AMBIENT-IT

+33 (0)1 48 10 09 46
formation@ambient-it.net
140 Avenue Jean Lolive 93500 Pantin - France

Updated on 04/22/2026

Sign up

# Apptainer (Singularity) Training

2 days (14 hours)

## Overview

Apptainer allows you to run containerized applications in a reproducible manner on Linux servers and HPC environments, without relying on a root daemon. Ideal for packaging scientific pipelines, securing multi-user executions, and facilitating deployment on clusters.

This training aims to make your workflows portable: image creation, execution on compute nodes, dependency and data management. You will learn how to choose between SIF and sandbox, use existing images, and build your own recipes.

The approach is 100% hands-on: guided workshops, build and execution demos, and troubleshooting common errors (permissions, mounts, network). Deliverables include ready-to-use definition files, a best-practices checklist, and sample commands for integrating Apptainer into your scripts.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

## Objectives

- Install and configure Apptainer on a Linux workstation or server.
- Run containers and manage volumes, variables, and working directories.
- Build an image using a definition file and version the recipe.
- Publish, sign, and verify images to secure the execution chain.
- Integrate Apptainer into batch jobs (Slurm) and automate pipelines.

## Target Audience

- DevOps / SRE Engineers
- Data scientists and ML engineers
- HPC users (research, scientific computing)
- Linux system administrators

# Prerequisites

- Proficiency with the Linux terminal and permissions (users, groups, chmod)
- Basic understanding of containerization (images, runtime, volumes)
- Basics of Bash scripting
- Understanding of Python/R environments and dependencies

# Technical prerequisites

- 64-bit Linux (Ubuntu/Debian/RHEL/Fedora), sudo access if installing locally
- Minimum 8 GB of RAM, 16 GB recommended for builds
- Internet access to download images (optional if using an internal registry)
- Apptainer, code editor, Git, access to a scheduler (Slurm) if in an HPC environment

# Apptainer training program

## [Day 1 - Morning]

### Apptainer fundamentals and getting started

- Context: HPC containers and constraints (multi-user, security, performance)
- Key differences vs. Docker
- Installation and checks
- Lifecycle: pull, build, inspect, run, exec, shell
- Hands-on workshop: Run an SIF container, inspect its contents, and execute an application command.

## [Day 1 - Afternoon]

### Building images with definition files

- Structure of a definition file
- Choosing a base: library, Docker/OCI, local image, reproducibility constraints
- Dependency management: system packages, Python/pip, environment variables, cleanup
- Best practices: fixed versions, image size, build logs, naming conventions
- Hands-on workshop: Write a definition file and build an executable SIF image (runscript) for a CLI tool.

[Day 2 - Morning]

### Advanced execution: I/O, networking, GPU, and HPC integration

- Mounts and file access: bind, working directories, path and permission management
- Isolation and execution options: namespaces, contain mode, variables, and environment propagation
- Hardware access: host libraries, compatibility best practices
- Batch integration: execution via Slurm (srun/sbatch), resource management, and logs
- Hands-on workshop: Launching a Slurm job that runs a container with data bind and result output.

[Day 2 - Afternoon]

### Security, deployment, and industrialization

- Security model: unprivileged execution, common risks, hardening options
- Signature and trust: sign, verify, key management, and internal policies
- Registries and sharing: Apptainer Library, OCI registries, versioning strategy, and traceability
- Automation: Reproducible CI builds, smoke tests, publishing SIF artifacts
- Hands-on workshop: Set up an image build/test pipeline and publish a signed version.

## Target Audience

This training is intended for both individuals and companies, large or small, wishing to train their teams in a new advanced IT technology or to acquire specific business knowledge or modern methods.

## Assessment upon enrollment

The pre-training assessment complies with Qualiopi quality standards. Upon final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency in various types of technologies, as well as their expectations and personal goals regarding the upcoming training, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could pose challenges for monitoring and ensuring the smooth running of the training session.

## Teaching Methods

Practical Course: 60% Practical, 40% Theory. Training materials distributed in digital format to all participants.

## Organization

The course alternates between theoretical input from the trainer, supported by examples and

reflection sessions and group work.

## Certification

At the end of the session, a multiple-choice quiz is used to verify that the skills have been properly acquired.

## Certification

A certificate will be issued to each trainee who has completed the entire training program.