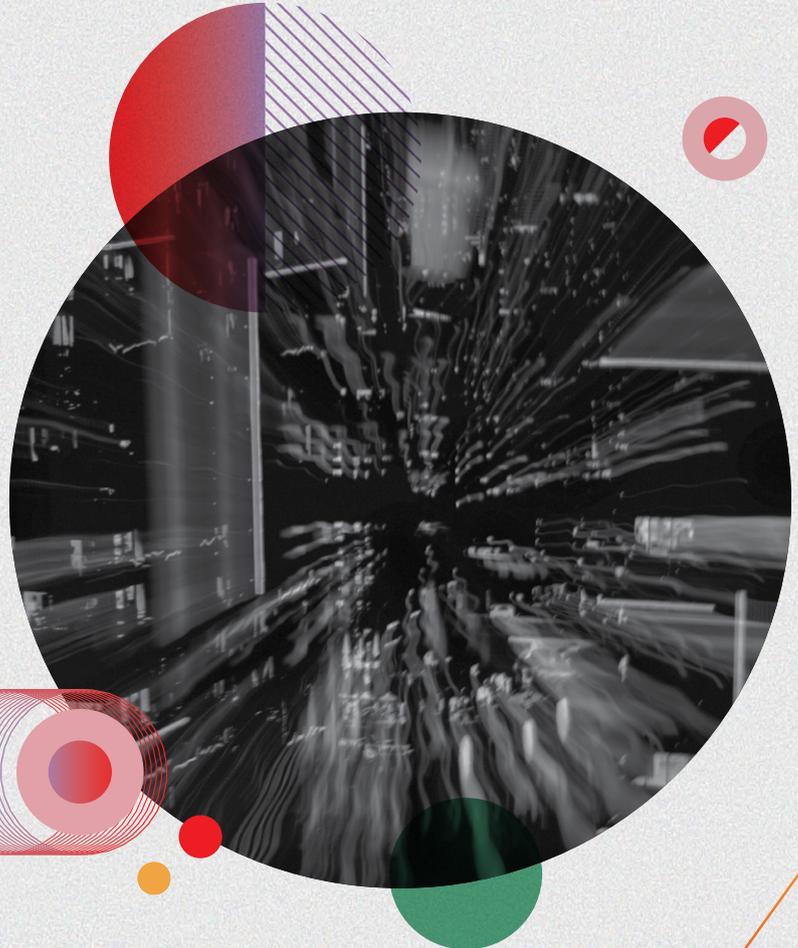


WHITE PAPER

Telecom Gets Smart

How AI-Native Networks Are Redefining the Industry



● KEY TAKEAWAYS

Transforming Telecom with AI-Native Networks

From siloed infrastructure into intelligent, autonomous platforms capable of self-management and real-time adaptation.

Building strong data foundations

Breaking silos, standardising formats, and integrating hybrid storage architectures to power advanced AI operations.

Enabling step-by-step autonomy for CSPs

From building data readiness to unlocking full ecosystem monetisation, with trusted partners like Celfocus guiding the journey.

● WHO SHOULD READ THIS DOCUMENT

- Head of Network Engineering
- Head of Network Operations
- Head of Orchestration
- Head of Network Domain
- Head of Network Virtual Infra
- Network Automation Managers
- Network Operations Managers
- Production/OSS Architecture Teams
- Chief Technology Officer
- Chief Information Officer

CELFOCUS

Why AI-Native Networks Matter



The telecom industry stands at a pivotal moment in its evolution. Traditional network systems are increasingly unable to keep pace with the rapid advancements in digital technology, user expectations, and service demands. With the widespread rollout of 5G and the emergence of 6G on the horizon, Communication Service Providers (CSPs) face mounting pressure to modernise their network strategies.

The solution lies in building **AI-native networks** - intelligent systems that not only support but actively enhance network performance by learning, adapting, and functioning autonomously in real time. Without this shift, CSPs risk losing market share, missing monetisation opportunities, and becoming obsolete in the face of more agile competitors.



Innovation Connected Experiences

Driving influence on the future digital society by providing connected experiences well beyond connectivity



- User Experience hyper-personalisation
- Network-as-a-Service
- TTM & Agility for business innovation



Differentiation Embedded Services

Proactive impact on the industry by providing functionality and data to be embedded in other player's products



- B2B2x Capabilities
- High-Value Partnerships
- Connectivity Excellence

HBO NETFLIX UBER Bundling Over-the-top differentiating services (eg.: HBO, Netflix, Spotify, Uber)



Commodisation Efficient Connectivity

Restricted influence and visibility on the industry by solely providing data transport & connectivity services



- Efficient & Sustainable Operations
- Automation Culture
- Open Network Architecture



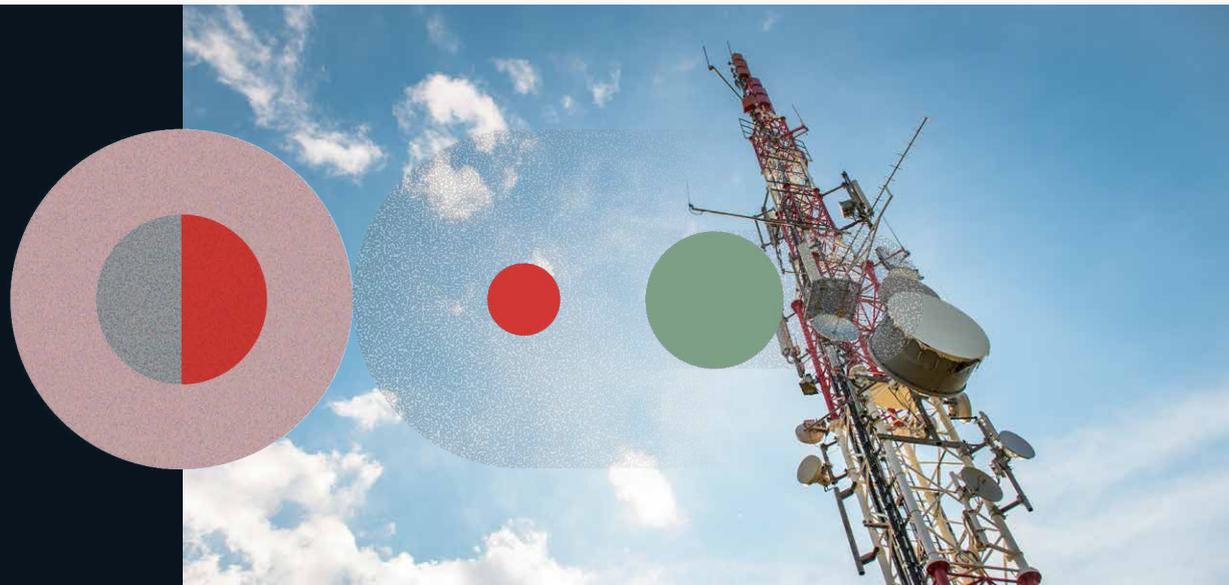
AI-native networks are not merely an upgrade. They represent a fundamental transformation in how networks are architected, deployed, and operated. They enable CSPs to evolve from being utility providers to becoming dynamic platforms for connected experiences, digital services, and real-time intelligence.

Today's **Challenges** for **Telecoms**

Telecom leaders - especially those managing network operations, architecture, and strategy - grapple with a growing list of challenges. These include constant firefighting across fragmented network domains, siloed data repositories that make cross-domain visibility nearly impossible, and underwhelming returns from early AI implementations. Many CSPs have invested heavily in 5G infrastructure, only to find that monetising the new capabilities remains elusive.

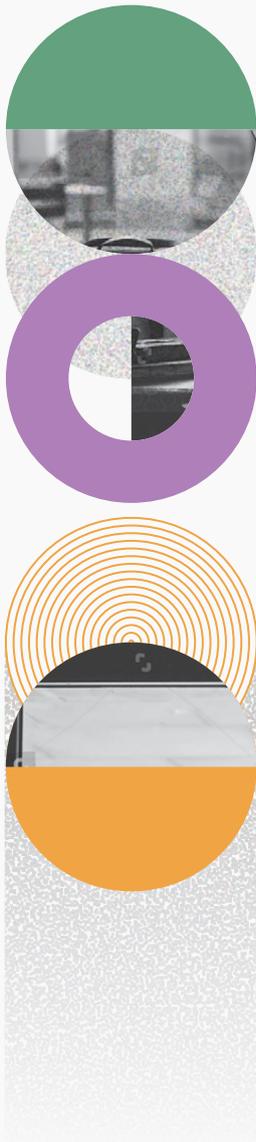
User demands continue to evolve, with expectations around instant service delivery, customised experiences, and always-on performance. Meanwhile, legacy systems struggle to accommodate these shifts. Innovation slows, costs rise, and customer satisfaction declines.

This white paper introduces a **unified, AI-driven path to address these growing pains and future-proof network operations.**



The Problem with Siloed Networks

A longstanding challenge in network engineering and operations is that critical data - vital for planning, optimisation, troubleshooting, resolution, and many other tasks - often resides in multiple, uncoordinated, and segregated systems. In many instances, data from different network and service domains - such as RAN, transport (TX), voice, and data - is further siloed, amplifying this complexity. Most existing telecom networks are highly compartmentalised. Each domain operates its own management systems, which often leads to fragmented operations and inconsistent data flows.



As a result, we encounter several persistent issues:

Disconnected Data Domains

Information related to topology, fault management, performance, configuration, incidents, and change logs is difficult to correlate or extract, as each domain is managed independently.

Multiple Network Ecosystems

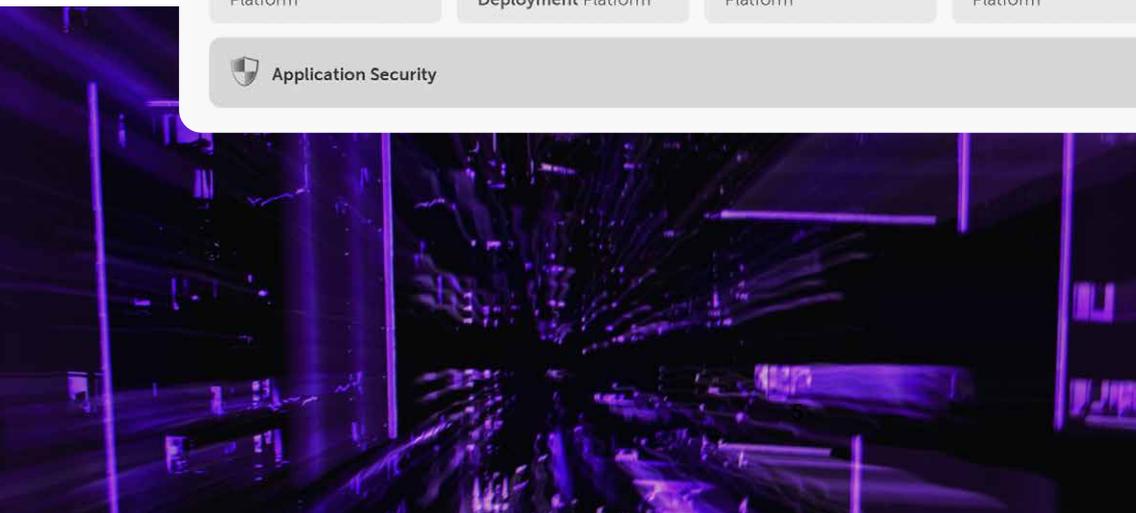
Individual domains maintain isolated data repositories, creating additional barriers to unified visibility and comprehensive analysis.

Even when AI is introduced into these environments, it typically functions in siloed or limited scopes, lacking access to cross-domain datasets. This severely hampers the potential for meaningful insights, automated orchestration, and effective root-cause analysis across the network. The outcome is slower time-to-market, increased operational costs, redundant diagnostics, and an underutilisation of valuable network telemetry.

The AI-Native Platform

Building Intelligence Into the Network

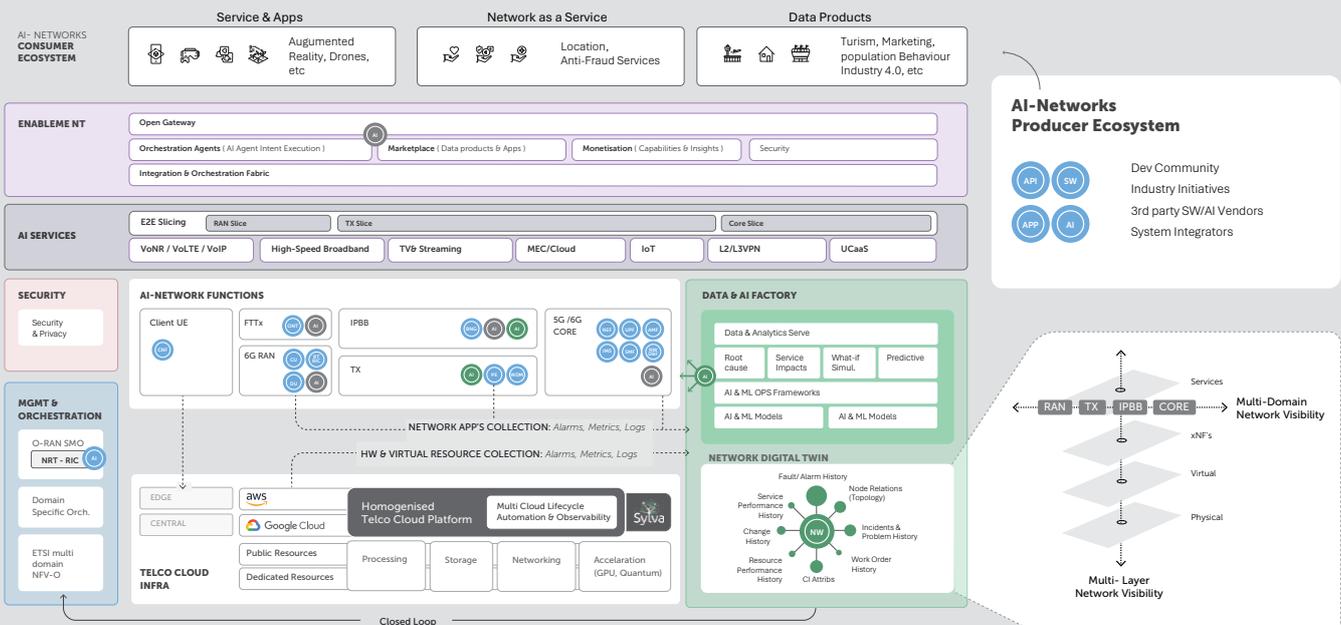
An AI-native platform addresses these challenges by turning the telecom network into a cohesive, intelligent system. It is built from the ground up with **AI, cloud-native infrastructure, and automation at its core**. Every layer - from the hardware and infrastructure to the orchestration tools and application services - is designed to collaborate, share data, and act based on real-time insights.





Critically, the platform supports a two-sided digital ecosystem. On one side, consumers - including individual users, businesses, and network functions - can subscribe to tailored services, insights, and capabilities.

On the other, producers - including CSP developers, industry partners, and third-party vendors - can publish services, deploy applications, contribute algorithms, and expose APIs. This forms a marketplace for network intelligence, where innovation flows freely across a federated, intelligent fabric.



Key AI-Networks Platform Components

Telco Cloud

Automated infrastructure provisioning using GitOps workflows. Supports containerised environments, distributed edge nodes, and accelerates AI performance using specialised hardware such as GPUs and FPGAs.

Network Functions

Integrates AI-enhanced network functions across all key domains, including 5G/6G core (UPF, SMF, AMF), RAN (CU/DU/RIC), broadband (FTTx), and transport (IP, optical). These components communicate seamlessly to provide end-to-end intelligent services.

Management & Orchestration Layer

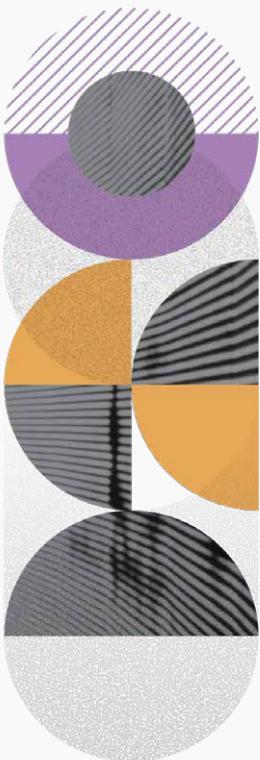
Uses policy-driven, intent-based orchestration to coordinate service delivery. Real-time closed-loop automation enables on-the-fly service reconfiguration based on current network state and business objectives.

Data & AI Factory Layer

Supports advanced analytics, observability, and real-time digital twin modelling. Facilitates lifecycle management of ML models, including continuous training, monitoring, and retraining.

App Enablement Layer

Abstracts the complexity of the network by offering developers easy-to-use APIs, toolkits, SDKs, and support for intelligent agents. These agents act as digital intermediaries that automate tasks, facilitate data exchange, and bridge applications with underlying network functions. This layer enables faster creation, testing, and integration of third-party services, streamlining the development of intelligent, real-time telecom applications.



Applications Layer

Hosts real-time consumer and enterprise applications that leverage AI-native capabilities such as hyper-personalisation, network slicing, and contextual optimisation.

The Journey to AI-Native Steps Towards Autonomy

Transforming into an AI-native network is not an all-at-once overhaul - it is an iterative process. CSPs progress through maturity phases, each delivering incremental value:

Step 1 - Data Readiness

Establish consistent, secure, and scalable data pipelines across network domains. This includes governance, transformation, and normalisation.

Step 2 - Digital Twin Creation

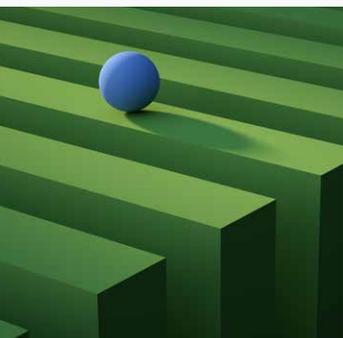
Build a real-time, data-rich virtual replica of the network that can simulate conditions, test changes, and predict failures.

Step 3 - Closed-Loop Automation

Implement intelligent agents that use AI-driven insights to take action automatically, optimising the network in real time.

Step 4 - Intent-Based Operations

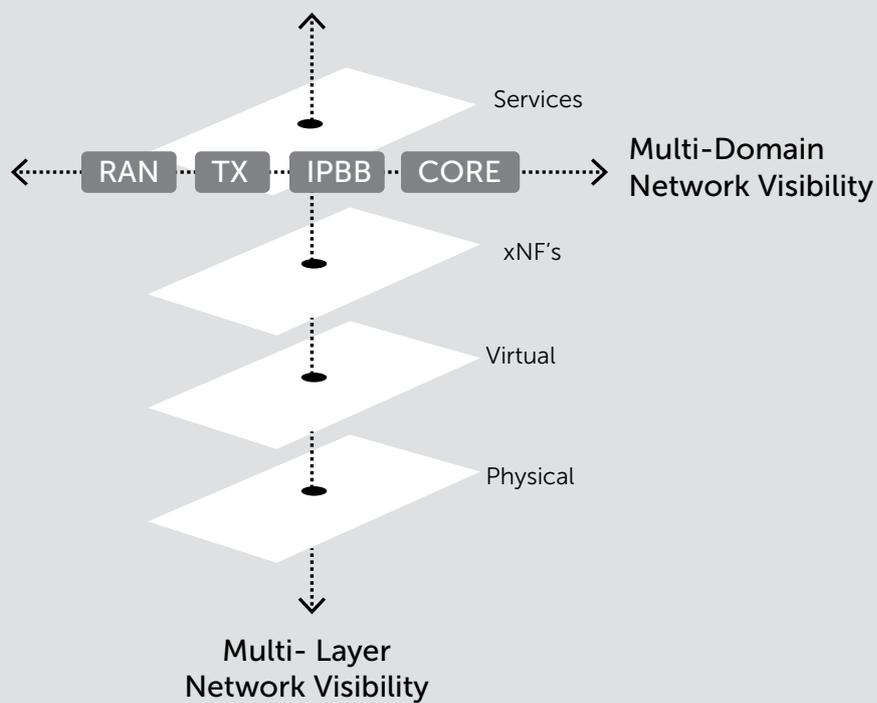
Replace static configuration with goal-oriented commands. Engineers define the "what," and the system determines the "how."



Step 5: Ecosystem Enablement & Monetisation

Open up platform APIs, marketplaces, and developer portals to allow partners and third parties to create new applications and services.

Each phase brings the organisation closer to full autonomous networking capabilities and allows for flexible scaling aligned with business priorities.



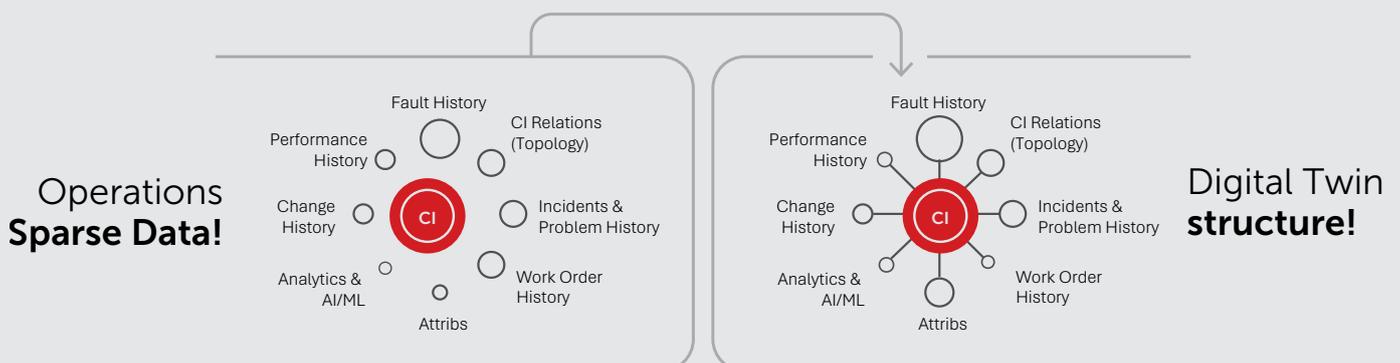
The Role of Digital Twins in AI-Native Evolution

Any AI application’s success relies on speed, quality, and relevance of the data feeding its algorithms. A robust data readiness framework and well-structured data model are therefore essential for rapidly querying, extracting, and processing information tailored to each targeted use case.

A well-designed data readiness framework addresses these issues by breaking down data silos, standardising formats, and enabling correlation across diverse sources.

This paves the way for more efficient troubleshooting, streamlined operations, and ultimately, better AI outcomes and restoring visibility to virtualised networks from an E2E perspective.

Implementing a robust data framework involves adopting a hybrid data ecosystem that strategically employs different storage technologies to balance cost efficiency and performance. Each component plays a specific role:





Data Lakes

Ideal for handling large-scale, unstructured data with minimal overhead, enabling both cost-effective storage and broad availability for analytics.

Graph Databases

Mimic real network topologies by capturing relationships between entities, making it simpler to visualise and analyse how components interact in a live environment.

Time-Series Databases

Designed to store recurrent, historical alarms, metrics and more - crucial for monitoring trends and events over time at high ingestion speeds. These metrics always stay connected to components of the topology within the graph.

Relational Databases

Best suited for well-structured attributes and transactional use cases that require consistency, referential integrity, and precise data relationships.

By combining specialised storage technologies, we ensure data from diverse sources is efficiently ingested, readily retrievable, and accurately reflects the complex interdependencies of modern network systems, both physical and virtual.

Insights

Model Abstraction



Advanced Analytics Apply algorithmics for anomaly and predictive algorithms on top of the unified data ecosystem

Graph-based Topology



Model and store network elements and their relationships in a graphDB architecture capturing topology

E.g: Network Topology

Time-Series Metrics



Store metrics and KPI's History relatively to each node and relation within the network topology

E.g: service & network performance

Relational Attributes



Store simple attributes, relatively to each node and relation within the network topology

E.g: CMDB Cis & Attributes



Efficient Storage Leveraging data lakes for scale and cost benefits storing large amounts of data

Beyond Generations: Embracing Continuous Evolution and Autonomy

AI-native networks break free from the rigid upgrade cycles that defined previous generations (2G to 3G to 4G, and so on). Instead, they embrace a model of continuous evolution. As AI capabilities expand and orchestration systems mature, CSPs transition through increasingly autonomous operations.

This journey aligns with industry frameworks like TM Forum’s Autonomous Networks model, moving from basic automation (Level 1) through assisted operations (Level 2–3) to full autonomy (Level 4–5). At the highest levels, networks can self-diagnose, self-heal, and self-optimize with minimal human intervention. This enables operators to launch services faster, minimize outages, and drastically reduce operational costs.



Take Action Now

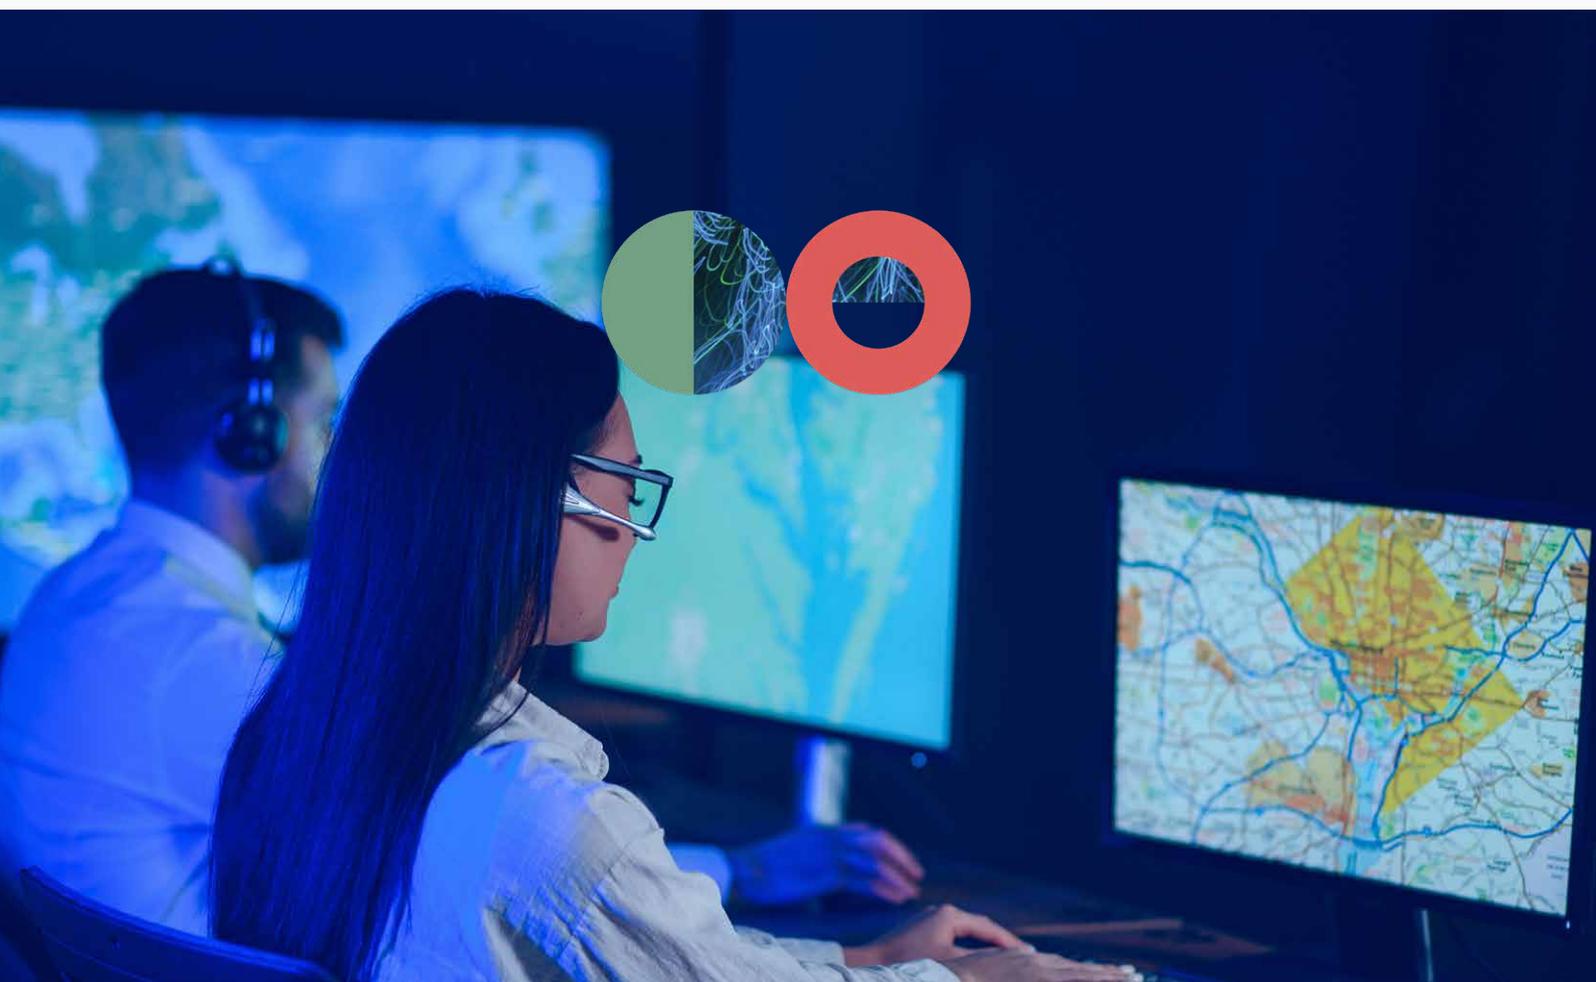
The journey to AI-native starts with today’s infrastructure. CSPs don’t need to wait for 6G - they can begin by making their 5G networks smarter. Start with scalable data strategies, introduce modular orchestration tools, deploy AI agents where insights are most needed, and expose APIs for external integration.

Each successful step builds confidence and proves ROI, making the case for broader transformation. Delaying the transition widens the gap between early movers and those left struggling with legacy systems.

Conclusion

AI-native networking represents the most significant architectural and strategic transformation in telecom since the shift to IP-based systems. The foundational tools - cloud-native platforms, AI engines, open APIs, and intelligent orchestration - are already available and mature. What remains is bold leadership and an aligned execution strategy.

Telecom leaders must recognise that AI-native is not a distant future - it is the strategic imperative of today. The networks that embrace this shift will lead the industry in innovation, efficiency, and profitability.



Why Celfocus?

Remove uncertainty from your path. Celfocus is a systems integrator with over 20 years of experience in CSP transformations, guiding clients from strategy to implementation and operation. We help our clients solve some of the most complex challenges in telecom by answering critical questions such as:

- **How to define and execute a roadmap for AI-native evolution?**
- **How to implement scalable, cloud-native platforms that support automation across network domains?**
- **How to set up effective data governance frameworks and unlock the full value of existing infrastructure?**
- **How to ensure end-to-end visibility and service assurance in hybrid, virtualised environments?**

We leverage the existing ecosystem and strategic CSP investments to deliver integrated, future-ready solutions. Our platforms are based on proven, pre-built assets that align with open standards and are designed to work seamlessly in complex technology landscapes.

Celfocus draws insights from large-scale data operations and open-source innovation, as well as hyperscaler frameworks, to accelerate delivery, reduce complexity, and maximise impact. Our reference projects demonstrate a consistent ability to deliver outcomes at scale, enabling telecom leaders to turn strategy into execution, and vision into value.



CELFOCUS

For more information about CELFOCUS,
please visit our website
www.celfocus.com

Follow us on: 

Copyright © CELFOCUS. All RIGHTS RESERVED TO CELFOCUS and its Licensors under Law. The disclosure, copying, adaptation, citation, transcription, translation, and/or any other form of total or partial use of the content, layout and graphic design, images or sound of this document or accessible herein, by any means of using any format (physical or virtual) without the respective authorization or licensing by CELFOCUS or its Licensors is prohibited and offenders shall be prosecuted. The user or licensee of all or part of this document may only use the document under the terms and conditions agreed upon with CELFOCUS and/or its Licensors for the purposes determined, otherwise subject to civil and/or criminal prosecution of the offenders.

CELFOCUS, S.A..
Avenida D. João II, Lote 34
Parque das Nações
1998-031 Lisboa, Portugal
Tel. +351 213 836 300 . Fax +351 213 836 301
www.celfocus.com