



AI agent trends 2026

Telecommunications

 Interactive report

About this report

This report provides key insights for business leaders to shape their AI agent strategy for 2026 and beyond. Within each trend, you will find real-life examples, technical resources, and customer stories to share with your teams for deeper learning.

These trends were identified using a blend of qualitative and quantitative data, including internal Google Cloud and Google DeepMind interviews with AI leaders, customer case studies, and insights from [The ROI of AI in telecommunications 2025](#) report (based on a global survey of 489 executives in telecommunications), with analysis using NotebookLM and Google AI Studio.

Ask the report

Want to dig deeper into the trends? NotebookLM is ready to help—simply click the icon for an interactive version of this report.





AI agents are redefining business value in 2026

The ceiling for human achievement has been lifted.

While headlines focus on the distant future of Artificial General Intelligence (AGI), the decisive shift for the telecommunications industry is happening right now. That shift is agentic AI. This is AI that moves beyond answering questions to understanding a goal, making a plan, and taking actions across applications to achieve it with extensive human guidance and oversight.

Due to this capability, AI is one of the first technologies that applies to every single person, in both our personal and professional lives. Its power lies in its ability to augment human capacity with better recall, faster data processing, and enhanced reasoning across the technical, commercial, and operations teams. However, the true transformation will unfold when telecommunications providers use these agents to bridge the gap to the end-users, gathering the feedback needed to drive continuous value of their products and services.

In this report, we explore five key AI agent trends shaping the telecommunications industry in 2026. Unlocking the value of these trends requires more than simply adopting new tools. It also demands that leaders question old assumptions and drive the cultural change necessary to thrive in this new, agentic AI era.

What are AI agents?

Agents are systems that combine the intelligence of advanced AI models with access to tools so they can take actions on your behalf, under your control.

| I/O, May 2025

| Keynote by Sundar Pichai, CEO, Google

5 AI trends shaping telecommunications in 2026


1 Agents for every employee

Empowering individuals to achieve peak productivity




2 Agents for every workflow

Running your business with grounded agentic systems




3 Agents for your customers

Delighting customers with concierge-like experiences




4 Agents for security

Advancing security from alerts to action



5 Agents for scale

Upskilling talent will be the ultimate driver of business value





Agents for
every employee

Empowering individuals to achieve peak productivity






The most significant business shift in 2026 is a move from simple efficiency to a fundamental, employee-centric transformation.

This new agentic model is designed to expand the potential of every individual, turning them into the primary engine for innovation and growth.

For example, employees will be increasingly able to state a desired outcome, and the computer—using LLMs and agents—determines how to deliver it.



56%

of telecommunications executives in organizations that use generative AI have AI agents in production,¹ deploying them across a wide range of use cases.²

Of those:

- 39% use them for network or equipment configuration and automation
- 34% use them for network remediation
- 32% have deployed network agents

¹ Google Cloud, [The ROI of AI in telecommunications](#), 2025 (Telco executives: n=489; Questions: (1) How are AI agents leveraged across the enterprise? (2) How many AI agents does your company currently have deployed in production across your organization?)
² Google Cloud, [The ROI of AI in telecommunications](#), 2025 (Telco executives whose organization is leveraging agentic AI: n=274; Question: What use cases has your company deployed AI agents for?)



“

We are witnessing a fundamental shift where every telecommunications employee evolves from a task executor into a strategic orchestrator. By creating specialized agent systems, teams can focus on solving the problems that matter—giving every person the space to do their best, most impactful work.”



Angelo Libertucci

Global Head of Industry – Telecom
Global Strategic Industries, Google Cloud





A new integrated working model

In this new model, every employee—from the field technician to the network engineer—becomes a human supervisor of agents.

Their primary job is no longer to perform every mundane task personally, but rather to orchestrate a team of specialized AI agents to achieve a goal. This model is about more than just delegation; it's about augmentation. The real power comes from giving every employee agents grounded in the company's own enterprise context—its internal systems, knowledge bases, customer data, and past work—to elevate the impact of their efforts.

The employee's core function becomes providing strategic direction. Their new responsibilities are to:

Delegate mundane or repetitive tasks

Identify which operational tasks are best suited for an agent and assign them.

Set goals

Clearly define the desired outcome for the agent.

Outline strategy

Use their human judgment to guide the agents and make the final, nuanced decisions that AI can't.

Verify quality

Act as the final checkpoint for quality and performance.

📋 What is grounding in AI?

It is the process of anchoring an AI model's responses to a specific, verifiable set of facts—its “ground truth.” For an enterprise, this ground truth is its own internal data.

[Learn more](#)

How it works

The 10x network operations specialist

A network operations specialist’s job used to be burdened by the constant scramble of repetitive ticket routing and reactive troubleshooting. In 2026, they can orchestrate a system of specialized AI agents to achieve their goals, rather than performing every task personally.

With agents dedicated to specific tasks, the network operations specialist can multiply their output by focusing on designing more resilient, self-healing systems that can fix outages and issues before they ever impact consumers.

Analyze real-time network telemetry data to predict expected traffic patterns and issues in supporting tonight's sold-out concert in Toronto

Thinking...

◆ Data sources gathered ▾

◆ Compiling data and trends...

I now have raw data. Now I need to process it to create the table structure and calculate variances.

Processing raw data

Processing DataFrame

◆ Executing calculations...

The table is now generated with all the required calculations. The percentage columns are still showing full precision, I need to round them to one decimal place. Also, the output format for percentages is not quite right as it is not standardized across the sources. I will fix this in the next step. I also need to extract raw data for '.csv' generation.

◆ Re-run calculations...

🛡️ Ask anything, search your data, @mention, or /tools

+

⌂

📄

◆

➤

May display inaccurate information, so double-check responses

An example of a Gemini Enterprise app home page (agent shown is built by a user)

Google Cloud

9

The network operations specialist's specialized agent system

Their new role involves orchestrating four specialized agents that bridge the gap between human intent and network outcomes:

These agents, in turn, will use horizontal agents such as the ones for translating natural language into SQL or GQL, or an automation script, for calling AI/ML models and so on.

Forecasting agent

Predicts capacity demand, defects, and other potential changes.

Anomaly detection agent

Identifies anomalies in the network behavior and compresses them by correlating ones with the same cause.

Root cause analysis agent

Identifies the most probable root causes. For example, it can use subagents that retrieve insights from multiple structured or unstructured data sources.

Remediation agent

Takes closed-loop automated remedial actions or creates tickets (or sends messages) for human intervention. For example, it could connect to information technology service management (ITSM) or field service management (FSM) tools such as ServiceNow, Remedy, or JIRA.



Data tip

Teams can use data agents as an intelligence multiplier.

[Learn how](#)



“

This surge in AI [agent] adoption reflects a fundamental mindset shift to recognizing it as a productivity instrument available 24/7. At TELUS, we’ve seen this firsthand—over 57,000 team members regularly use AI and save 40 minutes per AI interaction.”

Jaime Tatis


Chief AI Officer, TELUS



Ready to 10x your team?

Gemini Enterprise Agent Platform enables employees to build and manage their own specialized AI agents.

Try now





Agents for
every workflow

Running your business with grounded agentic systems

An abstract, flowing, wavy pattern in shades of orange and yellow, resembling a liquid or smoke effect. It is composed of many small, glowing dots and lines, creating a sense of movement and energy. The pattern starts from the bottom right and flows upwards and to the left, partially overlapping the text.

An agentic system is a digital assembly line—a human-guided, multi-step workflow that orchestrates multiple agents to run complex business workflows.

Telecommunications use cases can include coordinating technical operations to run complex sequences like autonomous network self-healing or site-maintenance optimization.

The true value in 2026 begins with augmenting the capabilities of individuals and teams, and continues to grow by making the entire business run more intelligently and efficiently, 24/7, at scale.



43%

of telecommunications executives say their organizations have launched more than 10 AI agents³

“The real value will lie in using agentic workflows to integrate siloed functions, such as network operations, field services, and customer contact centers in telecommunications. Agents could autonomously remediate network anomalies, proactively open a ticket with the field service systems, and alert contact centers to inform customers of a technician dispatch, all in one integrated sequence.”



**Angelo
Libertucci**

Global Head of Industry – Telecom
Global Strategic Industries, Google Cloud

³ Google Cloud, [The ROI of AI in telecommunications](#), 2025 (Telco executives: n=489; Question text: How many AI agents does your company currently have deployed in production across your organization?)


Digital assembly line: Orchestrating agentic systems

A digital assembly line is made possible by the Agent2Agent (A2A) protocol.

This open standard enables seamless integration and orchestration between AI agents, allowing them to work together even if they are from different developers, built on different frameworks, or owned by different organizations.

While LLMs are the “brains” of these agents, they have two major limitations: their knowledge is frozen at the time of their training, and they can’t interact with the outside world to access real-time data or perform actions.

The Model Context Protocol (MCP) solves this. It creates a standardized, two-way connection for AI applications, allowing LLMs to easily connect with various data sources and tools, such as managed databases (e.g., [Cloud SQL](#), [Spanner](#)) and data platforms (e.g., [BigQuery](#)).





Vodafone is using Google Cloud and the autonomous framework to establish a shared agentic architecture and A2A workflow to enhance investment planning, customer experience and operational efficiency.



Data tip

Discover how your data teams can ground AI in your business data.

[Get started](#)



How it works

An agentic network in action

Today's network systems assume a human initiates every order.

This poses a challenge for autonomous networks: How do we secure a system where an agent makes financial and operational decisions, under the guidance of a network engineer?

In 2026, network provisioning is treated as a micro-transaction. An agent initiating a configuration change utilizes the Agent Payments Protocol (AP2) to provide a cryptographically signed “Proof of Authorization.” This allows the network to verify that the agent is not only authorized to make technical changes but is also operating within a specific verified budget—creating an immutable audit trail for every micro-decision.





Example of the agentic workflow

Instead of writing static rules, a network engineer defines a strategic intent for their autonomous agent:

“Preserve ‘Gold Tier’ latency standards for the downtown Toronto cluster during a national hockey game. You have an authorized budget of \$5,000 for dynamic slice acquisition and cloud bursting. Optimize for lowest cost while maintaining service levels.”

The outcome

The network agent does not just react; it negotiates. It communicates with a finance agent (via the A2A protocol) to forecast cost-per-minute and dynamically provisions capacity only when the cost-benefit analysis aligns with the engineer’s intent. The agent digitally signs the resource purchase using AP2, executing the fix in milliseconds as part of the pre-approved strategy and guidance set by a network engineer.



**Try Gemini Enterprise
Agent Platform**

Get started



Deutsche Telekom uses a network AI agent to analyze network behavior, detect performance issues, and implement corrective actions to improve network reliability, reduce operational costs, and enhance customer experiences.

1

2

Trend 3


4

5



Agents for
your customers

Delighting customers with concierge-like experiences

A large, abstract, three-dimensional structure composed of numerous small, glowing green and blue particles. The structure has a jagged, crystalline appearance, resembling a mountain or a complex molecular model. It is set against a dark background, with the particles themselves providing the primary light source, creating a vibrant, ethereal glow.



For the last decade, customer service automation meant pre-programmed chatbots to answer simple questions and deflect support tickets.

They were efficient, but they lacked the ability to understand more nuanced and complex questions. This was especially true for home subscribers. Without a view into the network, these chatbots couldn't tell if a connection was lagging due to a neighborhood outage or a faulty home router—forcing the customer to act as their own technician.

With advances in LLMs and A2A, 2026 will deliver more helpful concierge-style agents. These AI agents will connect enterprises and customers by remembering preferences and past conversations to offer truly one-to-one experiences.



45%

of executives at telecommunications organizations with AI agents in production report adopting agents for customer service and experience⁴

“When AI handles the technical diagnostics and groundwork, your contact center teams can transform every interaction from a resolved ticket into a deepened relationship, driving growth through higher customer loyalty and long-term lifetime value.”



Angelo Libertucci

Global Head of Industry – Telecom
Global Strategic Industries, Google Cloud

⁴ Google Cloud, [The ROI of AI in telecommunications](#), 2025 (Telco executives whose organization is leveraging agentic AI: n=274; Question text: What use cases has your company deployed AI agents for?)

Personalization at scale

Your customer no longer has to start every conversation by proving who they are or re-explaining their problem.

The difference isn't just the AI, it's the data. The agentic concierge can succeed because it is grounded in an enterprise context you choose to share for a particular customer—from purchase history in the CRM to package tracking in the logistics database.

Chatbot

“Thanks for contacting support about slow internet speed. Please enter your 12-digit user ID.”

Agentic concierge

“Hi, Elizaveta.
This is your internet provider's automated assistant. Thanks for contacting support about your slow internet speed. The network is healthy, but you are hitting your bandwidth cap. I can upgrade your plan immediately within the support chat. Please respond “Upgrade” or “Talk to a human.” Thank you.”



verizon

Verizon partnered with Google Cloud to integrate Gemini-powered AI into customer service operations, achieving a 40% increase in sales and a 95% answer rate for customer inquiries.

How it works

Insightful and always-on help


Customer service is defined by its helpfulness.

An agentic concierge doesn't wait for a complaint. It monitors systems for triggers and resolves problems using real-time data to provide insights and take actions with human guidance and oversight.



Try Gemini Enterprise for Customer Experience

Get started






Building trust through seamless global service.

Imagine this scenario: A customer arrives in Japan. Instead of receiving a traditional text message saying “Welcome to Japan” that outlines the rates for calls and data, the concierge-style agent orchestrates the following:

- 1 **Checks** the network to confirm the customer landed in Japan.
- 2 **Accesses** the customer’s data usage history and determines they are a “High Data User.”
- 3 **Notifies** the customer via text message: “Welcome to Tokyo! Based on your typical data usage, the standard roaming pass will run out in 2 hours. I can activate an ‘Unlimited Data Pass’ for \$5/day that covers unlimited video streaming. Shall I turn that on for you?”

The problem is resolved quickly. The agent is designed to execute a smart handoff with a full summary for complex or emotionally charged issues; human staff can step in if the case becomes much more complex.



Bell Canada has built customizable contact center solutions for its business customers that offer AI agents to address callers, and Agent Assist, which listens when a human agent is on, offering suggestions and sentiment analysis. AI has contributed \$20 million in savings across customer operations.

1

2

3

Trend 4

5



Agents
for security

Advancing
security
from alerts
to action





In a modern security operations center (SOC), human analysts face a constant stream of data and alerts from across global networks, with 82% concerned or very concerned that they may be missing real threats or incidents due to alert fatigue.⁵

This alert fatigue is the attacker's greatest advantage; they only need to be right once, while the defender has to be right every single time.

While security orchestration automation and remediation (SOAR) solutions deliver some automation, they may offer only incremental benefits. But with their ability to reason, act, observe, and adjust actions across core and edge network infrastructures under the guidance of a security operations center analyst, AI agents have the potential to help security teams identify and respond to threats more effectively.

Expanded [Secure AI Framework 2.0](#) can help address the rapidly emerging risks posed by autonomous AI agents. Agentic chaos engineering can play an important role, too, by simulating complex “network-wide” attacks to identify hidden vulnerabilities in 5G slices and virtualized infrastructure before they can be exploited.

In 2026, AI agents will increasingly help with tasks like vulnerability discovery as well as alert triage and investigation.



47%

of executives at telecommunications organizations with AI agents in production report adopting agents for security operations and cybersecurity⁶

⁵ Forrester, [Threat Intelligence Benchmark: Stop reacting; start anticipating](#), 2025

⁶ Google Cloud, [The ROI of AI in telecommunications](#), 2025 (Telco executives whose organization is leveraging agentic AI: n=274; Question text: What use cases has your company deployed AI agents for?)



Today's CISO is laser-focused on achieving the greatest decrease of risk per dollar spent. Agents are essential to this, as they detect and respond faster to enterprise risks. More importantly, they elevate our SOC analysts from tactical responders to strategic defenders."




**Jon
Ramsey**

Vice President & General Manager,
Security, Google Cloud

The semi-autonomous security operations cycle

An agentic SOC orchestrates a system of task-based AI agents, each with a specific role, to achieve a common security outcome.

After receiving a security alert, the agentic SOC cycles through a process, engaging various agents:



This dynamic process of evaluating, acting, and re-evaluating enables the system to adapt to a changing security environment in real time, while freeing up time for human analysts to focus on higher-value work.

Multiple SOC agents need common enterprise context and can share the same security data sources (e.g., security telemetry data), regularly communicate, and adapt their actions through technologies like A2A and MCP. Agents should also be trained on continuously evolving real-world insights from security experts.



“

AI is already being used to find zero-days, identify malicious code, and uplift the work of defenders. As threat actors incorporate the technology into their operations, it will be our best tool to meet this new challenge.”



Sandra Joyce

Vice President of Threat Intelligence,
Google Cloud



One New Zealand used Google Security Operations to improve its ability to defend against attacks in real-time, resulting in a 48% operating cost reduction.

How it works

Elevating the security analyst

With the addition of agentic systems acting as force multipliers, human analyst roles are expected to change for the better. AI agents can take on the draining, reactive work of “alert-watching,” shifting the human analyst to a more strategic level, engaging in activities such as:

Threat hunting

Using their intuition and experience to guide the agents. Example: “Agent, I have a hunch about this server. Hunt for any unusual outbound data transfers.”

Supervising agents

Fine-tuning the agents’ “rules of engagement” and “performance-reviewing” their automated responses.

Defending

Focusing on long-term security posture, architecting better defenses and anticipating the next wave of attacks, rather than just chasing the current ones.



Explore Google
Security Operations

Get started

1

2

3

4

Trend 5



Agents
for scale

Upskilling
talent will be
the ultimate
driver of
business
value

A large, abstract, three-dimensional cloud-like structure composed of numerous small, glowing blue and purple spheres. It occupies the right side of the image, partially overlapping the text.



It is tempting to focus on the technology—the models, the platforms, and prompts—but this misses the most critical element: the people.

As AI is evolving, the skills gap is widening and it's getting harder for individuals and organizations to keep up. Skills themselves expire faster than ever: The “half-life” of a professional skill is now four years—and in tech, as short as two years.⁷

Both practitioners and decision-makers see the importance of closing this gap. Skills increase the ability to get hired, get promoted, and grow careers—plus, they have a positive impact on productivity, innovation, and revenue.

“In 2026, the primary constraint on growth is no longer technology, but the speed at which we can close the skills gap to wield it effectively. Success depends on building a culture where every employee can continuously learn to harness these agentic systems for massive operational scale.”



Angelo
Libertucci

Global Head of Industry – Telecom
Global Strategic Industries, Google Cloud

⁷ Forbes, [AI Puts The Squeeze On The Shrinking Half-life Of Skills](#), 2024




What executives are saying:⁸



What employees are saying:⁹



⁸ Google/Ipsos, Cloud Learning Services Market Pulse, fielded Sept.-Nov. 2024 (US, UK, FR, DE, IN, BR, MX, JP, AU/NZ; Decision-makers (n=902). Findings reflect the opinion of survey respondents only)
⁹ Google, [Beyond AI Optimism: Five ways to move your business from saving time to sparking innovation](#), 2025





AI offers an unprecedented opportunity for employees to harness the data and context around them. 2026 will be the year when every employee can go from guessing to knowing—but only if their organizations invest in the skills to make it possible.”



**Andrew
Milo**

Global Director, Customer Training,
Cloud Learning Services,
Google Cloud



How it works

The 5 pillars of AI learning

To thrive, organizations must move beyond simply buying technology and focus on building an AI-ready workforce. This requires a holistic strategy built on five key pillars.

01

Establish goals

Figure out what is most important and what you can actually measure. For example, perhaps your business goal is 100% AI tool adoption within your organization. So, every team member—from field technicians to network engineers— would use a Gemini Enterprise app at some point in their workflow, making the flow better than what they could have done on their own. Make sure your specific business goal aligns to the bigger picture of what the organization needs—like reducing time-to-resolution—and that it's measurable.





02

Secure sponsorship

Gathering a team together ensures that there is constant communication and momentum around your AI initiative. A good place to start is with a team of three primary stakeholders: an **executive sponsor**, the **groundswell lead**, and the **AI accelerator**. Let's break these roles down.

The **executive sponsor** provides the necessary funding, high-level backing, and consistent messaging on AI's importance.

The **groundswell lead** acts as the "AI megaphone" by managing grassroots campaigns, generating excitement, and collecting employee ideas.

The **AI accelerator** is the technical expert responsible for transforming those prioritized ideas into functional solutions in areas such as autonomous network healing, agentic customer support, and predictive infrastructure maintenance.

This structure ensures the program has the support to be funded, the momentum to engage employees, and the technical capability to deliver tangible AI tools.

03

Sustain momentum and reward innovation

To sustain engagement, the program should use a layered strategy centered on interactive platforms and consistent communication. Think of the program as a "network innovation hub"—a gamified idea exchange where field technicians, engineers, marketers, and customer support representatives can submit AI use cases and new ideas for service improvement. Peer-to-peer knowledge should be highlighted in a way that showcases successful AI applications across various roles.

Maintain momentum through regular communication, including executive updates, AI-focused team meetings, and a quarterly awards program to recognize top innovators.



04

Integrate AI into daily workflows

Keep the AI education going and tap into your team's collective expertise by hosting internal hackathons. Pair cross-functional teams, such as network engineers with customer support leads or field technicians with data scientists, to develop innovative AI solutions. Winning concepts could be integrated into the pilot programs.

Consider hosting Field Days, which are structured, challenge-based events that encourage teams to practice using new agentic AI tools in a collaborative, practical setting. These can be both in-person and virtual.

05

Prepare for increasing risks with trusted frameworks

As agent-accelerated cyber security risks become ever more sophisticated, security becomes everyone's responsibility. It's critical to ensure that employees understand their role—including training on what data can (and cannot) be used in AI tools, and how to recognize sophisticated threats like social engineering that uses AI.

Due to the heavy reliance on a global ecosystem of equipment vendors and cloud-native software providers, telecommunications employees must also be vigilant regarding vulnerabilities in decentralized 5G infrastructure or unauthorized access points. Mitigating these risks requires robust training alongside close partnerships with security teams to enforce zero-trust protocols and ensure end-to-end network integrity across all vendor interfaces.



Try Google Skills

Find AI training courses and lessons on a wide range of topics.

Get started



“

AI adoption should be an ongoing story, instead of just one chapter in a book. We had found that 96% of team members reported that they had increased confidence using AI tools. And 96% committed to applying these tools in their work. Our Google Skills training program has actually doubled in impact from February all the way till September 2025.”



Gina Narvaez


Director of Digital Talent Team, TELUS





The 2026 opportunity: A path to business growth





Although the 2026 opportunity can seem technical—and the adoption curve steep—it is fundamentally about innovation and business growth.

Telecommunications providers that experiment today are not just building tools; they are building the critical, in-house expertise to manage, govern, and scale this new capability.

For telecommunications providers, the move towards agentic AI presents a clear path to growth in 2026.

The true value of agentic systems for providers lies in freeing their teams to focus on the creative, strategic, and empathetic work that only they can do, fundamentally redefining how great products and services are designed, built, and delivered.

“ Access to agentic AI capabilities will democratize insights, innovation, creativity, and business growth, bringing value to consumers, employees, and organizations. This opportunity is truly significant, but comes with tremendous responsibility to ensure that the promise of AI delivers secure, ethical, and fair outcomes for all.”



Anil
Jain

Global Managing Director,
Global Strategic Industries,
Google Cloud

Google Cloud

See what
these trends
mean for your
business.

Get in touch

