

Updated on 19/08/2025

Sign up

ZFS training with Linux

2 days (14 hours)

Overview

ZFS with Linux is an advanced file system combining high performance, resilience and flexibility. Thanks to its unique design, it offers snapshots, replication, compression and self-healing features ideal for DevOps and cloud environments.

Our ZFS with Linux training course will teach you how to install, configure and administer ZFS in a Linux environment. You'll discover how to manage storage pools, create datasets and volumes, use snapshots and automate replication.

You'll also learn how to secure your data, optimize performance and integrate ZFS into your DevOps workflows and Kubernetes or cloud infrastructures.

By the end of the course, you'll be able to deploy a complete ZFS environment, ensure the resilience and performance of your systems and fully exploit the capabilities of OpenZFS.

Like all our training courses, this one uses the latest [OpenZFS](#) stable [v2.3.3](#).

Objectives

- Understand key ZFS concepts
- Create and administer zpools, datasets and volumes
- Use snapshots, clones and replication
- Securing and optimizing performance with ZFS
- Integrate ZFS in DevOps and container environments
- Deploy ZFS in cloud and Kubernetes environments

Target audience

- DevOps engineers
- System administrators
- Cloud and infrastructure architects
- Operations and support teams

Prerequisites

- General knowledge of Linux
- Notions of DevOps

ZFS with Linux training program

Introduction to ZFS and fundamental concepts

- ZFS history: Sun, OpenZFS, Linux integration
- Key features: checksumming, snapshots, copy-on-write
- Comparison with ext4, XFS, btrfs
- Notion of storage pools (zpools)
- DevOps use cases and cloud infrastructures
- Workshop: Installing ZFS on Linux and creating a first pool

Managing storage pools (zpools)

- Creating and administering zpools
- vdev types: RAID-Z, mirroring, striping
- Adding, replacing, extending and resizing disks
- Capacity monitoring and management
- Key commands: zpool create/status/scrub
- Workshop: Creating different types of pools (RAID-Z, mirror)

Managing datasets and volumes

- Notion of datasets, volumes and clones
- Compression, deduplication, atimes, recordsize
- Quota and reservation management
- Creating block volumes and file systems
- Sharing via NFS/SMB/iSCSI
- Workshop: Creating a dataset with quotas and compression

Snapshots, clones and replication

- Creating and managing snapshots
- Using clones for Dev/Test
- Backup/restore strategies
- Local and remote replication (zfs send/receive)

- Automating replication jobs
- Workshop: Creating a snapshot and restoring a dataset

Security, integrity and optimization

- Integrity via scrubbing and self-healing
- Optimization: ARC/L2ARC/ZIL, cache, SLOG
- Best practices in administration and monitoring
- Permission management and native encryption
- High availability (cluster concepts)
- Workshop: Configuring replication and testing restoration

Advanced ZFS and DevOps integration

- ZFS and containers: Docker, Kubernetes, LXD
- CI/CD usage and pipelines
- Infrastructure as Code (Terraform, Ansible) to manage ZFS
- ZFS in the cloud (Proxmox, TrueNAS, OpenZFS on AWS/Azure)
- OpenZFS roadmap and future developments
- Workshop: Deploying ZFS persistent storage for Kubernetes

Companies involved

This course is aimed at both individuals and companies, large or small, wishing to train their teams in a new advanced computer technology, or to acquire specific business knowledge or modern methods.

Positioning on entry to training

Positioning at the start of training complies with Qualiopi quality criteria. As soon as registration is finalized, the learner receives a self-assessment questionnaire which enables us to assess his or her estimated level of proficiency in different types of technology, as well as his or her expectations and personal objectives for the forthcoming training course, within the limits imposed by the selected format. This questionnaire also enables us to anticipate any connection or security difficulties within the company (intra-company or virtual classroom) which could be problematic for the follow-up and smooth running of the training session.

Teaching methods

Practical training: 60% hands-on, 40% theory. Training material distributed in digital format to all participants.

Organization

The course alternates theoretical input from the trainer, supported by examples, with brainstorming sessions and group work.

Validation

At the end of the session, a multiple-choice questionnaire is used to check that skills have been correctly acquired.

Certification

A certificate will be awarded to each trainee who completes the training course.