Omdia
by informa techtarget

The AI-Native Telco

New business outcomes rely on more than just AI technology investment

Author: Dean Ramsey, Chief Analyst

September 2026

In partnership with:

TATA CONSULTANCY SERVICES

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

Contents

Introduction: An AI-native telco is an entirely different entity 3
Section 1: What is an AI-native telco? 4
Section 2: Progress – Why now? 6
Section 3: Making the shift from AI bolted-on to AI-native 7
Section 4: The transformational path from today to AI-native 8
Section 5: Potential business outcomes 9
TCS is guiding telcos from today to success in the AI-native era 10
Appendix 11
Author 11

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

Introduction: An AI-native telco is an entirely different entity

The telecoms industry has spent several years bolting artificial intelligence (AI) onto networks and IT operations one use case at a time. In that time, there have been real successes, such as predictive maintenance in a radio access network (RAN), generative assistants in customer care, or anomaly detection in assurance. However, investments still outweigh returns, as shown in **Figure 1**.

Figure 1: Bolt-on AI is still producing average efficiencies vs investment

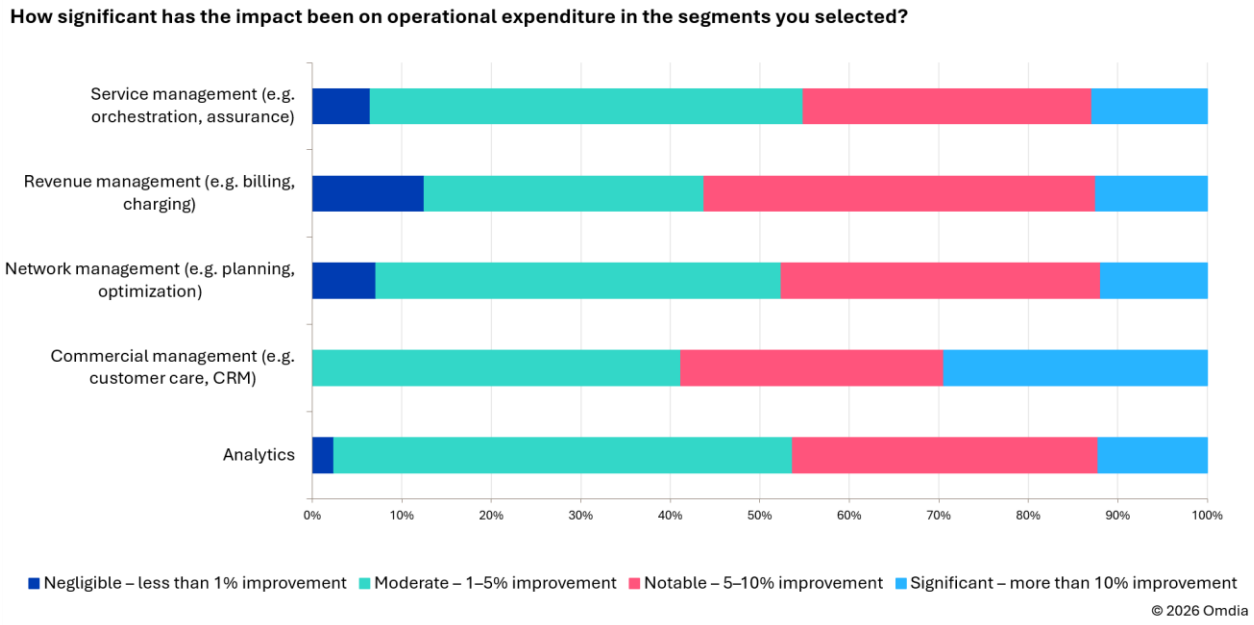

Source: Omdia

The reason is primarily structural rather than technological: a use case bolted onto a fragmented operational model inherently compromises the effectiveness of the AI.

Conceptually, an AI-native telco is a different proposition: it is not a telecom operator with AI use cases; it is a software-defined, data-product-driven, model-operated, agent-orchestrated, intent-aware enterprise in which intelligence is part of the operating fabric across a network, IT, customer, and business domains. The distinction sounds semantic, but it is the difference between adding intelligence to processes designed for human operators and rebuilding those processes around machine cognition.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This paper delves into what will separate an AI-native telco from a telco that has simply deployed AI tools, the scale of the change involved, and the potential business outcomes. Our conclusions draw on Omdia's telecoms research programs and ongoing interviews with telcos and technology suppliers around the world.

Section 1: What is an AI-native telco?

An AI-native telco is not an operator that has deployed AI. It is an enterprise whose systems, data, and processes are built on the assumption that machines, rather than people, are the primary decision-makers in day-to-day operations.

The industry's working definition of "AI-native" has become loose, to the point of being unhelpful, applied equally to a chatbot in a care center and to a network that reconfigures itself to meet a business objective. We define it narrowly, through five mutually dependent attributes, each describing how the enterprise will be built, rather than what it has bought, as shown in **Figure 2**.

Figure 2: The AI-native telco operating fabric

Source: Omdia

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

- **1. Software-defined and open environments:** Anything a model must reason about has to first be addressable by software, such as virtualized network functions, programmable transport, cloud-native IT, and inventory exposed through open APIs, rather than locked inside silos. This characteristic is a precondition rather than a differentiator; an AI model cannot act on data and infrastructure it cannot reach.

- **2. Data-product-driven:** Data is a managed product that is governed, versioned, discoverable, and owned, not a historical artifact to be stored and forgotten. While traditional telecoms data models describe what the estate contains, AI-native operations require a representation of how everything relates, connecting customers, services, resources, faults, and revenue in one traversable structure. Graph databases are not widespread in the industry today, but will be essential to AI-native operations.

- **3. Model-operated:** Decisions in the operational core will be made by models, with humans governing the overarching controls rather than approving each action. This demands model lifecycle management, carried out with the rigor that operators usually reserve for network elements. Architectures should avoid the sprawl that the bolt-on era has encouraged, where there is a model per use case, each with its own ungoverned pipeline. AI-native models will actively avoid this kind of accumulation of technical debt.

- **4. Agent-orchestrated:** Execution will be carried out by agents that plan and sequence actions across systems, rather than by scripts following a path a human once took. The key design question is not where agents can be used but rather where determinism, or giving the same output for the same input, must still rule. AI's reasoning and diagnosis are usually probabilistic, using sampling to create varied and creative results. Any action affecting a live customer, however, must be entirely predictable.

- **5. Intent-aware:** An operator declares an outcome, which could be a service level threshold, a coverage commitment, or a margin target, and the fabric resolves it into a configuration, ready for execution. This requires a semantic layer of ontologies that define what a service, a commitment, and a network capability mean, above a data plane able to serve that meaning in real time. Neither exists at scale today, which is why intent is the most discussed and least deployed concept of the five. **Figure 3** shows that the industry is already convinced that pursuing an intent-based operating model is the right approach.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

Figure 3: Survey – Which best characterizes your organization's plans for an intent-based operating model?

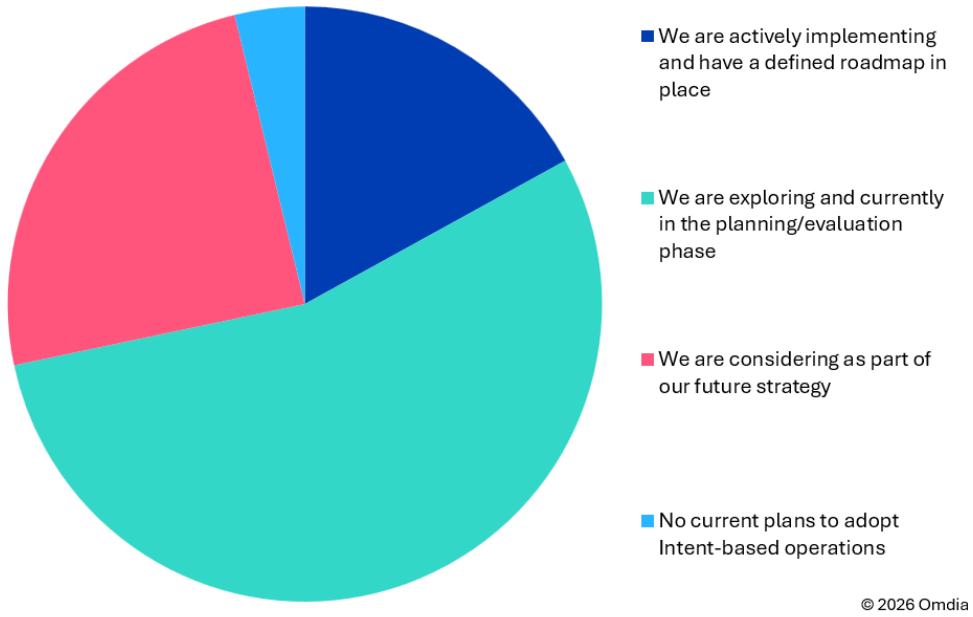

Source: Omdia

Section 2: Progress – Why now?

What is forcing this critical shift in how we manage telecoms? Three pressures have converged: (1) a monetization gap left by 5G, (2) an AI-value focus that is biased towards cost rather than growth, and (3) an operational complexity problem that human-led processes can no longer absorb. A fourth pressure, emerging 6G capabilities, sets the deadline.

Learnings from the 5G era: The industry invested heavily in 5G network deployment, but monetization rested on enterprise services that largely did not arrive at the scale forecast. AI is now being asked to deliver the returns 5G did not, but requests for fresh budgets to deliver unquantified results are being met with hesitance.

The economics: Where AI has provided value is narrowly concentrated: operational efficiency first, a better customer experience second, and net-new revenue a distant third. The industry has proved that AI can make existing business costs cheaper, but it has not proved that AI can make businesses larger and more profitable. That second proof is where AI-nativity earns its keep.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

Complexity is compounding: Multi-vendor, multi-domain estates are already at the edge of what human-scale operations can manage. Future scaling requires a new approach.

6G is the deadline, not the enabler: 6G will be AI-native by design, a direction embedded in the IMT-2030 framework and early 3GPP work, with the first networks to launch around 2030. That leaves less than four years to prove AI layers work in existing environments before they are deployed in entirely new ones. AI-nativity is not something to wait for 6G to deliver; it is the requirement needed for 6G to deliver.

Section 3: Making the shift from AI bolted-on to AI-native

Becoming AI-native requires an architectural reset, replacing data stores, the semantic model, the workflow engines, and the compute substrate that telecoms operations have run on for decades.

There is a useful precedent for comparison. Operational and business support systems (OSS and BSS) entered telecoms as bolt-on tools, then became baked-in processes, and eventually reshaped ways of working and organizational design. AI is retracing that arc, only faster and deeper. Predictive AI brought efficiency, and generative AI brought productivity, but both were additive. An operator could deploy them without disturbing the architecture underneath. To realize the true value of agentic AI, telcos need to stop thinking about agents as an additive technology but as part of the fabric. Four things need to happen to kickstart this fundamental transition:

- **1. The data layer becomes entirely available to AI:** Relational systems of record and warehouse-centric analytics describe the estate in tables built for human queries and periodic reporting. Agents need a knowledge graph spanning a network, services, customers, and commercial domains that is updated in real time. Graph technology is mature, but its adoption in telecoms is not.
- **2. The semantic and knowledge layers become a single source of truth:** Intent-based operations are impossible without shared meaning, yet ontologies defining these concepts in machine-resolvable terms consistently across domains do not exist in most operations today. A new semantic layer above a new data plane is the hard architectural work of the next five years, but almost none of it can currently be bought as a product.
- **3. The workflow engine becomes intent-based:** Today's orchestration encodes a procedure, or the sequence of steps a competent human would follow, and hardens it

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

into a deterministic flow. Intent-based operations invert this, declaring the end state before computing the path. Every process built on the old assumption is in scope: order fulfillment, OSS/BSS provisioning, activation, assurance, and optimization. Systems of record must become systems of reasoning.

- **4. The compute substrate:** Underneath, there must be an increase in AI density. Inference must run close to where operational decisions are made, pushing accelerated compute into an operator's own estate in an "AI factory" pattern, with the power, cooling, and sovereignty configurations to match. This is the one part of the transition where existing telecom assets are an advantage rather than a liability.

AI-native is not a modernization program with an AI budget; it is the replacement of the operational architecture that all telcos currently use. That is an uncomfortable message in a market where the accepted transformation methodology is usually incremental, and it is why we expect the gap between leaders and the rest to widen sharply. Some telcos, like Orange in France, have established research labs to specifically develop new paradigms for semantic communications in an AI-native telco, and to consider how all of these changes can be brought about in the second half of the 2020s.

Section 4: The transformational path from today to AI-native

For telcos to start making changes and investments to drive towards AI-native operations, at least five key moves are required:

- **1. Finish the substrate while following standards:** AI's work within various operational layers needs to be able to function without a human process bottleneck or a data roadblock. Enabling governed AI autonomy is essential. New architectures need to adhere to standards like the Agent2Agent protocol (A2A) and the Model Context Protocol (MCP).
- **2. Manage data as products and use knowledge graphs:** Telcos can intelligently connect network events to business impacts by using a single, federated data architecture with clear ownership of data products.
- **3. Treat model lifecycle as an operations discipline, not a data science one:** A "model as a service" posture keeps models trained on a telco's own ecosystem data, orchestrated centrally and governed as formally as network elements. The alternative risks model sprawl. Many standards development organizations, such as TM Forum and ETSI, are currently converging on the same conclusion.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

- **4. Draw strategic lines between deterministic and probabilistic approaches:** AI-native does not mean a probabilistic approach everywhere. Guardrails need to be placed between instances where deterministic or probabilistic approaches are used to ensure traceability for every agentic action, rollback paths, and documentation when an agent acts outside its envelope. Digital twins may prove to be a valued asset in simulating the real-world outcomes of complex agent actions before they are executed.
- **5. Start the semantic layer now, not after the data work:** The semantic layer has the longest lead time and the least off-the-shelf tooling of the five actions.

Section 5: Potential business outcomes

AI-native operators stand a lot to gain by pursuing AI-native operations. There are obvious operational efficiency and customer experience benefits, but these could be achieved through bolt-on AI. Omdia telco surveys have shown that process efficiency and staffing savings are delivering good results.

Figure 4: Bolt-on AI is delivering some efficiency benefits

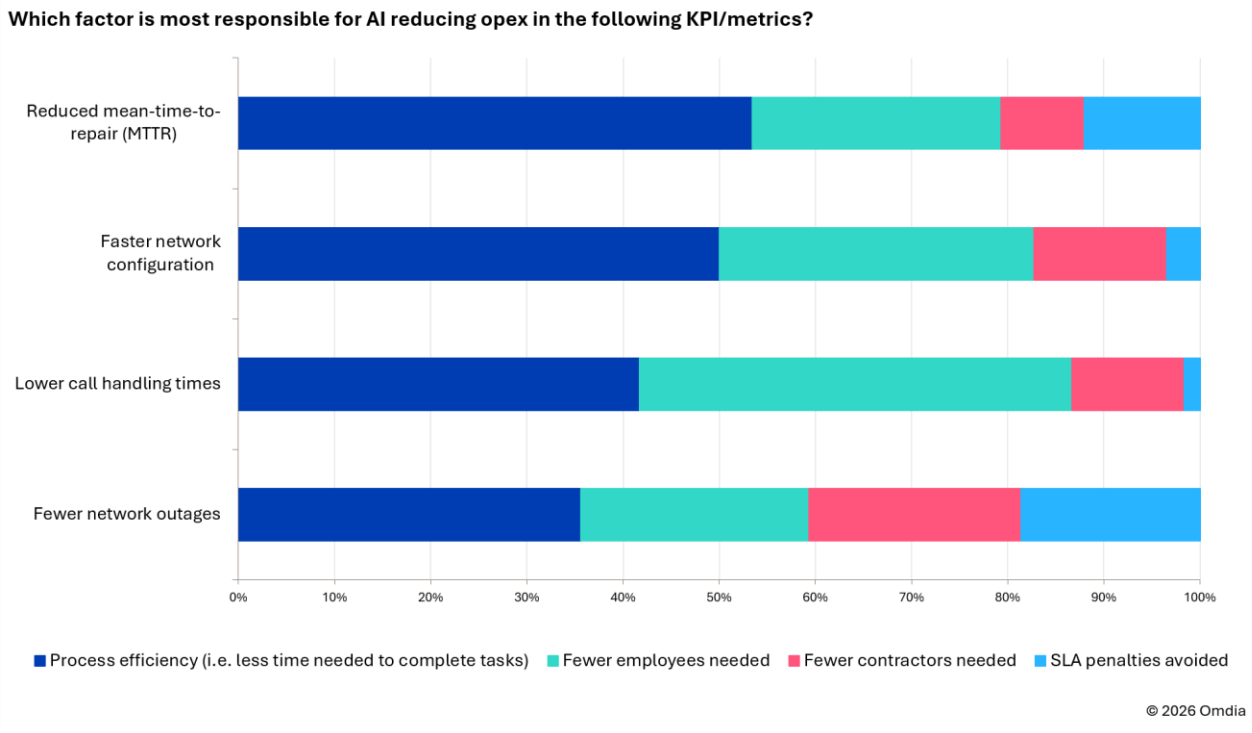

Source: Omdia

However, the largest opportunity, and the one least reflected in current telco planning, is the AI services market itself. Demand for inference capacity is growing faster than the

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

centralized data center estate can economically serve it, particularly where latency, data residency, or sovereignty constraints apply.

Telcos already have distributed real estate, robust power grid access (albeit in need of upgrades for running an AI factory), transport, regulatory standing, and local trust. These are the assets a distributed inference market needs and that hyperscalers and neoclouds cannot easily replicate. An AI-native telco is positioned to sell AI services rather than merely consume them, and to meet the hyperscale ecosystem as a peer rather than as the connectivity beneath it.

That opportunity is accessible to an AI-native telco, but it is not exclusively reserved for it. Unless telcos build an architecture capable of exposing, assuring, orchestrating, and monetizing intelligent services, they risk repeating the patterns of 3G, 4G, 5G, and IoT: heavy infrastructure investment, weak monetization, and the value accruing to the OTT layer above. The difference this time is that the transformation window is shorter, and competitors already operate in that market.

TCS is guiding telcos from today to success in the AI-native era

The obstacle for most operators is not conviction but sequencing—executing the data, developing the semantic layer, creating the workflow, and making infrastructure changes in the right order, in a live network, without interrupting the business. That is an integration and transformation problem, not an AI problem. Tata Consultancy Services (TCS) works across the full network services lifecycle, from plan and design through build and operate, which is the span this transition requires. For more information on their network solutions and services, please visit TCS's network solutions and services page or contact NSS.Marketing@tcs.com to discuss the topics covered in this paper.

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

This Omdia White Paper was commissioned by Tata Consultancy Services.

Appendix

Author

Dean Ramsay, Chief Analyst, Service Provider Transformation
dean.ramsay@omdia.com

Omdia consulting

Omdia is a market-leading data, research, and consulting business focused on helping digital service providers, technology companies, and enterprise decision makers thrive in the connected digital economy. Through our global base of analysts, we offer expert analysis and strategic insight across the IT, telecoms, and media industries.

We create business advantage for our customers by providing actionable insight to support business planning, product development, and go-to-market initiatives.

Our unique combination of authoritative data, market analysis, and vertical industry expertise is designed to empower decision-making, helping our clients profit from new technologies and capitalize on evolving business models.

Omdia is part of Informa TechTarget, a B2B information services business serving the technology, media, and telecoms sector. The Informa group is listed on the London Stock Exchange.

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help your company identify future trends and opportunities.

Get in touch

www.omdia.com
askananalyst@omdia.com

© 2026 TechTarget, Inc. All rights reserved. Unauthorized reproduction prohibited.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.