AMBIENT-IT

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

Updated on 08/26/2026

Sign Up

Longhorn.io Training

3 days (21 hours)

Overview

Longhorn.io is a distributed block storage solution for Kubernetes, designed to provide resilient persistent volumes for containerized applications. With synchronous replication, snapshots, remote backups, and its management interface, Longhorn.io enables you to build a reliable and operational cloud-native storage platform.

Our Longhorn.io training will enable you to master the essential mechanisms of persistent storage in Kubernetes and manage volumes tailored to stateful applications.

You'll learn how to deploy and configure Longhorn.io, prepare nodes and disks, and use StorageClasses, PVCs, and replication settings according to availability and performance requirements.

By the end of the training, you will be able to provision volumes, manage snapshots, backups, and restores, and apply best practices for capacity planning, scheduling, and business continuity.

This training also covers monitoring, incident diagnosis, replica evacuation, node maintenance, and the architectural choices required to operate Longhorn.io in a production Kubernetes environment.

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

Objectives

- Understand the architecture and components of Longhorn.io in a Kubernetes cluster.

- Deploy and configure Longhorn.io using Helm.
- Provision and manage persistent volumes using StorageClasses and PVCs.
- Configure replication, replica scheduling, and capacity management.
- Implement snapshots, recurring backups, and restores.
- Diagnose storage issues and perform routine maintenance tasks.

Target Audience

- Kubernetes Administrator
- DevOps Engineer
- Cloud Engineer
- Linux System Administrator
- Infrastructure Architect

Prerequisites

- Knowledge of container and orchestration concepts using Kubernetes.
- Practical experience with Pods, Deployments, StatefulSets, PersistentVolumes, and PersistentVolumeClaims.
- Proficiency with the Linux command line and kubectl.
- Basic understanding of storage administration, file systems, and block volumes.

Technical Prerequisites

- A computer running Linux, macOS, or Windows 11 with installation privileges.
- A hands-on lab environment with an accessible Kubernetes cluster containing at least three worker nodes.
- kubectl and Helm 3 installed on the workstation.
- A machine with at least 16 GB of RAM and a stable Internet connection, without any filtering that blocks container images and the required repositories.

Longhorn.io Training Program

[Day 1 - Morning]

Fundamentals of Kubernetes Distributed Storage

- Challenges of persistent storage for containerized and stateful applications.
- Review of PersistentVolume, PersistentVolumeClaim, StorageClass, and CSI.
- Longhorn.io’s position within the Kubernetes storage solutions ecosystem.
- Longhorn architecture, control plane components, volume engines, replicas, and management interface.

- Principles of synchronous replication, high availability, volume attachment, and fault tolerance.
- Understanding data flows between the application, CSI driver, volume engine, instance manager, and
  replicas.

[Day 1 - Afternoon]

Installation and preparation of the platform

- Analysis of installation prerequisites, Linux nodes, disks, file systems, connectivity, and resources.
- Preparing the Kubernetes cluster and verifying environment compatibility.
- Deploying Longhorn.io with Helm and reviewing the created resources.
- Exploring the dashboard, global settings, CRDs, and Longhorn namespaces.
- Initial configuration of nodes, disk paths, capacity reserves, and scheduling.
- Hands-on workshop: Install Longhorn.io on a lab cluster, check the status of components, and prepare storage disks.

[Day 2 - Morning]

Provisioning and managing volumes

- Creating and configuring Longhorn StorageClasses for application workloads.
- Dynamic provisioning using PVCs and association with Deployment and StatefulSet workloads.
- Selecting the number of replicas, affinities, disk tags, and scheduling rules.
- Managing data locality, balancing replicas, and balancing resilience against resource consumption.
- Common volume operations: attaching, detaching, expanding, cloning, and viewing events.
- Analysis of status metrics for a volume, its replicas, and its actual capacity usage.

[Day 2 - Afternoon]

Snapshots, Backups, and Restore

- Differences between snapshots, remote backups, and Longhorn configuration backups.
- Configuring an object-compatible or network share backup target.
- Creating snapshots, managing their lifecycle, and controlling their impact on capacity.
- Setting up recurring jobs to automate snapshots and backups.

- Restoring a volume from a backup and attaching it to a Kubernetes application.
- Hands-on workshop: Protecting a stateful application, performing a recurring backup, and then restoring the data to a new volume.

[Day 3 - Morning]

Availability, Monitoring, and Maintenance

- Designing a high-availability strategy tailored to critical nodes, disks, and workloads.
- Monitoring volumes, nodes, disks, replicas, resource usage, and capacity alerts.
- Interpreting degraded states, scheduling errors, and rebuild failures.
- Managing disk space, overprovisioning, minimum capacity thresholds, and snapshot cleanup.
- Controlled removal of replicas during disk or node maintenance.
- Best practices for non-disruptive updates and preparing for operational activities.

[Day 3 - Afternoon]

Diagnostics and Production Scenarios

- Troubleshooting methodology using the dashboard, Kubernetes events, logs, and Longhorn resources.
- Resolving common incidents, such as a full disk, a failed replica, a degraded volume, and unable to schedule.
- Analysis of the consequences of node unavailability and mechanisms for rebuilding replicas.
- Validation of a recovery strategy, restoration testing, and data integrity verification.
- Developing a checklist for operations, capacity, backups, monitoring, maintenance, and documentation.
- Hands-on workshop: Simulate a node failure or disk full condition, evacuate replicas, restore service, and validate the recovery plan.

Target Audience

This training is intended for both individuals and companies—large or small—that wish to train their teams in new, advanced IT technologies or to acquire specific industry 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 with various types of technologies, as well as their expectations and personal goals for 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 problems for the monitoring and smooth running of the training session.

Teaching Methods

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

Organization

The course alternates between theoretical input from the instructor—supported by examples and reflection sessions—and group work.

Assessment

At the end of the session, a multiple-choice quiz is administered to verify that participants have successfully acquired the skills.

Certification

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