tm forum | Sponsored by motive

November 2025
inform.tmforum.org

# Agents of change: integrating AI into customer experience operations

**Authors:**
Dawn Bushaus, Contributing Analyst
Joanne Taaffe, Editor in Chief, Insight

**Editor:**
Ian Kemp, Managing Editor

# Contents



We hope you enjoy the report and, most importantly, find ways to use the ideas, concepts and recommendations detailed within. You can send your feedback to the editorial team at TM Forum via **editor@tmforum.org**

# The big picture

**In September 2024 we published a comprehensive report about AI in customer experience (CX), but nowhere in it is the phrase "agentic AI" used. What a difference a year makes. Today, it's impossible to discuss customer experience – or any other topic in telecoms – without using the term.**

Agentic AI stands to revolutionize customer experience, going well beyond the capabilities of generative AI (GenAI) chatbots. As TM Forum Chief Analyst Mark Newman explains in our recent agentic AI Benchmark report: "Rather than putting GenAI into the hands of contact center agents, agentic AI can move beyond the 'chat interface' trigger point to one that is activated by, for example, a network error, an email or a conversation with a chatbot." This can set in motion a chain of events to fulfill customers' requests, or intents, and resolve their problems autonomously – often before they even realize there is an issue.

Leading communications service providers (CSPs) like Deutsche Telekom and SK Telecom are going beyond chatbots in the consumer market to develop sophisticated AI personal assistants that aim to accomplish a wide range of tasks for customers, such as planning and booking travel autonomously. They aren't the only companies doing this, of course. OpenAI's ChatGPT and Google Gemini already offer sophisticated travel-booking capabilities – using plug-ins and integrations in ChatGPT's case, and via Google services – with a goal of expanding these services and making them fully autonomous.

But telcos have two important competitive advantages: their mobile networks and devices, and their long-standing relationships with customers built over years of providing highly reliable services. Unlike general AI platforms that operate through apps or web interfaces, telcos can embed AI assistants directly into smartphones and SIM cards. This allows for real-time, context-aware assistance that doesn't rely on third-party apps. Deutsche Telekom's AI-phone, for example, is designed to replace traditional apps entirely, offering a seamless, voice-driven experience.

## Identifying the challenges

Agents, in the context of agentic AI, are essentially pieces of software that can act autonomously – in essence taking the place of a human agent making decisions and acting upon them. Full autonomy in customer experience operations is possible only if CSPs can securely integrate customer-facing AI agents with other AI agents and with existing network functions and operations and business support systems (OSS/BSS), many of which are aging and siloed. It will not be easy.

**Read the reports:**

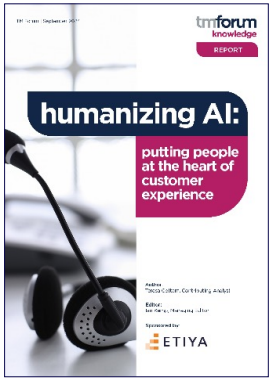


Legacy OSS/BSS weren't designed to work with modern AI tools. Many of the systems are hard to integrate, lack real-time access to data and don't support cloud-native technologies. Access to high-quality data is one of the biggest obstacles, if not the biggest. AI requires unified, data-as-a-product, but within telco environments data is fragmented and spread across multiple platforms. Addressing this challenge was the focus of our recent report Unlocking AI potential through a modern data architecture.

On top of the technical challenges, CSPs' limited AI expertise and cultural resistance to change make systems integration even more difficult. And then there is the cost. The business value of generative and agentic AI has not been proven yet, and it's hard to get investors to buy into potential savings and revenue when the cost of adopting AI is high.

The graphic oppostite shows how CSP respondents to our agentic AI Benchmark survey ranked these and other "major" challenges to becoming an AI-native telco.

"There are some challenges that we need to address before we are ready to scale [agentic AI], because we are talking about agents – effectively software – that take decisions," says Dr. Nektaria Efthymiou, Network Platform & Security Director, BT Group. "We need to make sure that there are guardrails. We need to make sure that there are security and cyber policies in place to make sure that the network is secure. Because it's not about AI sitting over the top of the network anymore – it is actually sitting in the heart of the network."

## Major challenges to becoming an AI-native telco

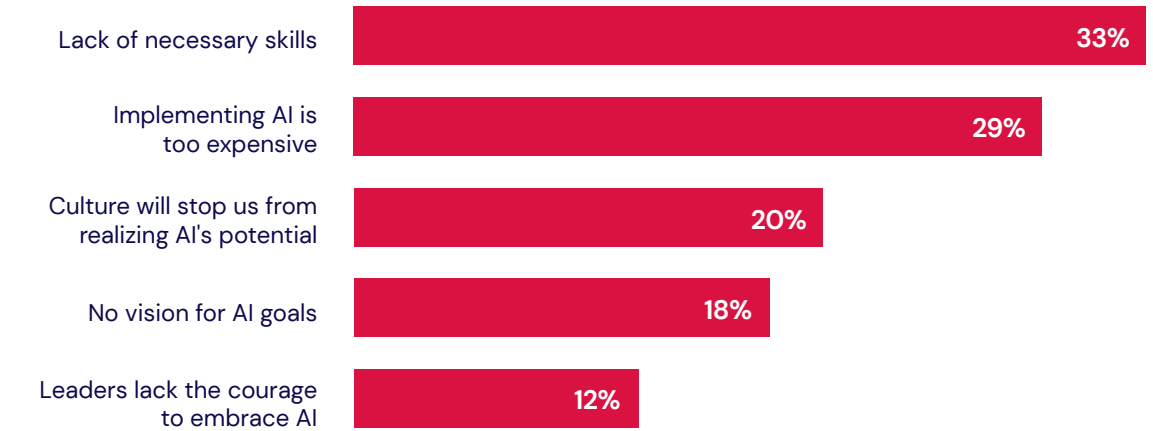


TM Forum, 2025

**Read the report:**

## What's in the report?

But even in the face of difficult challenges, operators are optimistic they will be successful in deploying agentic AI. This report looks at how CSPs intend to implement the technology by starting with customer experience use cases. We focus on some of the biggest integration challenges and how to solve them.

Read the report to understand:

- How customer-facing systems are evolving with AI
- How CSPs believe agentic AI will transform customer operations
- Why standards are needed to solve the biggest integration challenges as operators deploy agentic AI
- How CSPs expect to use Model Context Protocol (MCP) and Agent2Agent Protocol (A2A) in multi-agent architectures
- Why cost is a concern, and why it's difficult to prove the value of agentic AI.

> "We need to make sure that there are guardrails... Because it's not about AI sitting over the top of the network anymore – it is actually sitting in the heart of the network."
>
> **Dr. Nektaria Efthymiou, BT**

# Key report findings

In 2024 CSPs spent $90 billion on labor related to customer experience, making automation through agentic AI a major opportunity to cut costs.

Legacy OSS/BSS systems, fragmented data and cultural resistance are major obstacles to scaling agentic AI in customer operations.

High-quality, unified data and cloud-native infrastructure are prerequisites for scaling agentic AI.

MCP and A2A enable interoperability and intent-based orchestration across systems and domains, but security concerns must be addressed.

Agentic AI could cause telcos' costs to soar due to token-based LLM pricing and compute requirements.

CSPs should take a step-by-step approach to deploying agentic AI by focusing on high-value, high-priority scenarios that prioritize customer experience.

Section 1

# Customer engagement evolves with generative AI (GenAI)

## Improving customer experience is the top GenAI and agentic AI use case. The reason is simple: Satisfied customers are more likely to stay, upgrade services and recommend your company to others, leading to more revenue.

Reducing costs is also a major driver for using AI to improve CX. "Customer care is an enormous cost to the business," says Eoin Coughlan, Global CTO for Telco, Media and Entertainment at IBM.

Indeed, in 2024 CSPs collectively spent about $90 billion on labor related to customer experience, according to research firm Omdia. This includes staff supporting customer service and sales and marketing.

### Implementing GenAI chatbots

TM Forum's Generative AI Maturity Interactive Tool (GAMIT) survey, conducted at the end of 2024 and early in 2025, shows that CSPs are putting AI chatbots into production at a faster pace than other types of GenAI applications. A full 80% of CSPs said they are either developing proofs of concept for GenAI chatbots or putting them into production, with more than half claiming the latter (see graphic).

### Which GenAI use cases are CSPs developing?

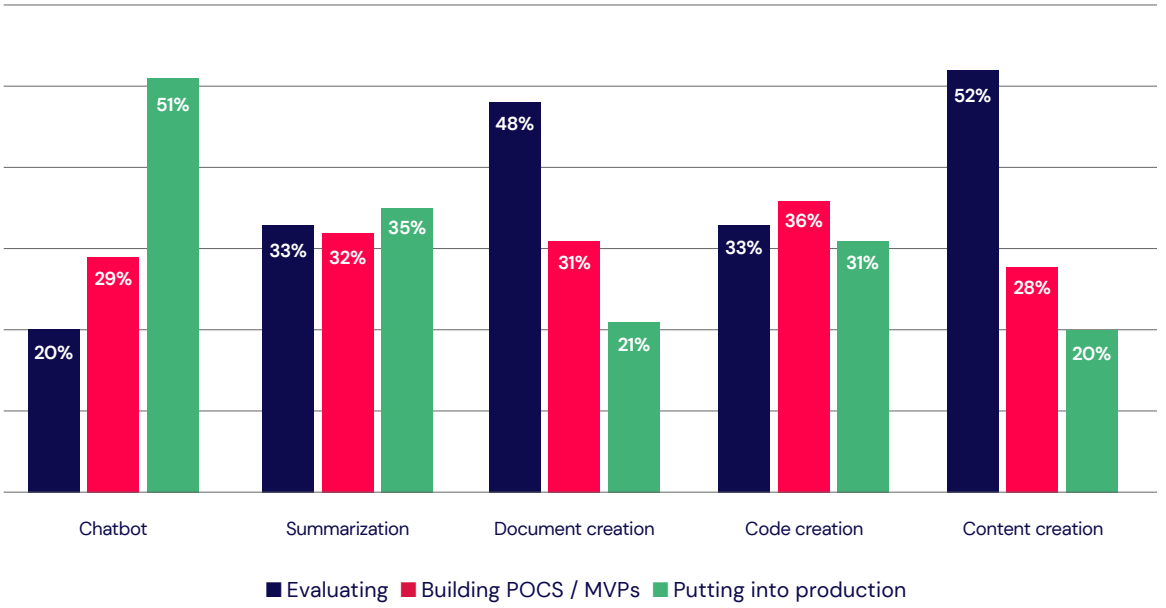


TM Forum, 2025

IBM's recent research supports these findings. The company's Institute for Business Value surveyed 106 CSP executives globally in 2025 about their AI strategy, initiatives and outcomes and found that 69% are using GenAI in customer care, with 44% reporting use of agentic AI capable of autonomous decision-making (see graphic).

### Agentic AI in CX

Software agents are not new in telecoms, but generative and agentic AI systems are. In our Benchmark report, we define agentic AI as an AI service that interacts with a large language model (LLM) and is capable of autonomous decision-making to achieve specific goals. AI agents leverage tools and contextual information, applying advanced reasoning and iterative planning to solve complex, multi-step problems independently. It's their ability to reason and take action that sets them apart from other software agents and GenAI systems.

In customer experience, this definition highlights a key distinction between agentic AI and robotic process automation (RPA), which emerged in the early 2000s and led to the development of simple rules-based chatbots (see graphic on page 11). While RPA uses AI to automate predefined processes, agentic AI goes further, determining the best course of action based on reasoning rather than following a fixed sequence.

## CSPs' deployment of AI in customer service

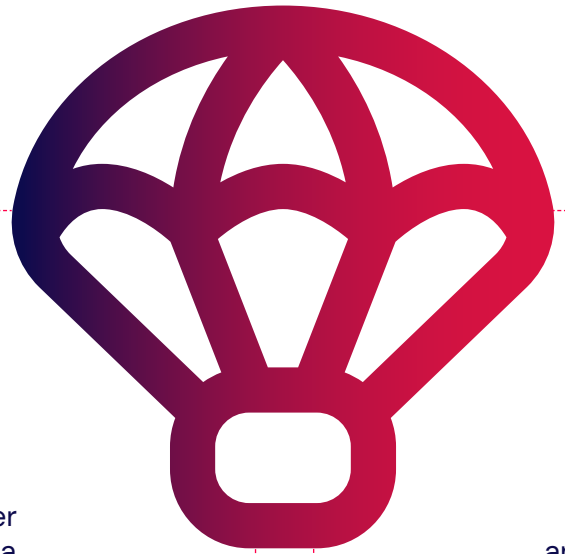

**55%**

implemented AI-powered customer service in 2024, and a full **87%** expect to be doing so in 2027

**66%**

are tracking the impact of AI in CX, and they have seen an **8%** improvement in customer satisfaction in the past year

**69%**

are using GenAI in customer care

**44%**

are implementing agentic AI in customer-facing systems

TM Forum, 2025 (source: IBM)

## Evolution of chatbots

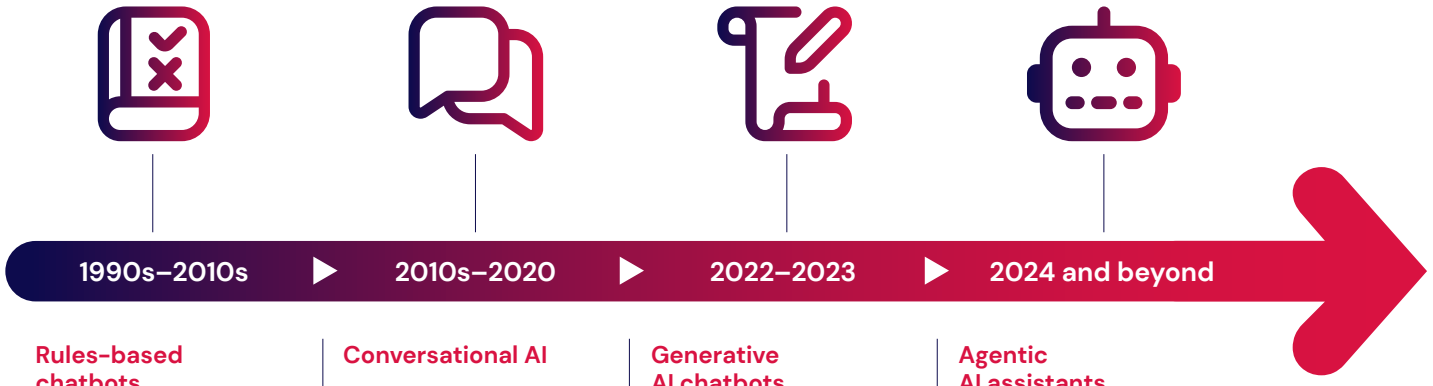


### Rules-based chatbots

Relies on rigid decision trees and keyword matching

Narrow conversations and poor handling of complex or multi-part queries

Supports customer service tasks like FAQs and billing inquiries

### Conversational AI

Introduces natural language processing (NLP) and machine learning

More dynamic conversations and better context handling

Broadens support tasks, but still limited in reasoning and adaptability

### Generative AI chatbots

Introduces LLMs for more natural interactions

Personalized responses and multilingual support

Enhances customer engagement, product recommendations and basic troubleshooting

### Agentic AI assistants

Introduces reasoning, learning and ability to evolve

Handles complex, multi-part problems; works with human agents to share insights and draft responses but can take these steps autonomously

Enables end-to-end customer service, network diagnostics, and proactive and predictive support

In a call center environment, an RPA chatbot might handle a billing issue by automatically retrieving the customer's account, checking for discrepancies, creating a support ticket, sending a confirmation email and updating internal systems. But all the steps are based on a fixed set of rules – the chatbot makes no decisions beyond its programming.

Agentic AI understands the customer's billing issue through natural conversation and then analyzes account details and other data to identify potential causes of the problem and evaluate possible solutions. Instead of following a fixed script, the agentic AI system uses knowledge and reasoning to decide the best course of action and then takes it.

**Read the GAMIT report:**

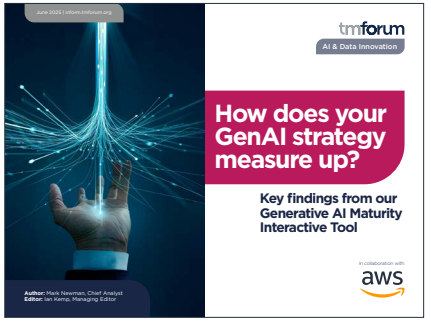

For example, if the best solution is to issue a refund, the AI agent can initiate the credit through the billing system autonomously. As we'll see in the next section, TM Forum members are working on an end-to-end agentic AI use case like this as part of a new Innovation Hub pioneer project that aims to make the Open Digital Architecture (ODA) AI-native.

This autonomous network (AN) approach is articulated in TM Forum's Autonomous Network Reference Architecture, which is built around intent-based operations where the desired outcome – what the customer wants – is the focus rather than the underlying technology. The intent ontology translates the outcome into service characteristics. Then, at the intelligent orchestration layer those characteristics guide the selection of appropriate network technologies. This approach leverages closed-loop controls, self-healing domains, Open APIs – and AI agents.

Within the CAMARA Project, for example, some operators are proposing to build AI agents around the GSMA and CAMARA APIs so that the interaction of a customer is with a local AI agent that talks to other AI agents rather than directly to network APIs, says TM Forum's Dave Milham, Chief Architect, Service Provider Engagement. "The thinking has always been that with APIs if we expose it, customers will want it. But the other point of view is, how can we give customers what they actually want, and how can we put together a solution that meets their needs? Agents could be quite a good way of doing this."

## Telstra empowers customers and eyes agentic AI

Telstra is laying the groundwork for using AI agents to empower customers with its new Adaptive Networks Product Experience (AN PEX). Through a set of canvas-like tools, AN PEX allows enterprise customers to design, configure, implement and change their network solutions through real-time visualization and intuitive drag-and-drop interfaces.

Telstra's teams drew on the example of hyperscaler service offerings to create AN PEX, describing it as "an inversion of control" because it places network product decisions firmly in the hands of enterprise customers and the operator's partners. Telstra is now working on adding agentic AI interfaces to its version of ODA components.

## Early days for agentic AI

The executives interviewed for this report and our agentic AI Benchmark stressed that it's very early days for agentic AI, even in customer experience. "I believe that as an industry we are still in the early stages of this journey," says BT's Efthymiou. "We have moved from the hype we were seeing, when everybody was looking at something new in AI every six months, to rationalizing what we need to do to scale."

In a recent interview with TM Forum Insight, Philippe Ensarguet, VP of Software Engineering at Orange, said his company is treading carefully with agentic AI. He stressed that multiple foundations must be in place before telcos can deploy agentic AI in production environments. This includes transitioning to "true cloud-native infrastructure". Today, between 15% and 20% of Orange's network functions are cloud native.

"You can't be successful at agentic AI if you are not successful at automation. You are not able to be successful at automation if you're not successful at data. You cannot be successful at data if you're not successful at telco infrastructure," said Ensarguet. "So, I would say it's like a cake with multiple layers, and you don't have the [agentic AI] cherry on top if you don't manage the rest."

Next, we look more closely at how CSPs believe agentic AI will transform customer operations.

> "It's like a cake with multiple layers, and you don't have the [agentic AI] cherry on top if you don't manage the rest."
>
> **Philippe Ensarguet, Orange**

Section 2

# How will agentic AI transform customer operations?

**The long-term vision for agentic AI in telecoms is that all the network functions, customer-facing systems and other OSS/BSS that handle voice, data and video communications today will evolve into AI agents, or sets of agents, that perform the function. But telecoms is a long way away from making that vision reality.**

In the future, even something as simple as making a phone call will involve multiple AI agents coordinating to verify details like: "Is the caller on the network? Have they paid their bills? Have they been blocked by the recipient?", explains IBM's Coughlan. Each of these checks will be handled by a different agent, and the call setup itself – across the network's control and user planes – could involve dozens of AI agents handling functions from the customer's device to the access network to the core.

"So, you actually come across the same type of technology issues that we have right now that were solved but have to be solved for agents," Coughlan says. While existing standards will guide how AI agents are built and determine what data they exchange, their autonomous nature – "their ability to think, plan and act" – requires new policies to ensure they don't "overstep [their] own boundaries or start to make decisions and take actions that they shouldn't", he adds.

BT is starting with intent-based orchestration but is aiming for a future where autonomous agents replace much of the complexity in OSS/BSS, according to Andy Corston-Petrie, Senior Manager of Intelligent Service Orchestration Research at BT Group. He believes AI agents can play a role in an "evolved OSS/BSS stack that sits above our programmable networks and above the APIs that are exposed by the network".

"What you'll end up with, potentially, is a multi-agent system," Corston-Petrie says, explaining that BT is exploring what this architecture should look like, how much autonomy to give each agent and whether a supervisory agent might be needed to coordinate others.

TM Forum's Andy Tiller, EVP, Member Products and Services, stresses that achieving autonomy in customer operations through agentic AI will not happen overnight. "In an imaginary world in the future, you just deploy agents and you don't need the existing systems and processes that run the business and network today, because the agents can talk to each other in whatever language they decide. If the agents need to get something done, they write some code to do it on the fly," he says. "Today, this is science fiction. What we need is a way to integrate agents with existing systems and processes."

> "In an imaginary world in the future...if the agents need to get something done, they write some code to do it on the fly. Today, this is science fiction."
>
> **Andy Tiller, TM Forum**

## Catalysts tackle agentic AI's integration challenges

CSPs must address multiple integration challenges simultaneously to deploy agentic AI in production at scale. The most important issues to tackle first are improving access to high-quality data, modifying existing system interfaces to support MCP (see section 3) and addressing the tendency for LLMs to hallucinate, according to China Unicom's Jie Zhang, Principal Researcher, Digital Transformation Strategy Consultant.

China Unicom, China Mobile, China Telecom and Telekom Malaysia have been working together in an ongoing Catalyst project to solve these challenges and others. In 2024 the team won an award for developing an AI-native computing network operating system that enables CSPs to intelligently and autonomously schedule computing and network resources across a cloud-edge architecture.

In the first phase of the project the team leveraged LLMs, multi-agent collaboration and automated code generation to interpret a customer's business scenario. The goal was to autonomously generate a tailored computing and networking integration proposal for the customer, who could then refine the proposal through an intuitive chatbot interface that served as a replacement for the traditional product catalog.

The second phase focused on two key objectives: using AI to design an energy-efficient computing infrastructure; and using cross-domain orchestration and heterogeneous resource scheduling to seamlessly integrate public cloud, network services, platform services and AI. The idea was to give customers a comprehensive, one-stop experience.

"The most important lesson from our Catalyst projects so far is the value of a focused, incremental approach over large-scale, all-encompassing deployment," says Zhang. "Instead of applying agentic AI and generative AI across all scenarios, we've found greater success in selecting high-priority, high-value use cases – like business processing, pre-sales consulting and after-sales complaints – where these technologies deliver tangible results."

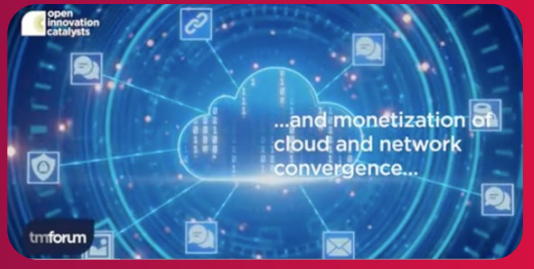

## Project ONE aims for agentic ODA

A new Innovation Hub pioneer project called Project ONE will bring together all three of TM Forum's Missions to demonstrate an end-to-end agentic AI fault resolution scenario at DTW-Ignite 2026. The graphic on the next page shows an example in which a CSP's small business customer experiences an outage.

In this case, the business customer's router has an embedded eSIM allowing it to switch over to 5G, but the CSP's backend systems detect that the customer's service level agreement (SLA) still cannot be met. So, the AI agent initiates a trouble ticket leading to autonomous fault resolution and delivery of compensation based on the SLA.

The idea is to highlight where agents may be a good replacement for humans throughout the journey from customer interaction to fixing the network fault. "We need to look for the parts of the picture where agents can do a better job than humans," says Tiller, noting that it can be expensive to deploy AI agents (see section 4).

Project ONE aims to put the focus on customer experience rather than just technology, because while a CSP's network operations team might be narrowly focused on fixing a network issue, the effect of a fault or outage flows all the way to the end user.

The problem with having an isolated view is that "I don't know who the customer is", explains Dr. Lester Thomas, Head of New Technologies and Innovation at Vodafone Group, and a leader in TM Forum's Innovation Hub. "If I have ten problems, how do I know which one to address first? The network will have its view of what's most important, but there might be a critical customer who should come first," he adds.

## Flexibility is key with agents

It's important to note that the Project ONE team is not setting out to define a standard set of AI agents. Rather, it is focusing on how to build agentic applications on top of ODA.

"We've got no idea what agents are going to do. They could do a low-level task, or they could manage a whole complex process. They could appear and disappear ephemerally," Tiller explains. "So, we're not concerned about trying to create a standard set of agents. But what we do care about is how you manage and operate them."

In the fault management scenario, when the AI agent decides to give the customer compensation it must communicate with a billing component to accomplish this. It needs to do that through the only language the billing system understands which is the API, says Tiller.

The Project ONE team plans to use MCP and A2A standards in their demo to enable an end-to-end experience for the customer. We'll look more closely at these new protocols in the next section.

> Project ONE is not setting out to define a standard set of AI agents. It aims to build agentic applications on top of ODA.

# Project ONE fault resolution scenario

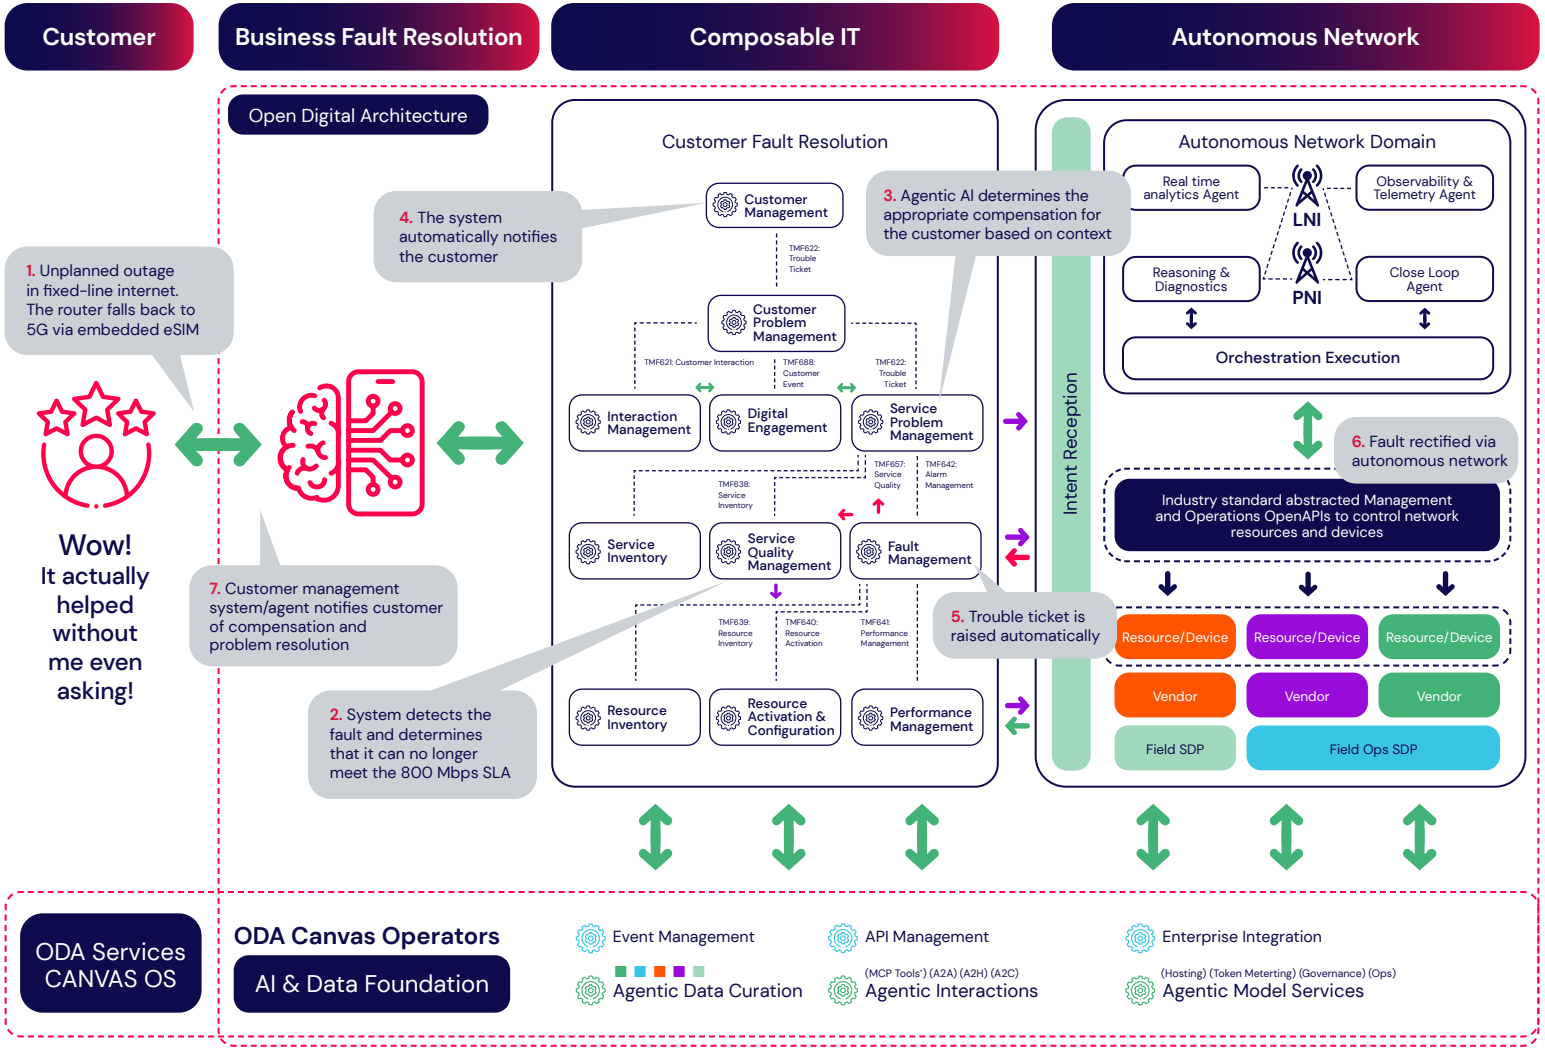


TM Forum, 2025

Section 3

# What are the roles for MCP and A2A?

**CSPs are experimenting with two (very) new standards – Model Context Protocol (MCP) and Agent2Agent (A2A) Protocol – to address agentic AI integration challenges. While debate is intensifying within the telecoms industry about the role for these standards in future network architectures, one thing is clear: The functions they perform are essential in AI-enabled autonomous networks.**

In customer experience, intent is the foundation for how agentic AI interprets and acts on users' goals. Instead of executing predefined tasks, agentic AI systems aim to understand why a customer is interacting (their intent) and then autonomously orchestrate actions across multiple systems in multiple domains to fulfill that goal. Together, MCP and A2A make intent actionable by embedding it in context and enabling inter-agent coordination.

Vodafone is using both MCP and A2A in its AI Booster platform. "MCP is very well adopted. It is almost a de facto standard," says Thomas. "We've got MCP servers live in production, doing real work." But he adds that MCP introduces security risks which must be addressed (see box on page 22).

Vodafone is using A2A to develop a multi-agent architecture. "One of the things we've learned is that smaller, tightly scoped agents are more effective. If you make them too big, they fail because they can't stay on track," says Thomas. "So for bigger problems you use multiple agents and they talk to each other through A2A. That's what we're trialing now."

## What are MCP and A2A?

MCP is an open standard introduced in November 2024 by Anthropic, creator of the Claude LLM, to give AI agents a unified way to interact with tools, services and data regardless of how those systems are built. MCP allows AI agents to autonomously carry out complex, multi-step tasks such as retrieving information, summarizing documents or saving content. Without a standard like MCP, these actions require separate APIs, custom logic and considerable integration effort by developers.

Google introduced the Agent2Agent (A2A) protocol in April this year to enable seamless interoperability between autonomous agents. A2A facilitates secure information exchange, coordinated actions and dynamic collaboration among agents without requiring them to share memory or internal tools. By abstracting these interactions, A2A opens the door to scalable multi-agent frameworks. The initiative has since been adopted as an open-source project under the Linux Foundation, ensuring broad accessibility and community-driven development.

One of the Catalysts Vodafone is championing, along with a dozen other CSPs, is called Agent Fabric. The first phase of the Catalyst introduced a copilot assistant that helps engineers resolve customer-affecting network issues faster. The second phase extended this capability by having multiple agents working together. The Agent Fabric Catalyst will feed into Project ONE in the Innovation Hub, according to Thomas.

## What's the debate?

Given that MCP and A2A are already being used in telecoms networks, the debate is not so much about whether the protocols are useful but rather how they will be standardized, secured and integrated to deliver on the promise of intent-based autonomous networks – and how quickly this can happen.

TM Forum released a position paper on MCP in August, arguing that the protocol is complementary to the Open API Program, essentially as an abstraction protocol that enhances the value of existing APIs. The document emphasizes the importance of semantic alignment, interoperability and structured governance, and proposes a roadmap that includes working groups, reference implementations and formal profiles to ensure that MCP adoption is consistent, secure and aligned with the ODA.

"One of the key arguments about why it makes sense to build agentic applications on top of ODA is that once somebody has written a tool for the API, everyone can use it. If everybody's got different APIs, you need to write the tool thousands of times," says TM Forum's Tiller. "ODA gives you a foundation on which you can build agentic AI-driven use cases, because you've got standards."

He adds: "There's a common misperception that MCP replaces APIs when you're talking to existing systems, but APIs is their language… MCP is enabled by APIs."

## Can telcos move quickly enough?

But some vendors and industry analysts view MCP as a disruptive force that could fundamentally reshape how telcos think about APIs. They believe the traditional model of API standardization in this case will be too slow for the pace of AI-driven innovation.

TM Forum members acknowledge the need for the telecoms industry to move quickly or risk being left behind in the AI innovation race, but some of them don't believe that MCP and A2A are suitable in their current forms for use in telecoms networks, largely because of security concerns (see box). As a result, they are looking at whether MCP-T and A2A-T versions of the protocols – where the "T" stands for telecoms – can strengthen security and governance.

**Read more about the Agent Fabric Catalyst:**

Vodafone's Thomas is not convinced that a telecoms-specific version of MCP or A2A is necessary. "If there's a deficiency, let's contribute to MCP and fix it in open source," he says. "I've always been a strong advocate of don't reinvent stuff… With MCP, we've done a lot of work to show how it links with TM Forum Open APIs. But I'm skeptical when someone says, 'We need a telco-specific standard'. What's so unique that other industries don't have?"

Next, we look at the cost of deploying agentic AI to improve customer experience.

## Beefing up security for agentic AI

Security is top of mind for all the executives we interviewed, and they stress that the risks must be addressed before MCP and A2A can be scaled in production.

"Agentic AI… is a direct entry point… into the very deep core of our network asset," says Orange's Ensarguet. "As a result, security in agent AI time will be even more important than today."

MCP introduces potential vulnerabilities, such as malicious clients intercepting authorization codes or compromised servers stealing user tokens, explains China Unicom's Zhao Yongjian, Senior Researcher, Network Intelligent Operation Expert. "These risks call for reinforced security mechanisms between the client and server sides, with particular emphasis on protecting user tokens and strengthening authentication and authorization processes," he says.

Accuracy is a related issue. Since AI models autonomously decide which tool or API to invoke, hallucinations can lead to inaccurate or inappropriate tool selection, Yongjian explains.

"This can be mitigated by improving the model's decision-making capability and ensuring that the descriptions of APIs and tools provided by the MCP service are precise and well-structured, he says.

Jeffrey Spiess, Head of R&D at Motive, a service management platform vendor and sponsor of this report, argues that CSPs should approach security "not as a late-stage implementation detail, but as a foundational element of the solution's architecture, deployment practices and operational governance". This requires a comprehensive, multi-layered approach that includes controls such as strong authentication and access; input validation and output guardrails; and encryption.

Vodafone is implementing MCP safeguards in its AI Booster platform. For example, the company only allows the use of verified and trusted MCP servers, which means maintaining an MCP catalog and an MCP gateway. Instead of an agent directly accessing an external MCP, it routes through the gateway, where access rules can be audited and enforced. Vodafone plans to contribute these learnings and capabilities to the ODA.

Section 4

# Where's the value in agentic AI for customer experience?

**Modernizing infrastructure and adding new AI systems are already expensive. Adopting agentic AI could cause telcos' costs to soar, and respondents to our surveys have acknowledged this. But overall, telco leaders are optimistic that agentic AI will reduce the cost of serving customers and improve their experience at the same time.**

Our agentic AI Benchmark report found that more than three quarters of CSP respondents view cost as at least a moderate challenge to becoming AI-native, with nearly a third believing it's a major challenge (see graphic). Only the lack of necessary skills was ranked higher as a major challenge in the survey.

Fees for the use of commercial LLMs are based on tokens, a unit of data processed by an LLM. While the cost of using LLMs has dropped significantly in the past year because of competition, usage-based billing for agentic AI workflows could lead to unexpectedly high costs – especially if the operator's usage spikes or isn't managed well.

"Agents have the potential to massively increase the use, and therefore the cost, of using LLMs," TM Forum's Newman explains in the Benchmark report. "Every interaction between an agent and an LLM consumes tokens." For example, when an agent needs to fetch data from an existing system via MCP and APIs it needs to understand which tool to use, which inputs to provide, and then process the response data.

## Cost of implementing AI as a challenge

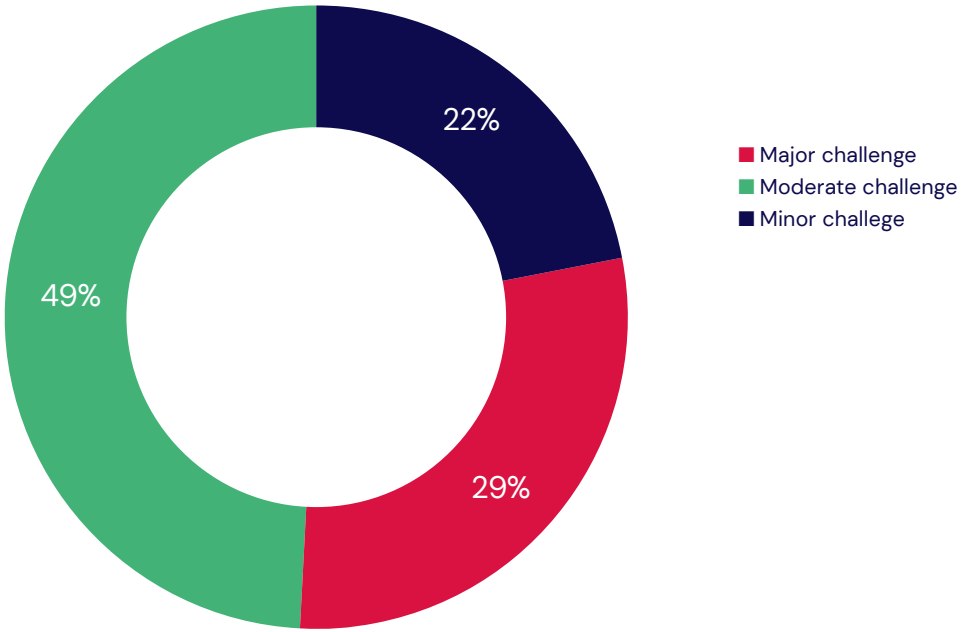


TM Forum, 2025

BT's Efthymiou says her team is tackling an open-ended challenge: designing and orchestrating the intelligence to scale AI agents for 20 million customers. "This is no longer just about proving the concept – it's about deploying in a real network," she explains. "We're asking critical questions: Where should compute be located? How do we manage and optimize energy to ensure efficiency at scale?"

Indeed, how quickly CSPs are able to implement agentic AI will depend greatly on the cost of rolling out the technology at scale. The "science fiction" scenario that TM Forum's Tiller imagines – where AI agents are interacting completely autonomously, without reference to the systems that run the business and network today – not only requires investment in AI models and technology but also in cloud-native network transformation.

"You don't necessarily need agentic AI to get just as good an experience," says Tiller. "Every place you put an agent making a decision, you've got to justify why it makes sense to do that rather than just use simpler automation. If you can do it with a simple coded business rule, then that may be the best way to do it."

## Agents running wild

During a panel at a recent Telecom TV event, TM Forum CTO George Glass explained that agentic AI is likely to be most useful and cost-effective for discovering the network domains that are available to deliver a service, but it doesn't need to be used everywhere.

**Watch the panel discussion:**

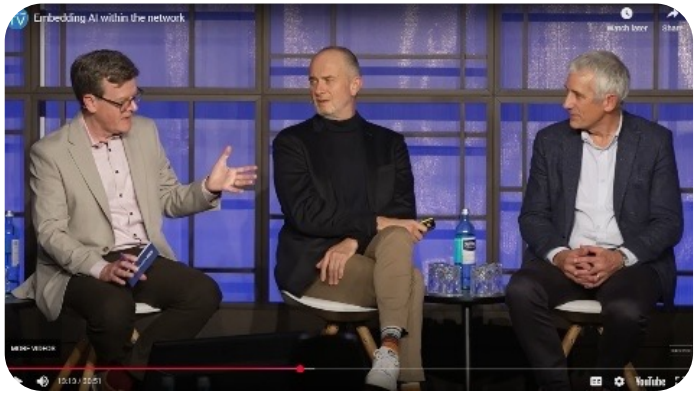


For example, GenAI may be a better choice for turning natural language into intent ontology, and predictive machine learning may be better at intelligent, closed-loop orchestration.

When agentic AI is used to fulfill a service using a specific domain, it is important to maintain awareness of the remaining domains in case the active one begins to degrade or fails. However, Glass cautioned that continuous communication between the orchestrator, the active domain and the other domains should be avoided in agentic AI deployment because excessive agent-to-agent "chatter" and MCP traffic can generate significant administrative overhead, potentially impairing network performance and increasing costs.

> "Every place you put an agent making a decision, you've got to justify why it makes sense to do that rather than just use simpler automation."
>
> **Andy Tiller, TM Forum**

"We are letting danger out of the box if we just let agents run wild over our network and over our IT estate without proper architecture and proper engineering," Glass emphasized. "You have been warned."

## Demonstrating value

In our agentic AI Benchmark report, an entire section is devoted to evaluating the economics of AI transformation. The report points out that while CSPs like AIS, Omantel and Swisscom are already reporting measurable gains from AI or AN initiatives, converting the improvements into a positive financial impact isn't easy because metrics like "time saved" don't always translate directly into cost reductions. "Indeed, there is a scenario where a CSP touts significant productivity gains from the use of AI but operating expenses continue to grow," writes Newman.

He suggests that a more effective approach may be to focus on outcomes like better customer experience and faster time-to-market. Telus, for example, has found that launching products 30-35 days sooner delivered far greater EBITDA gains than hours saved. Their AI teams now work with the finance department to track EBITDA impact for every use case.

To help CSPs demonstrate value, TM Forum's Autonomous Networks Project has identified 20 "high-value scenarios" that TM Forum members are addressing first to reach Level 4 in their network architectures. It's not surprising that half of them focus on improving customer experience.

"We are letting danger out of the box if we just let agents run wild over our network and over our IT estate without proper architecture and proper engineering."

**George Glass, TM Forum**

# Further reading

Read the AI in CX, agentic AI Benchmark, GAMIT and data architecture reports.

Read IBM's Telecommunications in the AI era report and another report on autonomous networks produced in conjunction with TM Forum.

Learn about autonomous networks high-value scenarios in this implementation guide.

**Check out these articles on TM Forum Insight:**

Orange's Philippe Ensarguet says his company is treading carefully with agentic AI.
Telstra lays the groundwork for using AI agents to empower customers with its new Adaptive Networks Product Experience (AN PEX).
Vodafone and Google demonstrate a production-grade reference implementation for agent-to-agent interaction.
AT&T is starting to realize value from scaling its Ask AT&T platform.
C-level executives call for enhancement of ODA to support agentic AI.

**Here are some Catalysts focusing on generative and agentic AI in CX:**

Agent Fabric
AI chat agent: The game-changer for telecoms
AI-driven proactive customer-centric O&M: Empowered by multi-agent and digital twin
Beyond Chatbots: Hybrid AI for fully automated proactive customer care
End-to-end service realization using intent-based networks
GenAI empowers computing force network
Predictive Intelligence for Optimized Networks & Enhanced Experience Resilience (PIONEER)

# Bridging the Gap: How AI-Enabled Integration Platforms Modernize Telecom Legacy Systems

**Abstract:**
Telecom operators are keen to adopt AI and have been playing with AI for years. However, the industry has been slower to adopt AI technology than other sectors due to decades of accumulation of legacy systems, fragmented customer data, and complex integration layers. An opportunity is emerging with new approaches that promise to accelerate AI adoption in Telcos without costly re-platforming. AI-enabled integration platforms provide automation, intelligence, and interoperability to overcome the legacy platform challenges.

**The Modernization Challenge**
Operators face a complex challenge given the fragmented and monolithic architectures that exist in their networks and support systems. OSS/BSS, CRM, billing, order management, and service assurance are often siloed, proprietary, and built on different technology stacks. Often there is a lack of RESTful APIs (or even any APIs), a lack of real-time or streaming data interfaces – some systems relying on file-based data exchange. There is a lack of cloud-native containerized services, and challenges with hybrid cloud setups.

There are challenges with data management; data exists in silos and with inconsistent schemas. Often there is a lack of IDs suitable for correlation across these different silos. There are difficulties with data quality, completeness, governance, and lineage. Operational and process gaps exist due to lack of automation hooks. Vendor systems are difficult to change due to strict availability and up-time constraints.

Organizational challenges abound due to this fractured and historical environment. There are often unclear ROI and ownership for transformation programs. Responsibilities are siloed; engineering, IT, and business units often work separately, making end-to-end AI solutions difficult. Operators have issues with vendor lock-in with resistance to interoperability. There is a cultural resistance to change with an old-school engineering mindset and lack of skills in the emerging AI technologies.

As an example of how these challenges are an impediment to AI adoption, an AI "next best offer" model may need to pull data from CRM, CDRs, network KPIs, and billing — all via separate, fragile APIs. This represents a challenging project on its own and scaling such an approach to support many AI use cases rapidly becomes too difficult.

AI-enabled integration platforms provide automation, intelligence, and interoperability to overcome the legacy platform challenges.

**The AI Integration Opportunity**
AI has the promise to transform the telecom industry by improving customer experience, reliability, efficiency, and profitability. To overcome the modernization challenge, an opportunity is emerging for AI integration solutions that simplify the integration of legacy systems and data. AI and LLMs can transform customer experience, network operations, and service automation — but only if connected to the right data and systems of record.

Traditional integration solutions (ETL, ESB) can't handle semantic, context-aware access or dynamic workflows. AI-enabled integration platforms go beyond data movement — they orchestrate intelligent actions across systems.

Introducing the AI-Enabled Integration Platform
An AI-enabled integration platform is a platform that combines:

- Traditional integration capabilities for API abstraction, composite services, multi-step processes, data consistency, and low-code tooling
- AI agents for reasoning and orchestration
- Model Context Protocol (MCP) for safe, standardized access to enterprise data and tools
- Agent-to-Agent (A2A) protocol for multi-agent collaboration
- AI sidecars and façades to incrementally connect legacy systems

Such a combination of functionalities enables semantic data access and context assembly across disparate systems. It can execute processes that consider policies for security, compliance and observability. It allows for incremental modernization and quick wins, building into a unified AI ecosystem for telecom customer experience and operations.

**Architecture Overview**
The diagram on the following page illustrates the architecture of an AI-enabled integration platform with integration with legacy systems.

The CX channel layer, or engagement layer consists of all the customer and/or internal operator channels, such as chatbots, voice, and apps. The channels interact with AI agents. These AI agents can reside within the AI-enabled integration platform; or they can be external AI agents which interact with the AI agents within the AI-enabled integration platform.

The AI agents are interconnected in a mesh using the A2A protocol. This "Agent-to-Agent" open protocol was launched April 2025 by Google to enable communication and collaboration between autonomous AI agents across different vendors, systems and frameworks. It is designed to address inter-agent interoperability rather than simply agent-to-tool or model-to-system integration.

The AI-enabled integration platform also includes domain adapters to interface with legacy systems using their native interfaces (e.g. APIs, files, UIs).

> AI-enabled integration platforms go beyond data movement — they orchestrate intelligent actions across systems.

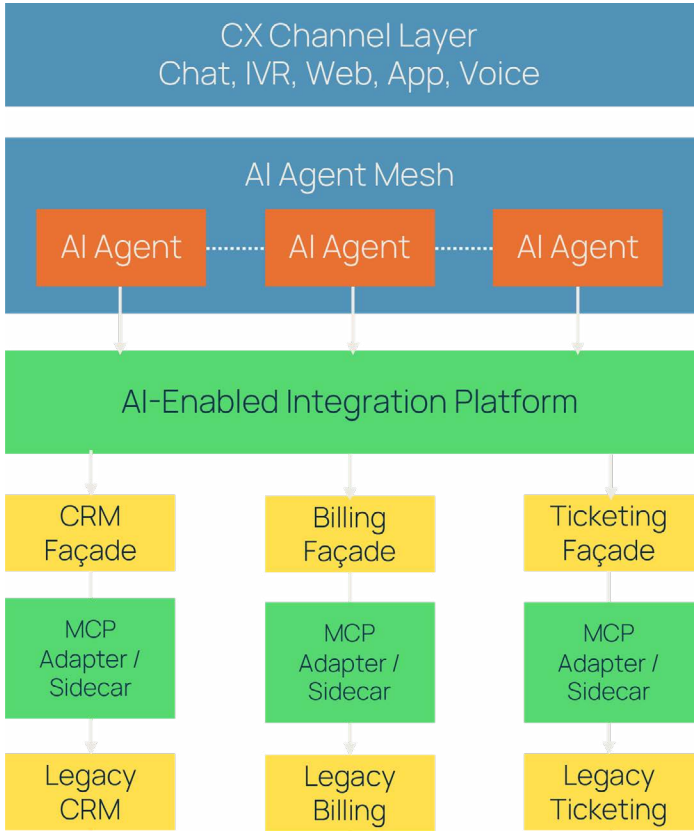


These adapters incorporate the Model Context Protocol (MCP). MCP was launched by Anthropic in November 2024 as a "universal AI interface" allowing large language models to communicate with external data sources, tools, and services. The AI-enabled integration platform places the MCP interface on top of the native system interfaces using the traditional integration platform capabilities augmented with MCP server capabilities.

The domain adapters can be implemented as façades and/or sidecars. With a façade, the adapter provides a proxy or front door to the legacy system without changing the underlying functionality of the system. It hides the complexity of the underlying legacy implementation. With a sidecar, the adapter is deployed alongside the legacy system, augmenting it with AI functionality without change to the underlying legacy implementation.

The AI-enabled integration platform also supports governance & observability. It has tools for security control and monitoring of all interfaces, enabling modern compliance tools to be used across the span of legacy systems.

**Incremental Integration Through Sidecars and Façades**
One of the challenges faced by operators is how to get started with AI. ROI models are often unclear, so a big bang AI transformation program is unrealistic. Sidecars and façades can provide a mechanism to incrementally introduce AI capabilities and use cases and gradually grow out AI functionality.

With the sidecar pattern, a domain adapter can be attached as an AI agent beside each legacy system for safe read/write access, without altering the existing interfaces to the legacy systems. In this way an AI use case can be introduced without risk of interference with existing functionality.

Over time, the capabilities of the sidecar can be improved, and more applications can move to the sidecar interfaces. At some point, all functionality is realized using the new interfaces, and the sidecar can be promoted to a façade pattern; the façade is now the front end for all interactions with the domain. In the future, when the legacy system is replaced, the MCP-based interfaces can be natively supported by the modern replacement for that domain.

For example, a ticketing sidecar agent may be introduced to enable AI analysis of customer tickets. Later the sidecar agent is enhanced to add ticket creation and management capabilities. Client systems that require access to ticketing are gradually migrated to utilize the AI agent interfaces. Once the migration completes, the sidecar becomes a ticketing Façade agent.

The benefits of this approach using sidecars and façades include fast rollout, minimal risk, and continuous improvement.

**Use Case Scenarios**

Consider an example use case. Customers may become frustrated when trying to subscribe to new services (say to an IPTV package) and are looking to confirm that their subscriptions have been enabled and they are being billed correctly. An AI agent can manage the customer interaction, verify the package is provisioned in their CRM record, check the billing status, and make corrections if necessary. The AI agent can open and close a ticket for the current care episode and look for historical tickets from this customer to optionally offer an incentive or commercial offer.

Another use case is predictive maintenance or predictive care for a customer's home network. Operators receive many complaints about Wi-Fi, IPTV, or other technical issues customers encounter in their home network. Predictive care works to analyze the data from the home network, predict the likelihood of future impact on customer service, execution of preventative actions to prevent impacts (self-healing), and continuous learning of data and outcomes to refine the models and improve results.

Over time, the capabilities of the sidecar can be improved, and more applications can move to the sidecar interfaces.

Predictive care utilizes multiple AI agents:

- Data analysis/cleansing agent that processes time-series data from the home network (through the ACS), Access network data (fiber/ONT, radio/FWA), CRM/ticketing data, etc.
- Issue prediction agent that detects anomalies in time-series data and predicts emergence of an impact to customer service.
- Next best action agent that determines the action most likely to resolve the issue.
- Issue remediation agent that executes the action in accordance with operator policies (pre-checks, scheduled time windows, etc.)
- Monitoring agent that confirms issue resolution.
- Optimization agent that compares results and adjusts prediction and other models accordingly using AI-ops methodology.

Without an AI-enabled integration platform, it is complex, expensive, and time-consuming to build such use cases. The result would be fragile and difficult to manage. These use cases would also be themselves siloed, without the ability to share components and achieve synergy. With an AI-enabled integration platform, the AI applications can be coordinated, automated, and contextual.

**Governance, Trust, and Compliance**
The AI-enabled integration platform plays a key role in data governance, trust and compliance. Adding a MCP layer on top of the legacy systems provides these key capabilities, that augment, rather than replace or bypass existing security controls.

MCP acts as a policy enforcement point (PEP) between the AI agent and the underlying data and tools. Access Control Lists (ACLs) define which agents or users can access which endpoints, datasets, or tool APIs. Attribute-Based Access Control (ABAC) implements policies that can depend on context (e.g., user role, purpose of use, sensitivity level, time, or location).

MCP enforces user and system-level consent constraints by binding context access to consent tokens or data agreements. Each data request carries a token that encodes user consent or legal basis (e.g., GDPR consent, service contract). The token specifies the intended use of data ("AI troubleshooting," "billing optimization," etc.) and can restrict further use or sharing. If consent is revoked, the MCP layer invalidates tokens or removes access automatically — downstream AI models are signaled to purge cached context. All consent validations are logged for compliance verification.

> Without an AI-enabled integration platform, it is complex, expensive, and time-consuming to build such use cases.

MCP provides end-to-end traceability for every piece of data that flows into an AI interaction. This enables auditable AI — critical for telecom operators. Each MCP context includes metadata tags: source system, dataset version, retrieval time, and transformation path. The MCP registry maintains a graph linking every AI interaction to its originating data sources and policies. Unique identifiers track data through adapters, transformations, and sidecars — so results can be traced back. Logs and telemetry allow compliance teams to inspect what data was accessed, by whom, when, and why.

**Business and Operational Benefits**
The introduction of an AI-enabled integration platform provides the following benefits across the operator's organization.

**Summary**
The emergence of new AI protocols (MCP, A2A) enables a new generation of Integration Platforms that are AI-enabled. These platforms let telecom operators modernize with AI without replacing their legacy systems.

Using a sidecar-first strategy, operators can incrementally add AI use cases to their existing network, operations, and customer service. Using MCP and A2A protocols provides standardization to future-proof interoperability and builds a foundation for enterprise-wide AI governance.

| Organization | Value |
|---|---|
| **Customer Experience** | Faster resolution, personalized offers |
| **Operations and Engineering** | AI-driven automation without re-engineering |
| **IT** | Lower TCO, fewer brittle integrations, faster experimentation and AI rollout |
| **Compliance** | Standardized audit and data governance |

# TM Forum Open Digital Architecture

# TM Forum Open Digital Architecture – A blueprint for intelligent operations

The TM Forum Open Digital Architecture (ODA) provides a migration path from legacy IT systems and processes to modular, cloud-native software orchestrated using AI.

ODA comprises tools, code, knowledge and standards (machine-readable assets, not just documents). It is delivering business value for TM Forum members today, accelerating concept-to-cash, eliminating IT & network costs, and enhancing digital customer experience.

Developed by TM Forum member organizations through our Collaboration Community and Catalyst proofs of concept, ODA is being used by leading service providers and software companies worldwide.

ODA includes:

- An architecture framework, common language, and design principles
- Open APIs exposing business services
- Standardized software components
- A reference implementation
- Guides to navigate digital transformation
- Tools to support the migration from legacy architecture to ODA
- Maturity models and readiness checks to baseline digital capabilities.

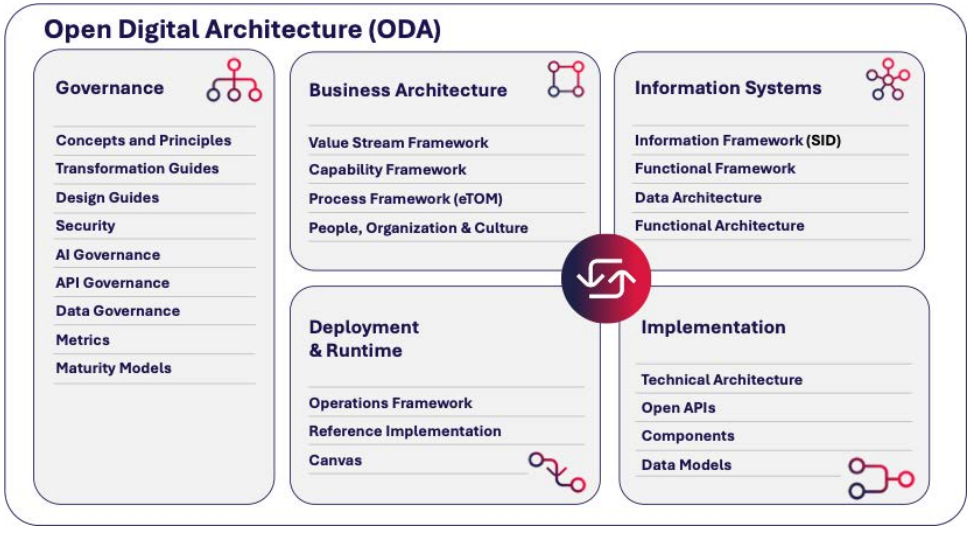


## Goals of the Open Digital Architecture

The aim is to transform business agility (accelerating concept-to-cash), enable simpler IT solutions that are easier and cheaper to deploy, integrate and upgrade, and to establish a standardized software model and market which benefits all parties (service providers, their suppliers and systems integrators).

### Learn more about collaboration

If you would like to learn more about the project or how to get involved in the TM Forum Collaboration Community, please contact **George Glass.**

# TM Forum research reports


Assessing the opportunities and challenges for AI-RAN strategies


ZTE
Optimizing governance to accelerate GenAI deployment

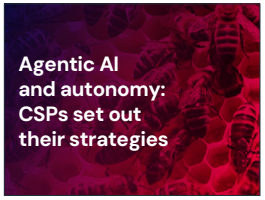
Agentic AI and autonomy: CSPs set out their strategies


csg
B2B2X business models: Making sense of partner platforms


How identity is driving the market for network APIs


CLOUDERA
Unlocking AI potential through a modern data architecture


tmforum
Monetizing the AI connectivity opportunity


tmforum
BENCHMARK REPORT
Composable IT & Ecosystems
assessing CSPs' progress
towards an Open Digital


tmforum
How does your GenAI strategy measure up?
Key findings from our Generative AI Maturity Interactive Tool
aws

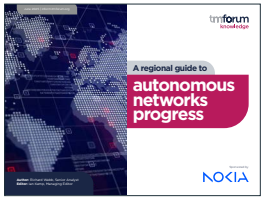
tmforum
A regional guide to
autonomous networks progress
NOKIA


tmforum
TREND REPORT
Composable IT & ecosystems
future BSS:
building new customer-focused strategies
ETIYA


tmforum
TREND REPORT
Autonomous network operations
the role of digital twins
in autonomous networks
CELFOCUS
ZTE


tmforum
BENCHMARK REPORT
Autonomous network operations
Autonomous networks:


tmforum
TREND REPORT
Data & AI Innovation
leveraging AI
for service and value creation
cerillion
PEGA
RedHat


tmforum
knowledge
how NaaS is driving the evolution
of telco operations
blueplanet


tmforum
TREND REPORT
Autonomous network operations
towards autonomy
in mobile private networks
ZTE


tmforum
BENCHMARK REPORT
Telco revenue growth:
Is B2B a beacon of light?


tmforum
REGIONAL REPORT
innovating in Asia:
exploring the key drivers for growth
csg


tmforum
REPORT
Meeting expanded expectations
for network charging
amdocs


tmforum
REGIONAL REPORT
innovating in the Americas:
exploring the key drivers for growth

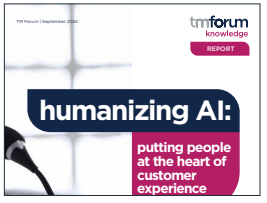
tmforum
knowledge
REPORT
humanizing AI:
putting people at the heart of customer experience


tmforum
knowledge
REPORT
in search of innovation:
the future for telecoms R&D


tmforum
knowledge
REPORT
Digital Transformation Tracker 8:
cloud adoption


tmforum
knowledge
REPORT
driving intelligence in network lifecycle automation:
focus on Asia-Pacific
servicenow


tmforum
knowledge
REPORT
finding skills for the future:
inside the telco talent revolution


tmforum
knowledge
REPORT
Managing the economics of hybrid cloud:
striking a


tmforum
knowledge
BENCHMARK
building an AI strategy:
telcos put the


tmforum
knowledge
REPORT
delayering the business:
is breaking up key to telco growth?

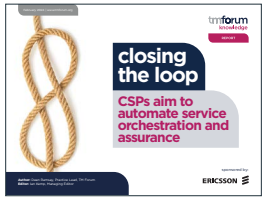
tmforum
knowledge
REPORT
closing the loop
CSPs aim to automate service orchestration and assurance
ERICSSON


tmforum
knowledge
REPORT
a numbers game:
exploiting data to drive future strategies
splunk>


tmforum
knowledge
REPORT
BSS for B2B
operators diverge on the path to cloud


tmforum
knowledge
REPORT
generative AI:
operators


tmforum
knowledge
REPORT
the sustainable telco:
navigating the


tmforum
knowledge
BENCHMARK
telco revenue growth:
time for


tmforum
knowledge
REPORT
Cybersecurity strategies:
risk management moves firmly into the telco spotlight
NOKIA


tmforum
knowledge
REPORT
leveling up:
achieving Level 3 autonomous networks and beyond


tmforum
knowledge
REPORT
wholesale changes:
rethinking support systems for new fiber operators
Prodapt


tmforum
Reigniting
telecoms growth
a Playbook for CEOs
TECH mahindra


tmforum
knowledge
REPORT
how software-as-a-service
is reshaping business support systems
Tecnotree


tmforum
knowledge
REPORT
Digital Transformation Tracker 7
cutting complexity
with automation and AI

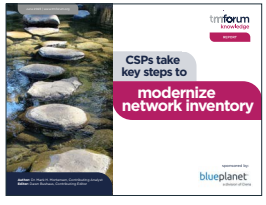
tmforum
knowledge
REPORT
CSPs take key steps to
modernize network inventory
blueplanet


tmforum
knowledge
REPORT
transforming BSS:
racing to a flexible, customer-focused future
qvantel

# Meet the Insight team

**Report Author:**
Dawn Bushaus
Contributing Analyst
dbushaus@tmforum.org

**Report Editor:**
Ian Kemp
Managing Editor
ikemp@tmforum.org

**Senior Analyst:**
Richard Webb
rwebb@tmforum.org

**Practice Lead:**
Dean Ramsay
dramsay@tmforum.org

**Head of Operations:**
Ali Groves
agroves@tmforum.org

**Commercial Manager:**
Tim Edwards
tedwards@tmforum.org

**Report Author:**
Joanne Taaffe
Editor in Chief, Insight
jtaaffe@tmforum.org

**Chief Analyst:**
Mark Newman
mnewman@tmforum.org

**Global Account Director:**
Carine Vandevelde
cvandevelde@tmforum.org

**Customer Success Project Manager:**
Maureen Adong
madong@tmforum.org

**Customer Success Project Manager:**
Amanda Alexander
aalexander@tmforum.org



**Published by:**
TM Forum

European Office
25 Worship St,
London EC2A 2DX
United Kingdom

US Office
181 New Road, Suite 304, Parsippany,
NJ 07054 USA
Phone: +1 862-227-1648

ISBN: 978-1-969387-04-3

**Report Design:**
Paul Martin