

Updated on 09/23/2026

Sign Up

IBM Storage Scale Training

3 days (21 hours)

Overview

IBM Storage Scale, formerly IBM Spectrum Scale and based on GPFS technology, is a distributed storage solution designed to provide high-performance, resilient, and scalable access to data in enterprise, hybrid cloud, AI, and HPC environments.

Our IBM Storage Scale: Fundamentals training course will teach you how to install, configure, and manage a distributed storage cluster. You'll learn about its architecture, how Network Shared Disks (NSDs) work, and the mechanisms for organizing and protecting data.

You will learn how to create and manage file systems, storage pools, and filesets; configure quotas; and implement high-availability and fault-tolerance mechanisms. You will also be able to monitor performance and identify key bottlenecks.

The training covers IBM Storage Scale's modern administration interfaces, including the GUI, the Native REST API, `scalectl`, and RBAC mechanisms. You'll discover how to automate routine operations while securing administrative access.

You will also explore the integration of Storage Scale into hybrid cloud and containerized architectures, particularly with Kubernetes and Red Hat OpenShift, as well as its use for workloads requiring high storage performance, such as AI, machine learning, HPC, and analytics.

Upon completion of this training, you will be able to administer an IBM Storage Scale infrastructure, ensure its availability, monitor its performance, and prepare for its integration into a modern production architecture.

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 IBM Storage Scale
- Install and administer a cluster and its file systems
- Manage NSDs, storage pools, filesets, and quotas
- Implement high availability and data protection mechanisms
- Monitor and optimize storage performance
- Automate administration using the Native REST API and modern interfaces
- Integrate Storage Scale with cloud, Kubernetes, and OpenShift environments

Target Audience

- System and storage administrators
- Infrastructure engineers
- Linux and AIX administrators
- Infrastructure and cloud architects
- Engineers working in AI, HPC, or analytics environments

Prerequisites

- General knowledge of Linux and system administration
- Knowledge of the fundamentals of storage and file systems
- Basic understanding of networking and high availability
- Some experience in infrastructure administration is a plus

Curriculum for our IBM Storage Scale: Fundamentals training course

[Day 1 - Morning]

Introduction to IBM Storage Scale and its architecture

- Understand the evolution from IBM Spectrum Scale to IBM Storage Scale and the role of GPFS technology
- Identify use cases for HPC, AI, analytics, and distributed storage environments
- Understand the architecture of a cluster: nodes, quorum, metadata, data, and networks
- Learn about Network Shared Disks (NSDs) and their role in distributed data access
- Understand the principles of scalability, parallelism, resilience, and high availability
- Hands-on workshop: Install and configure a test IBM Storage Scale cluster, then verify its status.

[Day 1 - Afternoon]

Create and manage file systems

- Create and manage IBM Storage Scale NSDs and file systems
- Configure storage pools and understand data placement policies
- Create and manage filesets to organize and isolate data
- Implement user, group, and fileset quotas
- Use essential administration commands and understand the logic behind the mm* tools
- Hands-on workshop: Create a file system, configure a fileset, and apply quotas to resources.

[Day 2 - Morning]

Manage data availability, security, and protection

- Understand replication, mirroring, and fault tolerance mechanisms
- Analyze cluster behavior when a disk, NSD, or node fails
- Implement backup, restore, and disaster recovery strategies
- Manage permissions, ACLs, and data access control mechanisms
- Understand encryption, auditing, and security mechanisms for a Storage Scale environment
- Hands-on workshop: Simulate a failure, observe recovery mechanisms, and verify availability of the data.

[Day 2 – Afternoon]

Monitor and optimize cluster performance

- Understand how parallel I/O, caching, and data placement work
- Identify factors influencing the performance of AI, analytics, and HPC workloads
- Use administration tools and the IBM Storage Scale GUI to monitor the infrastructure
- Collect and interpret metrics to identify bottlenecks
- Apply best practices for sizing, load balancing, tuning, and capacity management
- Hands-on workshop: Measure cluster performance, identify a bottleneck, and propose appropriate optimizations.

[Day 3 – Morning]

Automate and manage IBM Storage Scale using modern interfaces

- Automate administrative operations using commands, scripts, and programmable interfaces
- Discover the Native REST API for managing and integrating IBM Storage Scale with external tools
- Understand the role of scaleadm and use scalectl for remote administration
- Implement Role-Based Access Control (RBAC) principles to secure administrative operations
- Integrate administrative tasks into automation workflows and infrastructure tools
- Hands-on workshop: Automate an administrative operation using the Native REST API and manage associated permissions.

[Day 3 - Afternoon]

Integrating Storage Scale into cloud and containerized environments

- Understand the role of IBM Storage Scale in hybrid and distributed architectures
- Identify the principles of storage integration across data centers, the cloud, and remote environments
- Understand integration with Kubernetes and Red Hat OpenShift via CSI mechanisms
- Provide persistent storage for containerized applications and understand the volume lifecycle
- Identify architectures suited for AI, machine learning, HPC, and analytics workloads
- Hands-on workshop: Define an IBM Storage Scale architecture for a containerized workload and prepare a plan for production deployment.

Target Audience

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

Pre-training Assessment

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 level across 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 challenges for monitoring and ensuring the smooth running of the training session.

Teaching Methods

Hands-On Training: 60% Practical, 40% Theoretical. 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.

