Alcatel·Lucent
Enterprise

ALE DC Initiative

Why Data Centers Are ALE Strategic New Focus

Portfolio Diversificatio

Shift from legacy technologies to strategic, future-proof growth areas

Stronger Economics

Larger deals, longer customer lifetime value, multi-year engagements

Strategic Relevance

Core to digital transformation: AI, cloud, sovereignty initiatives

DC Solutions Landscape Today

Traditional & Enterprise DCs

On-premises facilities, owned and operated by a single organization,

For IT workloads,

Involve high CAPEX and limited scalability,

Cloud & Hyperscale DCs

Hosted by major public cloud providers,

Offer on-demand, scalable compute and storage,

Backbone for SaaS, hybrid cloud, and distributed workloads,

Colocation & Managed Services

Third-party facilities where companies lease space, power, and connectivity,

Offer flexibility without the need for direct infrastructure ownership and management.

Edge & Hybrid DCs

Smaller, decentralized facilities located close to end-users.

They reduce latency and support localized processing for IoT, real-time analytics, and 5G applications.

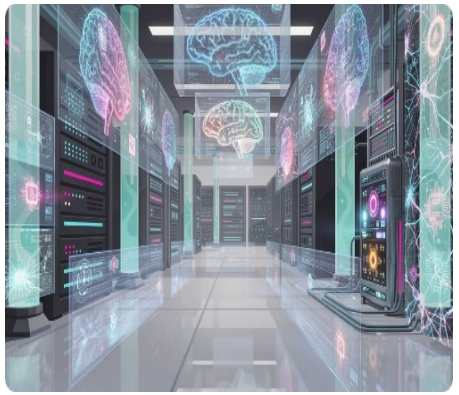

AI & High-Performance DCs

Purpose-built for compute-intensive AI workloads using GPUs/accelerators.

These require advanced cooling, high power density, and high-speed networking capabilities.

Data Center Investments Are Shifting Toward AI

Highlight: "While traditional Data Centres remain essential, nearly all new growth in Data Centre investment is now driven by AI."

- **AI-Driven Infrastructure:** GPU compute, high-speed networking, and advanced cooling are now the primary investment drivers, supporting cutting-edge AI workloads.
- **Future Growth in AI:** Over the next five years, most incremental data center spending will be AI-related, fueled by Generative AI, analytics, automation, and sovereign AI initiatives.
- **Strategic Capital Allocation:** Traditional data centre investments remain stable, while new capital is strategically directed towards accelerating AI capabilities and infrastructure.

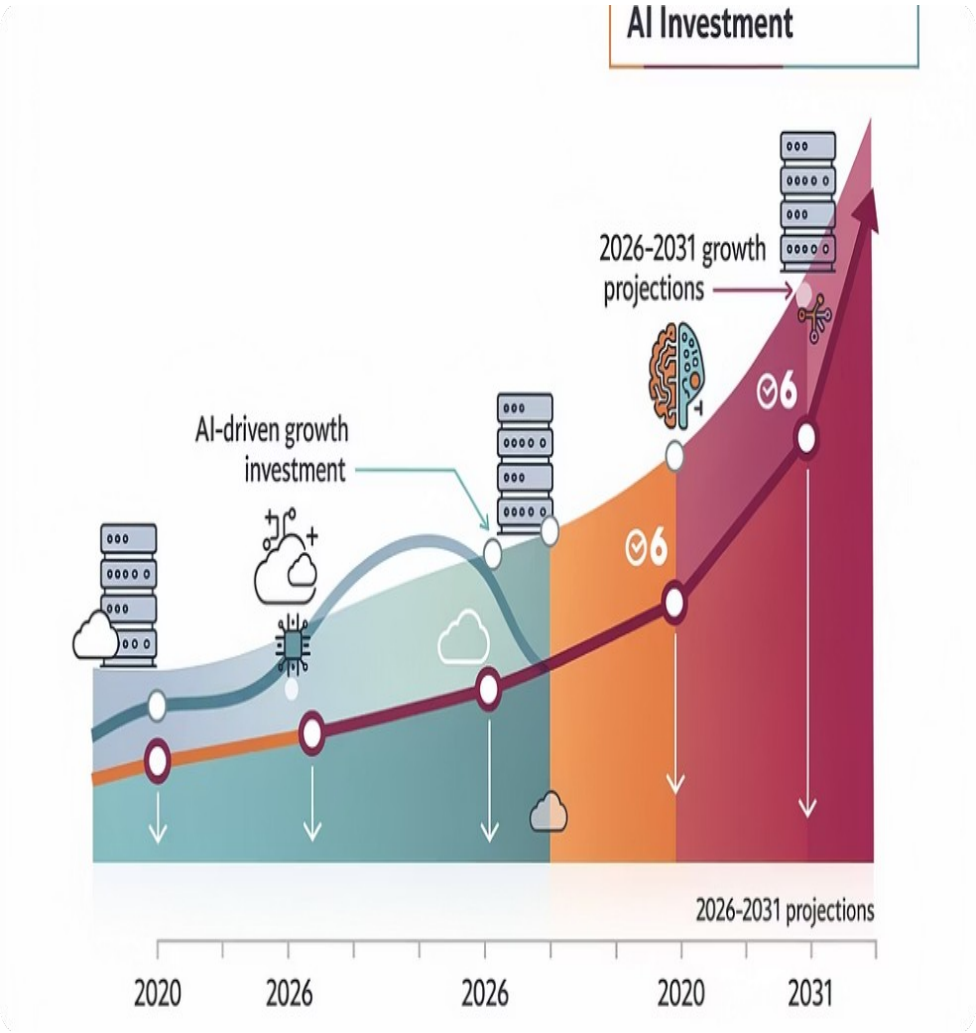

Why Everyone Is Using AI Today

Why Everyone Is Using AI Today

Artificial Intelligence (AI)

Perform tasks requiring human intelligence:
- Reasoning
- Learning
- Perception
- Decision making
- Interaction

Machine Learning (ML)

Learn patterns from historical data to make predictions or decisions, used for classification, regression, and anomaly detection.

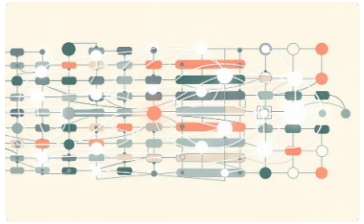

Deep Learning (DL)

multi-layer neural networks to learn highly complex patterns: In image recognition, speech recognition, and video analysis.

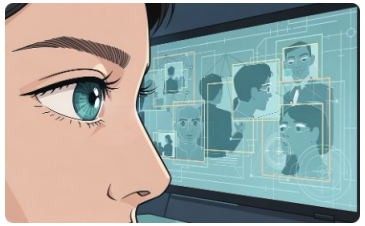

Computer Vision

Interprets visual data from cameras and images to enable object detection, facial recognition, image classification, video analytics, and quality inspection.

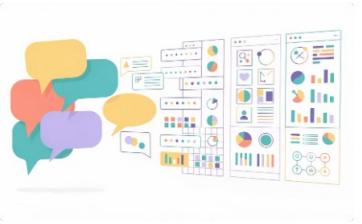

Natural Language Processing (NLP)

Processing and understanding human language, facilitating text understanding, translation, sentiment analysis, and speech-to-text conversion.

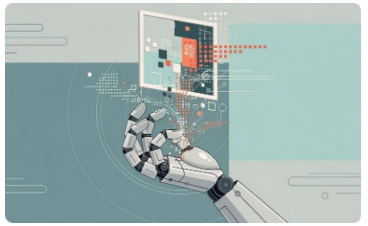

Generative AI (GenAI)

Generates original content based on learned patterns, ranging from text and images to video, music, and even code.

Large Language Models (LLMs)

Models trained on massive datasets to predict the next word, powering conversational AI, question answering, summarization, and coding assistance.

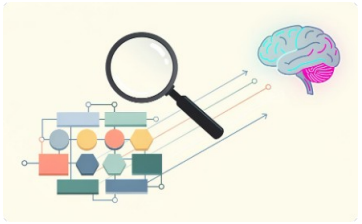

Retrieval-Augmented Generation (RAG)

Combines LLMs with enterprise or external data sources to provide context-aware, up-to-date, and accurate responses for knowledge assistants and enterprise Q&A.

Multimodal AI

Designed to understand and generate content across multiple modalities like text, images, audio, and video, enabling comprehensive cross-media reasoning.

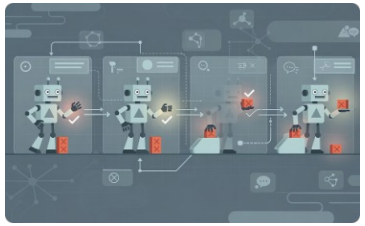

Reinforcement Learning (RL)

Learns optimal actions through feedback and rewards, primarily used in robotics, game-playing AI, and dynamic optimization for autonomous systems.

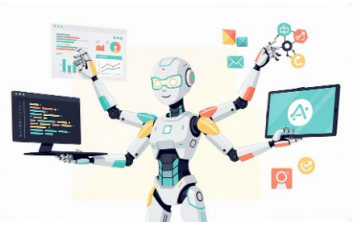

Autonomous & Agentic AI

Refers to AI systems that can plan, decide, and execute complex tasks with minimal human intervention, driving workflow automation and multi-step reasoning.

AI Delivering Tangible Business Outcomes

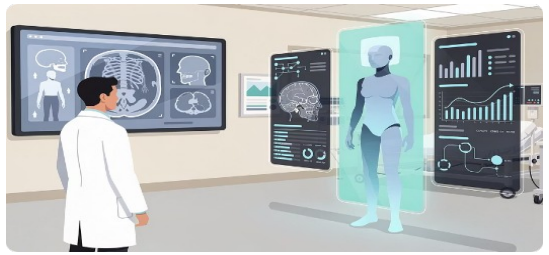

Healthcare

- Assist doctors in diagnosing diseases from medical images and test results
- Automatically prioritize patients based on urgency
- Support medical research by analyzing large clinical datasets
- Optimize hospital operations such as scheduling and bed usage
- Analyze population health data to predict outbreaks and risks

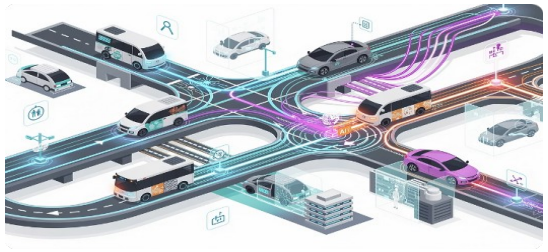

Transportation

- Predict and detect incidents and accidents in real time using video analytics
- Optimize train schedules based on passenger demand and traffic patterns
- Automate traffic management, passenger flow and signaling optimization
- Improve customer experience with AI-powered passenger information and chatbots
- Simulate rail operations using digital twins to test scenarios before deployment

Government & Cities

- Optimize traffic lights and reduce congestion in real time
- Improve public safety, incidents detection and emergency response times
- Optimize waste collection and city services
- Monitor air quality, noise, and pollution
- AI chatbots for citizen services and requests
- Automate document processing and public administration

Education

- Digital Education platforms and personalized learning (adaptive learning platforms, AI tutors)
- Secure assessment and academic integrity (automated grading, plagiarism detection, secure online exams, and AI-based proctoring)
- Research and AI skills development (AI research labs, student AI sandboxes, AI training programs, and cross-disciplinary innovation platforms)
- Operational efficiency and cost optimization (Enrollment forecasting, timetable

AI Delivering Tangible Business Outcomes

Finance

- Detect fraud and suspicious transactions in real time
- Improve risk assessment and credit scoring
- Optimize trading strategies using predictive models
- Provide instant customer support through AI chatbots
- Increase security while improving customer experience

Retail & eCommerce

- Predict product demand to optimize inventory
- Automate replenishment and supply chain decisions
- Generate product descriptions and marketing content
- Provide personalized shopping recommendations
- Enable conversational shopping assistants

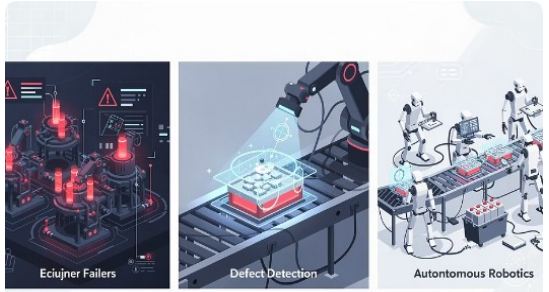

Manufacturing

- Predict equipment failures to reduce unplanned downtime
- Automatically detect defects during production
- Simulate manufacturing processes before physical deployment
- Enable smarter and more autonomous robotics systems
- Accelerate product design through AI-driven simulations

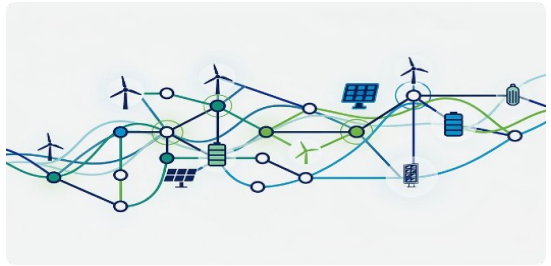

Energy

- Predict energy demand to balance supply and consumption
- Optimize smart grid operations in real time
- Maximize renewable energy production and usage
- Decide the best time to store or release energy
- Reduce energy waste through predictive analytics

Training vs Inference in an AI Data Centre

AI models undergo two critical phases: **training** and **inference**, each with distinct networking requirements.

Training : involves feeding vast datasets to GPU clusters, generating extremely heavy east-west network traffic between GPUs and servers. This demands high-bandwidth, low-latency, and lossless networks for optimal GPU utilization.

Inference : conversely, deploys trained models for real-time predictions. It is typically more latency-sensitive but less bandwidth-intensive than training, focusing on rapid responses to users.

AI Needs Specialized Infrastructure

Traditional IT infrastructure cannot meet the unique demands of modern AI workloads.

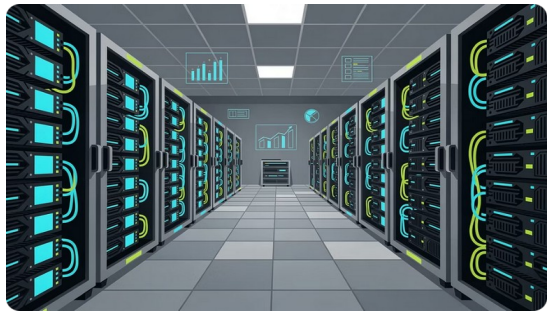

Massive Compute

Training large AI models demands **millions of GPU hours** and cost tens of millions of dollars in compute resources.

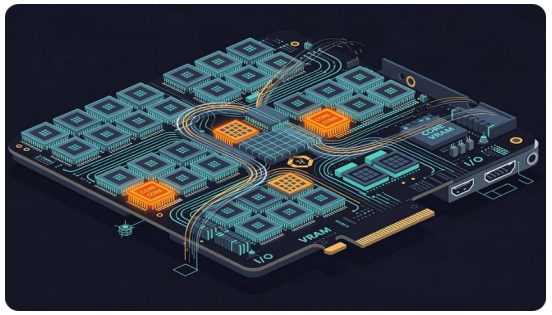

Parallel Processing

GPUs accelerate neural network training **by 100-1000x** compared to CPUs.

Extreme Data Scale

Models like GPT-5 train on **trillions of tokens** from hundreds of terabytes of text, requiring specialized data pipelines.

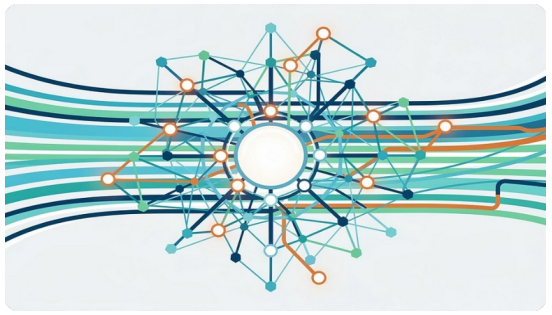

Low-Latency Inference

Production AI systems need to deliver predictions in **milliseconds** while handling millions of requests simultaneously

.

The Layered View of AI Infrastructure

- **Operations:** Transform experimental AI into production-grade systems with enterprise reliability
  - ✓ **CI/CD Pipelines:** Automated testing and deployment
  - ✓ **Monitoring & Observability:** Real-time metrics on model performance, drift detection, and infrastructure utilization
  - ✓ **Security & Compliance:** IAM Roles, encryption at rest/in transit, audit logs
- **Software:**
  - ✓ **Frameworks**: PyTorch, TensorFlow, JAX
  - ✓ **Acceleration Libraries:** ROCm, MIOpen, RCCL, CUDA, cuDNN, NCCL, oneDNN
  - ✓ **Containers & Orchestration:** Docker, Kubernetes
  - ✓ **MLOps Tools:** MLflow, Kubeflow, Weights & Biases
- **Hardware:**
  - ✓ **Compute** GPUs for matrix operations and parallel processing
  - ✓ **Storage** High-speed SSDs, distributed object stores, feature stores optimized for AI
  - ✓ **Networking** RDMAv2 lossless Eternet

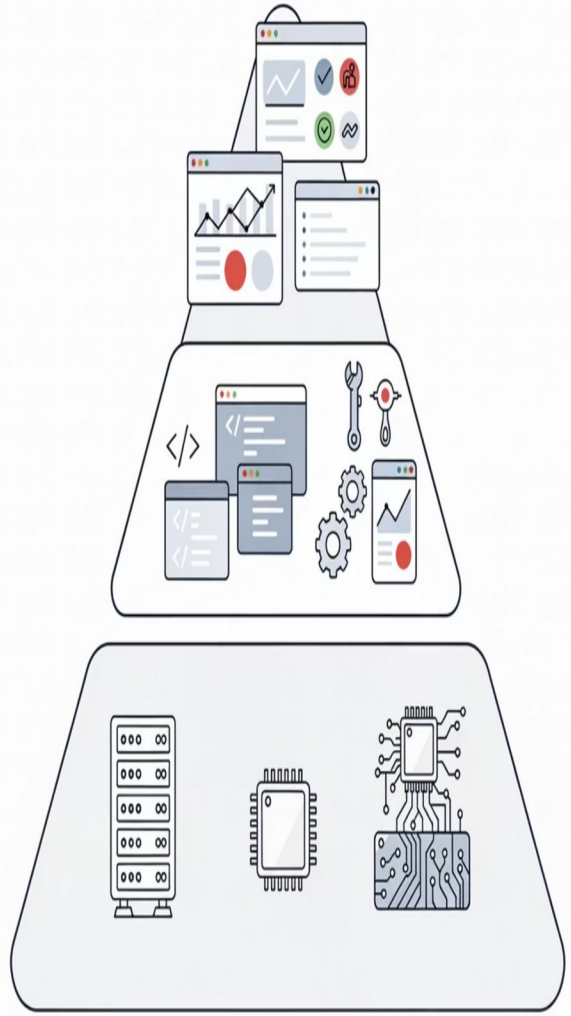

ALE AI-DC Offer

ALE AI Data Center offer

Computing

AI server for training Large Language Models and inference

Networking

Network infrastructure for AI clusters running at full speed, without delays or bottlenecks

ALE AI DATACENTER RACK COMPONENTS

Mgmt. Switch
OS6900

AI Fabric Switch
OS7900 (EC rebranded)
OS6900
OS6920-D32

Management / Storage Frontend Network

General Server/ Storage Server
OT2290 (future)

GPU Server
OmniCompute

Power: 10kW per GPU Server
(actual working consumption is 6~8kW)

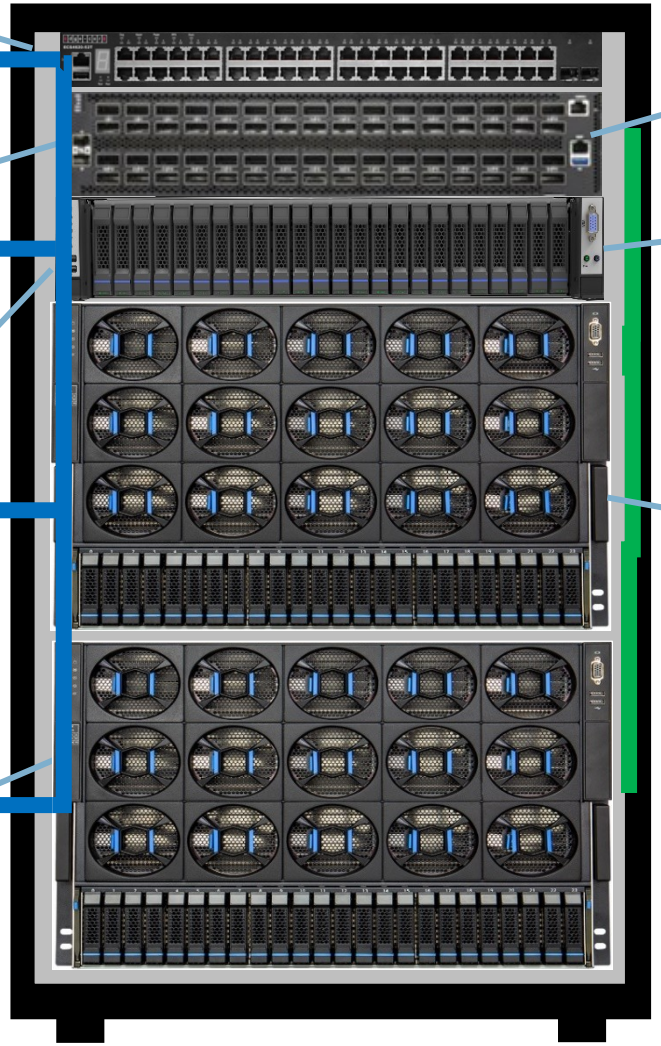

Switch NOS
SONiC/AOS-X

SW Orchestrator
Infrastructure Mgmt.
Service Mgmt.

AI Computing Backend Network

GPU Framework
AMD ROCm

Transceivers and Cables


AI-Data Center Networking [ALE-COMPONENTS]

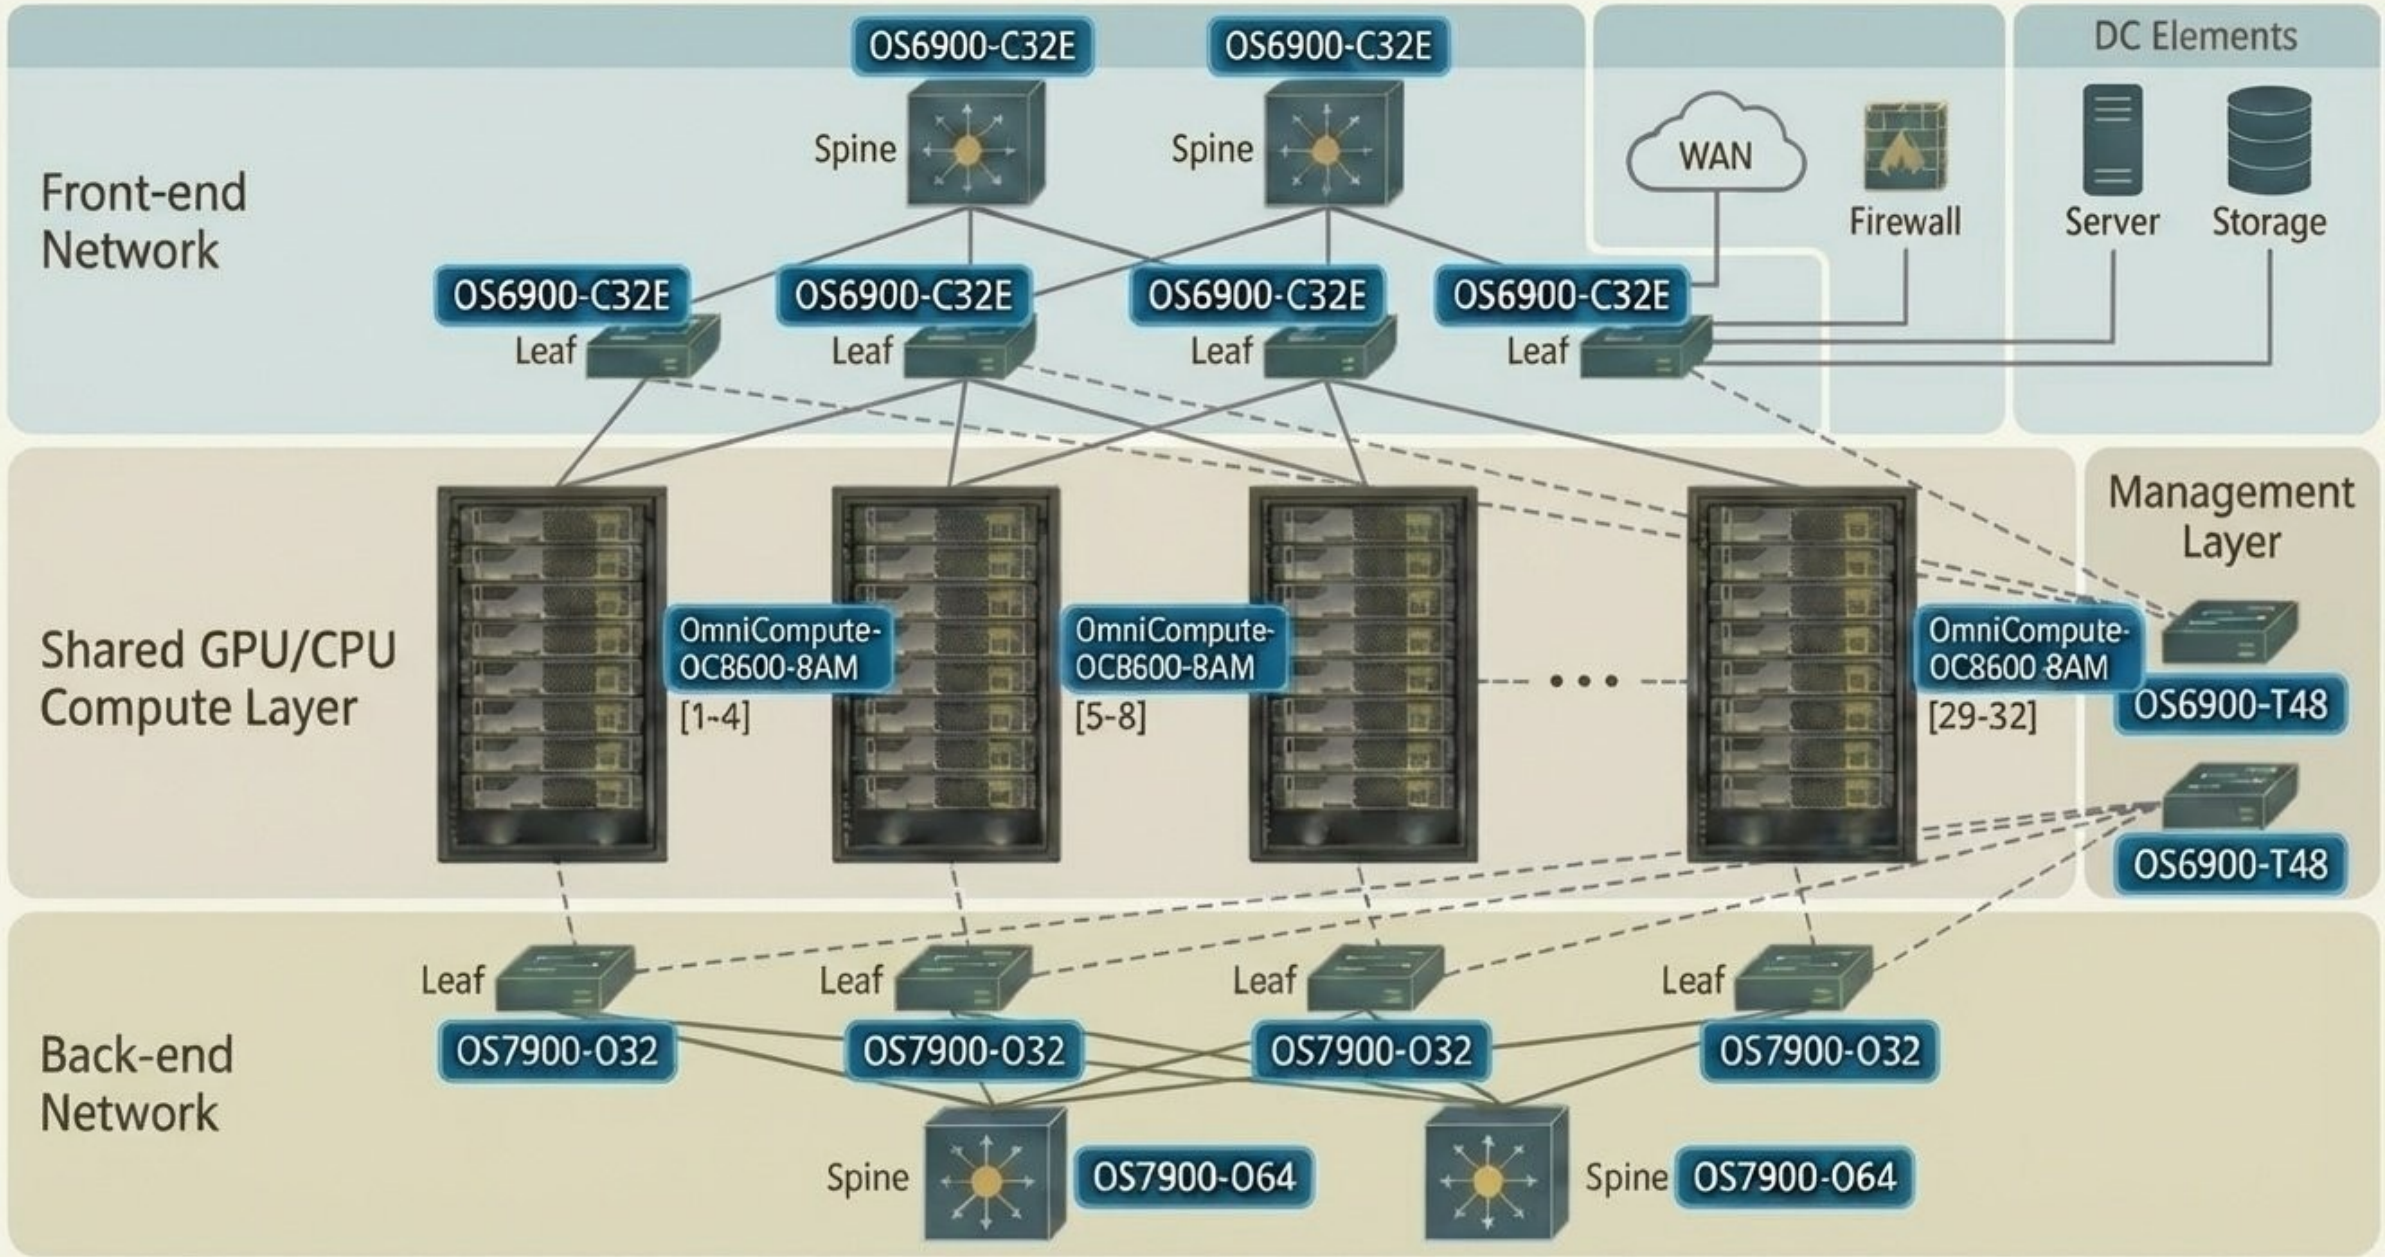

OS7900-064 & OS7900-032 FOR BACKEND FABRIC

OS7900-O64 (OSFP800 Ports)

OS7900-O32 (OSFP800 Ports

650 GROUP
MARKET INTELLIGENCE RESEARCH

- Switching capability of **51.2 Tbps** & **25.6 Tbps**.
  (H**igh bandwidth**, **low latency** & **Advanced load-balance**.)
- 8 SerDes lanes per port @100Gbps each.
- Broadcom Tomahawk5(**TH5**) ideal for **AI workloads.**
- **LPO**(Linear Pluggable optics) with **low power consumption**, **reduced cost** & **lower latency.**
- OSFP **migration** to higher speeds of **1.6Tbps**.

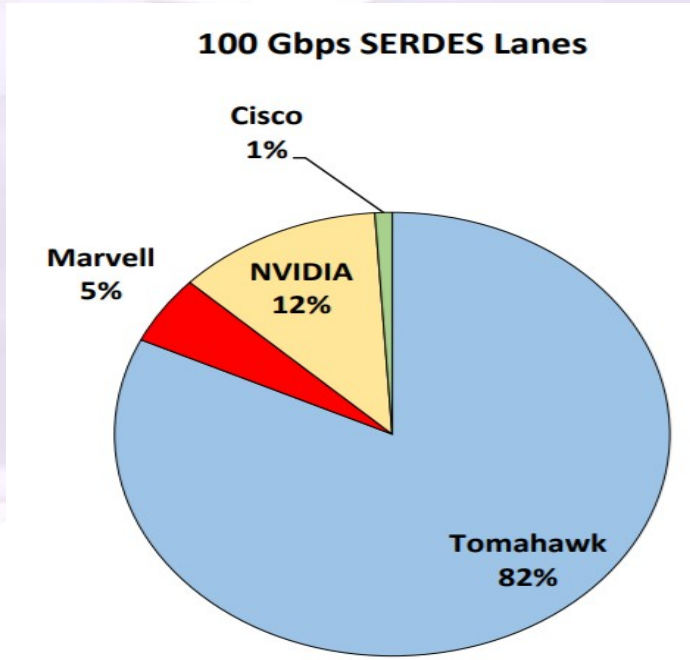

OS7900-D64 & OS7900-D32 FOR BACKEND FABRIC

OS7900-64D (QSFPD400 Ports)

OS7900-32D (QSFPD400 Ports)

650 GROUP
MARKET INTELLIGENCE RESEARCH

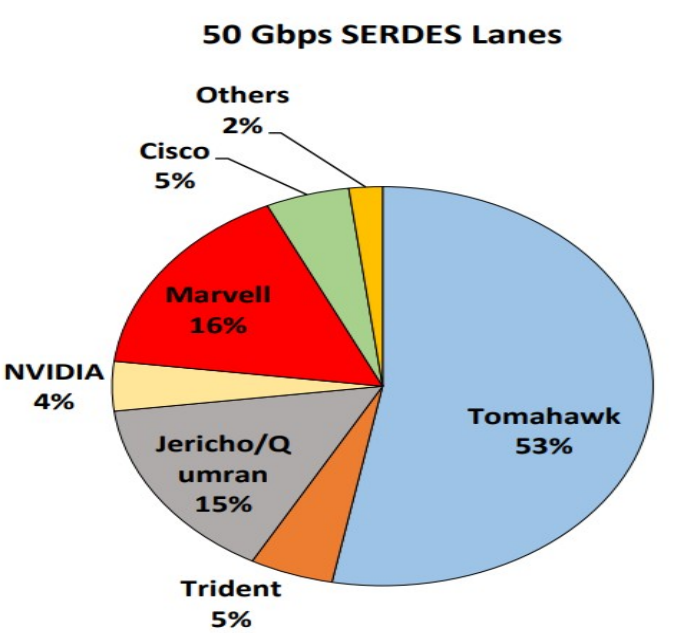

- Switching capability of **25.6 Tbps & 12.8 Tbps** 8 SerDes lanes per port @50Gbps each..
- Broadcom Tomahawk4(**TH5**) ideal for **AI workloads**
- **Dynamic load-balance**.
- Deployments for HPC/GPUs with lower bandwidth requirement**.**

ALE OMNICOMPUTE [OC8600]

AMD MI325X OAM CDNA3™ Architecture

- **Specifications**:
- GPU: **AMD 8 x MI325X GPUs** with Infinity Fabric
- CPU: AMD 2 x EPYC™ 9005/Turin Series Processor
- CPU NIC: 2x BCM957608-P2200GQF00 **Dual-Port 200GbE**
- GPU NIC: 8x BCM957608-P1400GDF00 **Single-Port 400G**
- Storage: 2x1.92TB U.2 2.5 NVME4 1DWPD (non SED)
- Form Factor / Cooling: 8U Air Cooling
- FP32 TFLOPS: 163.4 FP16 TFLOPS: 1300 FP8 TFLOPS: 2610
- Compute Memory per GPU: **256GB** HBM3e
- Memory Bandwidth: **6 TB/s**
- Peak Power Consumption per GPU: 1KW
- Actual Working Consumption per Server: 6-8KW
- Compute Fabric Interface: Addition External NIC

MI325X vs H200 (Key Specs)

NVIDIA H200 SXM	AMD MI325X OAM
HBM3 Memory: 141 GB	256 GB
Memory Bandwidth: 4.8 TB/s	6 TB/s
TF32 Peak: 494.7 TFLOPs	658.7 TFLOPs
FP16/BF16 (Tensor/Matrix) Peak: 989.4 TFLOPs	1207.4 TFLOPs
FP8 Peak: 1978.9 TFLOPs	2614.9 TFLOPs

www.amd.com

ALE SOFTWARE FOR SWITCHES & NETWORKS

SONiC

AOS-X

Dual NOS Strategy:

SONiC : ALE SONiC

- For large scale & cloud-native enterprise for **full control of stack**.

AOS-X :

- Provides **feature depth & maturity**.
- **Operational simplicity** for smaller teams/enterprises.

Key Features for NOS :

- **Modern NOS** principles.
- **Cut-through switching**.
- **Adaptive routing** & **dynamic load balancing**.
- **PFC**(Priority flow control),
- **RoCEv2**(RDMA over converged ethernet V2),
- **ECN**(explicit congestion notification).
- **Advanced telemetry.**
- **VXLAN-EVPN** for multi-tenancy.

Key Features for Orchestrator :

- Rack orchestration for backend, frontend & management network.
- Provide Industry standard Day-0, Day-1 & Day-2 operations.
- Intent based network provisioning
- Network Source of truth(NSOT) to maintain the integrity of the network.
- Integrated digital twin
- Monitoring & Maintenance using advanced telemetry.

GPU BACKEND FABRIC: [16 GPU] [LAB-BLUEPRINT-1]

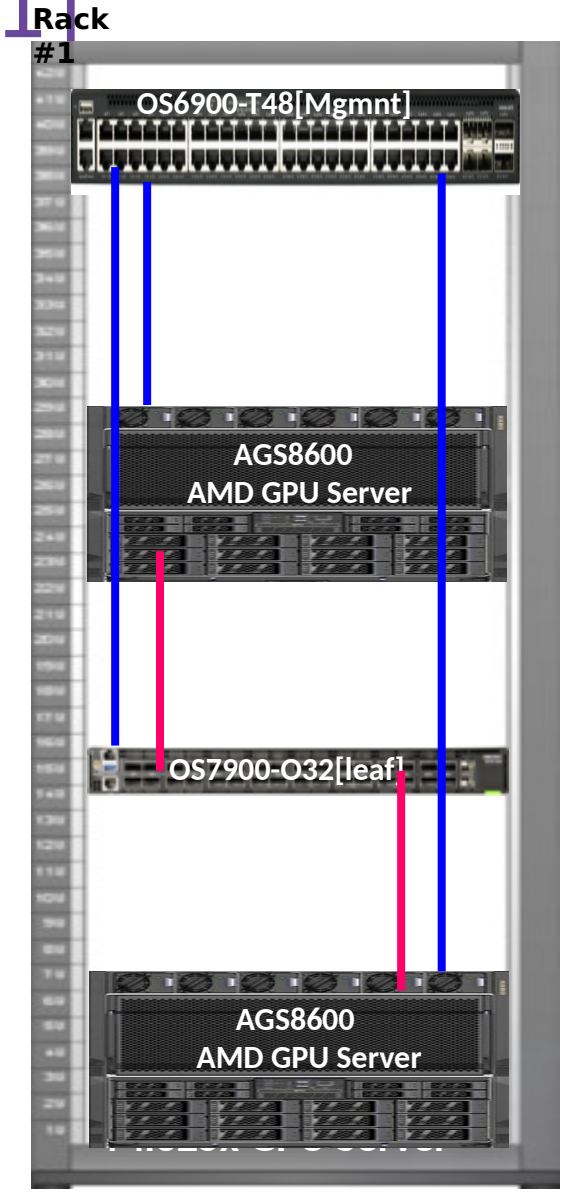

IN RACK AI FABRIC: 400G/800G DAC

Mngmt fabric: 1G/10G

#GPUs: ~16
#GPU Servers: 2
#Rack : 1
#Switches: 1

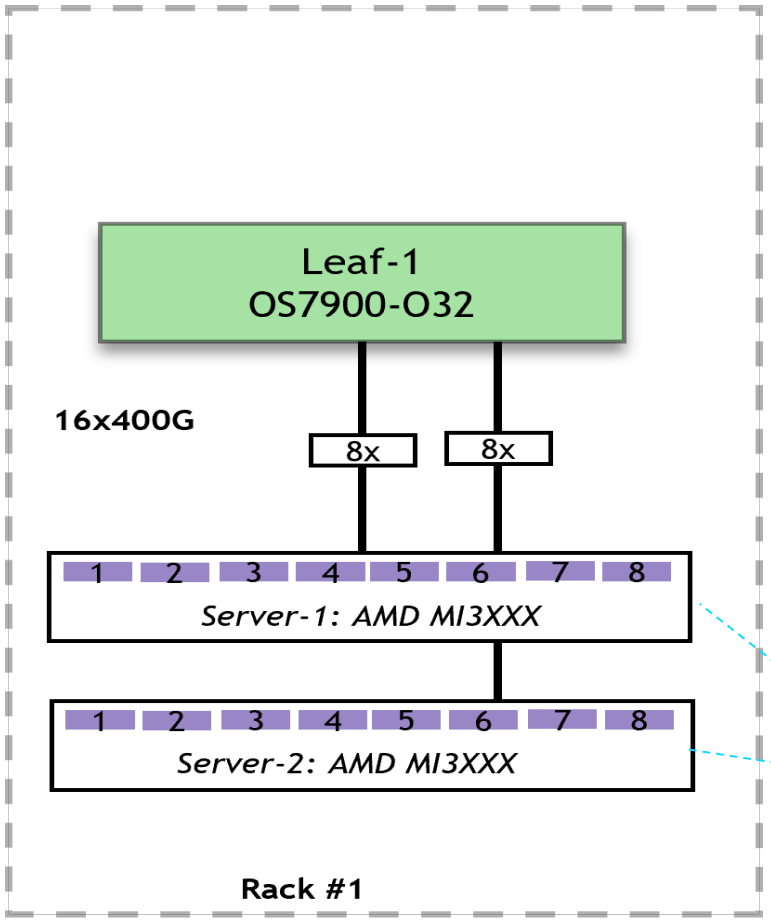

1. Turnkey Single-Rack AIDC solution with two compute servers.
2. Servers, backend network & management share the same rack.

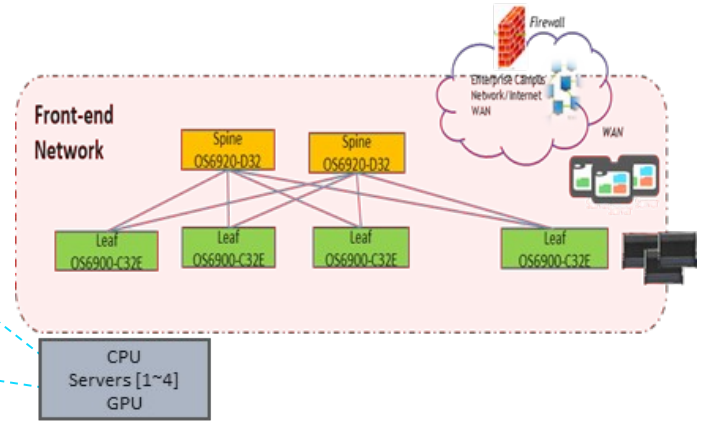

CPU of server connects to front-end network.

GPU BACKEND FABRIC: [16 GPU][LAB-BLUEPRINT-2]

#GPUs: 16
#GPU Servers: 2
#Rack : 3
#Switches: 4

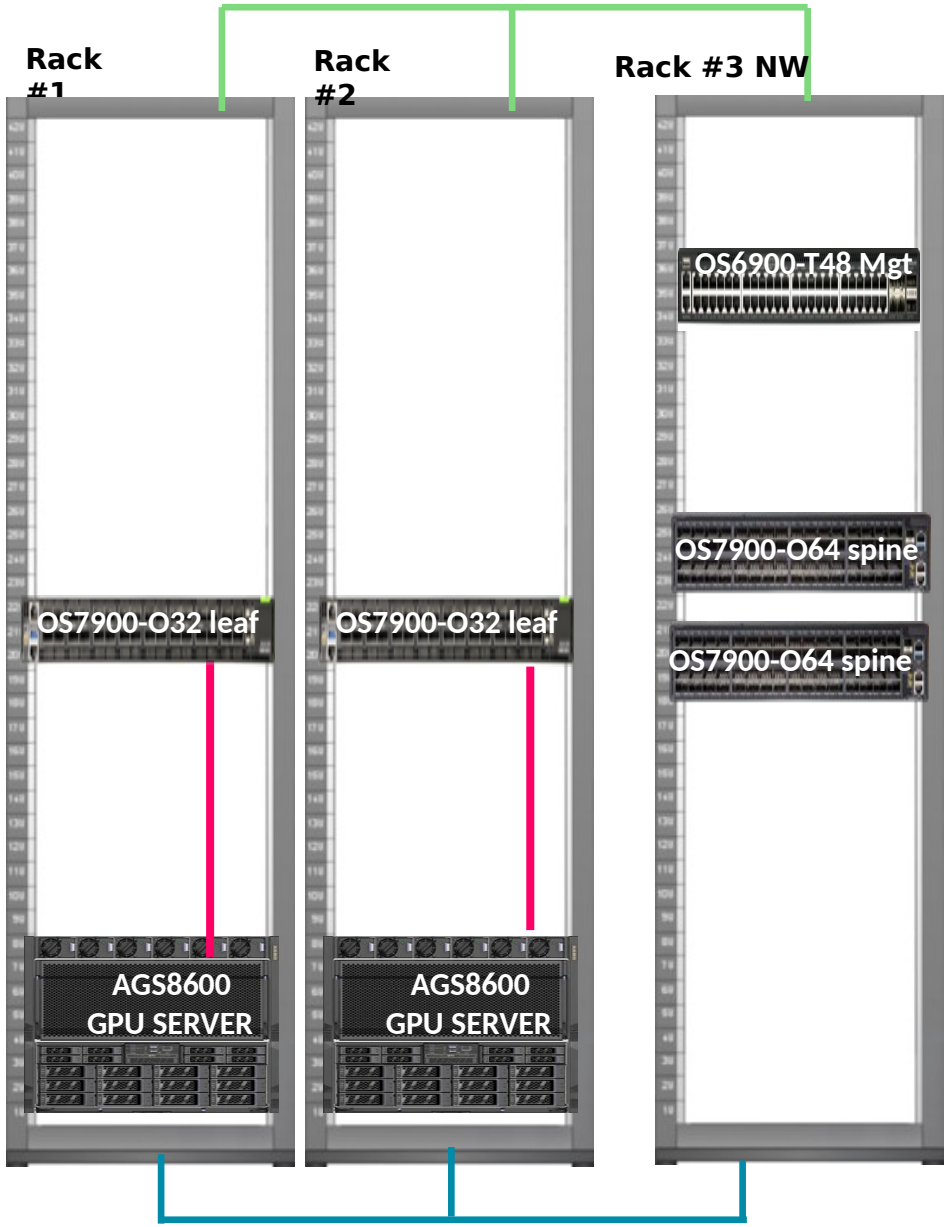

IN RACK AI FABRIC: 400G/800G DAC

INTER RACK AI FABRIC: 400G/800G AOC/Optical Cables

Mngmt fabric: 1G/10G

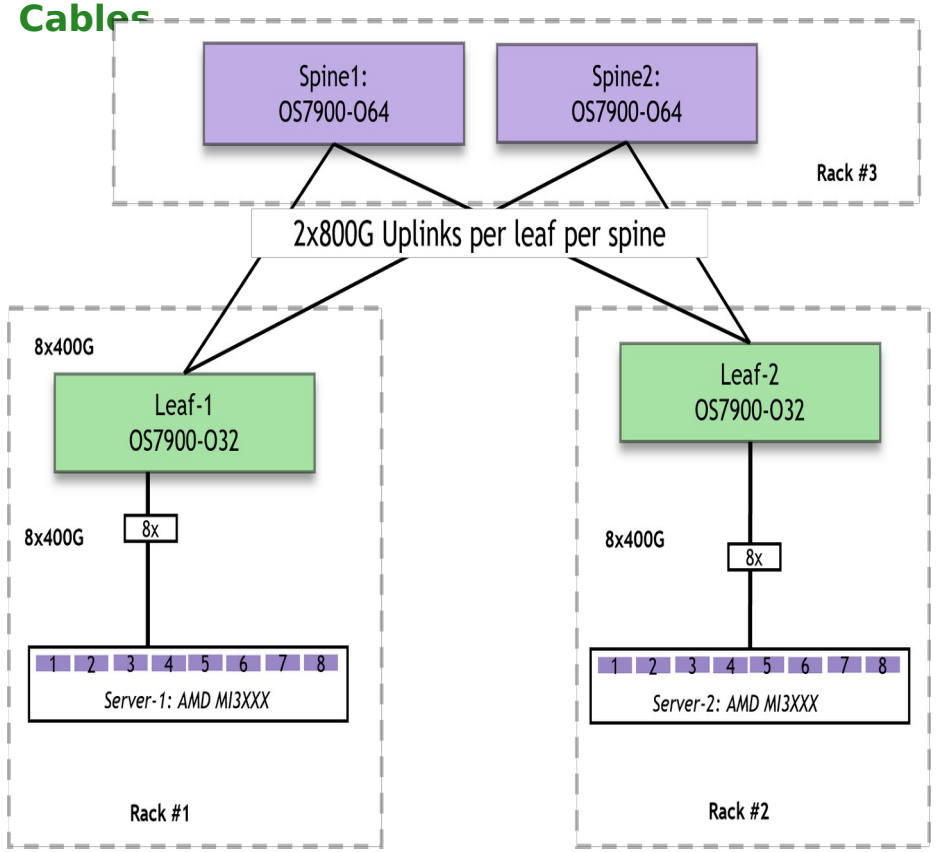

1. Dual-rack AIDC solution scalable to 8 racks.
2. Separate Network Rack for Mgmt/Spine.
3. DAC for intra-rack.
4. AOC/Optical cables for inter rack.

GPU BACKEND FABRIC: CLOS DESIGN[256 GPU

#GPUs: 256
#GPU Servers: 32
#Rack : 9
#Switches: 10

Scales to 8 racks.

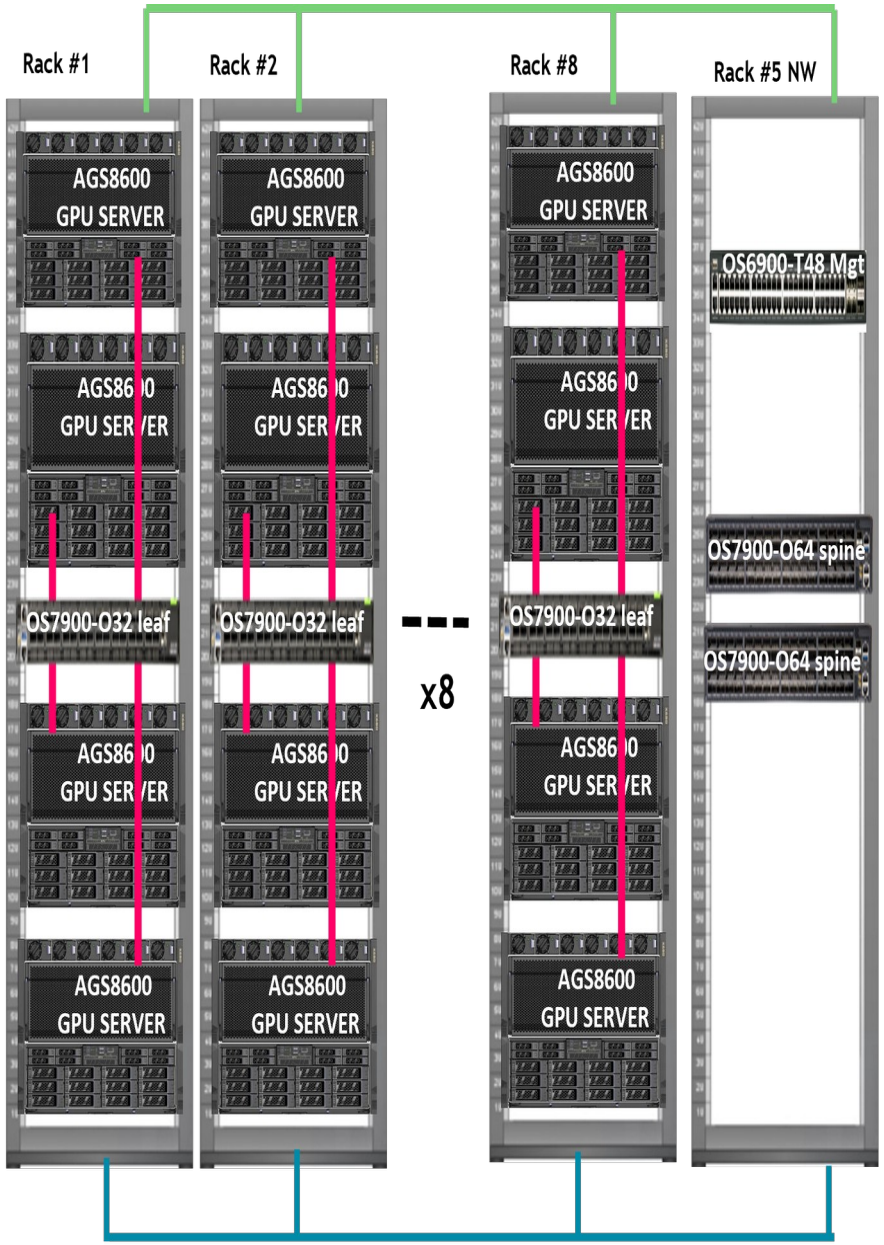

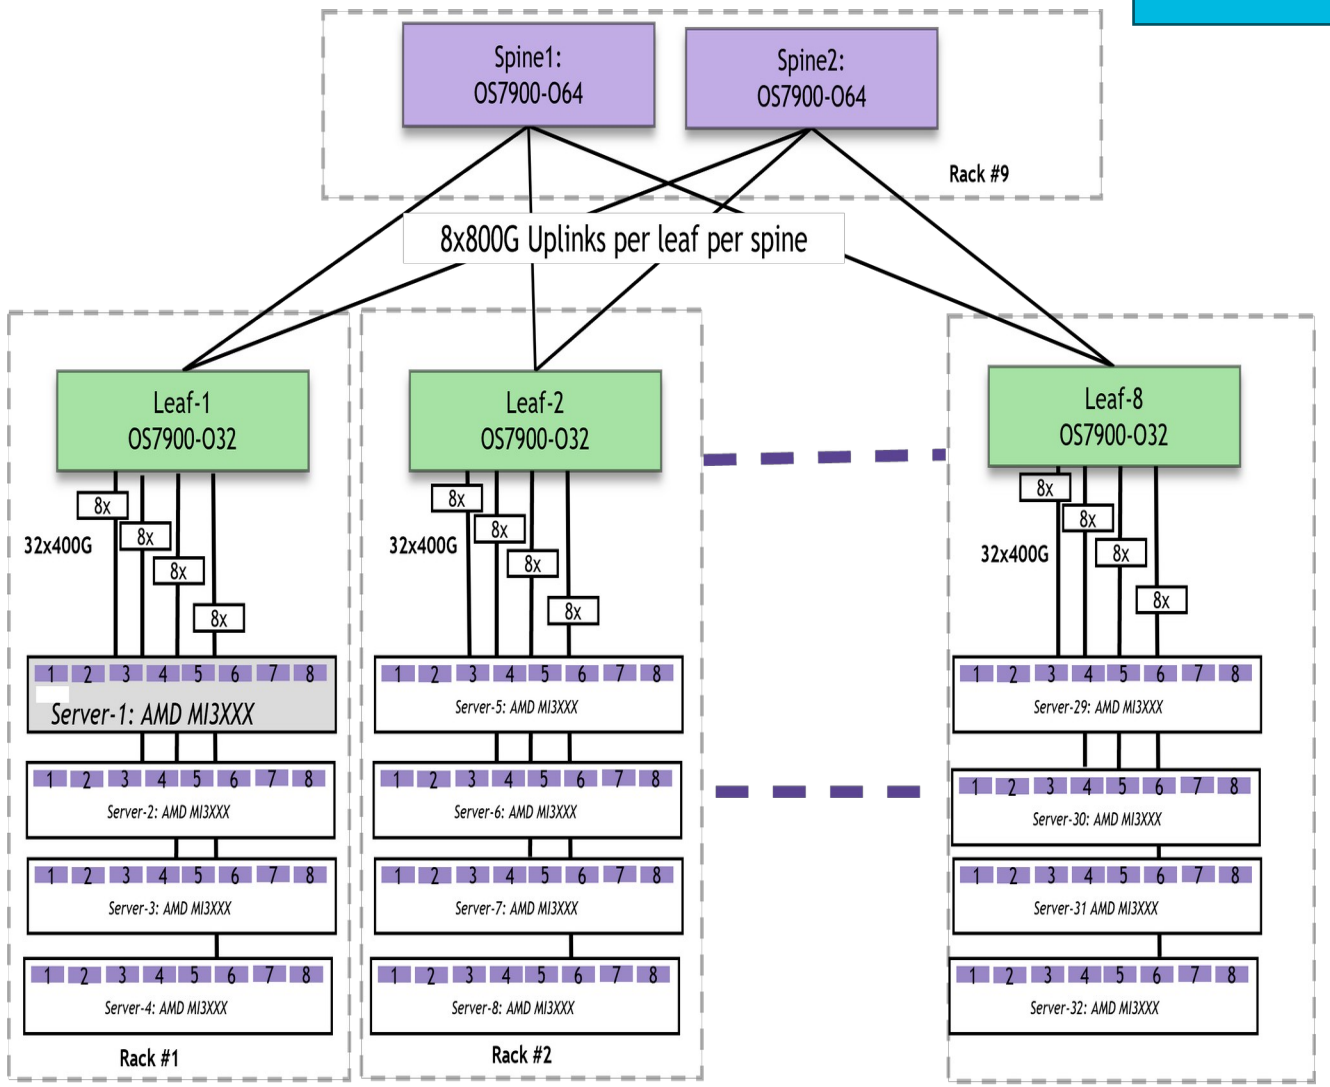

BOM AND PRICE-LIST

Product	Description	List Price
GPU Server	8x MI325x CPUs, NIC Card included	$300,000
800G Switch	64x 800G OSFP Backend switch as spine.	$132,000
800G Switch	32x 800G OSFP Backend switch as leaf	$110,000
800G to 800G Optics	800G SR8 Transceiver (50m)	$879
800G to 2x400G DAC	800G to 2x400G dac for leaf to server, 3m.	$300
Management switch	6900-T48 to connect to backend nw devices.	$34,275
NOS	SONiC(start zero but go to license in future).	$0
Orchestrator	Device Management software	$7000

Blueprint type	ASP Price
Blueprint-1(1-rack)	$753,675
Blueprint-2(2-rack)	$1,141,739

ALE AI-DC SOLUTION MILESTONES / PHASES

ACCOMPLISHMENTS and MILESTONES

- **PHASE 1**
  - Scope:
    - OEM offer with minimal Branding on switches.
      - ALE logo on overlay, Model name
      - ALE Login banner in NOS.
      - ALE logo on silk screen in Mgt. software
    - Documentation
      - HW/SW/CLI/Datasheet modified to ALE.
  - Progress:
    - Requirement established.
    - DR0: Completed on 26th Jan
    - DR1: Planned this week
    - DR4: Planned on 26th Feb
- **PHASE 2**
  - Scope:
    - Transition to ALE-NOS in front-end with AOS-X
    - Transition to ALE SONiC in back-end
  - Progress:
    - DR1: Completed on Dec 2025
    - DR4: Planned July 2026
- **PHASE 3**
  - Scope:
    - Dual NOS on back-end: AOS-X & ALE SONiC
    - ALE brand optics for 800/400G
    - Transition to ALE mgmt. software

PHASE-WISE ROADMAP

Feb 2026

PHASE 1: Technical Proof of Concept / MVP

- Back-end network : OEM 800G / 400G, OEM EC SONiC
- Front-end network : ALE 100G/25G, OEM EC SONiC
- Management network: ALE 10G,
- Optics : ALE (100G, 400G) OEM (400G, 800G)
- Compute : Edgecore AMD Server (MI325x GPU)
- Rack mgmt. SW : OEM EDM v2.3

July 2026

PHASE 2: Solution Launch

- Front-end network : +ALE 400G, AOS 10.1R2
- Back-end network: ALE SONiC
- Rack mgmt. SW: OEM EDM v2.4

Dec 2026

PHASE 3: SCALE

- Front-end network : AOS 10.1R3
- Back-end network: AOS 10.1R3, ALE SONiC
- Optics : ALE optics 10/100/400/800G
- Rack mgmt. SW: ALE OVCompass

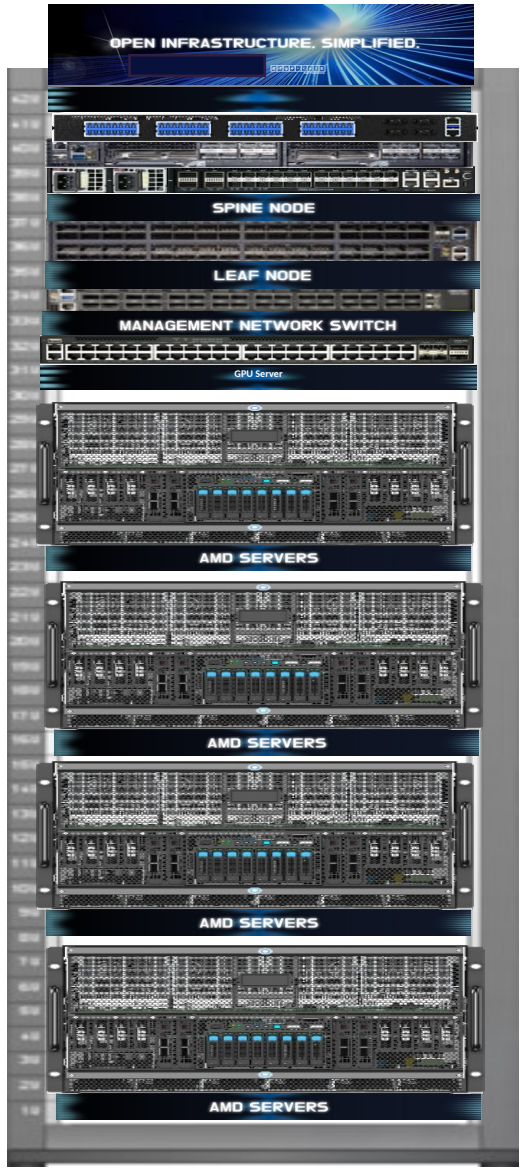

GO TO MARKET

Go-to-Market Strategy

Focus Countries:

- **Tier 1:** *France, Germany, UK, Nordics, UAE, Saudi* Arabia
- **Tier 2**: Rest of EUSO/EUNO/MEITA Regions and APAC
- **Tier 3**: USA/CANADA and CALA

Target Customer:

Enterprise, Private, and Regional Cloud AI DC opportunities (Tiers 2-3–4)

Sales & Channel Strategy:

- **Phase 1**: Indirect-only model through DC specialized ALE existing DR and VAD
- **Phase 2**: Direct sales to strategic customers where permitted by regulation and no channel conflict
- **Partner Recruitment:** Proven DC expertise and active Government, DCSP, and Telco relationships

Customer Profile

Sovereign / National AI Cloud Projects:

- Ministries
- Government-managed - cloud initiatives
- National AI clusters
- Sovereign data centers

Enterprise AI Transformation DCs

- Banks
- Energy/Utilities
- Telcos
- Aviation
- Manufacturing
- Healthcare
- Universities/ Research centers
- Transportation

Tier 2–3 Cloud Providers

- Regional cloud companies
- Managed service providers
- Emerging AI cloud platforms
- Colocation + cloud hybrid players

AI-First Companies

- AI startups
- HPC/
- Data analytics firms
- GenAI platform companies

Enterprise AI DC (ME)

	Vertical	Enterprise Example Type	Why Enterprise AI DC
UAE	Energy	ADNOC	Reservoir & seismic modeling AI sovereign
UAE	Aviation	Emirates Group	Ops AI + maintenance advantage
UAE	National AI	G42 Gov workloads	Sovereign AI DC
UAE	Healthcare	M42 Genome	Genomics AI must remain sovereign
UAE	Banking	FAB	AML + risk models remain private
SAUDI	Energy	Aramco	Reservoir modeling IP strategic national asset
SAUDI	Petrochemicals	SABIC	Material innovation AI IP protected
SAUDI	Gov Data	DGA/NDMO	National data sovereign AI
SAUDI	Smart City	NEOM	Sector AI DC for next gen smart region
SAUDI	Banking	Saudi National Bank	Financial AI private by design

Sovereign / National AI Cloud DC

	Example Sovereign Actor(s)	Why Sovereign / National AI Cloud DC
FR	GENCI / CEA / SecNumCloud certified providers (3DS Outscale / OVH Sovereign)	National HPC + AI model capacity under French jurisdiction. Strategic science AI + national data residency by law.
GER	JUPITER Exascale / Federal AI Research Centers (FZJ Jülich)	Exascale AI with national control — national industry AI + research compute protected.
UK	UKAI Compute (AI Research Resource) / National AI Research Resource (NIRR)	National AI cluster for AI R&D capacity – UK regulated, sovereign model training environment not public cloud multi-tenant.
UAE	National AI Authority programs + Core42 national compute + Khazna sovereign partitions	National AI model training + national data security priority. Sovereign DC strategy for next gen AI economy.
SAUDI	PIF backed national AI DCs (Humain initiative) / SDAIA national data compute	National strategic compute = sovereign asset. AI cluster capacity aligned to national transformation programs Vision 2030.

Tier 3 AI Cloud DC

	Tier-3 Local / National Private Cloud Provider Example Type	Why Tier-3 (Local Private Cloud)
FR	OVHcloud Enterprise Private Cloud instances (non-SecNumCloud) / Scaleway	Country & EU focused cloud. Limited global footprint. Commercial private cloud hosting for regulated enterprise workloads.
GER	IONOS Enterprise Cloud / PlusServer	National + EU compliance private cloud — but not AI specialized Tier-2 and not sovereign HPC state controlled.
UK	UK national telco private hosted cloud (BT Business / Vodafone Business Cloud)	Local private cloud sold commercially to UK enterprises — not global hyperscale, not sovereign designated workload.
UAE	Etisalat/DU Business Cloud private cloud offers	Country-bounded hosting for UAE enterprise workloads. Not national AI authority and not regional Tier-2 GPU capacity builder.
SAUDI	STC Cloud Enterprise hosting / Zain KSA Enterprise Cloud	National telco cloud DC used commercially by enterprises in the country. Not sovereign mandated DC and not hyperscale / Tier-2 GPU provider.

#WhereEverythingConnects

Why Everyone Is Using AI Today

AI Domain	PyTorch	TensorFlow	Industry Reality
Classical ML	✅ (via PyTorch + scikit-learn)	✅ (TF + Keras + TFX)	Both widely used
Deep Learning (DL)	✅	✅	Both core frameworks
Computer Vision	✅ (torchvision, Detectron2, YOLO, OpenCV)	✅ (TF Vision, KerasCV)	PyTorch dominant in research & SOTA
NLP	✅ (Hugging Face, spaCy, Fairseq)	✅ (TF Text, HF TF backend)	PyTorch dominant
GenAI	✅	⚠ Limited	PyTorch is the de-facto standard
LLMs	✅ (Transformers, vLLM, DeepSpeed, Megatron)	⚠ Rare	PyTorch overwhelmingly standard
RAG	✅ (LangChain, LlamaIndex, Haystack)	⚠ Indirect	PyTorch ecosystem dominant
Multimodal AI	✅ (CLIP, LLaVA, Flamingo-like stacks)	⚠ Limited	PyTorch dominant
Reinforcement Learning (RL)	✅ (RLlib, Stable-Baselines3, CleanRL)	⚠ Limited	PyTorch dominant
Agentic AI	✅ (LangGraph, AutoGen, CrewAI, OpenAI SDKs + PyTorch models)	⚠ Rare	PyTorch dominant

Data Center Investments Are Shifting Toward AI

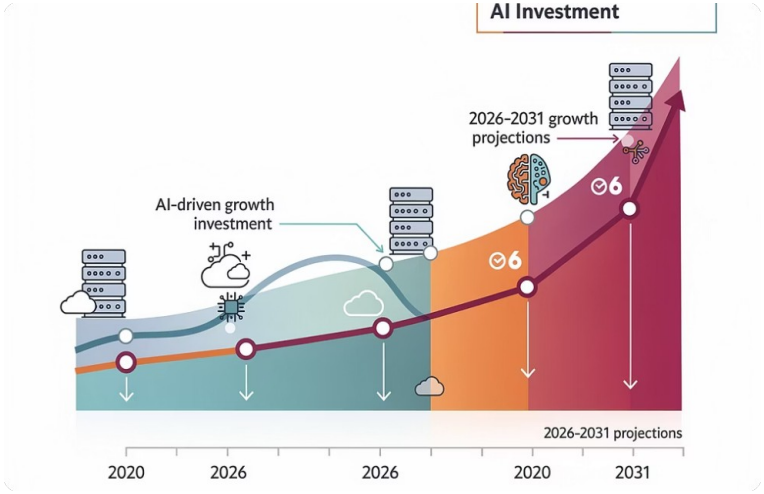

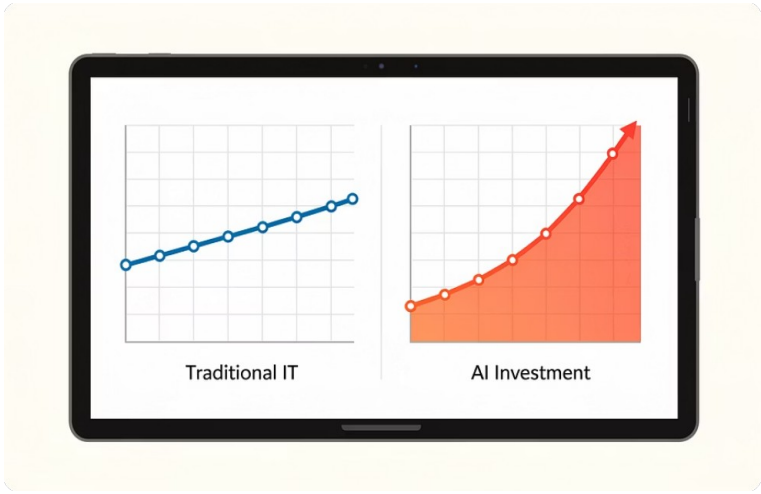

AI-Driven Infrastructure

GPU compute, high-speed networking, and advanced cooling are now the primary investment drivers, supporting cutting-edge AI workloads.

Future Growth in AI

Over the next five years, most incremental data centre spending will be AI-related, fuelled by Generative AI, analytics, automation, and sovereign AI initiatives.

Strategic Capital Allocation

Traditional data centre investments remain stable, while new capital is strategically directed towards accelerating AI capabilities and infrastructure.

Highlight: "While traditional Data Centres remain essential, nearly all new growth in Data Centre investment is now c

Why Data Centers Are ALE Strategic New Focus

Portfolio diversification and business transformation:
Expanding into Data Centers allows us to diversify away from technologies that are becoming legacy, commoditized, and lower-value, and reposition the company toward strategic, future-proof growth areas.

Stronger economics and long-term value creation:
The Data Center segment offers structurally larger deal sizes, longer customer lifetime value, and multi-year engagements, supporting more predictable and sustainable revenue.

Increased strategic relevance with enterprise customers:
Data Centers sit at the core of digital transformation initiatives (AI, cloud, sovereignty), enabling earlier engagement with C-level decision makers and positioning us as a long-term strategic partner.

Enterprise AI DC (Europe)

	Vertical	Enterprise Example Type	Why Enterprise AI DC
FR	Banking	BNP Paribas / Société Générale	Sovereign + AML + Risk AI must be private
FR	Aerospace	Airbus	Safety critical simulation + IP protection
FR	Automotive	Renault Group	AI for manufacturing + digital twin
FR	Pharma	Sanofi	R&D AI drug discovery data sensitivity
FR	Energy	EDF	Critical national infra + sovereign AI
GER	Automotive	BMW / Mercedes-Benz / VW	Autonomous & simulation IP must remain sovereign
GER	Industrial	Siemens	Industrial AI + factory digital twin
GER	Pharma	Bayer	Drug discovery AI is core R&D IP
GER	Energy	E.ON / RWE	Critical infrastructure optimization AI
GER	Banking	Deutsche Bank	Internal fraud/risk AI must be internal
UK	Pharma	GSK / AstraZeneca	Drug discovery AI IP protected
UK	Telecom	BT	Network AI optimization private
UK	Defense	MoD/DSTL	National security AI
UK	Media	BBC R&D	Generative Media AI IP advantage

Tier 2 AI Cloud DC

	Tier-2 Specialized Provider Example Type	Why Tier-2 (Specialized GPU AI Cloud)
FR	OVHcloud AI GPU Clusters / NexGen Cloud (EU)	multi-regional GPU cloud expansion model, AI compute specialization, faster provisioning vs hyperscalers
GER	Northern Data (EU HQ)	EU scale GPU compute provider; not national telco, not hyperscaler; AI compute capacity sellers
UK	Vultr GPU PoPs (London) / DigitalOcean (Paperspace-origin)	GPU cloud builder expansion into UK; specialized capacity provider not full cloud hyperscaler
UAE	Vultr GPU Dubai / regional Tier-2 GPU startup footprints	Tier-2 providers expanding into GCC region as GPU capacity sellers for enterprises needing faster access
SAUDI	Tier-2 GPU entrants (regional PoP builds) / EU Tier-2 rented into KSA	“bring-in” Tier-2 GPU clusters into Saudi via regional PoPs → specialized GPU supply, not hyperscaler, not sovereign