

CASE STUDY

MARCH 2026

The Infrastructure Bet That Changes Everything



madlen



INFERCOM

Table of Contents

03	Executive Summary
04	The Challenge: Three Problems Every Expanding EdTech Faces
07	The Solution: Infercom Sovereign AI Infrastructure
11	What the Partnership Delivered
14	Why This Matters: The Broader Context
17	What Comes Next
19	About the Companies

Executive Summary

While most EdTech platforms quietly route their AI through servers based in the United States — subject to US law, US oversight, and the ever-present shadow of the CLOUD Act — Madlen made a different call.

Madlen is an AI-native school operating system built for teachers. With over 50,000 students, 20,000 teachers, and more than 300,000 hours of administrative time saved, it is one of the fastest-growing EdTech platforms in Turkey and across the MENA region. But growth alone was never the goal. The goal was to build something schools could trust — not just to work, but to safeguard the people inside them.

That is why Madlen partnered with Infercom: Europe's first sovereign AI inference provider, headquartered in Luxembourg, operating out of certified data centers in Munich. All data processed through Madlen's AI tools — student performance records, teacher content, assessment data — stays within European jurisdiction. GDPR compliant. EU AI Act ready. ISO 27001:2022 certified. Zero exposure to foreign surveillance law.

This case study explains what that means in practice: for school directors making compliance decisions, for education ministries evaluating vendor risk, and for every school that has ever asked, "Where exactly does our data go?"



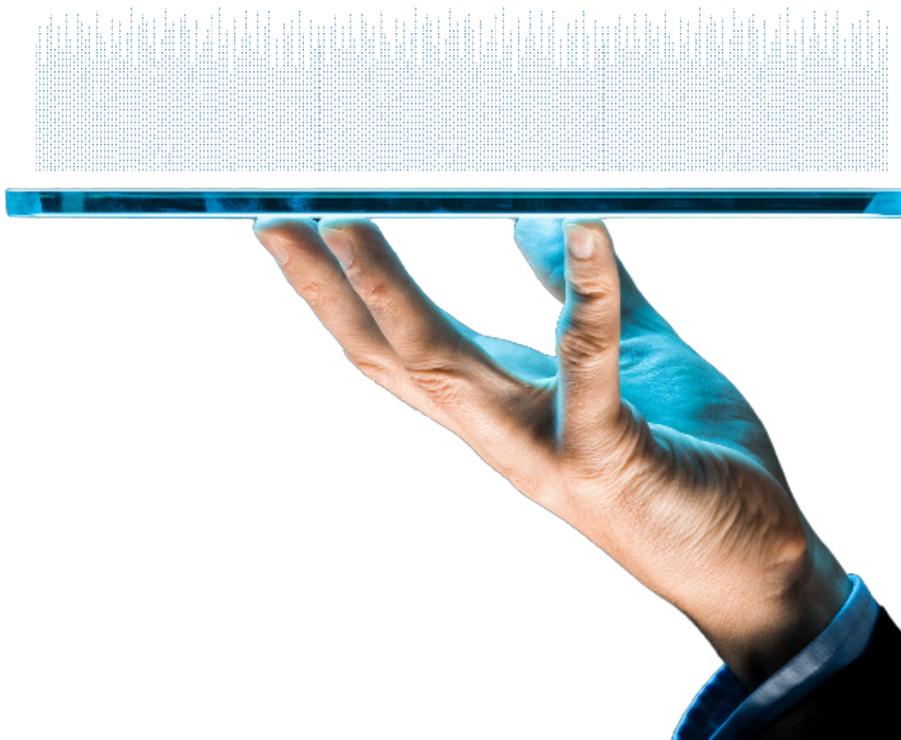
The Challenge: Three Problems Every Expanding EdTech Faces

1

Data Sovereignty and Regulatory Compliance

Entering European markets is not simply a matter of translating an interface. The EU General Data Protection Regulation (GDPR) and the newly enacted EU AI Act — which classifies education as a high-risk AI application domain — impose strict requirements on how student data is collected, stored, processed, and transferred. Platforms that route sensitive educational data through US-based cloud providers expose their institutional clients to CLOUD Act and PATRIOT Act risk: the possibility that a foreign government agency could, under certain legal conditions, access that data without the knowledge of the school or the students involved.

For school directors and education ministry procurement teams, this is not a theoretical concern. It is a dealbreaker. When a Turkish EdTech platform wants to serve a German school network, a French public institution, or a Scandinavian university, it cannot afford ambiguity on data residency. It needs proof.



The Challenge: Three Problems Every Expanding EdTech Faces

2

Performance Under Peak Load

Madlen's platform includes over 40 AI-powered tools: a lesson plan assistant, an exam generator, a composition evaluator, a differentiation engine, and more. These tools are used simultaneously by thousands of teachers — and not at random intervals. Usage spikes sharply at the start of the school day, during exam periods, and around assignment deadlines. The AI infrastructure beneath those tools must absorb these peaks without degradation.

Traditional GPU-based cloud inference systems handle this in one of two ways: they either over-provision capacity (expensive) or they queue requests and introduce latency (unacceptable for a teacher mid-lesson). Neither option fits an EdTech platform trying to compete on quality and scale simultaneously.



The Challenge: Three Problems Every Expanding EdTech Faces

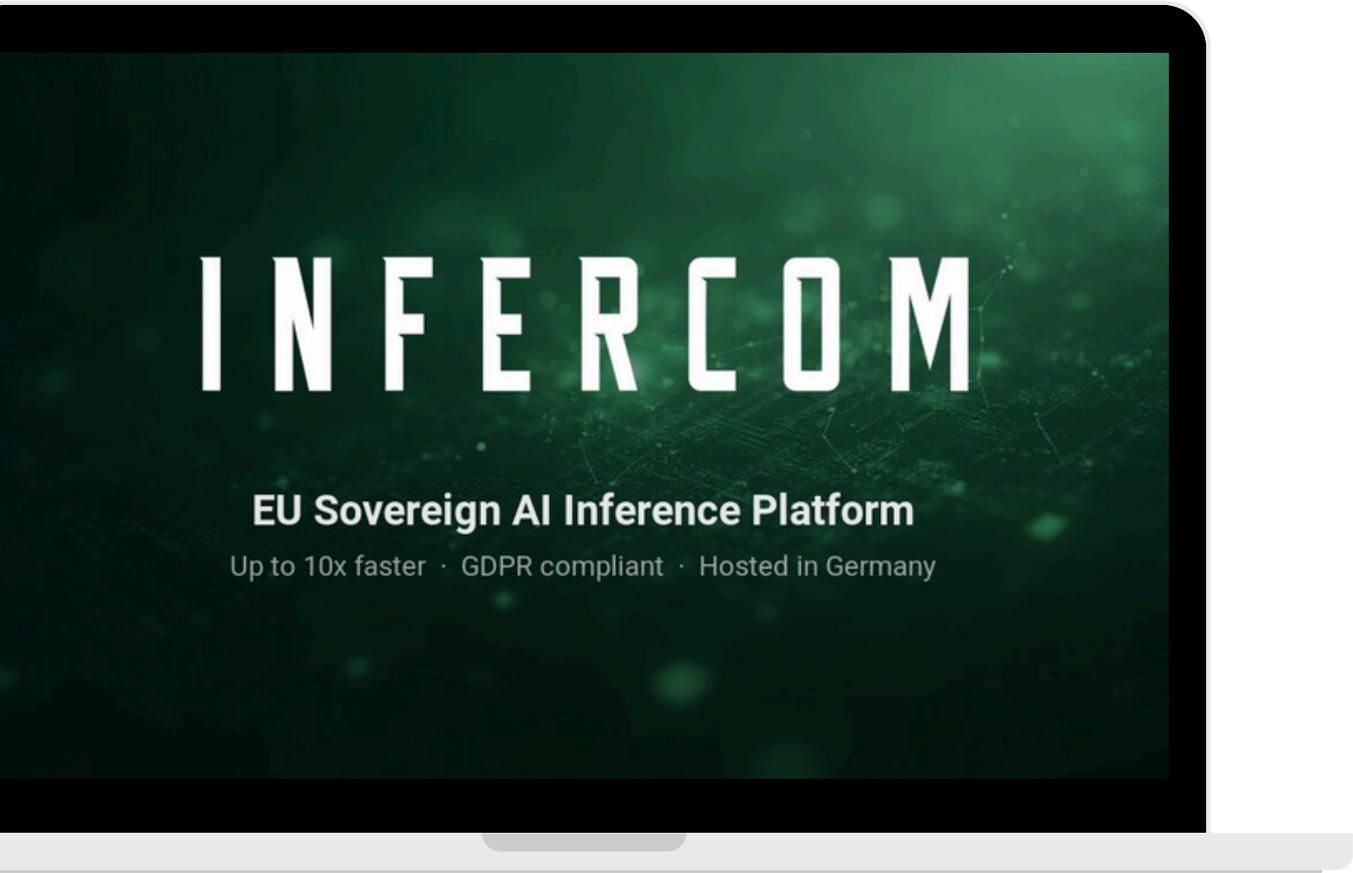
3 Unit Economics and Market Accessibility

Education is not enterprise software. Schools have defined per-student budgets. The cost of AI inference is not an abstract infrastructure line — it directly determines how much Madlen can charge per seat, which markets it can enter, and whether it can build a business model that sustains long-term access rather than pricing out the institutions that need it most. High inference costs are a structural barrier to expansion into developing markets, including MENA and Southeast Europe, where Madlen has identified significant demand.



The Solution: Infercom Sovereign AI Infrastructure

Madlen resolved all three challenges — compliance, performance, and cost — through a single architectural decision: migrating its AI inference layer to Infercom's European-native platform.



INFERCOM

EU Sovereign AI Inference Platform

Up to 10x faster · GDPR compliant · Hosted in Germany

The Solution: Infercom Sovereign AI Infrastructure

What Infercom Is

Infercom is Europe's first sovereign AI inference provider. Founded in Luxembourg, it operates GPU-free AI inference out of Equinix Tier III+ data centers in Munich, Germany — entirely within EU jurisdiction. Its infrastructure is built on SambaNova's Reconfigurable Dataflow Unit (RDU) technology, an architecture designed specifically for large language model inference at scale. The result is a platform that is simultaneously faster, more energy efficient, and more compliant than conventional GPU cloud alternatives.

Think of Infercom as what AWS would look like if it were built from scratch in Europe, for European institutions, with European law as a first-order design constraint — not an afterthought.



GDPR Compliant by Design

EU-hosted models run on dedicated hardware at Equinix Munich, Germany. Your data never leaves the European Union. Full GDPR compliance is built into the infrastructure — not bolted on after the fact.



Zero Data Retention Policy

Your prompts and outputs are never stored, logged, or used for model training. Unlike major AI providers, we don't learn from your data. What you send stays yours.



No CLOUD Act Exposure

This is where compliance and sovereignty diverge. US-parented providers — even with EU data centers — remain subject to the CLOUD Act. Our infrastructure is operated under EU jurisdiction only.



AI Act Ready

Built for EU AI Act requirements from day one. Transparent operations, auditability, and the documentation you need for regulated AI use cases.

The Solution: Infercom Sovereign AI Infrastructure

The Integration: One Endpoint Change

Because Infercom's APIs are fully OpenAI-compatible, the technical migration was remarkably straightforward. Madlen's existing AI pipeline — content generation for teachers, adaptive learning for students, analytics for administrators — was redirected to Infercom infrastructure with a single API endpoint change. No retraining of models. No rebuilding of workflows. No vendor lock-in risk.

This simplicity matters. It means that Madlen's engineers spent their time on product, not on infrastructure archaeology. And it means that other platforms considering a similar move face a much lower technical barrier than they might expect.

SambaNova Dataflow Architecture

The performance advantages of Infercom's infrastructure are not marginal. Compared to GPU-based inference systems, SambaNova's RDU architecture delivers up to 10x faster inference and 5x better energy efficiency. The three-tier memory system and native BF16 precision support allow large language models to run at full quality — no quantization, no hidden accuracy tradeoffs. For Madlen's use case, this translates directly to teacher experience.

The exam generator returns results in seconds, not minutes. The lesson plan assistant responds before the teacher has second-guessed the prompt. At peak times — Monday morning, the week before national exams — the platform holds steady. The infrastructure earns its invisibility.

The Solution: Infercom Sovereign AI Infrastructure

Data Sovereignty: The Architecture of Trust

This is the differentiator that matters most to institutional buyers. All data processed through Madlen's AI tools is handled exclusively in Infercom's Munich data centers, under German and EU law. There is no path — technical or legal — by which a US government request could compel access to that data.

The exposure to the US CLOUD Act and PATRIOT Act is zero.

For a school director signing a contract that commits them to protecting the data of hundreds or thousands of minors, this is not a footnote. It is the entire conversation.

Technical Infrastructure Summary

Inference Engine	SambaNova RDU (Dataflow Architecture)
Data Center	Equinix Munich, Tier III+
Legal Jurisdiction	German / EU Law
API Compatibility	OpenAI-compatible endpoints
Certifications	ISO 27001:2022, GDPR, EU AI Act Ready
Model Support	DeepSeek V3.1, Minimax M2.5, GPT-OSS 120B
Deployment Options	Managed / Dedicated / On-Premises
Performance	10x faster than GPU, 5x energy efficient

Results: What the Partnership Delivered



Performance

The transition to Infercom's infrastructure produced measurable improvements across Madlen's most-used AI tools. Latency dropped during peak hours — the moments that matter most for usability. Concurrent user capacity increased. The exam generator, the lesson plan assistant, and the composition evaluator all became noticeably more responsive. Teachers who had developed workarounds for slow response times found they no longer needed them.

The energy efficiency gains matter too, not only as a sustainability metric but as a long-term cost signal. A 5x improvement in energy efficiency at the infrastructure level creates compounding economics as usage scales.



Results: What the Partnership Delivered



Regulatory Compliance as a Competitive Asset

Before the Infercom partnership, Madlen's compliance posture in European markets was aspirational. After it, compliance became infrastructural — baked into the architecture, not added as a layer of policy documentation on top of a non-compliant system.

This distinction matters to procurement teams. An ISO 27001:2022 certification backed by data residency in certified German facilities is a different kind of guarantee than a privacy policy that promises data protection while routing data through US servers. In public procurement processes, education ministry evaluations, and school board due diligence reviews, the Infercom partnership gives Madlen a response to the most important question — "Where is our data?" — that no competitor currently offers.



Results: What the Partnership Delivered



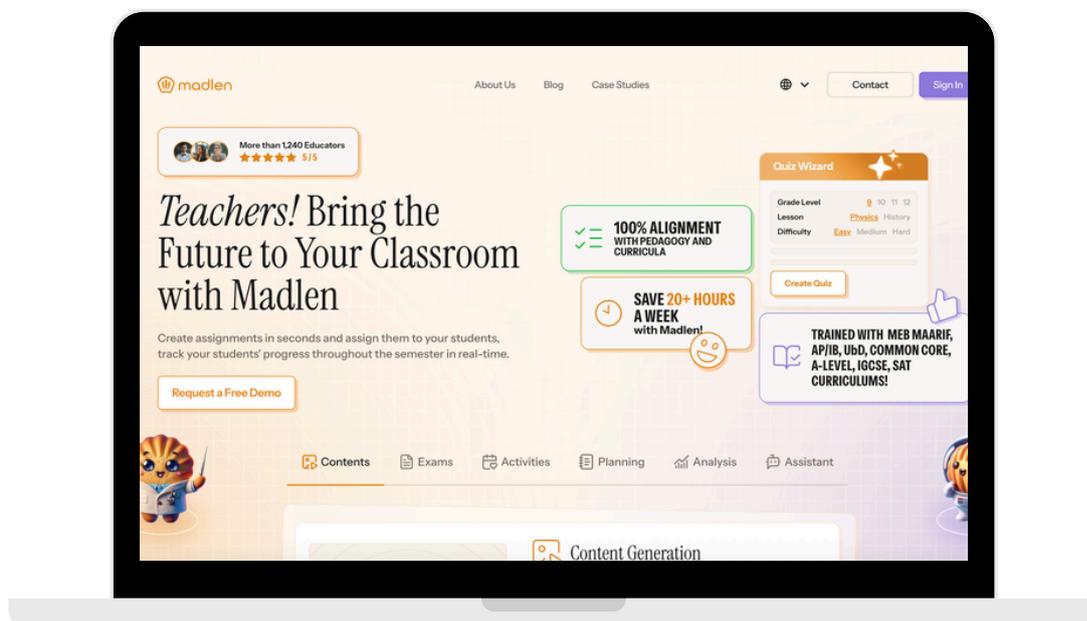
Strategic Repositioning

Perhaps the most significant result of the Infercom partnership is not quantitative. It is the shift in how Madlen is understood — by clients, by investors, and by the market.

Madlen is no longer simply an AI application for teachers. It is an education infrastructure platform: **sovereign, compliant, scalable, and built on an architecture that treats student data as a protected resource rather than a commodity.** That repositioning changes the sales conversation, the partnership landscape, and the long-term valuation story.

"When we enter European markets, data sovereignty is not a luxury. It is a prerequisite. With Infercom, we meet that prerequisite without compromising on performance."

— Atakan Özkaya, Founder & CEO, Madlen

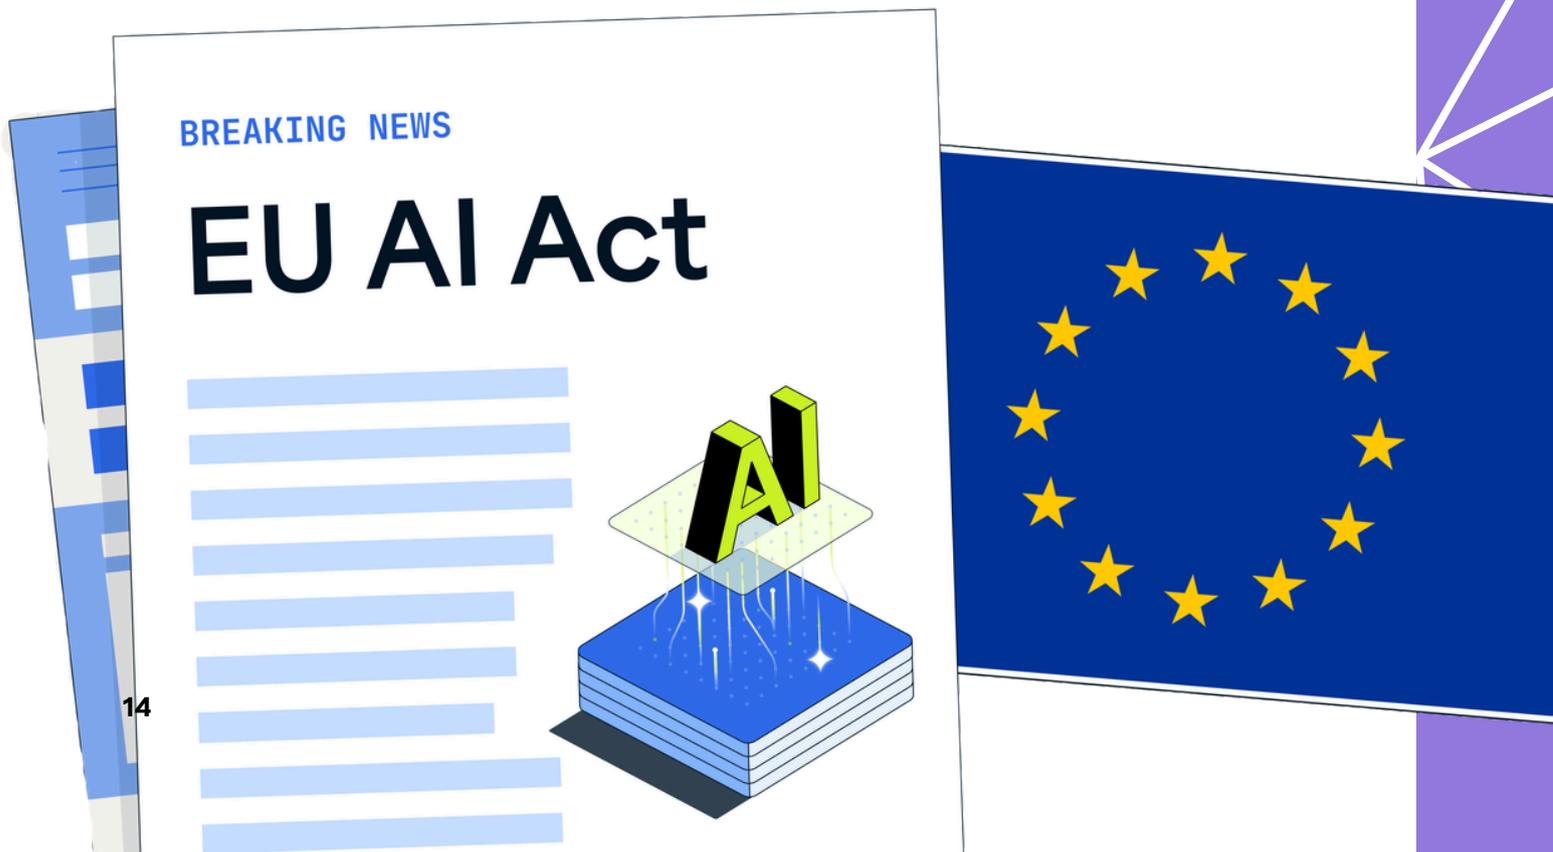


Why This Matters: The Broader Context

The EU AI Act Changes the Rules

The EU AI Act, which came into force at the end of 2024, classifies education as a high-risk AI application domain. This means EdTech platforms operating in Europe are subject to stricter oversight, more rigorous audit requirements, and higher standards for transparency and accountability than software in lower-risk categories. Platforms that process student data outside EU jurisdiction will carry an elevated risk profile for institutional buyers — and that risk will only grow as enforcement ramps up.

The Madlen-Infercom partnership is a proactive response to this regulatory environment. It is not a compliance patch applied after the fact. It is an architectural choice made before the regulation fully bites — which means Madlen arrives at the table already in compliance, while competitors are still calculating the cost of catching up.



Why This Matters: The Broader Context

The Turkish EdTech Ecosystem Goes Global

Turkey's EdTech market reached \$2.2 billion in 2024 and is projected to exceed \$6 billion by 2033. The companies driving that growth are the ones successfully translating local success into international scale. **Madlen's recognition as the number one EdTech tool in MENA at GES Awards and its Times Higher Education Asia award** demonstrate that this translation is already underway.

But international scale in education requires more than a good product and an aggressive go-to-market strategy. It requires the infrastructure credibility to walk into a ministry of education, a national school network, or a university consortium and answer every compliance question with documentation rather than reassurances. The Infercom partnership provides that credibility.



Why This Matters: The Broader Context

Infrastructure as Philosophy

There is a deeper alignment between Madlen and Infercom that goes beyond technical compatibility. Both companies have made the same bet: **that the most important thing you can build in AI is not the model — it is the trust layer around the model.**

For Infercom, that means sovereign infrastructure. For Madlen, that means designing every AI tool around the real needs of teachers and the real obligations of schools. The partnership is a natural extension of a shared philosophy.



What Comes Next

2026 - 2027

Medium Term

Near Term

Long Term

Near Term (6-12 Months)

- Full operational deployment in European pilot schools on Infercom infrastructure
- Transition to dedicated capacity model with guaranteed performance SLAs
- Regulatory compliance documentation package for education ministries and public institutions

Medium Term (12-24 Months)

- Fine-tuning of Madlen's proprietary education data graph on Infercom infrastructure, creating a compounding model advantage
- Regional inference node deployment for MENA markets
- Joint go-to-market strategy leveraging Infercom's European partner ecosystem for access to institutional buyers

What Comes Next

Long-Term Vision

Madlen's "**teacher intelligence premium**" thesis holds that the platform's value compounds over time as it accumulates richer, more nuanced understanding of how teachers teach and students learn.

That data — the most sensitive kind — belongs in sovereign infrastructure. Infercom's on-premises deployment option represents the furthest expression of this vision: a future in which a school or ministry can run Madlen's AI tools on their own infrastructure, with complete control over their data, and no dependency on any external cloud.

That future is not a product pitch. It is a commitment to the institutions that trust Madlen with the data of the next generation.

"Education is the highest-leverage thing we can fix. Doing it on the right infrastructure is our responsibility to future generations."

— Atakan Özkaya, Founder & CEO, Madlen



About the Companies

Madlen is an AI-native school operating system designed for teachers. It offers over 40 AI-powered tools across content creation, activity planning, assessment, differentiation, and lesson planning. With \$1M+ ARR, 50,000+ students, 20,000+ teachers, and 300,000+ hours saved, Madlen operates across Turkey, MENA, and Europe. It is the GES Awards MENA number one EdTech tool and a Times Higher Education Asia award recipient.

Infercom is Europe's first sovereign AI inference provider. Headquartered in Luxembourg, the company delivers managed, dedicated, and on-premises AI inference built on SambaNova's dataflow architecture. Operating from Equinix Tier III+ data centers in Munich, Infercom provides ISO 27001:2022 certified, GDPR and EU AI Act compliant infrastructure that enables European organizations to run AI at full performance and full compliance – simultaneously.