

Cluster Computing Administration and Performance Training Course

3 – 7 October 2027
Toronto

AGILE LEADERS
Training Center

AGILE LEADERS
Training Center

Cluster Computing Administration and Performance Training Course

Ref: 1739_230736 **Date:** 3 – 7 October 2027 **Location:** Toronto **Fees:** 12000 Euro

Overview

Cluster Computing Administration and Performance Training Course is a one-week foundation-to-intermediate course for infrastructure, research computing, data, and scientific computing functions who leave with a Cluster Workload Deployment and Tuning Plan. Participants connect cluster architecture, Linux environments, networking, storage, scheduling, parallel execution, observability, security, and capacity with measurable workload behavior. The course turns shared compute resources into a controlled operating workflow. Agile Leaders Training Center develops cluster computing administration and performance capability.

Who Should Attend

- Infrastructure functions responsible for Linux compute nodes, networking, storage, and access
- Research computing functions responsible for supporting scientific and engineering workloads
- Data functions responsible for running scheduled processing across shared compute resources
- Application functions responsible for parallel execution, scaling, and software environments
- Operations functions responsible for monitoring, faults, security, and capacity

The course assumes participants use Linux commands or support technical workloads and leaves out cloud service architecture, data center facilities, advanced parallel algorithm development, and vendor certification preparation.

Departments and Industries

The course supports departments and industries that operate or use shared compute clusters.

- Infrastructure, platform operations, research computing, data engineering, and application support
- Research, education, engineering, and scientific services
- Energy, manufacturing, healthcare, and life sciences
- Technology, telecommunications, finance, and media
- Public services, weather, transport, and analytics functions

Learning Objectives

By the end of this course, participants will be able to:

- Analyze cluster architecture, node roles, networks, and storage paths
- Apply Lmod software environments and controlled user access
- Build Slurm job scripts, resource requests, and job arrays
- Compare MPI and shared-memory parallel execution approaches
- Evaluate monitoring, benchmarking, profiling, scaling, and efficiency evidence
- Prioritize fault, security, and capacity actions for cluster workloads

Course Agenda

Day 1: Cluster Architecture and Linux Environment

- Head, Login, Compute, and Storage Node Architecture Map
- Linux Access, Identity, Permission, and Shell Workflow
- Cluster Network and Shared Filesystem Path Diagram
- Lmod Module Discovery, Load, Swap, and Purge Procedure
- Software Stack and Dependency Compatibility Record

Day 2: Slurm Scheduling and Resource Control

- Slurm Controller, Node, Partition, Job, and Job-Step Model
- Batch Script and Resource Request Template
- Queue State, Priority, Dependency, and Time-Limit Analysis
- Job Array and Repeated Workload Submission Pattern
- CPU, Memory, GPU, and Task Binding Checklist

Day 3: Parallel Workloads and Containers

- Serial, Shared-Memory, and Distributed-Memory Comparison
- MPI Process Placement and Communication Map
- Thread Count, Affinity, and Hybrid Execution Plan
- Scheduled Container Image and Bind-Mount Workflow
- Parallel Input, Output, Checkpoint, and Restart Design

Day 4: Performance and Cluster Operations

- Baseline Benchmark and Reproducibility Record
- Runtime, Utilization, Throughput, and Queue-Wait Dashboard
- Profiling Bottleneck and Scaling Efficiency Analysis
- Node Health, Alert, Fault, and Recovery Runbook
- Cluster Security and Capacity Planning Register

AGILE LEADERS
Training Center

Day 5: Cluster Workload Deployment Practice

- Exercise: Map Cluster Nodes, Networks, Storage, and Software
- Exercise: Submit, Monitor, and Diagnose a Slurm Job Array
- Exercise: Run and Compare Parallel Workload Configurations
- Exercise: Benchmark, Profile, and Tune Resource Requests
- Capstone Exercise: Cluster Workload Deployment and Tuning Plan

Practical Exercises

The course uses suggested activities to connect cluster controls with workload behavior.

- Suggested activity: map login, compute, storage, network, identity, and software dependencies for a research cluster.
- Suggested activity: write and diagnose Slurm jobs for data processing and simulation cases.
- Suggested activity: compare MPI, threaded, hybrid, and containerized execution patterns.
- Suggested activity: assemble benchmark, profiling, fault, security, and capacity evidence in a deployment and tuning plan.

FAQs

Who suits cluster computing administration training, and what does it assume?

Infrastructure, research computing, data, application, and operations functions suit the training. It assumes basic Linux command use or responsibility for technical workloads.

How does cluster computing training differ from general cloud computing training?

Cluster computing training focuses on nodes, shared filesystems, workload schedulers, parallel jobs, software modules, observability, and resource tuning, while general cloud training focuses on service models, virtual resources, managed services, and cloud governance.

What does Slurm manage in a compute cluster?

Slurm allocates compute resources, starts and monitors work, and manages queued jobs through concepts such as nodes, partitions, jobs, job steps, priorities, limits, and accounting.

How do MPI and shared-memory computing differ?

MPI coordinates explicit communication among processes, including processes on different nodes, while shared-memory approaches coordinate threads or processes that access memory on the same node.

AGILE LEADERS
Training Center

How is cluster workload performance evaluated?

Cluster workload performance is evaluated through reproducible baselines, runtime, utilization, throughput, queue wait, scaling efficiency, profiling evidence, input and output behavior, and resource-request accuracy.

Conclusion

Participants take a Cluster Workload Deployment and Tuning Plan back to work. It connects architecture, access, software environments, scheduling, parallel execution, containers, observability, performance, faults, security, and capacity. It changes an ad hoc cluster workload into a controlled and measurable operating workflow with documented resource and tuning decisions.