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

Updated on 09/10/2025

Sign up

IBM Spectrum Scale Essentials Training

3 days (21 hours)

Overview

IBM Spectrum Scale (formerly GPFS) is a high-performance distributed storage solution designed to manage large volumes of data in hybrid and multi-cloud environments. It combines scalability, resilience and flexibility to meet the needs of DevOps, modern data centers and mission-critical infrastructures.

Our IBM Spectrum Scale Essentials training course will teach you how to install, configure and administer a Spectrum Scale cluster.

You'll discover how to manage files and storage pools, ensure high availability, optimize performance and integrate the platform with your DevOps tools (Ansible, Terraform, Kubernetes).

You'll also learn how to secure your data, automate administration and operate Spectrum Scale in cloud or containerized environments.

At the end of the course, you'll be able to deploy a complete Spectrum Scale cluster, ensure its resilience and integrate it efficiently into your IT infrastructures.

Like all our training courses, this one uses the latest stable version of IBM Storage Scale.

Objectives

- Understand the key concepts of IBM Spectrum Scale
- Install and configure a distributed storage cluster
- Manage pools, filesets, high availability and replication
- Optimize performance and ensure compliance
- Automate administration with API and IaC
- Integrate Spectrum Scale with Kubernetes and the cloud

Target audience

- DevOps engineers
- Systems and storage administrators
- Cloud and infrastructure architects
- Big Data teams

Prerequisites

- General knowledge of Linux and system administration
- Notions of storage and high availability
- Notions of DevOps

IBM Spectrum Scale Essentials training program

Introduction to IBM Spectrum Scale

- Introduction to IBM Spectrum Scale (GPFS) and use cases
- Fundamental concepts: distributed storage clusters, performance, scalability
- Differences with other solutions (NFS, Ceph, Lustre)
- General architecture: nodes, files, metadata, data placement
- Positioning in the DevOps and cloud ecosystems
- Workshop: Installation and initial configuration of a test cluster

Components and internal architecture

- Cluster node management and roles
- How the metadata manager works
- Storage pools, filesets and policy rules
- Internal communication: protocols and network
- Integrated resilience and high availability
- Workshop: Creating and configuring a fileset on a cluster

Basic administration

- Essential management commands (mmcli, mmchconfig, mmcrfs, mmlsfs)
- User management and permissions
- Integration with authentication services (LDAP, AD)
- Resource monitoring and quota management
- Key administration best practices
- Workshop: Creating a filesystem and setting user quotas

High availability and fault tolerance

- Replication and mirroring concepts in Spectrum Scale
- Failure management and automatic recovery

- Disaster recovery methods
- Backup and restore strategies
- Integration with external solutions (TSM, Commvault, Veeam)
- Workshop: Simulate a disk failure and observe recovery

Performance and optimization

- Cache, I/O and data placement
- Optimization for Big Data and HPC workloads
- Monitoring and profiling tools (mmperfmon, Grafana, Prometheus)
- Strategies for balancing performance and resilience
- Cluster tuning and sizing best practices
- Workshop: Measuring and analyzing cluster performance

Security and compliance

- Role-based access control (RBAC)
- Encryption management at rest and in transit
- Integration with SIEM and audit logging
- Regulatory compliance (ISO, RGPD, HIPAA)
- Best practices in data security
- Workshop: Configuring an encrypted filesystem and enabling auditing

DevOps automation and integration

- Automation with scripts and REST APIs
- Integration with Ansible, Terraform and Puppet
- Automated filesystem provisioning
- CI/CD use cases with shared storage
- Centralized supervision in multi-cloud environments
- Workshop: Deploying and managing a cluster via Ansible

Integration with Cloud and Containers

- Deploying Spectrum Scale in AWS, Azure, GCP
- Integration with Kubernetes and OpenShift (CSI drivers)
- Persistent volume management for containers
- IA/ML use cases and massively distributed storage
- Scale-out demonstration
- Workshop: Creating persistent storage for Kubernetes with Spectrum Scale

Summary and outlook

- Review of key concepts and best practices
- Industrial use cases: HPC, data analytics, AI, DevOps
- Limitations and common pitfalls

- What's new in the IBM Spectrum Scale roadmap
- Conclusion and industrialization roadmap
- Workshop: Drawing up a production deployment roadmap

Companies concerned

This training course is aimed at 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.

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 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 verifies the correct acquisition of skills.

Certification

A certificate will be awarded to each trainee who has completed the entire course.