

Updated on 09/30/2026

Sign Up

SeaweedFS Training

2 days (14 hours)

Overview

SeaweedFS is a distributed storage platform designed to efficiently manage files, objects, and data at scale using Master, Volume, and Filer components, along with an S3-compatible gateway

Our SeaweedFS training will enable you to master the essential mechanisms of this distributed storage solution so you can design, deploy, and administer a platform tailored to the needs of data and application teams.

You'll learn to understand the SeaweedFS architecture, organize volumes, collections, replication, and data placement, and then expose storage spaces as files or S3 objects.

By the end of the training, you'll be able to install a SeaweedFS environment, configure a Filer, publish an S3 endpoint, and manage access, buckets, and security policies tailored to business needs.

This training also covers high availability, data resilience, observability, metrics, maintenance operations, and best practices for preparing for deployment in a production 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 distributed architecture and components of SeaweedFS
- Deploy a SeaweedFS cluster and configure its core services

- Organize storage using volumes, collections, replicas, and storage tiers
- Expose and use an S3-compatible API
- Secure, monitor, and prepare a SeaweedFS platform for operation

Target Audience

- Data Engineer
- System Administrator
- DevOps Engineer
- Cloud Architect

Prerequisites

- Experience with Linux system administration
- Knowledge of Docker containers
- Basic understanding of object storage and the S3 API
- Knowledge of TCP/IP networks and HTTP services

Technical Prerequisites

- Computer running Linux, macOS, or Windows 11 with installation privileges
- At least 8 GB of RAM and 20 GB of available disk space
- Docker Desktop 4.x or Docker Engine installed and running
- Terminal, a recent web browser, and a stable Internet connection without any filters blocking downloads

SeaweedFS Training Course Outline

[Day 1 - Morning]

Architecture and Fundamentals of Distributed Storage

- SeaweedFS's Position in File, Object, and Data Storage Architectures
- Role of the Master, Volume Server, Filer, and S3 gateway components
- Write and read cycles, file identifiers, volumes, and volume maps
- Organizing data with collections, directories, and metadata
- Principles of Capacity, Topology, Placement, and Disk Types
- Hands-on Workshop: Setting up a local SeaweedFS environment and exploring services, interfaces, and administration commands.

[Day 1 - Afternoon]

Data Deployment and Management

- Installing SeaweedFS with Docker and structuring persistent directories
- Configuring a Master, a Volume Server, and a Filer in a distributed environment
- Managing volumes, size limits, collections, and placement policies
- Data replication, fault tolerance, and an introduction to erasure coding
- Accessing Files via HTTP, FUSE, WebDAV, and Interfaces Provided by the Filer
- Hands-on Workshop: Deploying a multi-service cluster, creating a collection, and storing files with replication.

[Day 2 - Morning]

Object storage and S3 integration

- How the S3 gateway works and the relationship between buckets, collections, and namespaces
- Creating and managing buckets, objects, versions, tags, and lifecycle rules
- Using S3-compatible clients and SDKs for upload and download operations
- Access management with users, access keys, IAM policies, and pre-signed URLs
- Best practices for integration with data pipelines, backup tools, and business applications
- Hands-on workshop: publishing an S3 endpoint, creating a bucket, and automating object transfers with a compatible client.

[Day 2 - Afternoon]

Security, Operations, and Production Readiness

- Securing communications, authentication, access control, and secret management
- Data protection using encryption, retention, versioning, and Object Lock
- High availability of services, metadata replication, and recovery strategies
- Monitoring with Prometheus metrics, logs, capacity indicators, and alerts
- Volume maintenance, compaction, load balancing, incident diagnostics, and sizing principles
- Hands-on workshop: Securing an S3 bucket, monitoring cluster health, and applying a guided maintenance scenario.

Target Audience

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

Prerequisites

The assessment conducted upon enrollment 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 technology, as well as their personal expectations and objectives

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 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 instruction from the trainer—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.