

ALE High-Performance AI & HPC Platform

AI & HPC Rack – Ready to Go

Accelerate AI training, inference, and simulation with a fully pre-integrated, high-performance AI & HPC rack. Built with AMD GPU servers, ultra-low-latency Ethernet fabrics, and validated leaf-spine architecture,optimized for power efficiency. ALE delivers predictable performance from day one — without deployment complexity.

Brochure

Alcatel·Lucent Enterprise

Scaling AI & HPC Without the Pain

AI and HPC infrastructures are under unprecedented pressure:

AI workloads are exploding
Training, inference, simulation, analytics

Performance demands push **latency and bandwidth** to the limit

Clusters grow fast, but **operational complexity grows faster**

Security, isolation, and reliability are no longer optional

Time-to-value & TCO matter more than ever

Enterprises need performance at scale — without building everything from scratch.



Pre-integrated GPU compute and high-performance Ethernet fabric — validated, scalable, and ready to deploy.

- High-performance **compute and network switching**
- AI-optimized **Ethernet fabric**
- **Unified orchestration** with ML-Ops and AIOps
- Embedded **security and micro-segmentation**
- Scalable, **ready-to-deploy**
- **End-to-end ALE services**
- **Lower energy impact**

Brochure
ALE High-performance AI & HPC Platform



AI-ready fabric

400G/800G backend switching with lossless Ethernet (RoCE-ready) for **predictable performance**.

GPU compute

AMD MI325X-based GPU server building blocks designed for **dense racks and cluster scaling**.

Operate with confidence

SONiC-based or AOS NOS and rack Orchestration for **provisioning, monitoring, and lifecycle management.**

AI training, inference & advanced simulation workloads demand high-throughput compute, predictable low-latency networking, and simplified operations at rack and cluster scale.

Designed to maintain peak performance while maximizing energy efficiency

What ALE delivers

- **Backend leaf-spine switching** (400G/800G)
- **Frontend connectivity** for management and storage
- **Telemetry and automation-ready interfaces**
- **GPU servers** based on AMD MI325X (8 GPUs per node)
- **High-speed NICs** for CPU and GPU traffic profiles
- **Dense rack building blocks** for scale-out interfaces
- **SONiC-based or AOS NOS** for backend/front-end fabrics
- **Rack orchestration** for inventory, provisioning, and monitoring
- Validated **blueprints**, runbooks, and services



Architecture at a glance

An AI rack typically uses three physical networks:

- **Frontend** (management/storage/client)
- **Backend** (GPU-to-GPU / internode)
- **Out-of-Band** (server and switch management)

Typical workloads

- **Large language model training** and fine-tuning
- Enterprise **inference clusters and RAG pipelines**
- **High-performance simulation** and engineering workloads
- AI-assisted analytics and accelerated **data processing**

Key benefits

- **Faster time to value** with a validated rack design and defined BOM
- **Predictable performance** with an AI-optimized Ethernet fabric
- **Simplified operations** through orchestration, telemetry, and lifecycle tooling
- **Scale** from a single rack to multi-rack clusters **with repeatable designs**

Inside the rack

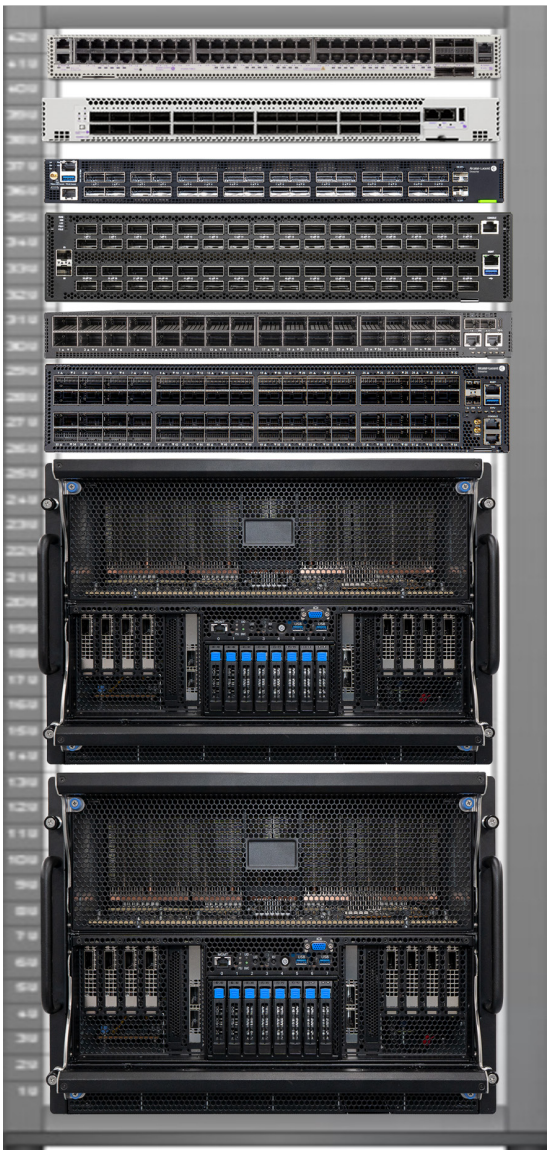

Management Switch
OmniSwitch 6900 T24-T48

AI Fabric Switch
OmniSwitch 6920 Series
OmniSwitch 7900 Series

Switch NOS
AOS X or SONiC

Orchestrator

GPU Server

AMD ROCm

GPU Server

Example AI Data Center rack composition (compute, fabric, management, NOS and orchestration)
Note: final rack population, power/cooling profile, and optics vary by validated design and customer requirements.

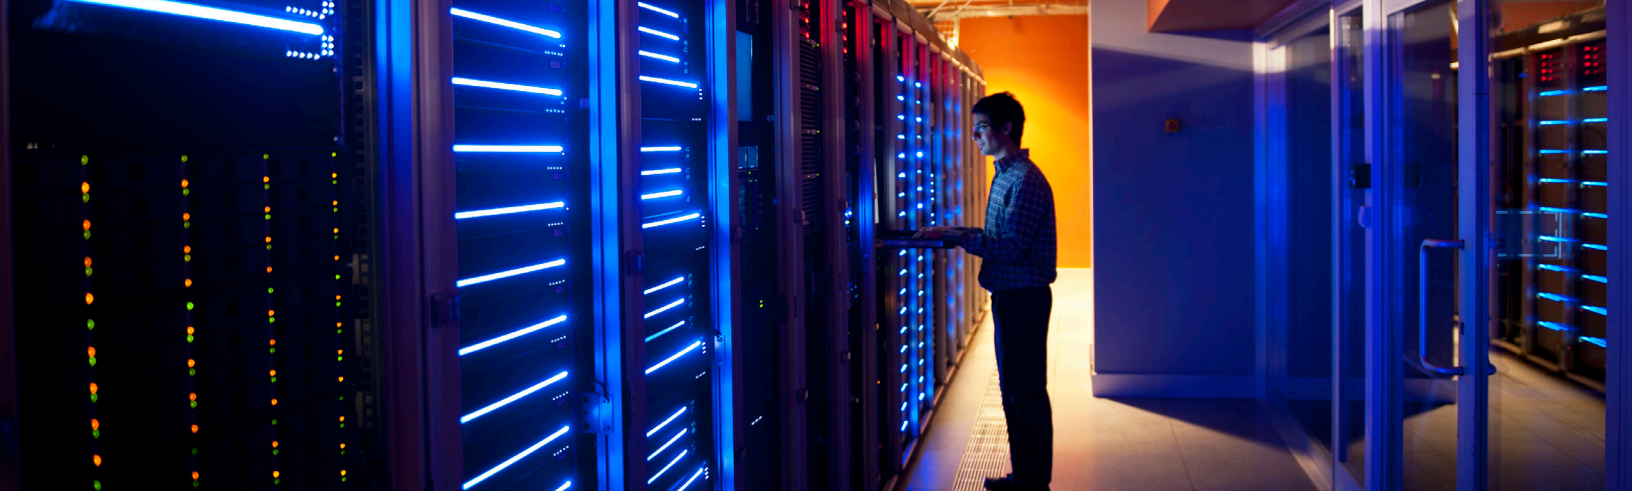

Core building blocks

Backend switches (leaf/spine)	Alcatel-Lucent **OmniSwitch 7900 series**: • OmniSwitch 7900-D32: 32×400G (leaf) • OmniSwitch 7900-D64: 64×400G (spine) • OmniSwitch 7900-O32: 32×800G (leaf) • OmniSwitch 7900-O64: 64×800G (spine)
Front-end switches	Alcatel-Lucent **OmniSwitch 6900 / OmniSwitch 6920 series** for 10G/25G/100G/400G front-end
Management (OOB) switches	Alcatel-Lucent **OmniSwitch 6900-T24/T48** for out-of-band management and device access
GPU server	• AMD MI325X (8 GPUs) with Infinity Fabric • CPU: AMD EPYC 9005 (Turin) • CPU NIC: 2×200GbE • GPU NIC: 8×400G • Storage: 2×1.92TB NVMe (U.2)
Optics & cables	Validated transceivers and cabling options for 100G/400G/800G (DAC/AOC/optics) depending on rack and cluster design
GPU software stack	ROCm-based stack with validated container/runtime approach and integration options for leading AI/HPC frameworks
Network operating systems	• **ALE SONiC** (ASON) or **AOS-X** for AI fabrics
Orchestration & lifecycle	Rack orchestration for inventory, topology discovery, configuration templates, telemetry, and software image management

Scale and operate from rack to cluster

Scalability reference

Validated backend fabric designs scale from 1 rack to 8 racks, mapping to 32 to 256 GPUs (4 GPU servers per rack; 8 GPUs per server).

Rack	GPU servers	GPUs	Typical use
1	4	32	Pilot / lab / first production pod
2	8	64	Department-scale training and inference
4	16	128	Shared AI cluster / private cloud
8	32	256	Large-scale training, multi-team cluster



Fabric capabilities Design targets

- Lossless Ethernet for RoCE-ready workloads (PFC/ECN/priority handling as required by the design)
- Leaf-spine architectures with ECMP for scalable bandwidth
- EVPN-VXLAN options for front-end segmentation and multi-tenant designs
- Streaming telemetry and automation interfaces for fast operations



Services and delivery

- Validated bill of materials (BOM) and reference architecture
- Pre-cabling guidance and deployment runbooks (Day 0/Day 1)
- Acceptance testing and handover documentation
- Lifecycle services: upgrades, configuration governance, and support

Ready to discuss a pilot rack or a scale-out cluster?

Contact your ALE representative to plan a proof of concept and validated design.

Alcatel·Lucent Enterprise

© 2026 ALE International, ALE USA Inc. All rights reserved in all countries. The Alcatel-Lucent name and logo are trademarks of Nokia used under license by ALE. To view a list of proprietary ALE trademarks, visit: www.al-enterprise.com/en/legal/trademarks-copyright. February 2026