCLUSIVE	INSECURE
SylabsCloud	cloud.sylabs.io	NO	YES	NO	NO
company-remote	enterprise.example.com	YES	YES	YES	NO
myremote	enterprise.example.com	NO	NO	NO	NO

```

Keyservers
=====

URI                GLOBAL  INSECURE  ORDER
https://keys.example.com  YES    NO        1*

* Active cloud services keyserver
```

Insecure (HTTP) Endpoints

From SingularityCE 3.9, if you are using a endpoint that exposes its service discovery file over an insecure HTTP connection only, it can be added by specifying the `--insecure` flag:

```
$ sudo singularity remote add --global --insecure test http://test.example.com
INFO: Remote "test" added.
INFO: Global option detected. Will not automatically log into remote.
```

This flag controls HTTP vs HTTPS for service discovery only. The protocol used to access individual library, build and keyserver URLs is set by the service discovery file.

Additional Information

For more details on the remote command group and managing remote endpoints, please check the [Remote Userdocs](#).

3.7.2 Keyserver Configuration

By default, SingularityCE will use the keyserver correlated to the active cloud service endpoint. This behavior can be changed or supplemented via the `add-keyserver` and `remove-keyserver` commands. These commands allow an administrator to create a global list of key servers used to verify container signatures by default.

For more details on the `remote` command group and managing keyservers, please check the [Remote Userdocs](#).

USER NAMESPACES & FAKEROOT

User namespaces are an isolation feature that allow processes to run with different user identifiers and/or privileges inside that namespace than are permitted outside. A user may have a `uid` of `1001` on a system outside of a user namespace, but run programs with a different `uid` with different privileges inside the namespace.

User namespaces are used with containers to make it possible to set up a container without privileged operations, and so that a normal user can act as root inside a container to perform administrative tasks, without being root on the host outside.

SingularityCE uses user namespaces in the following situations:

- When the `setuid` workflow is disabled or SingularityCE was installed without root.
- When a container is run with the `--userns` option.
- When `--fakeroot` is used to impersonate a root user when building or running a container.
- When the `-oci` runtime mode is used.

4.1 User Namespace Requirements

To allow unprivileged creation of user namespaces a kernel ≥ 3.8 is required, with ≥ 3.18 being recommended due to security fixes for user namespaces.

To use a persistent overlay directory with `--overlay` when running unprivileged, a kernel ≥ 5.11 is required.

Additionally, some Linux distributions require that unprivileged user namespace creation is enabled using a `sysctl` or kernel command line parameter. Please consult your distribution documentation or vendor to confirm the steps necessary to 'enable unprivileged user namespace creation'.

4.1.1 Debian

```
sudo sh -c 'echo kernel.unprivileged_userns_clone=1 \  
>/etc/sysctl.d/90-unprivileged_userns.conf'  
sudo sysctl -p /etc/sysctl.d /etc/sysctl.d/90-unprivileged_userns.conf
```


4.1.2 RHEL/CentOS 7

From 7.4, kernel support is included but must be enabled with:

```
sudo sh -c 'echo user.max_user_namespaces=15000 \
  >/etc/sysctl.d/90-max_net_namespaces.conf'
sudo sysctl -p /etc/sysctl.d /etc/sysctl.d/90-max_net_namespaces.conf
```

4.2 Unprivileged Installations

As detailed in the *non-setuid installation* (page 14) section, SingularityCE can be compiled or configured with the `allow_setuid = no` option in `singularity.conf` to not perform privileged operations using the `starter-setuid` binary.

When SingularityCE does not use `setuid` all container execution will use a user namespace. In this mode of operation, some features are not available, and there are impacts to the security/integrity guarantees when running SIF container images:

- Unless using the experimental `--sif-fuse` function, all containers must be run from sandbox directories. SIF images are extracted to a sandbox directory on the fly, preventing verification at runtime, and potentially allowing external modification of the container at runtime.
- Filesystem image, and SIF-embedded persistent overlays cannot be used. Directory overlays require kernel ≥ 5.11 .
- Encrypted containers cannot be used. SingularityCE mounts encrypted containers directly through the kernel, so that encrypted content is not extracted to disk. This requires the `setuid` workflow.
- Fakedroot functionality will rely on external `setuid` root `newuidmap` and `newgidmap` binaries which may be provided by the distribution.

4.3 `--usersns` option

The `--usersns` option to `singularity run/exec/shell` will start a container using a user namespace, avoiding the `setuid` privileged workflow for container setup even if SingularityCE was compiled and configured to use `setuid` by default.

The same limitations apply as in an unprivileged installation.

4.4 `--sif-fuse` option

If `squashfuse` $\geq 0.1.100$ is installed on the system, and available on the `$PATH`, SingularityCE can use it to mount a container filesystem from a SIF file in unprivileged / user namespace flows.

This is an experimental feature, and does not currently apply to `--fakedroot`, or support additional overlays etc.

To always attempt a `squashfuse` based mount, set `sif fuse = yes` in `singularity.conf`. Otherwise, use the `--sif-fuse` flag. E.g.:

```
$ singularity run -u --sif-fuse ~/ubuntu_latest.sif
INFO:   Mounting SIF with FUSE...
Singularity>
```

(continues on next page)

(continued from previous page)

```
exit
INFO:   Unmounting SIF with FUSE...
INFO:   Removing image tempDir /tmp/rootfs-33363059
```

SingularityCE will call `squashfuse` before container startup, to mount the container filesystem read-only, from the SIF to a temporary location. When the container exits, `fusermount` is used to unmount the SIF filesystem, and the temporary location is cleaned up.

If a `squashfuse` mount cannot be performed successfully, SingularityCE will fall back to extracting the SIF image to a temporary directory, with a warning as this extraction is performed.

4.5 Fakeroot feature

Fakeroot (or commonly referred as rootless mode) allows an unprivileged user to run a container as a “**fake root**” user by leveraging user namespaces with [user namespace UID/GID mapping](#).

User namespace UID/GID mapping allows a user to act as a different UID/GID in the container than they are on the host. A user can access a configured range of UIDs/GIDs in the container, which map back to (generally) unprivileged user UIDs/GIDs on the host. This allows a user to be `root` (`uid 0`) in a container, install packages etc., but have no privilege on the host.

4.5.1 Requirements

In addition to user namespace support, SingularityCE must manipulate `subuid` and `subgid` maps for the user namespace it creates. By default this happens transparently in the `setuid` workflow. With unprivileged installations of SingularityCE or where `allow setuid = no` is set in `singularity.conf`, SingularityCE attempts to use external `setuid` binaries `newuidmap` and `newgidmap`, so you need to install those binaries on your system.

4.5.2 Basics

Fakeroot relies on `/etc/subuid` and `/etc/subgid` files to find configured mappings from real user and group IDs, to a range of otherwise vacant IDs for each user on the host system that can be remapped in the user namespace. A user must have an entry in these system configuration files to use the fakeroot feature. SingularityCE provides a `config fakeroot` (page 41) command to assist in managing these files, but it is important to understand how they work.

For user `foo` an entry in `/etc/subuid` might be:

```
foo:100000:65536
```

where `foo` is the username, `100000` is the start of the UID range that can be used by `foo` in a user namespace uid mapping, and `65536` number of UIDs available for mapping.

Same for `/etc/subgid`:

```
foo:100000:65536
```

Note: Some distributions add users to these files on installation, or when `useradd`, `adduser`, etc. utilities are used to manage local users.

The glibc nss name service switch mechanism does not currently support managing subuid and subgid mappings with external directory services such as LDAP. You must manage or provision mapping files direct to systems where fakeroot will be used.

Warning: SingularityCE requires that a range of at least 65536 IDs is used for each mapping. Larger ranges may be defined without error.

It is also important to ensure that the subuid and subgid ranges defined in these files don't overlap with each other, or any real UIDs and GIDs on the host system.

So if you want to add another user bar, /etc/subuid and /etc/subgid will look like:

```
foo:100000:65536
bar:165536:65536
```

Resulting in the following allocation:

User	Host UID	Sub UID/GID range
foo	1000	100000 to 165535
bar	1001	165536 to 231071

Inside a user namespace / container, foo and bar can now act as any UID/GID between 0 and 65536, but these UIDs are confined to the container. For foo UID 0 in the container will map to the host foo UID 1000 and 1 to 65536 will map to 100000-165535 outside of the container etc. This impacts the ownership of files, which will have different IDs inside and outside of the container.

Note: If you are managing large numbers of fakeroot mappings you may wish to specify users by UID rather than username in the /etc/subuid and /etc/subgid files. The man page for subuid advises:

“When large number of entries (10000-100000 or more) are defined in /etc/subuid, parsing performance penalty will become noticeable. In this case it is recommended to use UIDs instead of login names. Benchmarks have shown speed-ups up to 20x.”

4.5.3 Filesystem considerations

Based on the above range, here we can see what happens when the user foo create files with --fakeroot feature:

Create file with container UID	Created host file owned by UID
0 (default)	1000
1 (daemon)	100000
2 (bin)	100001

Outside of the fakeroot container the user may not be able to remove directories and files created with a subuid, as they do not match with the user's UID on the host. The user can remove these files by using a container shell running with fakeroot.

4.5.4 Network configuration

With fakeroot, users can request a container network named `fakeroot`, other networks are restricted and can only be used by the real host root user. By default the `fakeroot` network is configured to use a network veth pair.

Warning: Do not change the `fakeroot` network type in `etc/singularity/network/40_fakeroot.conflict` without considering the security implications.

Note: Unprivileged installations of SingularityCE cannot use `fakeroot` network as it requires privilege during container creation to set up the network.

4.5.5 Configuration with `config fakeroot`

SingularityCE 3.5 and above provides a `config fakeroot` command that can be used by a root user to administer local system `/etc/subuid` and `/etc/subgid` files in a simple manner. This allows users to be granted the ability to use Singularity's `fakeroot` functionality without editing the files manually. The `config fakeroot` command will automatically ensure that generated subuid/subgid ranges are an appropriate size, and do not overlap.

`config fakeroot` must be run as the root user, or via `sudo singularity config fakeroot` as the `/etc/subuid` and `/etc/subgid` files form part of the system configuration, and are security sensitive. You may `--add` or `--remove` user subuid/subgid mappings. You can also `--enable` or `--disable` existing mappings.

Note: If you deploy SingularityCE to a cluster you will need to make arrangements to synchronize `/etc/subuid` and `/etc/subgid` mapping files to all nodes.

At this time, the `glibc` name service switch functionality does not support subuid or subgid mappings, so they cannot be defined in a central directory such as LDAP.

Adding a `fakeroot` mapping

Use the `-a/--add <user>` option to `config fakeroot` to create new mapping entries so that `<user>` can use the `fakeroot` feature of Singularity:

```
$ sudo singularity config fakeroot --add dave
# Show generated `/etc/subuid`
$ cat /etc/subuid
1000:4294836224:65536
# Show generated `/etc/subgid`
$ cat /etc/subgid
1000:4294836224:65536
```

The first subuid range will be set to the top of the 32-bit UID space. Subsequent subuid ranges for additional users will be created working down from this value. This minimizes the change of overlap with real UIDs on most systems.

Note: The `config fakeroot` command generates mappings specified using the user's uid, rather than their username. This is the preferred format for faster lookups when configuring a large number of mappings, and the command can be

used to manipulate these by username.

Deleting, disabling, enabling mappings

Use the `-r/--remove <user>` option to `config fakeroot` to completely remove mapping entries. The `<user>` will no longer be able to use the fakeroot feature of Singularity:

```
$ sudo singularity config fakeroot --remove dave
```

Warning: If a fakeroot mapping is removed, the subuid/subgid range may be assigned to another user via `--add`. Any remaining files from the prior user that were created with this mapping will be accessible to the new user via fakeroot.

The `-d/--disable` and `-e/--enable` options will comment and uncomment entries in the mapping files, to temporarily disable and subsequently re-enable fakeroot functionality for a user. This can be useful to disable fakeroot for a user, but ensure the subuid/subgid range assigned to them is reserved, and not re-assigned to a different user.

```
# Disable dave
$ sudo singularity config fakeroot --disable dave

# Entry is commented
$ cat /etc/subuid
!1000:4294836224:65536

# Enable dave
$ sudo singularity config fakeroot --enable dave

# Entry is active
$ cat /etc/subuid
1000:4294836224:65536
```

4.6 Unprivileged Builds Without User Namespaces

Where local container builds need to be performed unprivileged, but user namespaces and / or subuid mapping cannot be enabled, limited support is provided via the use of `proot`. This functionality was introduced in {SingularityCE} 3.11.

`proot` is an optional dependency of SingularityCE that can be installed from community distribution repositories, or a static binary available from proot-me.github.io. The `proot` executable should be on the `PATH` in order for SingularityCE to use it.

When `singularity build` is run against a definition file by a non-root user, and without the `--fakeroot` option, SingularityCE will search the `PATH` for `proot`. If it is found, then the `%post` section of the build will run as an emulated root user. Commands run as the user who invoked `singularity build`, but `proot` will intercept system calls, so that the commands appear to be running as root.

Unprivileged builds with `proot` have limitations, as the emulation of the root user is not complete. These builds:

- Do not support `arch` / `debootstrap` / `yum` / `zypper` bootstraps. Use `localimage`, `library`, `oras`, or one of the `docker/oci` sources.

- Do not support %pre and %setup sections.
- Run the %post sections of a build in the container as an emulated root user.
- Run the %test section of a build as the non-root user, like singularity test.
- Are subject to any restrictions imposed in singularity.conf.
- Incur a performance penalty due to proot's ptrace based interception of syscalls.
- May fail if the %post script requires privileged operations that proot cannot emulate.

SECURITY IN SINGULARITYCE

5.1 Security Policy

If you suspect you have found a vulnerability in SingularityCE we want to work with you so that it can be investigated, fixed, and disclosed in a responsible manner. Please follow the steps in our published [Security Policy](#), which begins with contacting us privately via security@sylabs.io

Sylabs discloses vulnerabilities found in SingularityCE through public CVE reports, and notifications on our community channels. We encourage all users to monitor new releases of SingularityCE for security information. Security patches are applied to the latest open-source release.

SingularityPRO is a professionally curated and licensed version of SingularityCE that provides added security, stability, and support beyond that offered by the open source project. Security and bug-fix patches are backported to select versions of SingularityPRO, so that they can be deployed long-term where required. PRO users receive security fixes as detailed in the [Sylabs Security Policy](#).

5.2 Background

SingularityCE grew out of the need to implement a container platform that was suitable for use on shared systems, such as HPC clusters. In these environments multiple people access a shared resource. User accounts, groups, and standard file permissions limit their access to data, devices, and prevent them from disrupting or accessing others' work.

To provide security in these environments a container needs to run as the user who starts it on the system. Before the widespread adoption of the Linux user namespace, only a privileged user could perform the operations which are needed to run a container. A default Docker installation uses a root-owned daemon to start containers. Users can request that the daemon starts a container on their behalf. However, coordinating a daemon with other schedulers is difficult and, since the daemon is privileged, users can ask it to carry out actions that they wouldn't normally have permission to do.

When a user runs a container with SingularityCE, it is started as a normal process running under the user's account. Standard file permissions and other security controls based on user accounts, groups, and processes apply. In a default installation SingularityCE uses a `setuid` starter binary to perform only the specific tasks needed to setup the container.

5.3 Setuid & User Namespaces

Using a setuid binary to run container setup operations is essential to support containers on older Linux distributions, such as CentOS 6, that were previously common in HPC and enterprise. Newer distributions have support for ‘unprivileged user namespace creation’. This means a normal user can create a user namespace, in which most setup operations needed to run a container can be run unprivileged.

5.3.1 Security Implications of Unprivileged User Namespaces

Warning: If you rely on the ECL or other container execution limits, you must disable unprivileged user namespace creation on your systems.

When unprivileged user namespace creation is allowed on a system, a user can supply and use their own unprivileged installation of Singularity or another container runtime. They may also be able to use standard system tools such as `unshare`, `nsenter`, and FUSE mounts to access / execute arbitrary containers without installing any runtime. Both of these approaches will allow users to bypass any restrictions that have been set in a system-wide installation of SingularityCE. These include:

- The `allow container` and `limit container` directives in `singularity.conf`.
- The Execution Control List, which restricts execution of SIF container images via signature checks.

Note also that SingularityCE’s `-oci` mode is an unprivileged runtime that requires unprivileged user namespace creation. It does not implement the container restrictions that cannot be effectively enforced when unprivileged user namespaces are available.

If your primary security concern is that of restricting the containers which users can execute, you should use singularity in setuid mode, and ensure unprivileged user namespace creation is disabled on the host.

5.3.2 Configuration and Limitations of User Namespace Mode

SingularityCE supports running containers without setuid, using user namespaces. It can be compiled with the `--without-setuid` option, or `allow setuid = no` can be set in `singularity.conf` to enable this. In this mode *all* operations run as the user who starts the `singularity` program. However, there are some disadvantages:

- SIF and other single file container images cannot be mounted directly, except via the experimental `--sif-fuse` feature. The container image must be extracted to a directory on disk to run. This impact the speed of execution. Workloads accessing large numbers of small files (such as python application startup) do not benefit from the reduced metadata load on the filesystem an image file provides.
- Replacing direct kernel mounts with the experimental `--sif-fuse` FUSE mount approach can cause a significant reduction in performance.
- The effectiveness of signing and verifying container images is reduced as, when extracted to a directory, modification is possible and verification of the image’s original signature cannot be performed.
- Encryption is not supported. SingularityCE leverages kernel LUKS2 mounts to run encrypted containers without decrypting their content to disk.
- Some sites hold the opinion that vulnerabilities in kernel user namespace code could have greater impact than vulnerabilities confined to a single piece of setuid software. Therefore they are reluctant to enable unprivileged user namespace creation.
- Limitations on container execution by location, valid signatures, user/group cannot be enforced.

Because of the points above, the default mode of operation of SingularityCE uses a `setuid` binary. Sylabs aims to reduce the circumstances that require this as new functionality is developed and reaches commonly deployed Linux distributions.

5.4 Runtime & User Privilege Model

While other runtimes have aimed to safely sandbox containers executing as the `root` user, so that they cannot affect the host system, SingularityCE has adopted an alternative security model:

- Containers should be run as an unprivileged user.
- The user should never be able to elevate their privileges inside the container to gain control over the host.
- All permission restrictions on the user outside of a container should apply inside the container.
- Favor integration over isolation. Allow a user to use host resources such as GPUs, network file systems, high speed interconnects easily. The process ID space, network etc. are not isolated in separate namespaces by default.

To accomplish this, SingularityCE uses a number of Linux kernel features. The container file system is mounted using the `nosuid` option, and processes are started with the `PR_NO_NEW_PRIVS` flag set. This means that even if you run `sudo` inside your container, you won't be able to change to another user, or gain root privileges by other means.

If you do require the additional isolation of the network, devices, PIDs etc. provided by other runtimes, SingularityCE can make use of additional namespaces and functionality such as `seccomp` and `cgroups`.

5.5 Singularity Image Format (SIF)

SingularityCE uses SIF as its default container format. A SIF container is a single file, which makes it easy to manage and distribute. Inside the SIF file, the container filesystem is held in a SquashFS object. By default, we mount the container filesystem directly using SquashFS. On a network filesystem this means that reads from the container are data-only. Metadata operations happen locally, speeding up workloads with many small files.

Holding the container image in a single file also enable unique security features. The container filesystem is immutable, and can be signed. The signature travels in the SIF image itself so that it is always possible to verify that the image has not been tampered with or corrupted.

We use private PGP keys to create a container signature, and the public key in order to verify the container. Verification of signed containers happens automatically in `singularity pull` commands against the Sylabs Cloud Container Library. A Keystore in the Sylabs Cloud makes it easier to share and obtain public keys for container verification.

A container may be signed once, by a trusted individual who approves its use. It could also be signed with multiple keys to signify it has passed each step in a CI/CD QA & Security process. In `setuid` mode, SingularityCE can be configured with an execution control list (ECL). The ECL requires the presence of one or more valid signatures, to limit execution to approved containers on systems that have unprivileged user namespace creation disabled.

In SingularityCE 3.4 and above, the root filesystem of a container (stored in the squashFS partition of SIF) can be encrypted. As a result, everything inside the container becomes inaccessible without the correct key or passphrase. The content of the container is private, even if the SIF file is shared in public.

Encryption and decryption are performed using the Linux kernel's LUKS2 feature. This is the same technology routinely used for full disk encryption. The encrypted container is mounted directly through the kernel. Unlike other container formats, an encrypted container is not decrypted to disk in order to run it.

5.6 Plugins

As discussed in the SingularityCE User Guide, [plugins](#) provide a way to augment Singularity with additional functionality. Before using the `singularity plugin compile` or `singularity plugin install` commands to compile or add a new plugin to your SingularityCE installation, make sure that you trust the origin of the plugin, and that you are certain it does not contain any malicious code.

For further information on verifying the contents of SIF files using cryptographic signatures, see the “[Sign and Verify](#)” section of the SingularityCE User Guide.

5.7 Configuration & Runtime Options

System administrators who manage SingularityCE can use configuration files to set security restrictions, grant or revoke a user’s capabilities, manage resources and authorize containers etc.

Configuration files and their parameters are *documented for administrators here* (page 20).

When running a container as root, SingularityCE can apply hardening rules using `seccomp` and `apparmor`. See the ‘Security Options’ section of the user guide.

Limits on resource usage by containers can be enforced using `cgroups`. On systems that use `cgroups v1`, only the root user can set resource limits. On systems that use `cgroups v2` and `systemd`, all users can apply resource limits as long as the system is configured for delegation.

By default, EL9, Ubuntu 22.04, Debian 11, Fedora 31 and newer use `cgroups v2` and are configured for delegation so that unprivileged users will be able to use the `--apply-cgroups` and other resource limit flags of SingularityCE without further configuration.

On EL8 and Ubuntu 20.04 it is possible to setup a compatible configuration by following the ‘Enabling `cgroup v2`’ and ‘Enabling CPU, CPuset, and I/O delegation’ steps at the [rootless containers website](#)

See the ‘Limiting Container Resources’ section of the user guide for more details of how to apply `cgroups` limits to containers at runtime.

INSTALLED FILES

An installation of SingularityCE 3.11.0, performed as root via `sudo make install` consists of the following files, with ownership and permissions required to use the *setuid* workflow:

```
# Main executables
bin root:root 755 (drwxr-xr-x)
bin/singularity root:root 755 (-rwxr-xr-x)
bin/run-singularity root:root 755 (-rwxr-xr-x)

# Configuration files
etc root:root 755 (drwxr-xr-x)
etc/bash_completion.d root:root 755 (drwxr-xr-x)
etc/bash_completion.d/singularity root:root 644 (-rw-r--r--)
etc/singularity root:root 755 (drwxr-xr-x)
etc/singularity/singularity.conf root:root 644 (-rw-r--r--)
etc/singularity/remote.yaml root:root 644 (-rw-r--r--)
etc/singularity/network root:root 755 (drwxr-xr-x)
etc/singularity/network/00_bridge.conflist root:root 644 (-rw-r--r--)
etc/singularity/network/10_ptp.conflist root:root 644 (-rw-r--r--)
etc/singularity/network/20_ipvlan.conflist root:root 644 (-rw-r--r--)
etc/singularity/network/30_macvlan.conflist root:root 644 (-rw-r--r--)
etc/singularity/network/40_fakeroot.conflist root:root 644 (-rw-r--r--)
etc/singularity/capability.json root:root 644 (-rw-r--r--)
etc/singularity/ecl.toml root:root 644 (-rw-r--r--)
etc/singularity/seccomp-profiles root:root 755 (drwxr-xr-x)
etc/singularity/seccomp-profiles/default.json root:root 644 (-rw-r--r--)
etc/singularity/nvliblist.conf root:root 644 (-rw-r--r--)
etc/singularity/rocmlliblist.conf root:root 644 (-rw-r--r--)
etc/singularity/cgroups root:root 755 (drwxr-xr-x)
etc/singularity/cgroups/cgroups.toml root:root 644 (-rw-r--r--)
etc/singularity/global-pgp-public root:root 644 (-rw-r--r--)

# Runtime executables
libexec root:root 755 (drwxr-xr-x)
libexec/singularity root:root 755 (drwxr-xr-x)
libexec/singularity/bin root:root 755 (drwxr-xr-x)
libexec/singularity/bin/common root:root 755 (-rwxr-xr-x)
libexec/singularity/bin/starter root:root 755 (-rwxr-xr-x)
libexec/singularity/bin/starter-suid root:root 4755 (-rwsr-xr-x)

# CNI network plugins
```

(continues on next page)

(continued from previous page)

```

libexec/singularity/cni root:root 755 (drwxr-xr-x)
libexec/singularity/cni/dhcp root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/host-local root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/static root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/bridge root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/host-device root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/ipvlan root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/loopback root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/macvlan root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/ptp root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/vlan root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/bandwidth root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/firewall root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/portmap root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/sbr root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/tuning root:root 755 (-rwxr-xr-x)
libexec/singularity/cni/vrf root:root 755 (-rwxr-xr-x)

# Documentation (man pages)
share root:root 755 (drwxr-xr-x)
share/man root:root 755 (drwxr-xr-x)
share/man/man1 root:root 755 (drwxr-xr-x)
share/man/man1/singularity.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-build.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-cache.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-cache-clean.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-cache-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-capability.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-capability-add.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-capability-avail.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-capability-drop.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-capability-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-config.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-config-fakeroot.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-config-global.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-delete.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-exec.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-inspect.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-instance.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-instance-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-instance-start.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-instance-stop.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-export.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-import.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-newpair.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-pull.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-push.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-remove.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-key-search.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci.1 root:root 644 (-rw-r--r--)

```

(continues on next page)

(continued from previous page)

```

share/man/man1/singularity-oci-attach.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-create.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-delete.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-exec.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-kill.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-mount.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-pause.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-resume.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-run.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-start.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-state.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-umount.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-oci-update.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-overlay.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-overlay-create.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-compile.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-create.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-disable.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-enable.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-inspect.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-install.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-plugin-uninstall.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-pull.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-push.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-add.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-add-keyserver.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-login.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-logout.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-remove.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-remove-keyserver.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-status.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-remote-use.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-run.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-run-help.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-search.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-shell.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-add.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-del.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-dump.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-header.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-info.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-list.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-new.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sif-setprim.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-sign.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-test.1 root:root 644 (-rw-r--r--)
share/man/man1/singularity-verify.1 root:root 644 (-rw-r--r--)

```

(continues on next page)

(continued from previous page)

```
share/man/man1/singularity-version.1 root:root 644 (-rw-r--r--)
```

```
# Container state directories
```

```
var root:root 755 (drwxr-xr-x)
```

```
var/singularity root:root 755 (drwxr-xr-x)
```

```
var/singularity/mnt root:root 755 (drwxr-xr-x)
```

```
var/singularity/mnt/session root:root 755 (drwxr-xr-x)
```

LICENSE

This documentation is subject to the following 3-clause BSD license:

Copyright (c) 2017, SingularityWare, LLC. All rights reserved.
Copyright (c) 2018-2023, Sylabs, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.