MARKET PERSPECTIVE

Understanding Twilio's Agentic Opportunity: Expanding from Communications Utility to Agent Control Plane

Arnal Dayaratna

EXECUTIVE SNAPSHOT

Twilio is poised to transform from a communications infrastructure provider to the agentic control plane, enabling enterprises to govern, orchestrate, and scale autonomous AI agents across all channels with compliance and reliability. As organizations shift toward communication-first agents, Twilio's robust APIs, identity management, and governance capabilities address critical needs such as authentication, observability, and auditability.

Key takeaways

- Twilio is strategically positioned to evolve from a communications infrastructure provider to the agentic control plane, enabling enterprises to govern, orchestrate, and scale autonomous AI agents safely and compliantly across all communication channels.
- As organizations shift from generative AI coding assistants to autonomous, communication-first agents, Twilio's robust APIs, identity management, and governance capabilities address critical needs such as agent authentication, observability, state management, and auditability, ensuring reliable and measurable business outcomes.
- Twilio's model-agnostic, neutral communications backbone allows enterprises to avoid vendor lock-in, maintain flexibility across AI models, and enforce consistent compliance and operational policies, making it an essential foundation for the next generation of enterprise AI and digital labor.

Recommended actions

- Prioritize the implementation of a centralized governance and orchestration layer for agentic systems to ensure compliance, identity verification, observability, state management, and auditability across all AI-driven customer interactions. This will mitigate regulatory risks, prevent operational chaos, and enable safe scaling of autonomous agents.
- Leverage Twilio's model-agnostic communications infrastructure as the foundational control plane for multi-agent, multi-model environments. This approach preserves enterprise flexibility, prevents vendor lock-in, and ensures

February 2026, IDC #US54252825

consistent policy enforcement and customer experience across diverse AI models and agent vendors.

- Invest in operational intelligence capabilities by aggregating and analyzing communication data across all agents and human handlers. Use key performance metrics — such as containment rates, handle times, sentiment shifts, and downstream business outcomes — to continuously optimize agent performance, refine automation-to-human support strategies, and drive measurable business impact.

NEW MARKET DEVELOPMENTS AND DYNAMICS

Introduction

This document argues that Twilio is positioned to expand from a leading provider of programmable communications infrastructure to the "agentic control plane" that provides the necessary governance, identity, and orchestration layer transforming experimental AI models into compliant enterprise employees, natively integrating these new capabilities with pervasive, bidirectional communications channels. This expansion is driven by a fundamental shift in software development as organizations move away from generative AI coding assistants that help developers write code to agentic developer technologies where autonomous agents execute tasks independently. As enterprises race to deploy this digital workforce, they face a fragmentation crisis where these agents lack a standardized way to interact with the real world safely. By establishing itself as the governance and orchestration backbone for agentic systems, which communicate over diverse channels with broad audiences, Twilio has the potential to enable enterprises to scale agent deployments with confidence, compliance, and measurable business impact across customer interactions and critical operations.

The developer pivot: From "connecting humans" to "managing fleets"

Twilio is well known for solving the problem of converting traditional telephony, email, and messaging into flexible programmable building blocks. By translating the complexity of global telecom networks into APIs, Twilio empowers developers to treat calls and texts as native software primitives. As a result, developers no longer need to operate their own telecom infrastructure or negotiate directly with carriers. This abstraction was transformative enough to create an entirely new category, namely, communications platform as a service (CPaaS), that positioned Twilio as the dominant player at the intersection where software meets end users.

While generative AI coding tools like Claude Code, GitHub Copilot, and Amazon Q Developer have amplified developer productivity, the industry now faces a more fundamental shift as organizations transition toward agentic developer technologies that execute tasks autonomously without human oversight. Agentic systems differ from code generation tools in a critical way by autonomously executing tasks in the real world

without human intervention or oversight, operating with genuine independence and real-world consequences. Communication-first AI agents are exemplary of this new era by scheduling appointments, confirming deliveries, verifying identities, or managing collections through seamless voice, messaging, and email conversations. For enterprises deploying these agents, the quality of the underlying communications substrate becomes a direct determinant of customer experience and business success.

For developers building such systems, the infrastructure requirements have shifted dramatically. When a developer deploys a support agent authorized to issue refunds or schedule technician visits, the risk profile changes immediately. Developers need far more than an API to send text messages. They require a runtime environment that guarantees four critical properties as follows:

- **Identity:** Developers must verify whether the agent making a call is legitimate or whether a rogue instance or competing service is spoofing their phone number.
- **Observability:** Developers need assurance that the agent stayed within its sentiment guardrails, is not hallucinating, and did not violate compliance constraints during the interaction.
- **State management:** If the agent crashes mid-call or needs human intervention, developers must know whether the customer gets seamlessly handed off with full context preserved, or whether they are starting from scratch.
- **Auditability:** Developers must be able to prove to a regulator or customer exactly what happened in that call and verify their agent behaved correctly.

To guarantee that agentic interactions adhere to these properties, developers must orchestrate the individual workflows of their agents and supporting ecosystem. This orchestration is low level, and removed from the business intent that a developer may be trying to accomplish. This only becomes more challenging as agents scale in complexity.

Twilio's APIs today provide the communication reliability, compliance, and observability of a proven infrastructure layer, which allows developers to focus on agent behavior and business logic as a new interface for interaction. Just as Kubernetes abstracted container orchestration by transforming complex distributed compute management into declarative APIs, Twilio is positioned to abstract the orchestration and management of communications-first agents, at scale. As a result, Twilio has the potential to transition from a provider of programmable communications primitives into a state management engine for communication-first AI agents. Without this abstraction, every new agent deployment requires rebuilding the entire telephony and messaging integration from scratch. Twilio's forward-looking, API-centric capabilities provide developers with a robust, resilient set of abstractions that they can drive to ensure that their businesses can navigate the tremendous operational shift that agentic workflows promise.

IDC'S POINT OF VIEW

The "shadow workforce": A quantifiable demand shock

To understand the business magnitude of this shift, consider the volume of digital labor currently entering the development pipeline. This labor represents a massive, unmodeled demand shock for communication infrastructure because — unlike previous waves of automation, which typically replaced human workers in back-office roles — agentic AI is creating a new class of extroverted workers: systems that must interact directly with customers, partners, and stakeholders in real time. These agents will handle customer support interactions, sales conversations, collections outreach, technical troubleshooting, appointment scheduling, and hundreds of other functions that currently require human availability. As a result, enterprises are not deploying one or two agents but architecting multi-agent systems where different specialized agents handle different customer engagement functions under the same business logic, customer relationships, and compliance frameworks.

IDC forecast scenarios

- **Total agent population:** IDC estimates that by 2029 the number of discrete, task-specific AI agents deployed in the enterprise will reach 1 billion.
- **The "extroverted" share (SAM):** Unlike back-end automation bots that process invoices or manage databases, our analysis suggests 20% of these agents are designed specifically for "external actuation." This creates a serviceable addressable market (SAM) of 200 million agents whose primary job is to talk, text, or email humans.
- **The Twilio " serviceable obtainable market (SOM):** Based on its incumbent status and developer ecosystem and the importance of communication channels to agentic workflows accomplishing business intent, we project Twilio can capture between 40% and 50% of this market due to the company's agentic road map and strength in CPaaS offerings. By 2029, Twilio will become the underlying infrastructure for 80 million to 100 million agents.

This forecast illustrates exponential concurrency growth because, unlike human agents restricted by 8-hour shifts, geographical constraints, and vacation time, AI agents operate constantly. An enterprise deploying the same number of autonomous agents as human employees would see their "contact center" capacity increase by a factor of three to five simply from the elimination of human scheduling constraints. This could create a volume multiplier on the network that is currently unmodeled by traditional financial analysis and represents a fundamental repricing of communication infrastructure demand.

Multi-agent orchestration and the network effect

Different agents will specialize in various customer engagement functions, each potentially leveraging different models or runtimes. Sales agents identify opportunities,

qualify leads, and guide customer conversations toward purchase decisions. Support agents handle troubleshooting, issue resolution, and knowledge base queries with human-like empathy. Collections agents manage payment outreach and account management with regulatory compliance. Technical troubleshooting agents diagnose and resolve complex issues, escalating when necessary.

For these systems to operate coherently across an enterprise, they must share access to the same customers and communication touch points. A customer calling to resolve a billing issue might be routed to a collections agent, who recognizes the root cause is technical and seamlessly hands it off to a technical agent, all without requiring the customer to repeat their situation. This requires more than simple call routing. It demands a unified orchestration layer understanding customer history, agent specialization, and business rules. It also requires the ability to identify trusted agents versus rogue agents. Twilio with its recent acquisition of Stytch provides the tools needed for identity and access management of these agents.

A consistent communication layer provides three critical capabilities that no individual agent can offer in isolation:

- **Orchestration:** This determines which agent should handle which interaction based on type, urgency, customer history, and current system load through dynamic, real-time decision-making.
- **Context preservation:** This ensures that when conversations move between agents or between automated and human support, context moves seamlessly without requiring customers to repeat information.
- **Extensibility:** This describes the ability to evolve the agent mix over time without disrupting how calls and messages reach users, allowing enterprises to swap agents without rebuilding infrastructure.

Twilio already abstracts the complexity of global telephony, SMS, and IP messaging, allowing application developers to treat calls and messages as simple programmable operations. As organizations introduce agents that must interact over these same channels, they plug agents into this familiar communication fabric. This reduces complexity, accelerates time to value, and ensures enterprise communication strategies remain coherent as agent architectures evolve. Twilio is embedded in many channels that conversation-driven systems require, including phone, text, email, and popular messaging apps like WhatsApp and Facebook Messenger. Its core value lies in translating heterogeneous networks into a coherent set of programmable capabilities that behave consistently from developers' and product teams' perspectives. As enterprises layer agents atop existing applications and workflows, they rely on Twilio to provide stable access to calls, texts, and chat endpoints while experimenting with different agent designs and orchestration patterns.

The "agent sprawl" crisis and the imperative of governance

As AI agent forecasts scale, CIOs face an urgent new threat that can be characterized as unmonitored agent interactions. Different agents will be optimized for different functions. Some may be powered by OpenAI's models and tuned for sales conversations. Others might use Anthropic's architecture, optimized for detailed troubleshooting. Still others might leverage open source models that are deployed on-premises for sensitive technical conversations. A large financial services firm might deploy dozens of distinct agents, each with different capabilities, risk profiles, and compliance requirements. When an enterprise deploys open source Agent A and proprietary Agent B, they cannot afford them operating in silos with different compliance standards. If Agent A inadvertently discloses customer personally identifiable information (PII) during a call while Agent B properly redacts it, the enterprise faces regulatory liability, customer trust erosion, and operational chaos.

Twilio could serve as the centralized control plane for the consistent enforcement of enterprise policies and governance about agentic communication. Regardless of whether agents are powered by OpenAI, Anthropic, NVIDIA Nemotron, or proprietary in-house models, traffic flows through a single governance layer that enforces consistent policies across consent management, recording controls, PII redaction, sentiment monitoring, and escalation rules. This would transform Twilio into the "system of record" for communications-first agent behavior. For enterprises, this is not a cost center but a compliance necessity that functions as a critical infrastructure control just as essential as firewall policies are for data security.

Twilio's structural advantage in this domain is significant. It has spent over a decade building authentication frameworks, call recording controls, regional data residency compliance, and consent management systems for human-led interactions. These same tools now become guardrails for agentic systems.

Rather than treating each deployment as a fresh compliance project, organizations leverage a common set of policies and technical controls already governing human-led interactions. This lowers friction and accelerates safe agent adoption at scale. As enterprises introduce agents interacting over channels used by humans, they plug agents into a familiar communication fabric rather than constructing new bespoke integrations for each deployment. This creates a stable foundation where channels, routing, and compliance controls remain consistent even as specific agents, models, and vendors evolve. The separation of concerns — with governance handled by Twilio and model innovation handled by development teams — is essential to scaling agentic systems safely.

Operational intelligence: Turning the pipe into a lens

Measuring and improving agent performance once they operate on real communication channels is another crucial dimension for agentic systems. It is insufficient for an agent to merely connect to a phone number or messaging endpoint. Enterprises must understand

what transpires in those conversations. Key performance metrics driving agent tuning and governance include four metrics, as follows:

- **Containment rates:** These measure how often agents resolve issues without human escalation, directly impacting cost and customer satisfaction.
- **Handle times:** These measure interaction duration and resolution efficiency, though customer satisfaction matters more than speed.
- **Sentiment shifts:** These capture the emotional trajectory of conversations and customer satisfaction signals, reflecting whether customers felt heard and perceived agents as competent.
- **Downstream business outcomes:** These measure impact on conversions, churn, retention, and revenue, recognizing that high containment can still damage long-term customer value.

These metrics form the basis for deciding which agents to expand, which to retrain, and when to adjust routing for the right automation-to-human support mix. A communications platform positioned in the path of every call and message is uniquely situated to surface operational insight. By aggregating transcripts, sentiment data, outcome signals, and customer profiles across all agents and human handlers sharing the same channels, Twilio provides a unified performance view cutting across different models and use cases. This feedback loop transforms the communication layer from a passive conduit into an active enabler of governance, experimentation, and continuous improvement.

Over time, combining programmable communications with rich operational insight creates reinforcing advantages. The greater the volume of conversations that traverse Twilio-powered channels, the more data enterprises can use refining routing logic, escalation policies, and channel strategies for both agents and humans. This transforms the communication layer into an intelligent control point rather than a simple pass-through, positioning Twilio to anchor an emerging generation of communication-centric applications where agents, people, and workflows depend on the same reliable and measurable fabric.

The independent communications backbone: Staying above the AI wars

Developers and the organizations they work for fear vendor lock-in. They are wary of building entire communication stacks directly inside OpenAI's platform or Google Gemini because doing so binds customer relationships, phone numbers, and communication history to a specific model provider. If that provider raises prices, changes terms, or loses competitive ground, the enterprise finds itself trapped in an expensive, difficult migration. Twilio offers the neutral substrate alternative that prevents this outcome. By building communication logic on Twilio rather than embedding it within a specific model provider, developers ensure their customer data, phone numbers, and compliance logs remain independent of the underlying AI model. If OpenAI raises prices or NVIDIA releases a

superior open source model, developers can swap the "brain" without ripping out the "nervous system," which is constituted by Twilio.

This neutrality carries significant weight. As enterprises adopt multi-model strategies that use OpenAI for sales conversations, Anthropic for detailed support interactions, and specialized models for domain-specific tasks, they require a model-agnostic communication substrate. Twilio's position as "Switzerland" of the AI stack becomes increasingly valuable precisely because it maintains enterprise optionality. In this fragmented agent ecosystem, Twilio serves as the stable communications backbone that every agent must traverse to reach customers. This enables organizations to compare outcomes across agents while keeping telephony and messaging architecture consistent. This positioning preserves flexibility at the agent layer and channels competitive pressure toward better performance without forcing repeated reinvestment in communications fundamentals. Each agent vendor competes on intelligence and capability, while Twilio maintains the integrity and governance of the channel layer beneath them.

Conclusion

Technology buyers would do well to appreciate Twilio's potential to serve as a strategic governance firewall that is essential to safe, compliant, and auditable agentic systems, as opposed to only programmable infrastructure provider for SMS, email, and voice. As enterprises deploy agentic systems at scale, the infrastructure requirements shift dramatically. While the ability to select, train, and optimize AI models remains strategically important, the ability to govern how those models interact with customers, data, and business processes becomes equally critical. The bottleneck in agentic systems is no longer simply acquiring intelligence but also the management, control, verification, and governance of that intelligence once deployed. Moreover, the value of orchestrating those interactions with reliability and compliance becomes paramount. Ensuring that intelligence stays within guardrails, respects consent, handles PII properly, and escalates appropriately becomes the core strategic imperative.

Twilio's move to position itself as the agentic control plane aligns perfectly with the developer's and enterprise's core needs in the next phase of the AI life cycle: safety, standardization, and independence from any single model provider. Organizations that fail to adopt a governance layer for agentic systems face regulatory risk, customer trust erosion, and operational chaos. Just like the more than 9,000 start-ups building in the AI space are already doing on Twilio, those that adopt Twilio as their communications control plane gain the confidence to scale agents into increasingly critical and higher-stakes customer journeys.

LEARN MORE

Related research

- *Microsoft Ignite 2025: Security Primitives in Action — Securing Agents Amid the AI Expansion* (IDC #IcUS54039225, December 2025)

- *CIO Peer Perspective — Agentic AI and the Emergence of Digital Labor* (IDC #US53384325, December 2025)
- *Dreamforce 2026: The Era of the Agentic Enterprise* (IDC #US53894225, November 2025)
- *Dreamforce 2025: Agentforce and the Journey to an Agentic Enterprise* (IDC #IcUS53884525, October 2025)
- *IDC PlanScape: Agentic AI Security* (IDC #US53731425, August 2025)
- *Ready or Not, Here Comes Agentic AI: Exploratory View from IT and Business Leaders at IDC Data Xchange 2025* (IDC #EUR153545725, June 2025)

Synopsis

This IDC Market Perspective provides insights into Twilio's agentic opportunity. Twilio is poised to evolve from a communications utility into the "agentic control plane" for enterprises deploying autonomous AI agents. As organizations shift from generative AI coding tools to agentic systems that independently execute tasks, the need for robust governance, identity, and orchestration becomes critical. Twilio's proven infrastructure, API-centric approach, and recent acquisitions position it to provide the compliance, observability, and operational intelligence required to safely scale digital workforces across diverse communication channels. With IDC projecting up to 100 million agents running on Twilio by 2029, the company's neutral, model-agnostic platform offers enterprises the flexibility to innovate while maintaining control and regulatory compliance. Twilio's strategic move addresses the emerging "agent sprawl" crisis, ensuring enterprises can govern, audit, and optimize agentic interactions at scale. In the coming era, Twilio's role as the agentic governance backbone will be essential for safe, compliant, and scalable AI-driven customer engagement.

"Enterprises deploying autonomous AI agents face a new infrastructure imperative around governance at scale. As digital workforces grow from dozens to thousands of agents, the ability to authenticate, monitor, and audit their interactions becomes as critical as the intelligence driving them. Twilio's evolution into an agentic control plane addresses this challenge by providing the neutral governance layer that allows enterprises to innovate with confidence while maintaining compliance and operational control across their entire agent ecosystem." — Arnal Dayaratna, research vice president, Software Development at IDC

ABOUT IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology, IT benchmarking and sourcing, and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives. Founded in 1964, IDC is a wholly owned subsidiary of International Data Group (IDG, Inc.).

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com