

# Kolo Partner Briefing: Secure, Scalable AI at the Edge



# Kolo

## Why Read This Guide?

If you're working with enterprise clients exploring AI, edge compute, or hybrid cloud infrastructure, this guide is for you. It will help you understand how Kolo's metro-edge infrastructure can solve your customers' most pressing challenges—and how you can lead those conversations with confidence.

### Use this guide to:

- Understand what sets Kolo apart in the AI infrastructure market
- Identify which customers are a fit for metro-edge deployments
- Learn how to qualify opportunities and handle common objections
- Access tools and resources to support your sales efforts



## Contents

### [3. The market opportunity](#)

### [4. Who This Helps and Their Challenges](#)

### [5. Why Customers Choose Kolo](#)

### [6. The Opportunities to Target: Metro Edge AI & Hybrid Cloud Use Cases](#)

### [7. Customer Use Cases](#)

### [12. How to Position Kolo](#)

### [13. Objections You Might Hear—and What to Say](#)

### [14. Get Started Today](#)

# The Market Opportunity

- AI and machine learning workloads are shifting to the edge to reduce latency and improve responsiveness.
- Enterprises need a secure, sovereign, and compliant environment to train and deploy AI models.
- Public cloud alone doesn't address cost, data sovereignty, or high-performance GPU needs.
- Kolo enables a hybrid strategy: run sensitive, high-density workloads in our metro-edge facilities while maintaining flexible connections to public cloud.

---

**>50%**

of enterprise data is now  
created and processed outside  
of traditional data centers.

---

**80%**

of enterprises plan to run AI  
workloads at the edge within  
the next 24 months.

---

# Who This Helps and Their Challenges



## Industries

Kolo solutions are relevant across a broad range of industries, but especially relevant for enterprises in:

- Healthcare
- Manufacturing
- Financial Services
- Public Sector
- SaaS / ISVs

These sectors benefit most from reduced latency, data sovereignty, and AI-ready infrastructure at the edge.

## Key Personas

- CTO / CIO / Chief Architect
- Head of Infrastructure / Digital Transformation
- VP AI / Data Science / Cloud Platform

## Common Challenges

### Ensuring regulatory compliance

Many are navigating strict requirements around data residency, GDPR, or sector-specific rules like those in healthcare, finance, and public sector AI deployments. AI training and inference often involves sensitive datasets that must remain within regional boundaries.

### Reducing the risk of cyberattacks.

Security is critical, especially for AI models that handle proprietary or sensitive data. Public internet-based cloud exposure increases the attack surface. Enterprises seek private, dedicated connectivity options.

### Minimizing the risks of migration

Shifting to a hybrid or edge infrastructure must be simple, fast, and without disruption. Enterprises need confidence that workloads can move between environments without compatibility or downtime concerns.

### Ensuring regulatory compliance

IT systems must support the business to meet local data compliance and governance legislation.

### Reducing time-to-market

AI innovation cycles require environments that can launch quickly. Delays in GPU provisioning from the public cloud or long deployment cycles from traditional DCs slow innovation.

### Lowering IT costs

Public cloud GPU usage introduces pricing volatility, especially under sustained workloads. Enterprises want predictable infrastructure costs for budgeting and cost optimization

### Delivering consistent user experiences

Whether it's real-time imaging in healthcare or analytics dashboards in finance, users expect consistent, low-latency access to applications and data. That's only possible with infrastructure placed closer to the users

# Why Customers Choose Kolo

Kolo enables enterprises to confidently run high-performance, AI-ready workloads in secure, sustainable, metro-edge environments. Our hybrid model supports real-world deployments where data gravity, compliance, and latency all matter.

## Here's how Kolo delivers:



**Metro-edge proximity** to major European industrial and metro hubs enables AI inference and hybrid cloud workloads with sub-5ms latency to users or public cloud regions.



**High-density, GPU-ready colocation** (up to 50 kW+ per rack) gives clients consistent access to compute.



**Seamless hybrid connectivity** with direct interconnects to AWS, Microsoft Azure, and Google Cloud makes it easy to train in the cloud and infer at the edge.



**Compliance and sovereignty-first architecture** ensures workloads stay in-region, with ISO certifications, GDPR adherence, and support for upcoming EU AI Act alignment.

In short: Kolo delivers the flexibility of the cloud with the performance and control of on-prem, in locations that meet today's AI and data sovereignty needs.

# The Opportunities to Target: Metro Edge AI & Hybrid Cloud Use Cases

The Kolo platform is uniquely positioned to address the core infrastructure challenges your customers face—from AI performance demands to cloud integration and compliance pressure. Our hybrid-ready, metro-edge model allows you to deliver fast, local AI compute that connects directly to the cloud, enabling more intelligent and secure operations across industries.

## Explore use cases

1. Real-time AI inference at the edge
2. Hybrid AI deployments
3. Sovereign, compliant compute for regulated workloads
4. Disaster recovery for AI and analytics platforms
5. Multi-cloud interconnect and workload portability

## How to Qualify an Opportunity

Ask questions like:

- Are you running or planning AI/ML workloads?
- Is latency or compliance a concern?
- Are you frustrated by cloud GPU pricing or delays?
- Do you need to connect local compute with public cloud?
- Are compliance or data residency concerns influencing your cloud or AI infrastructure decisions?

## USE CASE 1

# Real-time AI inference at the edge

## Customer challenges

Enterprises need to process data instantly for critical decisions—whether it's for medical imaging, industrial quality control, or smart logistics. Routing inference through cloud introduces unacceptable latency and data compliance risks.

## How Kolo helps

Kolo's metro-edge colocation enables AI models to run inference as close as possible to the data source. This reduces response times, limits cloud dependency, and ensures compliance with sensitive data processing requirements.

## Kolo Cloud Access

Kolo provides the continuous connectivity required for hybrid deployments between systems and cloud-hosted models. Through our low-latency interconnects to AWS, Azure, and Google Cloud, enterprises can seamlessly bridge edge infrastructure with public cloud AI pipelines—without compromising on privacy, performance, or compliance.

---

## Exploratory questions

- How do latency requirements impact your AI models and real-time decision-making?
  - Are there use cases that cannot tolerate roundtrip delays to the cloud?
  - Where does inference happen today—and what challenges does that create?
-

## USE CASE 2

# Hybrid AI deployments

## Customer challenges

Many organizations want to train AI models in the public cloud, but need low-latency, high-performance environments to run inference closer to end users—or to ensure data doesn't leave the region.

## How Kolo helps

Kolo provides seamless interconnection with AWS, Azure, and GCP while hosting inference workloads locally in sovereign environments. Our cloud access services deliver the continuous, low-latency connectivity required for hybrid AI deployments between cloud environments and local GPU infrastructure—ensuring high performance without compromising on privacy, security, or regional compliance.

## Kolo Cloud Access

Kolo provides the continuous connectivity required for hybrid deployments between systems and cloud-hosted models. Through our low-latency interconnects to AWS, Azure, and Google Cloud, enterprises can seamlessly bridge edge infrastructure with public cloud AI pipelines—without compromising on privacy, performance, or compliance.

---

## Exploratory questions

- Do you train in the cloud but need local performance for inference?
  - How critical is it for your data or models to stay in-region?
  - Are you seeing cloud GPU shortages or cost spikes affecting project rollouts?
-



## USE CASE 3

# Sovereign, compliant compute for regulated workloads

## Customer challenges

Governments, financial institutions, and healthcare providers require full control over data location, access, and processing due to regulations like GDPR, HIPAA, or the EU AI Act.

## How Kolo helps

Kolo delivers secure colocation in sovereign jurisdictions with full compliance alignment. Customers retain control of their infrastructure, supported by certified facilities and optional EMP-protected environments..

## Kolo Cloud Access

Kolo provides the continuous connectivity required for hybrid deployments between systems and cloud-hosted models. Through our low-latency interconnects to AWS, Azure, and Google Cloud, enterprises can seamlessly bridge edge infrastructure with public cloud AI pipelines—without compromising on privacy, performance, or compliance.

---

## Exploratory questions

- Are you subject to data sovereignty regulations like GDPR or the EU AI Act?
  - What level of control and auditability do you require over your infrastructure?
  - Do you need high-security or air-gapped environments for sensitive workloads?
-

## USE CASE 4

# Disaster recovery for AI and analytics platforms

## Customer challenges

AI and analytics platforms are increasingly business-critical. Downtime affects operations, compliance, and revenue. Public cloud DR models may not be cost-effective or compliant.

## How Kolo helps

Kolo supports disaster recovery strategies through regionally distributed data center sites with redundant network connections, high-availability configurations, and support for cold or warm DR environments. While we do not currently offer managed DRaaS as a bundled service, our platform enables enterprises and partners to build resilient GPU-based recovery solutions with flexible deployment models.

## Kolo Cloud Access

Kolo provides the continuous connectivity required for hybrid deployments between systems and cloud-hosted models. Through our low-latency interconnects to AWS, Azure, and Google Cloud, enterprises can seamlessly bridge edge infrastructure with public cloud AI pipelines—without compromising on privacy, performance, or compliance.

---

## Exploratory questions

- What is your current RTO/RPO for AI infrastructure?
  - Do you have DR for GPU workloads or high-performance analytics?
  - Are you concerned about cloud-only failover models meeting compliance or performance expectations?
-

## USE CASE 5

# Multi-cloud interconnect and workload portability

## Customer challenges

Enterprises running workloads across multiple cloud providers face rising egress costs, inconsistent performance, and operational complexity. They need a neutral platform to interconnect environments.

## How Kolo Helps

Kolo's direct interconnects to major cloud platforms support fast, secure data movement across environments. Customers can stage, route, and control workloads across clouds—and deploy local infrastructure as an anchor point.

## Kolo Cloud Access

Kolo provides the continuous connectivity required for hybrid deployments between systems and cloud-hosted models. Through our low-latency interconnects to AWS, Azure, and Google Cloud, enterprises can seamlessly bridge edge infrastructure with public cloud AI pipelines—without compromising on privacy, performance, or compliance

---

## Exploratory questions

- Do you run across multiple clouds today—or plan to?
  - Are you experiencing high costs or limitations moving data between cloud environments?
  - Would centralized interconnection or localized infrastructure improve your performance or compliance posture?
-

# How to Position Kolo

Kolo is built for the next wave of enterprise infrastructure, where AI, compliance, and hybrid strategies intersect.



## Built for AI workloads

Kolo provides consistent, high-density infrastructure optimized for GPU deployments—ideal for training, inference, and hybrid AI architectures.



## Cloud-adjacent metro-edge design

Our data centers are strategically located near major European industrial and metro hubs, minimizing latency to both end users and public cloud regions.



## Hybrid-ready by design

We integrate seamlessly with AWS, Azure, and GCP, enabling clients to connect local workloads with cloud-based pipelines.



## Compliance-first architecture

All locations are GDPR-compliant and aligned to ISO standards, with support for customers addressing NIS2 and EU AI Act requirements.



## Partner-centric flexibility

Kolo works with MSPs and integrators to deliver private cloud, DRaaS, and managed GPUaaS—extending your ability to deliver tailored solutions.



## Speed to deployment

With available space and retrofit capabilities, we help customers go live in weeks—not quarters.



## High-security offerings for critical workloads

Select Kolo locations are designed to meet the highest security demands of government, defense, and regulated industries. EMP-shielded suites, private cages, and 24/7 guarded access ensure physical and logical security far beyond standard colocation environments.

# Objections You Might Hear—and What to Say

## Objections and concerns you could face

## How to overcome them

We use the cloud for everything.



Kolo complements cloud. Use us for latency-sensitive, secure, or cost-sensitive workloads.

We don't want to manage infrastructure.



Kolo works with partners who offer managed services, GPUaaS, bare metal, and turnkey DRaaS.

Colocation seems more expensive than the cloud.



Compared to long-term GPU rentals and egress fees, Kolo provides better predictability and lower TCO.

Is it secure and compliant?



Yes. We meet GDPR, ISO, and EU AI Act standards—and never touch your data.

Why should we choose Kolo over large global providers or hyperscalers?



Kolo offers a tailored, high-touch alternative to hyperscale and global platforms. We prioritize partner relationships, regional focus, and rapid deployment—without the complexity or premium pricing of global incumbents. With direct access to decision-makers and faster turnarounds, Kolo helps your clients move faster, stay compliant, and get more value from their infrastructure.

# Get Started Today

Let's help your customers deploy smarter AI—securely, locally, and reliably. With Kolo.

---

Identify 2–3 accounts that need faster, smarter AI infrastructure

Use Kolo materials to initiate the conversation

Register the lead with your Kolo Sales Director or AE

Set up a joint workshop or solution design call

---

# Data centers



Kolo operates a growing platform of strategically located data centers across Northern Europe—offering low-latency access, sovereign control, and high-performance infrastructure tailored for AI and hybrid workloads.

## Sweden

- **Stockholm:** 2 high-security facilities with dedicated AI capacity
- **Piteå:** 1 high-security ideal for AI-deployments

## Denmark

- **Copenhagen:** Premium facility with AI capacity
- **Jutland Region:** AArhus, Skanderborg (with AI capacity), Viby (with AI capacity), Randers

## The Netherlands

- **Amsterdam:** Metro connectivity hub
- **Regional Data Centers:** Meppel, Dronten, Apeldoorn

All locations feature resilient interconnects, sovereignty-aligned hosting, and scalable infrastructure optimized for enterprise, and hybrid workloads.

## Custom Retrofit Solutions

Beyond our established network, Kolo specializes in rapidly delivering custom retrofit solutions for high-density workloads in new locations. Our expert team can transform facilities into state-of-the-art data centers optimized for your specific requirements—with an industry-leading Ready-for-Service (RFS) timeframe of just 6-8 months.

This accelerated deployment model enables enterprises to quickly establish sovereign, high-performance infrastructure in strategic locations not yet served by our existing footprint.

## About Kolo

Kolo is a next-generation colocation and infrastructure platform built for the AI era. We operate a growing network of secure, high-performance data centers across Northern Europe—strategically located near major industrial hubs and metro areas. Our mission is to enable enterprises and partners to deploy latency-sensitive, AI-powered, and compliant digital infrastructure without compromise.

### What sets Kolo apart:

- **Metro-edge focus:** Close to users, close to data, close to cloud.
- **GPU-ready capacity:** High-density racks, optimized for AI training and inference workloads.
- **Hybrid-ready interconnectivity:** Seamless links to AWS, Azure, and Google Cloud.
- **Sovereign infrastructure:** 100% European ownership, with facilities aligned to GDPR, ISO, NIS2, and EU AI Act requirements.
- **Secure by design:** Optional EMP shielding, private cages, and 24/7 physical protection for high-security deployments.
- **Partner-first approach:** We work closely with MSPs, integrators, and resellers to build tailored hybrid and edge solutions.

Whether your clients are building AI pipelines, modernizing legacy infrastructure, or securing sensitive workloads, Kolo is the platform that bridges cloud agility with physical control.





# Thank you!