

Whitepaper

April 2025

Role of Agentic AI in Telecommunications

Streamlining Digital and Network Operations

Author: Patrick Kelly, Principal Analyst Appledore Research



Appledore
RESEARCH

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: patrick.kelly@appledorerg.com • www.appledorerresearch.com

© Appledore Research LLC 2025. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means – electronic, mechanical, photocopying, recording or otherwise – without the prior written permission of the publisher.

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Appledore Research LLC recognizes that many terms appearing in this report is proprietary; all such trademarks are acknowledged, and every effort has been made to indicate them by the normal USA publishing standards. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Appledore Research LLC maintains that all reasonable care and skill have been used in the compilation of this publication. However, Appledore Research LLC shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising because of the use of this publication by the customer, his servants, agents or any third party.

Cover image by Patrick Kelly generated in chatgpt4.0 using prompt beehive box in Vincent van Gogh impressionist style

Contents

EXECUTIVE SUMMARY	1
INTRODUCTION.....	2
INDUSTRY OVERVIEW.....	4
KEY AREAS OF APPLICATION/USE CASES	5
Billing and Monetization.....	5
Customer Care.....	6
Service Assurance.....	6
Network Optimization	6
Network Security	7
CASE STUDIES	8
Enhancing AT&T's Customer Support with Generative AI	8
Lumen's Implementation of Microsoft Copilot for Sales	8
Vodafone Transforms Customer Engagement and Operations with Microsoft Generative AI.....	9
INNOVATION APPLIED TO NEW FRONTIERS OF AGENTIC AI IN TELECOM.....	11
Applying Agentic AI in Network Traffic Optimization	11
Service Assurance with RAN Slicing.....	12
Microsoft Expands Security Copilot with AI-Powered Autonomous Agents.....	13
MICROSOFT ISV ECOSYSTEM AND AI TRANSFORMATION	14
Nokia NetGuard solution.....	15
Norwood Systems.....	15
VALUE TO THE TELECOM C-SUITE	17
CONCLUSION	18

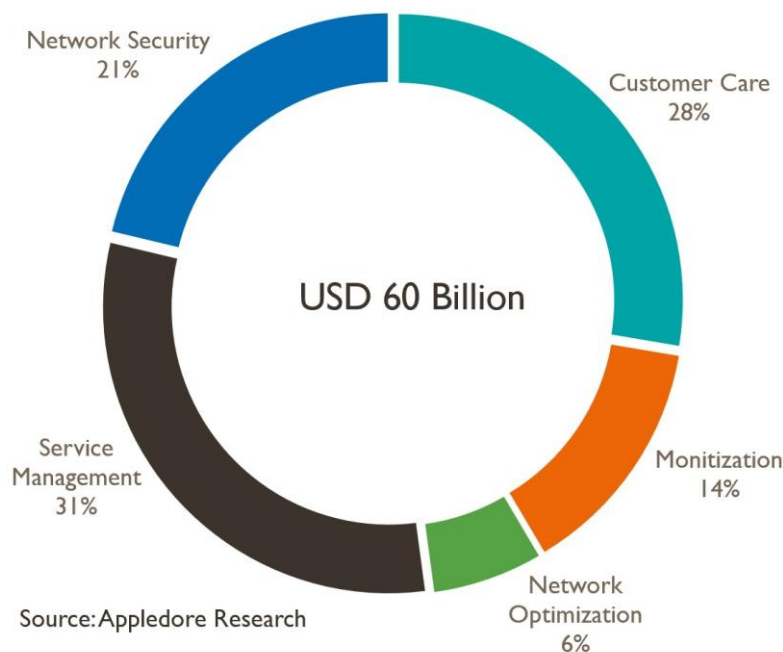
EXECUTIVE SUMMARY

In 2024, the combined global revenue of the top 100 telecommunications operators worldwide reached \$1.75 trillion. However, an analysis by Appledore Research revealed that the operational expenses for these operators amounted to a staggering \$1.38 trillion. This massive expenditure underscores a significant economic opportunity to enhance efficiency and automate numerous tasks that currently hinder Communication Service Providers (CSPs) from improving customer service, increasing revenue, and optimizing network performance.

Looking ahead to 2030, our investment thesis for Agentic AI assumes that allocating just 0.4 percent of the current operational expenses could dramatically transform digital enablement, operations, and management functions. By strategically applying these investments to the right areas of the business, there is immense potential to generate additional revenue, elevate customer satisfaction, and further optimize network performance.

Appledore Research forecast that AI Agent investment in the telecommunications sector is set to expand by a factor of 100 over the next five years, marking a profound shift in how operators approach automation and efficiency. A significant portion of this investment—approximately 60%—will be concentrated in Service Management and Customer Care, reflecting the industry's prioritization of customer-centric and operationally intensive domains for AI-led transformation. By 2030, Agentic AI is projected to unlock up to \$60 billion in operational cost savings, establishing itself as a critical lever for driving profitability and long-term competitiveness in the telecom market.

Figure 1: Appledore Research Forecast for AI OPEX Savings by Segment in 2030



Microsoft is pioneering a transformative shift in enterprise automation through its Agentic AI strategy, empowering Communication Service Providers (CSPs) to achieve significant levels of operational efficiency, productivity, and innovation. Agentic AI refers to autonomous AI agents capable of independently performing complex, multi-step tasks by dynamically interacting with structured and unstructured data, continuously learning and adapting to evolving conditions.

At the core of Microsoft's Agentic AI vision is the seamless integration of these intelligent agents into business processes across domains such as sales, customer service, finance, network operations, and cybersecurity. These agents are embedded into Microsoft 365, Dynamics 365, Copilot Studio, and Azure AI Foundry—providing low-code environments to customize, deploy, and manage agents at scale.

For CSPs, Agentic AI is rapidly redefining how networks are operated, monetized, and secured. AI agents can improve service assurance through real-time network anomaly detection, predictive maintenance, and application-level SLA enforcement. **Vodafone's** deployment of ODA Canvas, powered by Microsoft Azure and Agentic AI, demonstrated a 10x improvement in operational agility and enhanced app-level performance assurance for 5G services.

In customer experience, Microsoft has partnered with **AT&T** to deploy a generative AI chatbot that reduced unanswered queries by 50%, and with **Lumen** to implement Copilot for Sales—saving \$50M annually through productivity gains and automated CRM integration. These real-world implementations prove measurable impacts today of AI agents on customer engagement and revenue growth.

Agentic AI also enables CSPs to optimize network planning, reduce energy consumption, and enhance cybersecurity. Through real-time analysis and autonomous response mechanisms, AI agents support dynamic traffic routing, energy-efficient RAN management, and proactive defense against evolving threats. Microsoft's collaborations with **Nokia** (NetGuard Cybersecurity Dome) and **Norwood Systems** (CogVoice OpenSpan) showcase how AI-driven platforms are modernizing core telecom infrastructure.

From the C-suite perspective, Microsoft's Agentic AI empowers CEOs with strategic simulations, CFOs with improved forecasting, CIOs with streamlined governance, and CTOs with predictive network control. These capabilities are already helping Tier 1 CSPs improve forecast accuracy by 15–20%, reduce manual IT processes by 30%, and boost network resource efficiency by 25%.

Microsoft's Agentic AI strategy positions it as a key enabler of the telecom industry's journey toward autonomous networks, AI-native operations, and customer-centric innovation.

INTRODUCTION

Agentic AI performs specific tasks mostly independent of any human intervention. It is the next frontier of innovation in the AI race which allows agents to dynamically adapt to new situations. The promise of Agentic AI is its ability to manage complex, multi-step tasks by interacting with structured and unstructured data sets. The agent evolves via continuous learning, and in some cases

taking action in sequential order, much like a human being performing complex tasks. The promise of agentic AI is that many specialty agents can coordinate activities to complete highly complex tasks more efficiently than human experts.

The implementation of an AI agent begins with defining its role or task. For example, an agent can be set up to assist with software development. Before coding, this involves establishing the agent's responsibilities, such as defining application requirements, generating user stories, and creating a development plan. Once these foundational steps are in place, the agent can proceed with writing the code, followed by generating automated tests to ensure functionality. Additionally, another agent can be utilized to conduct security testing, identifying potential software vulnerabilities. The strength of agentic AI lies in its ability to seamlessly integrate multiple agents, accelerating task completion and scaling operations to industrialize routine tasks efficiently.

In the telecommunications market, AI agents will transform workflow tasks in all areas of business operations. AI agents will enable the personalization of service offerings, predict customer churn, optimize network performance, and identify the root causes of sophisticated infrastructure attacks. Additionally, AI agents will forecast network bandwidth demand, diagnose cell site quality issues, anticipate network outages, and enhance power efficiency in high-cost areas like the radio access network.

A key aspect of agentic AI is its ability to interact with a variety of data sources in real time. In Operations, timely and accurate data access empowers service providers to proactively detect network and security issues before they escalate into trouble tickets. In this context, the mean time to identify (MTTI) issues becomes a critical metric, often more significant than the mean time to resolve (MTTR), especially when AI is leveraged alongside relevant datasets. This proactive strategy reduces disruptions to business operations and ensures a seamless experience for end-users. The true power of agentic AI emerges when it can sift through enormous datasets, identifying patterns and extracting knowledge that human analysis might overlook. Agentic AI possesses the capability to autonomously achieve complex objectives with minimal human intervention. It dynamically adapts to evolving conditions and seamlessly integrates into diverse environments, optimizing productivity and efficiency.

Microsoft's strategy for agentic AI focuses on developing autonomous AI agents that can perform tasks and make decisions on behalf of users, aiming to enhance productivity and streamline business operations. This approach involves integrating these agents across various Microsoft products and services.

Microsoft Dynamics 365 is increasingly embedding agentic AI—AI that can act autonomously on behalf of users—across its suite of business applications. The goal is to streamline operations, reduce manual effort, and support decision-making across sales, service, finance, supply chain, and other functions.

Microsoft Copilot Studio is a low-code development platform that enables organizations to create, customize, and deploy AI-powered agents, known as "copilots," tailored to specific business needs. These agents can automate tasks, engage in natural language conversations, and integrate seamlessly with various Microsoft products and services. For additional info on this Microsoft solution see [MWC 2025 Microsoft Agentic AI](#).

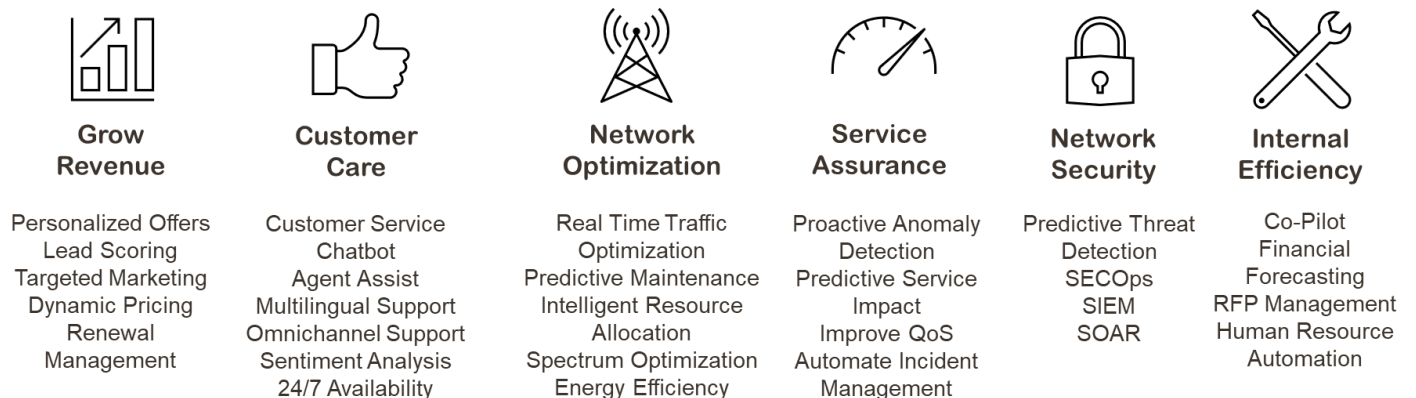
INDUSTRY OVERVIEW

Agentic AI has the potential to revolutionize every aspect of the business value chain for Communication Service Providers (CSPs). With access to vast network data and deep customer insights, CSPs are uniquely positioned to harness Agentic AI to drive innovation. By leveraging this technology alongside their proprietary datasets, CSPs can transform key areas such as monetization, customer care, network operations, and security.

In the telecom industry, Agentic AI's ability to convert natural language into SQL, REST, SOAP, and TM Forum Open APIs to execute complex queries unlocks new opportunities for efficiency and innovation. This capability enables users to interact with data effortlessly, translating natural language inputs into BSS and OSS commands, retrieving relevant information from databases, and delivering insights in an intuitive format. By integrating this functionality, telecom companies can empower teams to analyze data more effectively, make data-driven decisions, and enhance overall operational agility.

In our research, we have identified six major areas in which Agentic AI is being exploited in early pilot deployments.

Figure 2: Major segments of Telecom Agentic AI Investments



Source: Appledore Research

Microsoft is partnering with communications service providers (CSPs) to develop innovative, telecom-specific solutions tailored to their customers' needs. Two notable examples of this collaboration are AT&T and Lumen.

AT&T faced challenges with its customer-facing chatbot on att.com, which often struggled to provide relevant and accurate responses to user queries. To address this issue, AT&T partnered with Microsoft to implement a generative AI-powered solution known as *Ask AT&T*. This tool leverages retrieval-augmented generation (RAG) in combination with ChatGPT to deliver contextual and accurate answers. The implementation integrated *Ask AT&T* with AT&T's proprietary data to ensure secure and personalized interactions. Hosted within a dedicated Azure tenant, the deployment was designed to prevent data leakage and ensure compliance with internal security standards. As a result, AT&T achieved a 50% reduction in unanswered customer queries during the solution's initial launch, significantly improving the customer support experience.

Lumen, undergoing a transformation into a global communications provider, recognized the need to modernize its sales workflows to support its 3,000+ strong sales team. To drive efficiency and reduce manual workload, Lumen adopted Microsoft 365 Copilot, deploying both Copilot for Sales and GitHub Copilot across the organization. This implementation automated critical sales processes such as CRM updates, meeting summaries, and reporting, while also accelerating the onboarding of new team members. A notable improvement included reducing the time to analyze complex documents like 10-K reports from two days to just one hour. On average, each salesperson saved approximately four hours per week, which translated into over \$50 million in annual savings. The initiative not only improved productivity but also enhanced the work-life balance for Lumen's sales professionals.

KEY AREAS OF APPLICATION/USE CASES

Billing and Monetization

The concept-to-cash process in most CSPs remains largely manual, leading to inefficiencies, delays, and revenue leakage. The primary challenge extends beyond solution deployment to ensuring seamless data access across multiple silos, establishing efficient workflows, retraining employees, and redesigning customer engagement channels, all of which require significant time and resources. Consequently, many CSPs remain cautious about the short- and medium-term feasibility of concept-to-cash solutions.

However, CSPs recognize the need to modernize both consumer and enterprise digital engagement stacks. They are increasingly transitioning from fragmented legacy architectures with multiple manual touchpoints to automated, cloud-based, and unified frameworks. Productized, configurable solutions are gaining traction over highly customized offerings that require extensive integration and ongoing maintenance. A clear use case in this domain is order fallout detection analysis.

In sales, Agentic AI is emerging as a powerful tool for generating personalized outreach emails, sales proposals, and customer-specific content. By enhancing the relevance and effectiveness of sales interactions, AI streamlines the sales cycle and improves engagement. Lumen is actively implementing this approach today, as highlighted in the use cases referenced later in the document.

The future of Business Support Systems (BSS) lies in AI-native architectures, where AI is embedded at the core to manage complex business scenarios, optimize operational efficiencies, and unlock new revenue streams.

Customer Care

Consumer billing inquiries still account for 30% of all calls into the contact center every month. This is due to variability in usage and service plans including roaming charges.

Customer service remains a priority for CSPs, but there is a strong focus on leveraging technology to minimize support costs. AI-driven enhancements to self-service capabilities will be a key area of interest and investment in the coming years. Additionally, CSPs will invest in solutions aimed at improving call center efficiency, including advanced AI-based call deflection strategies, enhanced contextual information, and real-time recommendations for agents. These advancements will help reduce average call handling times, increase sales, and optimize overall customer support operations.

Agentic AI can be thought of as a virtual assistant for customer service agents, providing real-time recommendations, summarizing case details, and offering guidance on best responses. This improves agent productivity and ensures consistent quality across interactions. Agentic AI can also be used for agent training and quality assurance by scoring interactions and providing feedback.

Agentic AI can identify gaps in existing knowledge bases and automatically generate new recommendations based on successful past resolutions. This keeps knowledge bases up-to-date and comprehensive, benefiting both self-service and agent-assisted support.

Service Assurance

As 5G networks transition to Standalone (SA) to enable network slicing and ultra-low latency, real-time network status updates become critical for maintaining high-quality service. Instead of relying on periodic reports, CSPs can leverage Agentic AI to detect subscriber impacts and identify anomalies in real time. One particularly promising area of innovation is analysing how different devices perform under varying traffic conditions, allowing for more precise network optimizations.

The escalating complexity of modern network architectures has led to prolonged incident investigation and recovery times. Telecom carriers now expect their network engineers to manage, optimize, monitor, forecast, and troubleshoot multi-layered, multi-technology, and multi-vendor networks, all while meeting increasing demands for higher reliability and expedited issue resolution. As operational complexities continue to rise, the ability to transform vast amounts of network traffic data into actionable intelligence using AI agents has become a critical priority for Communication Service Providers (CSPs). This urgency is underscored by insights from numerous Tier 1 CSPs aiming to advance their operations beyond Level 2 autonomy towards more advanced stages. According to the TM Forum's six-level Autonomous Network Maturity Model, as many as 83% of CSPs are currently at Level 1 or 2 autonomy, highlighting the significant journey ahead.

Network Optimization

Network planning is pivotal in delivering efficient, reliable, and scalable connectivity services while adhering to regulatory mandates such as coverage obligations and service quality standards. This strategic process encompasses decisions on infrastructure deployment, resource allocation, and capacity management to meet current and future demands effectively. The integration of Agentic AI

into network planning introduces a transformative approach, enabling real-time adaptation to fluctuating traffic patterns and preventing potential capacity issues. By autonomously analyzing vast datasets, Agentic AI can optimize network configurations, forecast demand surges, and ensure compliance with regulatory requirements, thereby enhancing overall network performance and customer satisfaction.

Effective network planning impacts several key areas:

- **Capacity Management:** Ensures networks can handle peak demands without sacrificing service quality, by optimizing resources to meet user needs for data and bandwidth-intensive services.
- **Quality of Service (QoS):** Maintains high service quality for customer satisfaction and retention, avoiding bottlenecks and enhancing network reliability.
- **Cost Efficiency:** Optimizes infrastructure spending and operational costs, preventing over-provisioning and ensuring cost-effective network expansions.

Network energy efficiency is a critical priority for telecom operators, particularly in wireless infrastructure, which accounts for approximately 45% of total power consumption. A key challenge lies in the inability of networks to dynamically adjust power usage in response to fluctuating traffic volumes, leading to unnecessary energy expenditure during off-peak periods.

To tackle this, efficient energy management strategies such as "zero bits, zero watts" are being implemented. These approaches optimize power consumption based on real-time traffic demands across business districts, residential areas, critical infrastructure, and suburban regions.

Traditional energy-saving methods rely on manual data analysis and broad, predefined rules, often resulting in suboptimal efficiency or service degradation during peak usage. Furthermore, the widespread adoption of unlimited data plans has introduced new complexities, making traffic patterns increasingly unpredictable.

Agentic AI presents a transformative solution by leveraging advanced pattern recognition and predictive analytics to dynamically optimize energy consumption, ensuring both sustainability and network performance without compromising service quality.

Network Security

The cybersecurity landscape is facing unprecedented challenges, with over 100 million malware variants and thousands more emerging daily. The expansion of mobile devices, cloud storage, and IoT technologies has significantly widened attack surfaces, making traditional security measures inadequate. Conventional signature-based threat detection is struggling to keep pace with the sheer volume and complexity of emerging threats, highlighting the urgent need for autonomous, AI-driven cybersecurity solutions.

Agentic AI is redefining cybersecurity by enabling real-time, adaptive threat detection and response. Unlike traditional approaches that rely on predefined signatures, Agentic AI autonomously detects, predicts, and mitigates zero-day vulnerabilities and advanced persistent threats (APTs) by continuously analyzing patterns, anomalies, and evolving attack behaviors.

In today's networks, where ransomware, APTs, and unknown malware strains are escalating, AI-driven security systems can autonomously learn from new attack patterns, proactively adapt defences, and minimize response times, drastically reducing the impact of cyberattacks. By integrating Agentic AI into cybersecurity frameworks, organizations can shift from reactive to proactive responses, intelligent threat mitigation, ensuring stronger, more resilient protection against modern cyber threats.

CASE STUDIES

Enhancing AT&T's Customer Support with Generative AI

Challenge

AT&T aimed to improve the performance of its customer-facing chatbot on att.com, which often struggled to provide relevant responses.

Solution

To address this, AT&T collaborated with Microsoft to integrate retrieval-augmented generation (RAG) and boost its capabilities with generative AI. The tool is branded as *Ask AT&T*.

Implementation & Impact

- **Generative AI Integration:** The first iteration of *Ask AT&T* harnesses OpenAI's ChatGPT functionality while being designed for seamless integration with other emerging generative AI tools.
- **Secure Deployment:** To ensure security and data integrity, AT&T partnered with Microsoft to deploy *Ask AT&T* within a dedicated Azure tenant, rigorously tested to prevent data leakage. This allows AT&T employees to leverage company data confidently, knowing it remains protected and inaccessible to the public domain.
- **Customization:** What differentiates *Ask AT&T* from generic generative AI models is its deep integration of AT&T's proprietary knowledge and processes, ensuring responses are accurate, reliable, and responsibly generated.
- **Time and Cost Impact:** The impact was significant—AT&T reduced the number of unanswered customer queries by 50% in its initial commercial deployment.

Lumen's Implementation of Microsoft Copilot for Sales

Challenge

As Lumen transitioned from a traditional telecom company to a global communications provider, its 3,000+ sales team needed new tools to streamline workflows, improve efficiency, and adapt to evolving sales strategies. The company sought a solution that would reduce manual administrative tasks, provide deeper customer insights, and accelerate sales processes.

Solution

In 2023, Lumen strategically implemented **Microsoft 365 Copilot** through a phased rollout. The company designated internal “champions” to drive adoption, conducted informational sessions via Teams, and shared best practices to ensure a smooth transition. Following this, **Copilot for Sales** and **GitHub Copilot** were introduced, rapidly expanding the use of generative AI across the organization.

Implementation & Impact

- **Enhanced Sales Efficiency**

Lumen’s sales team leveraged **Copilot for Sales** throughout the sales process, from researching customer needs to summarizing meetings and crafting follow-up communications. A key improvement was the automation of syncing Outlook emails with customer relationship management (CRM) software, which eliminates a previous manual process.

“The company can capture a lot of the information that traditionally would stay in someone’s Outlook inbox,” says **Delvin Holman, Senior Lead Customer Service Enablement Manager**.

Additionally, Copilot’s AI-powered summarization significantly improved data analysis. For example, summarizing 10-K reports—a crucial but time-intensive tasks were reduced from **two days to just one hour**.

- **Accelerated Onboarding & Knowledge Transfer**

Copilot proved invaluable for onboarding new team members. **Kerrie Davis, Senior Director of Revenue Enablement Execution**, recalls a customer success manager who used Copilot to analyze 18 months of emails and data, organizing them into an Excel-based timeline: “It organized an entire account at a very pivotal customer moment in record time,” Davis explains. “It put together days and days of work that she said was probably not humanly possible otherwise.”

- **Time & Cost Savings**

On average, **Copilot is saving Lumen’s sales team four hours per week per seller**, translating to an estimated **\$50 million in annual savings**. This efficiency increase allows sales professionals to spend more time with customers while improving their work-life balance.

Vodafone Transforms Customer Engagement and Operations with Microsoft Generative AI

Challenge

Vodafone, one of the world’s largest telecommunications providers with over 330 million customers across 15 countries, faced growing pressure to deliver faster, more personalized customer service across digital and human-assisted channels. The company needed to address rising customer expectations, reduce response times, and empower call center agents to handle increasingly complex inquiries—while also fostering a more innovative, inclusive work environment.

Vodafone recognized that its customer service tools, while effective for routine inquiries, lacked the intelligence and flexibility to handle more nuanced interactions. Additionally, internal employees needed better resources to develop real-time solutions for operational challenges and improve customer satisfaction without increasing workloads or costs.

Solution & Implementation

To address these challenges, Vodafone signed a multi-year deal with Microsoft to implement the following Microsoft Azure AI tools:

- **Azure OpenAI Service:** To enhance the intelligence of its existing chatbot, TOBi, for more context-aware, personalized interactions.
- **Microsoft Azure AI Foundry:** A development platform that enabled internal experimentation and deployment of AI solutions.
- **Microsoft Copilot and Azure AI Search:** To support a new internal virtual assistant, "SuperAgent," used by live customer service agents for resolving complex queries.

The solution was piloted in Vodafone's Corso, Italy call center. Two core applications were deployed:

1. **TOBi (Enhanced Chatbot)** – TOBi was upgraded with Azure OpenAI Service to handle natural, dynamic conversations for common inquiries such as billing, network issues, and device orders. If TOBi couldn't resolve the issue, it escalated the interaction—along with a conversation summary—to a human agent.
2. **SuperAgent (Internal Assistant)** – This tool supported call center agents by surfacing relevant content from Vodafone's knowledge base in real-time, eliminating the need to manually search through long documents. SuperAgent also conducted compliance checks on call transcripts and provided feedback for ongoing training and quality assurance.

To drive innovation internally, Vodafone also began training senior operators at the Corso center to build AI-powered tools with Azure AI Foundry, further democratizing access to cutting-edge technology.

Impact

The introduction of generative AI resulted in measurable improvements in both operational efficiency and customer satisfaction:

- **Volume & Efficiency:** TOBi now handles approximately 45 million customer interactions per month, resolving 70% of inquiries within digital channels without human intervention.
- **Support Enhancement:** The remaining 30% of calls are assisted by SuperAgent, which helps human agents quickly find precise answers, reducing average call handling time by more than one minute.
- **Customer Satisfaction:** Vodafone saw an increase in customer satisfaction scores due to faster, more accurate, and more personalized support experiences.

- **Employee Enablement:** By using Azure AI Foundry, employees—particularly at the Corso call center—are now co-creators of their own AI solutions, supporting Vodafone’s goals of inclusion and innovation.
- **Scalability:** Encouraged by the success at Corso, Vodafone is planning a broader rollout of both TOBi and SuperAgent to other call centers across its footprint.

INNOVATION APPLIED TO NEW FRONTIERS OF AGENTIC AI IN TELECOM

Emerging applications of agentic AI are unlocking new opportunities to optimize telecom networks and improve operational efficiency. By integrating AI-powered tools and assistants, these solutions are empowering network engineers and field agents, enhancing both productivity and precision in day-to-day operations.

For network engineers, AI agents streamline complex tasks such as Methods of Procedures (MOPs) for configuration changes, deployment rollouts, and outage recoveries. Leveraging CSP-specific network data and resources, the agent enables engineers to locate intricate configurations, extract and summarize critical information, and retrieve past configurations with ease. This significantly reduces MOP cycle times, improves accuracy, and minimizes the risk of misconfigurations—supporting the CSP’s "Zero Outage" objective.

In the field, AI assistants enhance operational agility by integrating real-time data sources and APIs to deliver incident updates, provide quick access to technical procedures, and offer guided troubleshooting steps. This real-time, context-aware support allows field agents to resolve issues faster and reduces dependence on the Network Operations Center, creating a more autonomous and efficient field workforce.

Azure OpenAI Service, a specialized RAN (Radio Access Network) agent, supports network teams with instant access to incident history, resolution workflows, regulatory documentation, and key performance insights. Automating these routine yet critical tasks improves overall operational efficiency, shortens mean time to repair (MTTR), and streamline engineering workflows. This aligns with the CSP’s broader vision of transitioning to a Smart Network—one capable of self-configuration, self-diagnosis, and self-correction.

For more information on products making their way into the market we highlight some MWC25 announcements below:

1. [AI & Gen AI: Key to Level 5 Autonomous Networks | Nokia](#)
2. [MWC 2025 AI-driven Operations](#) (Nokia)
3. [MWC 2025 Intelligent NOC](#) (Accenture)

Applying Agentic AI in Network Traffic Optimization

Communication Service Providers (CSPs) have long leveraged automation to manage and reroute network traffic. However, Agentic AI is taking this capability to the next level by enabling real-time, intelligent traffic optimization. By continuously analyzing traffic patterns over time, AI-driven

networks empower Network Service Providers (NSPs) with the insights needed to optimize routing, capacity management, and resource allocation more dynamically than ever before.

With Agentic AI, networks can autonomously detect and respond to unexpected traffic spikes, proactively allocating temporary capacity to prevent congestion and maintain seamless user experiences. This ensures that CSPs can not only meet but exceed customer expectations for low latency, high reliability, and uninterrupted service quality.

Beyond peak-time management, AI-driven networks can intelligently regulate network components during off-peak hours. In mobile networks, for instance, Radio Access Networks (RANs) can be dynamically adjusted based on real-time demand—entering low-power mode or temporarily shutting down when traffic is minimal. This capability optimizes energy efficiency, reducing operational costs while enhancing 5G network sustainability.

By integrating Agentic AI into network operations, CSPs can achieve truly autonomous, self-optimizing networks, ensuring superior service delivery, reduced energy consumption, and unparalleled network resilience.

Microsoft is also supporting CSPs in adopting AI agents across key operational domains, including within the TM Forum ODA Canvas. By leveraging Azure Kubernetes Service to host ODA components, CSPs can deploy their own carrier-grade operational Canvases using commercially supported, pre-built modules—accelerating automation and agility. Vodafone, a founding member of TM Forum's Innovation Hub and an early leader in ODA adoption, is already operating an ODA Canvas in a live production environment in Greece.

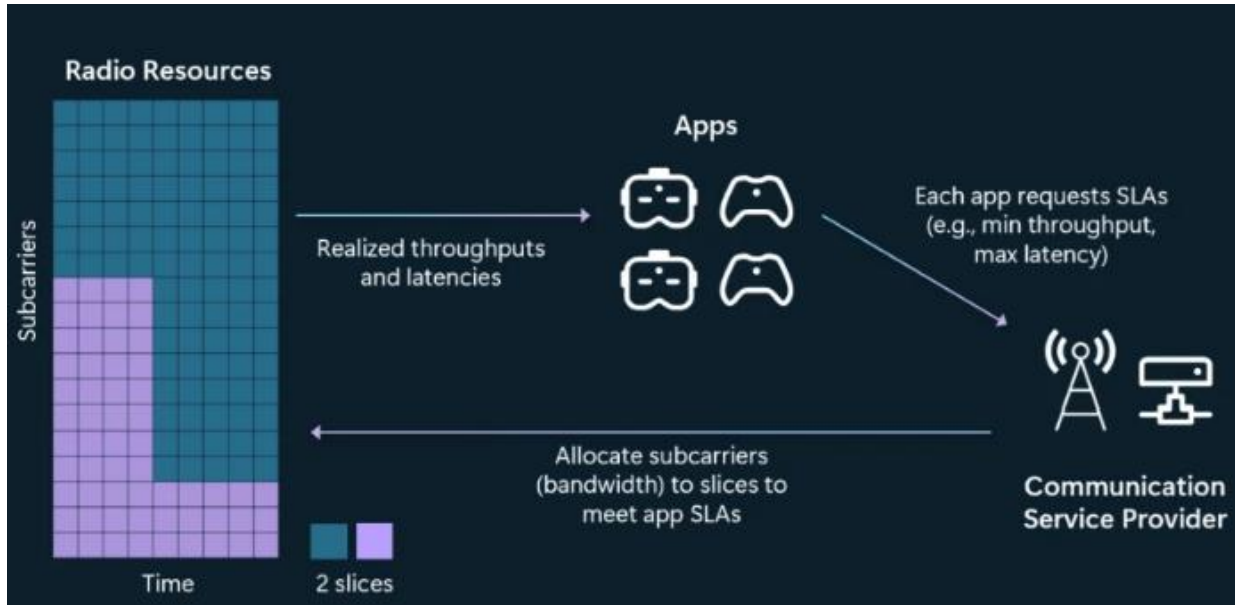
Service Assurance with RAN Slicing

Latency-sensitive mobile applications—such as Xbox Cloud Gaming, Microsoft Teams, Mixed Reality, telemedicine, and cloud robotics—require consistent and predictable network performance, particularly in terms of throughput and latency. Recognizing this need, 3GPP introduced network slicing, which allows operators to create multiple virtualized “slices” atop a single physical network, each tailored to different connectivity requirements. Radio Access Network (RAN) slicing is especially critical because the wireless “last mile” is often the primary bottleneck. However, traditional base station schedulers typically optimize for overall throughput or bundled application performance, rather than individual application needs, leading to suboptimal outcomes. While existing systems can assign physical resource blocks (PRBs) to network slices to meet high-level service-level agreements (SLAs), these do not guarantee performance for every application within a slice. Diverse channel conditions, asynchronous app connections, and dynamic traffic patterns make it difficult to provide reliable app-level performance using conventional methods.

To address this gap, Microsoft has developed a production-class, O-RAN-compatible RAN slicing system integrated into its end-to-end 5G platform. The RAN slicing system is part of Microsoft's **Project Janus**, a programmable RAN platform designed to enhance Open RAN networks through advanced programmability. This system enables app-level service assurance by dynamically bundling apps with similar SLA needs into slices and allocating PRBs based on model predictive control (MPC). MPC separates the network model from the control policy, allowing it to predict environmental variables (like wireless channel conditions) and compute optimal bandwidth

allocations in real time. To handle the state-space complexity, Microsoft's solution leverages deep neural networks trained on simulations to forecast PRB requirements and resource availability. The control algorithm can efficiently prune the allocation search space and adapt bandwidth allocations based on predicted signal quality, allocating more PRBs under poor channel conditions to mitigate packet loss.

Figure 3: 5G Slicing in RAN Network Illustrative Diagram



Source: Microsoft

Operators using Microsoft’s system can define slices for distinct services—such as Microsoft Teams or Xbox Cloud Gaming—and meet each app’s SLA for latency and throughput far more accurately than with slice-level guarantees alone. In testing, Microsoft’s approach significantly reduced SLA violations, measured as deviations from an app’s requested performance. Additionally, the system provides primitives that help admission controllers assess real-time capacity without interfering with the operator’s monetization or policy decisions. Overall, this implementation offers a scalable, intelligent, and field-ready framework for delivering predictable, app-aware network performance in modern 5G environments. For more details on this project see <https://www.microsoft.com/en-us/research/publication/application-level-service-assurance-with-5g-ran-slicing/>

Microsoft Expands Security Copilot with AI-Powered Autonomous Agents

Microsoft has significantly advanced its cybersecurity portfolio with the introduction of autonomous agents in *Security Copilot*, a generative AI solution designed to enhance the capabilities of security and IT teams. Originally launched to help defenders detect threats faster and respond more effectively, Security Copilot has shown measurable impact, with users reporting up to 30% reductions in mean time to resolution.

As cyber threats become increasingly complex and adversaries adopt AI, Microsoft is responding by evolving Security Copilot into a suite of autonomous agents. These agents integrate directly with Microsoft Security and partner solutions, enabling automated handling of high-volume security and IT operations tasks. Built on Microsoft's Zero Trust architecture, they continuously learn from feedback, adapt to workflows, and deliver enterprise-grade security with minimal manual intervention.

To operationalize this vision, Microsoft is addressing the market opportunity with agents tailored to a specific challenge:

- **Phishing Triage Agent:** Embedded in Microsoft Defender, it autonomously analyzes phishing alerts using multi-modal AI and learns from analyst feedback, reducing triage time and false positives.
- **Alert Triage Agents for DLP and IRM:** Found in Microsoft Purview, these agents prioritize data-related alerts based on risk and policy parameters, with natural language explanations and adaptive learning.
- **Conditional Access Optimization Agent:** Integrated with Microsoft Entra, it detects and resolves access policy gaps by monitoring new users and applications, streamlining updates through one-click recommendations.
- **Vulnerability Remediation Agent:** Available in Microsoft Intune, it uses Defender Vulnerability Management to assess and prioritize CVEs, providing guided remediation to minimize risk exposure.
- **Threat Intelligence Briefing Agent:** Leveraging Microsoft's threat intelligence ecosystem, this agent produces executive-ready briefings in minutes, tailored to the organization's threat landscape and attack surface.

MICROSOFT ISV ECOSYSTEM AND AI TRANSFORMATION

The AI supplier ecosystem is as varied as the numerous use cases it supports, encompassing both established industry leaders and specialized providers. Major players such as Amdocs, Nokia, Ericsson, CSG, ServiceNow, and Ciena, alongside niche suppliers like Mavenir, Cohere Technologies, Norwood, and Neuron7.ai, actively contribute to the market. Each of these companies leverages their domain expertise and proprietary high-value datasets to apply pattern recognition techniques and deliver contextual insights into complex problem areas.

Microsoft has strategically partnered with many of these suppliers to tap into their specialized knowledge across various facets of telecommunications operations. Its partnership approach in the telecommunications sector for Agentic AI is centered around several key priorities:

- **Enhancing Customer Experience:** By integrating AI-powered support systems, Microsoft helps telecom operators elevate customer interactions. For example, Vodafone has adopted Microsoft's Azure AI services to enhance its digital assistant, resulting in more personalized and efficient customer support.
- **Optimizing Business and Operations Support Systems:** Microsoft collaborates with partners such as Amdocs to modernize telecom operations. Through the integration of Amdocs' amAlz platform with Microsoft's Azure OpenAI Service, these partnerships streamline critical

processes such as billing and network management, driving efficiency and reducing operational costs.

- **Modernizing Network Infrastructure:** As the demand for AI workloads continues to grow, Microsoft has partnered with Lumen Technologies to expand data center capacity. This collaboration improves network performance and scalability, ensuring seamless service delivery for telecom providers.
- **Unlocking New Revenue Streams:** Leveraging AI and generative AI, Microsoft enables telecom operators to develop innovative services that extend beyond traditional offerings. This includes creating value-added services, strengthening B2B solutions, and increasing customer lifetime value.

Nokia NetGuard solution

Nokia is leveraging NetGuard Cybersecurity Dome software with Microsoft Azure OpenAI Service, integrating a telecom-focused GenAI assistant to enhance the ability of communication service providers (CSPs) and enterprises to detect and respond to cyber threats more swiftly and effectively.

The Telco GenAI assistant enhances Cybersecurity Dome's efficiency by rapidly analyzing vast amounts of cyber threat data. By incorporating large language models within Microsoft Azure OpenAI Service, the assistant benefits from extensive training on 5G network architecture, security practices, and Nokia's telecom expertise. This AI-driven approach strengthens threat identification and resolution, allowing CSPs to respond to attacks with greater accuracy and speed.

NetGuard Cybersecurity Dome serves as Nokia's extended detection and response (XDR) security platform, providing advanced 5G network defenses through AI and machine learning. With cyberattacks increasing in sophistication and frequency, the XDR security market is projected to grow by nearly 150%, from \$965 million in 2022 to approximately \$2.4 billion by 2027.

The GenAI assistant is equipped with a comprehensive knowledge base, including 3GPP and NIST network architecture specifications and 5G topology spanning RAN, Transport, and Core. Nokia anticipates that this AI-driven assistant will reduce threat identification and resolution times by up to 50%, depending on the nature of the attack. Additionally, it is expected to decrease false-positive alerts, improving incident handling precision and overall cybersecurity efficiency.

Microsoft is partnering with Nokia to enhance automated workflow creation by leveraging advanced tools and AI models. Through Microsoft Azure AI Foundry, integrated with Azure OpenAI Service's large language models (LLMs), Nokia can capture initial requirements and efficiently convert them into structured workflow templates or fully executable workflows. For more information see [Autonomous Network](#)

Norwood Systems

Norwood Systems a leader in AI-powered telecom voice technologies, has announced the launch of CogVoice OpenSpan™, an advanced AI-driven voice platform built on Microsoft Azure. OpenSpan enables Communication Service Providers (CSPs) to modernize and monetize their voice services by integrating AI-powered features and real-time intelligence.

Traditional CSP infrastructures are designed for reliability and scalability but often lack native AI integration, making it difficult to support emerging AI-powered applications. At the same time, AI-driven voice services frequently require seamless core network integration to function efficiently at scale. OpenSpan bridges this gap by leveraging Azure’s AI capabilities and Norwood’s expertise in real-time cloud-based media processing, allowing CSPs to deploy AI-powered voice services such as fraud detection, live translation, and automated call summarization.

OpenSpan offers several key features that enhance its value for CSPs. AI-powered intelligence, integrated with Azure AI Speech and Azure OpenAI Service, enables real-time in-call AI applications that enrich voice interactions. Seamless deployment is facilitated through its modular architecture, ensuring low-friction integration with 4G/5G SIP-IMS cores and Diameter signaling environments, minimizing complexity and implementation time. Additionally, flexible monetization options allow CSPs to generate revenue from both legacy and AI-driven services, supported by an open, service-based architecture with customizable pricing models. Designed for carrier-grade scalability, OpenSpan handles high concurrent call volumes while integrating effortlessly with Azure-based services.

Figure 4 highlights a selection of Microsoft partners that offer domain expertise and commercial solutions tailored for telecommunications providers.

Figure 4: Microsoft Telco AI Supplier Partner Ecosystem



Source: Microsoft

VALUE TO THE TELECOM C-SUITE

Microsoft's Agentic AI is increasingly delivering measurable value to telecom C-suite leaders by enabling **goal-driven, autonomous intelligence** across strategic, financial, technological, and operational domains.

For **CEOs**, Agentic AI enhances strategic responsiveness by continuously ingesting operational, customer, and market intelligence to surface actionable insights. In a recent TM Forum Catalyst project involving Microsoft, CSPs used generative AI agents to simulate new service offerings and received scenario-based projections on revenue, cost, and market readiness in minutes rather than weeks. This kind of AI-driven simulation supports faster go-to-market strategies and empowers CEOs to pivot based on real-time market dynamics. In parallel, Agentic AI supports ESG commitments by tracking energy efficiency across networks and modeling the impact of sustainability initiatives—capabilities already being piloted in Vodafone and Telefónica operations through Microsoft-Azure-based analytics.

For **CFOs**, Agentic AI offers autonomous forecasting and anomaly detection capabilities that enhance financial oversight. In recent telecom deployments, AI agents have been used to automate rolling forecasts, detect early deviations in CAPEX and OPEX, and simulate financial impacts of network rollout delays or service disruptions. One Tier 1 European CSP, using Microsoft's AI and Power BI Copilot tools, reported a **15–20% improvement in forecast accuracy** and significantly reduced the cycle time for monthly financial close. AI agents also monitor cloud spend and network utilization patterns, helping CFOs identify savings opportunities—particularly important as CSPs virtualize infrastructure and move toward cloud-native operations.

CIOs benefit from Agentic AI through automation of IT operations and governance of digital systems. In Microsoft-supported TM Forum Catalyst pilots, AI agents have been deployed to streamline IT service management (ITSM), reduce ticket resolution times, and orchestrate identity and access provisioning across hybrid cloud environments. One telecom operator using Microsoft Entra ID with Agentic AI saw a **30% reduction in manual access provisioning tasks** and improved security compliance through continuous access reviews. CIOs also gain transparency into AI usage across the enterprise through governance dashboards, aligning with responsible AI standards, especially relevant in regions where data sovereignty and AI regulation are tightening.

For **CTOs**, particularly in telecom, the impact of Agentic AI is most visible in **network operations and service assurance**. AI agents integrated into Azure and TM Forum's ODA Canvas architecture are enabling predictive maintenance, closed-loop automation, and autonomous network configuration. For example, Vodafone's early deployment of an ODA Canvas in Greece, with Azure Kubernetes and Agentic AI components, demonstrated up to a **10x improvement in operational agility**, allowing faster onboarding of services and automated responses to network events. In another deployment, AI agents were used to manage and optimize 5G network slicing, resulting in a **25% increase in resource utilization efficiency** during high-demand events.

CONCLUSION

Microsoft's **Agentic AI strategy** is fundamentally reshaping how businesses operate by deploying **autonomous AI agents** that automate routine tasks, enhance workforce productivity, and unlock innovation across industries. This transformative approach enables organizations to act with greater agility and efficiency, streamlining operations while simultaneously opening new pathways for growth.

At the heart of Microsoft's strategy are several key pillars. First, **enhanced productivity and operational efficiency** are achieved by automating repetitive and time-consuming activities, allowing employees to focus on strategic, high-value work. This shift not only improves output but also boosts employee engagement and job satisfaction. Second, Microsoft delivers **tailored AI solutions** via Copilot Studio, which enables organizations to build customized AI copilots aligned to their specific operational challenges. These AI agents integrate seamlessly with Microsoft 365 and the Power Platform, ensuring adaptability to a wide range of business functions.

Third, Microsoft offers **scalable AI agents** embedded within Microsoft 365 and Dynamics 365, purpose-built to optimize enterprise functions like sales, customer service, finance, and supply chain management. These AI-driven tools empower organizations to scale while maintaining consistency and precision in execution.

Fourth, **seamless integration with Microsoft's broader ecosystem** ensures interoperability with third-party platforms such as Salesforce and ServiceNow. Solutions like Copilot for Service and Dynamics 365 Contact Center come pre-integrated, allowing for rapid deployment with minimal disruption.

Fifth, Microsoft maintains a strong **commitment to responsible and secure AI implementation**. Its AI agents are governed by a robust responsible AI framework that ensures data privacy, compliance, and ethical use—critical for maintaining trust in autonomous decision-making.

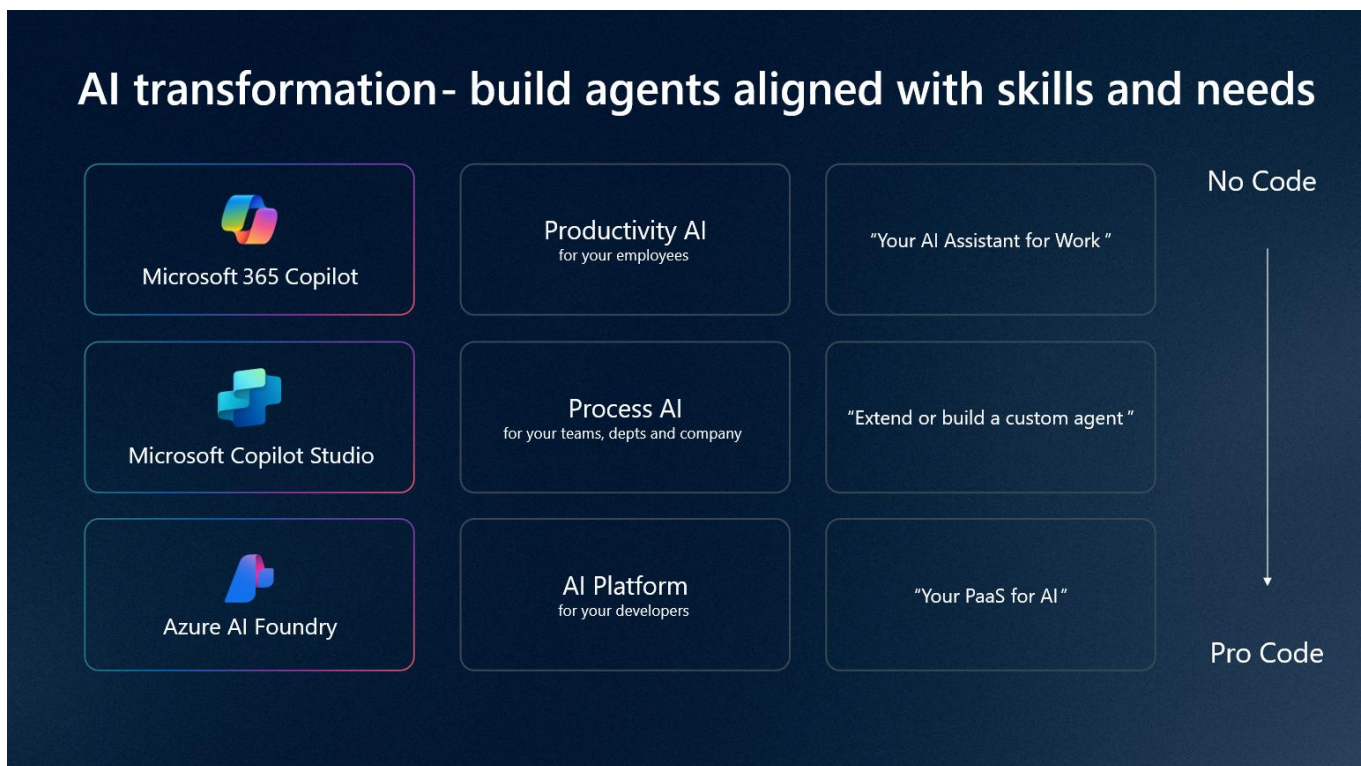
As part of its forward-looking vision, Microsoft is introducing a new wave of **autonomous AI agents** designed to help businesses transition from legacy systems to AI-first processes. Starting in late 2024 and into early 2025, ten pre-built agents will be available in public preview, representing the start of a broader AI transformation in enterprise automation. These agents will support core functions across **sales, operations, finance, and customer service**, driving measurable improvements in productivity and decision-making.

For instance, **Sales Agents** streamline lead qualification and automate order processing, enabling sales teams to focus on building relationships and closing deals faster. **Operations Agents** automate procurement, supply chain, and financial workflows to reduce manual workload and accelerate time-to-decision. **Finance Agents** handle reconciliation and close financial tasks, improving cash flow visibility and reporting accuracy. Meanwhile, **Service Agents** enhance customer experience by automating case management, dynamically updating knowledge bases, and optimizing field service schedules.

Designed for **continuous learning and self-improvement**, these AI agents adapt over time, improving performance and service quality. Microsoft also prioritizes **security and governance**, embedding privacy controls, IT oversight, and data protection measures to ensure safe adoption at scale.

The strategic vision is clear: **every employee will have a Copilot, and every business process will have its own AI agent**. Given the complexity and volume of business processes, particularly within Communication Service Providers (CSPs), Microsoft has established a robust platform ecosystem to support agent creation, deployment, and management. This includes **Microsoft 365 Copilot with Agent Builder, Copilot Studio, and Azure AI Foundry**, offering no-code, low-code, and pro-code environments respectively. Together, these platforms provide a flexible, secure, and centralized framework for designing and orchestrating the next generation of intelligent, autonomous enterprise systems.

Figure 5: Custom AI Agent Channels to Market



Source: Microsoft

Microsoft continues to advance AI-driven automation with new features in **Copilot Studio**, a platform designed to solve enterprise-scale business challenges using generative AI. Following recent enhancements—such as deep reasoning, Model Context Protocol (MCP) support, and general availability of agent flows—Microsoft has introduced a new capability called **"computer use"** in early access.

The **computer use** feature enables AI agents in Copilot Studio to interact with websites and desktop applications through their graphical user interfaces (GUIs), mimicking human actions like clicking buttons, navigating menus, and entering data. This unlocks automation for systems that lack API integration, significantly expanding the range of tasks that agents can perform.

Key highlights of the computer use capability:

- Agents can now operate across GUI-based applications and web browsers (Edge, Chrome, Firefox).
- Built-in reasoning allows agents to dynamically adapt to interface changes in real time.
- Runs on Microsoft-hosted infrastructure, minimizing overhead for enterprise IT teams.
- Complies with Microsoft's security and governance frameworks; data remains within Microsoft Cloud boundaries and is not used to train foundation models.

This advancement positions Microsoft Copilot Studio as a leader in AI-powered task automation, enabling broader use cases and improving operational efficiency for enterprise users. For further information on computer use see:

1. [Announcing new computer use in Microsoft Copilot Studio for UI automation | Microsoft Copilot Blog](#)
2. [Introducing Model Context Protocol \(MCP\) in Copilot Studio: Simplified Integration with AI Apps and Agents | Microsoft Copilot Blog](#)
3. [Introducing agent flows: Transforming automation with AI-first workflows | Microsoft Copilot Blog](#)



Insight and analysis for telecom transformation.

 @AppledoreVision

 Appledore Research

www.appledorerresearch.com

info@appledorerg.com

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA

© Appledore Research LLC 2025.