

AI-INFRA 100: AI Infrastructure

OEM: Mirantis • Code: AI-INFRA 100



COURSE DESCRIPTION

A GPU cluster built for AI workloads runs on infrastructure that does not behave like a conventional datacenter. This introductory-level 2-day course covers that stack across four modules: fundamental architecture, fleet hardware management, network fabrics, and bare-metal GPU provisioning from Kubernetes. The goal is a working mental model of the whole stack, and hands-on exposure to the key components. The four modules run in sequence, each building on the vocabulary and mental model established by the one before it. Three of the four include hands-on labs, on virtual environments that stand in for a real cluster.

COURSE CONTENT SUMMARY

- The curriculum encompasses four module domains:
- AI Infrastructure Foundations: A conceptual walkthrough of the GPU cluster stack, covering the GPU node, HBM and the checkpoint burst, NVLink and intra-node interconnect, the four fabrics and what each one carries, the management plane, and the failure domains at each layer. Introduces the vocabulary and architecture the remaining three modules assume, and gives participants enough grounding to follow a vendor conversation without nodding at words they cannot define. Conceptual; no lab.
- Redfish and Fleet Hardware Management: An introduction to hardware management and observability through one HTTPS API across a mixed-vendor fleet. Covers the Redfish resource model, health validation before provisioning, firmware compliance and drift, the scrape pipeline behind a fleet dashboard, and VirtualMedia provisioning boot. The multi-vendor problem is the practical heart of the module: every BMC vendor has its own interface and its own quirks, and clicking through four different web UIs does not survive contact with a real fleet. Hands-on labs.
- GPU Cluster Networking: RoCE v2, PFC/ECN: An introduction to lossless fabric, covering why RDMA cannot tolerate loss, what PFC and ECN do on both the switch and the NIC, fabric telemetry, and worked incidents traced from symptom to root cause. Participants leave able to reason about fabric behavior and read fabric telemetry, not to design a production fabric unaided. Hands-on labs.
- Bare-Metal GPU Provisioning with Kubernetes: An introduction to declarative bare-metal provisioning. Covers host enrollment, provisioning with k0rdent and Metal3/Ironic, GPU Operator and Network Operator deployment on a mixed fleet, an incremental deploy-and-validate methodology, and a provisioning failure traced back to its hardware root cause. Hands-on labs.

AUDIENCE PROFILE

- The course addresses infrastructure engineers, datacenter and platform operations staff, and field, solutions, and pre-sales engineers.
- Infrastructure engineers encountering GPU cluster work for the first time get the map and the vocabulary before the first outage. Platform and operations teams, inheriting a cluster somebody else stood up, get a model of what they have been handed. Field and pre-sales roles who need to follow a technical conversation about GPU topology, RDMA fabric, and cluster automation without implementing it get most of that value from the first two modules.

- Participants should possess practical datacenter experience with servers, switching, and BMC-based hardware management. No RDMA, DPU, Redfish, or GPU cluster experience is assumed.

PREREQUISITES

- Solid Linux command line proficiency
- TCP/IP networking, VLANs, and L2/L3 switching fundamentals
- Server hardware fundamentals, including exposure to BMC or IPMI
- Ability to read JSON and YAML, and use curl
- Kubernetes and kubectl basics, for the provisioning module
- Beneficial (not mandatory): shell scripting, and familiarity with Kubernetes operators

Lab Requirements

- WiFi-enabled laptop
- Current Chrome or Firefox browser
- SSH client
- Lab environments are hosted and preconfigured; no local installation or hardware is required

COURSE OBJECTIVES

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

- This section doesn't contain any content.

COURSE OUTLINE

Module 1: AI Infrastructure Foundations

- The GPU node: accelerators, HBM, and the checkpoint burst
- NVLink and intra-node interconnect
- The four fabrics and what each one carries
- The management plane
- Failure domains at each layer

Module 2: Redfish and Fleet Hardware Management

- The Redfish resource model and mixed-vendor BMCs
- Health validation before provisioning
- Firmware compliance and drift
- The scrape pipeline behind a fleet dashboard
- VirtualMedia provisioning boot
- Hands-on labs

Module 3: GPU Cluster Networking: RoCE v2, PFC/ECN

- Why RDMA cannot tolerate packet loss
- PFC and ECN on the switch and on the NIC
- Fabric telemetry and what it surfaces
- Worked incidents traced from symptom to root cause
- Hands-on labs

Module 4: Bare-Metal GPU Provisioning with Kubernetes

- Host enrollment and hardware inventory
- Declarative provisioning with k0rdent and Metal3/Ironic
- GPU Operator and Network Operator on a mixed fleet
- Incremental deploy-and-validate methodology
- Tracing a provisioning failure to its hardware root cause
- Hands-on labs

DELIVERY NOTE

- Instructor-led, two consecutive days, four sequential modules. Hands-on labs in modules 2, 3, and 4 on hosted, preconfigured environments. Reference materials are yours to keep. Available as a private delivery for your team.