

Twimbit AI Radar

Roundup of innovative enterprise
deployments and announcements in AI

Copyright © Twimbit 2024



Twimbit is a research and advisory firm driven by a singular mission: to empower businesses that are making a difference. We specialise in providing invaluable industry intelligence to executives and teams, acting as a catalyst for innovation and growth.




Contents

Summary	4
Introduction	6
SK Group and AWS collaborate on sovereign AI infrastructure in South Korea	7
Singtel builds sovereign AI cloud and GPU-as-a-Service ecosystem across Asia	8
Deutsche Telekom & Nvidia launch Europe's first industrial AI cloud	10
Huawei debuts AI Data Lake to ease data bottlenecks in AI workloads	12
Nokia unveils AI Network Fabric to fast-track telco automation	14
Shaping the Next Decade of Telecom	16

Summary

Twimbit AI Radar is a monthly report series that recaps innovative AI deployments and announcements in telecommunications, financial services, and customer experience fields of practice. It also offers insights into companies deploying AI, aimed at assisting business executives and technology leaders develop their own AI projects and long-term strategies.

Company(s)	Deployment/Initiative
	<p>SK Group and AWS collaborate on sovereign AI infrastructure in South Korea</p> <ul style="list-style-type: none">• Partners with AWS to build a 100 MW AI datacentre in Ulsan as part of SK Telecom’s AI superhighway strategy.• Deploys localized GPU-as-a-Service, ML orchestration, and GenAI tools in an AWS “AI zone” to process workloads domestically.• Strengthens South Korea’s data sovereignty, fuels economic growth, and positions the country as a major AI hub.
	<p>Singtel builds sovereign AI cloud and GPU-as-a-Service ecosystem across Asia</p> <ul style="list-style-type: none">• Teams up with Nvidia to roll out energy-efficient AI datacentres across Southeast Asia under the Nxera brand.• Offers GPU-as-a-Service with Hopper GPUs, integrated with Singtel’s Paragon platform for multi-cloud, multi-network orchestration.• Supports Singapore’s AI Strategy 2.0 and regional sovereign cloud goals, enabling enterprises to adopt AI with compliance and scale.

	<p>Deutsche Telekom and Nvidia build Europe's first industrial AI cloud</p> <ul style="list-style-type: none">• Collaborates with Nvidia to launch an "AI Gigafactory" with 10,000 GPUs to serve European manufacturers by 2026.• Provides a sovereign industrial AI cloud hosted entirely under EU data standards, supporting simulations and digital twin workloads.• Advances Europe's technological autonomy, reducing reliance on global hyperscalers while modernizing the manufacturing sector.
	<p>Huawei launches AI Data Lake to tackle AI data bottlenecks</p> <ul style="list-style-type: none">• Introduces an AI Data Lake combining OceanStor high-speed access, Pacific scale-out, and E8000 backup storage.• Integrates Omni-Dataverse for global data visibility and ModelEngine for low-code AI model development, plus xPU pooling and AI O&M.• Helps enterprises handle intensive AI training and inference workloads with optimized, energy-efficient infrastructure.
	<p>Nokia unveils Autonomous Network Fabric for telecom automation</p> <ul style="list-style-type: none">• Expands partnership with Google Cloud to deliver a telco-grade Autonomous Network Fabric across cloud, on-prem, and hybrid environments.• Combines data mesh, observability, explainable LLMs, and agentic workflows leveraging Vertex AI and BigQuery.• Enables operators to move from fragmented systems to adaptive, self-managing networks, improving automation and resilience.

Introduction

Artificial intelligence has rapidly evolved from an experimental promise to a strategic necessity, fundamentally reshaping the telecommunications and digital infrastructure landscape. As networks grow more complex and data volumes soar, AI is emerging as the critical pillar for long-term competitiveness. The global AI in telecom market, valued **at \$2.66 billion in 2025**, is expected to surpass **\$50 billion by 2034**—reflecting an extraordinary **CAGR of nearly 39%**. This trajectory underscores a clear warning: operators that fail to embed AI across their operations risk falling irreversibly behind.

A defining shift underway is the race to build sovereign AI infrastructure. More than **18 telecom operators** across five continents—including Orange, Telenor, Swisscom, Telefónica, Fastweb, Singtel, Indosat, TELUS, Telconet, and Cassava—are investing in **“AI factories,”** GPU-accelerated facilities designed to keep AI workloads and sensitive data within national borders. Governments are amplifying this push, channeling **over \$900 billion into sovereign AI initiatives across 70 countries**.

Across key markets, several operators and technology partners are already demonstrating what this transformation looks like in practice. In South Korea, SK Group and AWS are constructing a sovereign AI ecosystem anchored by a **100 MW** data center, local GPU-as-a-Service, and advanced generative AI platforms to safeguard national interests and drive economic growth. In Southeast Asia, Singtel is rolling out a regional **GPU-as-a-Service** backbone across Singapore, Thailand, Malaysia, Indonesia, and Japan, positioning itself as a linchpin for sovereign AI cloud services tailored to local governance requirements.

Europe’s Deutsche Telekom is collaborating with Nvidia to build the continent’s first industrial AI cloud—an expansive **“AI Gigafactory”** designed to help manufacturers tap into simulation-first, AI-driven production under robust EU data protection rules. Huawei is tackling a critical operational bottleneck with its **AI Data Lake**, integrating high-speed, large-scale storage and intelligent orchestration to support enterprise AI at scale. Meanwhile, Nokia is redefining network automation with its Autonomous Network Fabric, blending **telecom-trained AI models**, explainable intelligence, and agentic workflows delivered via Google Cloud to enable self-managing, adaptive networks.

Collectively, these efforts mark a decisive industry pivot. The networks of tomorrow will be defined not just by speed and scale, but by sovereign architectures, adaptive intelligence, and AI-driven resilience. This edition of the AI Telecom Radar explores the operators and partners leading this charge—and what their strategies reveal about the competitive landscape ahead.

SK Group and AWS collaborate on sovereign AI infrastructure in South Korea

SK Group, the parent of SK Telecom, is partnering with Amazon Web Services (AWS) to develop a large-scale AI-focused datacentre in Ulsan, South Korea. This initiative is a key component of SK Telecom's broader AI infrastructure superhighway strategy, designed to strengthen data sovereignty and support the country's ambitions to process AI workloads locally.

Under this partnership, SK Group will handle the construction of the 100 MW datacentre, while AWS will establish an "AI zone" inside the facility. This zone will provide enterprises and public sector organizations in South Korea with localized access to advanced AI capabilities, including GPU-as-a-Service (GPUaaS) for high-intensity computing, managed machine learning via Amazon SageMaker, multi-model orchestration through Amazon Bedrock, and enterprise GenAI tools like Amazon Q.



Sources: How SK Telecom is Building AI-Enable data centers (<https://sktelecom.com>)

Key Features & Capabilities

- **Localized AI processing** – Ensures AI data remains within South Korea, aligning with national data sovereignty priorities and reducing latency for local applications.
- **Multi-layered AI infrastructure** – Combines regional datacentres, edge AI deployments, and GPUaaS to support scalable AI training and inference.
- **Integrated ecosystem** – Draws on the combined strengths of SK's telecom, semiconductor, energy, and cloud businesses to create an efficient AI environment.
- **Advanced AI tooling** – Offers organizations access to AWS's managed AI services for building, training, and deploying generative AI and machine learning applications.
- **Economic uplift** – Expected to drive significant employment and contribute to South Korea's goal of becoming a major AI hub.

This collaboration demonstrates how telecom groups and hyperscalers can jointly build sovereign AI infrastructure that accelerates local innovation while supporting strategic goals like data localization, industry growth, and future-ready AI capabilities.

Singtel builds sovereign AI cloud and GPU-as-a-Service ecosystem across Asia

Singtel is actively positioning itself at the center of Asia's sovereign cloud movement by developing an extensive network of advanced AI datacentres and offering GPU-as-a-Service (GPUaaS) capabilities in partnership with Nvidia. This strategy aligns closely with the Singapore **government's National AI Strategy 2.0**, which allocates **US\$767** million over the next five years to support AI training, chip procurement, and the establishment of AI centers of excellence.



Sources: How Singtel develops AI Cloud and GPUaaS across Asia (www.singtel.com)

The approach goes beyond Singapore. Singtel, through its dedicated datacentre arm Nxera, is rolling out a series of energy-efficient AI-ready datacentres across Southeast Asia — including planned facilities in Thailand, Malaysia, Indonesia, and Japan. As an Nvidia Cloud Partner, Singtel integrates Nvidia's Hopper GPUs, which significantly accelerate AI model training and inference, along with Nvidia's AI Enterprise software suite, into these sites.

Singtel's infrastructure also capitalizes on its extensive fixed broadband and submarine cable networks, along with its 5G high-speed connectivity. This is further supported by Singtel's Paragon cloud platform, designed to help enterprises orchestrate AI workloads across multi-network, multi-cloud environments with carrier neutrality to maximize connectivity resilience.

Key Features & Capabilities

- **Sovereign AI cloud infrastructure** – Tailors datacentres and GPU resources to meet local data residency, regulatory, and cultural requirements, supporting sovereign AI objectives in Singapore and beyond.
- **GPU-as-a-Service with Nvidia** – Offers enterprise access to powerful Nvidia Hopper GPUs and AI Enterprise software hosted on Nxera's platform, enabling faster AI training and inference.
- **Multi-country AI footprint** – Establishes a regional AI infrastructure backbone with datacentres in Singapore, Thailand, Malaysia, Indonesia, and Japan.

- **Advanced orchestration & connectivity** – Uses Singtel’s Paragon platform along with submarine and 5G networks to allow flexible, multi-cloud, multi-network AI workload management.
- **Alignment with national AI policy** – Supports Singapore’s AI Strategy 2.0 by expanding AI access and accelerating enterprise adoption, while also positioning Singtel as a partner of choice for government-backed AI initiatives.

This use case illustrates how Singtel is moving beyond traditional telecom services to become a key enabler of sovereign AI ecosystems across Asia, combining specialized datacentre infrastructure, high-performance GPUs, and orchestration platforms to meet the emerging needs of both governments and enterprises.

Deutsche Telekom & Nvidia launch Europe's first industrial AI cloud

Deutsche Telekom (DT) has partnered with Nvidia to develop what it calls the first **industrial AI cloud for European manufacturers**, a sovereign infrastructure designed to accelerate AI-driven production and maintain Europe's competitive edge in global manufacturing. The project, described by DT as an "**AI Gigafactory**," is scheduled for implementation by 2026 and is part of a broader push by European telcos and technology providers to establish regional AI capabilities that align with local data sovereignty and security standards.

The initiative will leverage a massive scale of AI compute — featuring **10,000 GPUs via Nvidia DGX B200 systems and RTX Pro Servers**, running software stacks such as Nvidia CUDA-X, RTX, and Omniverse to support high-end simulation, design, and digital twin applications. Deutsche Telekom will provide the secure infrastructure and handle datacentre operations, sales, and security, ensuring strict compliance with European data protection norms.

This approach reflects a growing trend among European operators and policymakers to prioritize **technological sovereignty**, reducing dependence on non-European hyperscalers for sensitive industrial workloads. DT's strategy is reinforced by its recent launch of Nvidia H100 Tensor Core processors on the Open Telekom Cloud, marketed explicitly as a European alternative to global cloud giants.



Sources: How Deutsche Telekom develops AI Cloud for manufacturing ([www.deutschetelekom .com](http://www.deutschetelekom.com))

Key Features & Capabilities

- **Sovereign industrial AI cloud** – Built to host critical manufacturing AI workloads entirely within European jurisdiction, adhering strictly to EU data protection and cybersecurity requirements.
- **High-performance AI infrastructure** – Equipped with 10,000 Nvidia DGX B200 GPUs and RTX Pro Servers, enabling advanced AI model training, digital twin simulations, and AI-enhanced design workflows.

- **Secure European operations** – Deutsche Telekom will manage all datacentre operations, security, and service delivery, guaranteeing that industrial data is processed under European standards.
- **Sector-focused innovation** – Specifically targets Europe's manufacturing sector, providing the AI resources needed for simulation-first, AI-driven production processes.
- **Alternative to global hyperscalers** – Strengthens the role of European cloud ecosystems like Open Telekom Cloud as sovereign platforms for industrial transformation.


This use case highlights how Deutsche Telekom and Nvidia are responding to Europe's strategic need for autonomous AI infrastructure, enabling manufacturers to leverage next-generation AI capabilities while maintaining control over sensitive data and adhering to local regulatory frameworks.

Huawei debuts AI Data Lake to ease data bottlenecks in AI workloads

Huawei has introduced its **AI Data Lake** solution, designed to help enterprises overcome the intensive data storage and management demands of AI model training and real-time inference. Unveiled at the 4th Huawei Innovative Data Infrastructure (IDI) Forum in Munich, the solution targets a fundamental hurdle in enterprise AI adoption: preparing high-volume, high-speed, cost-efficient, and resilient data environments that can keep pace with evolving AI workloads.

Huawei's approach integrates three key storage products:

- **OceanStor A800** delivers high-performance data access, achieving up to **500GB/s throughput and 24 million IOPS**, and ranks first in the MLPerf™ AI benchmark. It also includes built-in key-value caching to reduce repetitive computations, boosting **GPU utilisation to around 70%**.
- **OceanStor Pacific** offers large-scale, energy-efficient storage, providing up to **4PB in a 2U rack** with an industry-low energy use of **0.25W/TB**. The system enables 2:1 compression via dedicated hardware, with use cases like a German university storing over 50PB while cutting **physical space by 70%** and energy **consumption by 46%**.
- **OceanProtect E8000** provides high-speed backup for AI training data (corpus) and vector databases, with data reduction up to 72:1 and backup rates of 255TB/hour. It also features ransomware detection with a claimed 99.99% detection rate.



Sources: How Huawei comes up with AI Data lakes (www.huawei.com)

The AI Data Lake also layers in:

- **Omni-Dataverse**, Huawei's data management framework, offering a unified global view of data assets across distributed storage. It accelerates data retrieval (locating from **100 billion files in seconds**) and enables intelligent data movement with local cache enhancements.
- **ModelEngine**, an AI tool chain that streamlines data ingestion, preparation, model development, and application deployment, supporting low-code environments and automated evaluations for quicker AI rollout.
- **Resource management and AI O&M**, which includes xPU pooling to share GPU and NPU resources across workloads and an AI Copilot for operations that can resolve up to **80% of common issues**, plus real-time anomaly detection for root cause analysis.

Key Features & Capabilities

- **Optimised AI storage architecture** – Integrates high-performance, scale-out, and backup storage tailored to intensive AI workloads, supporting faster training and more efficient inference.
- **Energy-efficient, space-saving design** – Delivers high density with significantly reduced power and physical footprint, making it viable for large-scale AI deployments.
- **Unified data oversight with Omni-Dataverse** – Offers cross-datacentre visibility and accelerated data movement, enhancing training data accessibility.
- **End-to-end AI enablement** – Through ModelEngine, enterprises get tools for ingestion, data preparation, low-code model building, and seamless deployment.
- **Advanced resource orchestration and AI-driven operations** – Enables GPU/NPU resource pooling across AI tasks, and incorporates AI assistants for monitoring, troubleshooting, and automating common maintenance workflows.

Huawei's AI Data Lake illustrates how infrastructure vendors are tackling one of the biggest constraints in enterprise AI — preparing data environments that can handle the speed, scale, and operational complexity of modern AI initiatives, ultimately smoothing the path to large-scale AI adoption.

Nokia unveils AI Network Fabric to fast-track telco automation

Nokia has introduced its **Autonomous Network Fabric**, a comprehensive solution designed to help telecom operators advance toward fully autonomous networks by unifying observability, analytics, security, and automation across all network domains. The platform integrates telco-specific AI models—including LLMs and machine learning algorithms—along with explainable AI capabilities and a robust data management layer, creating a foundation for next-generation, adaptive networks.



Sources: Nokia launches AI Autonomous network fabric (www.Nokia.com)

A key component of Nokia's approach is its **expanded collaboration with Google Cloud**, enabling the Autonomous Network Fabric to be deployed flexibly: as a SaaS offering on Google Cloud, on-premises through Google Distributed Cloud, or in hybrid environments. This integration allows operators to combine Nokia's telecom data expertise with Google's AI platforms, such as Vertex AI and BigQuery, to support agentic workflows that handle real-time anomaly detection, zero-touch remediation, elastic scale-out, and disaster recovery.

Key Features & Capabilities

- **Unified data management with data mesh architecture** – Aggregates, curates, and publishes network data into reusable products, with low-code/no-code tools for rapid development of automation-ready data assets.
- **360-degree observability & chain of custody** – Provides end-to-end monitoring of data use and quality across the organization, ensuring reliable automation outcomes.
- **Explainable telco-trained AI models** – Uses domain-specific LLMs that not only automate decisions but also provide transparent reasoning, helping operators understand why specific actions are taken.
- **Agentic workflows with Google Cloud AI** – Leverages Vertex AI and BigQuery for intelligent network operations, including anomaly detection, subscriber experience optimization, and automated remediation.

- **Flexible deployment** – Available on Google Cloud, on-premises, or in hybrid configurations, supporting evolving telecom infrastructure needs.

This use case highlights how Nokia is addressing the complexity of network automation by combining deep telecom expertise with advanced AI, data management, and observability tools—enabling operators to transition from fragmented systems to cohesive, adaptive networks capable of self-management and rapid innovation.

Shaping the Next Decade of Telecom

The momentum behind these initiatives signals more than incremental improvement—it marks a structural transformation in how telecom operators and technology partners approach growth, resilience, and market leadership. As the sector races toward sovereign AI infrastructures, advanced data orchestration, and deeply embedded automation, the contours of competition are being fundamentally redrawn.

What unites these diverse strategies—from SK Group and AWS’s sovereign AI superhighway in South Korea, to Singtel’s regional GPU-as-a-Service backbone, Deutsche Telekom’s industrial AI cloud, Huawei’s purpose-built AI data lakes, and Nokia’s agentic network automation—is a shared recognition: controlling the data, compute, and intelligence layers of AI is no longer optional. It is the foundation for securing national interests, meeting regulatory mandates, and unlocking entirely new business models rooted in predictive, self-optimizing capabilities.

For industry leaders, the strategic mandate is clear. Investing in AI-native infrastructure today is about far more than keeping pace with technological hype—it is about building the resilient, sovereign, and adaptive platforms that will determine who thrives in the next decade of telecom and digital transformation. Those that act decisively now are positioning themselves not just to navigate complexity, but to shape the very future of global connectivity and intelligent industry.