

## How Anthropic and OpenAI Are Building "Everywhere Ecosystems"

The Multi-Dimensional Go-to-Market Strategies Designed to Embed AI into Enterprises

Analyst(s): Alex Smith

Publication Date: July 1, 2026

Document #: AIOAS202606

### Key Points

- OpenAI and Anthropic are redefining enterprise software by creating "Everywhere Ecosystems," a multi-dimensional strategy that embeds cognitive reasoning engines directly into the operational fabric of global businesses.
- The new competitive frontier is structural entrenchment; securing dominance across physical compute, capital-backed deployment vehicles, proprietary marketplaces, and platform-native integrations.
- Foundational AI companies are bypassing traditional sales cycles by using aggressive, capital-funded vehicles to embed internal engineering cohorts directly into client organizations, while simultaneously bridging the "Implementation Gap" through GSI partnerships.
- Achieving full market penetration requires a pivotal transition from high-touch, bespoke enterprise deployments to standardized, transactional channels capable of efficiently reaching fragmented SME and mid-market sectors.

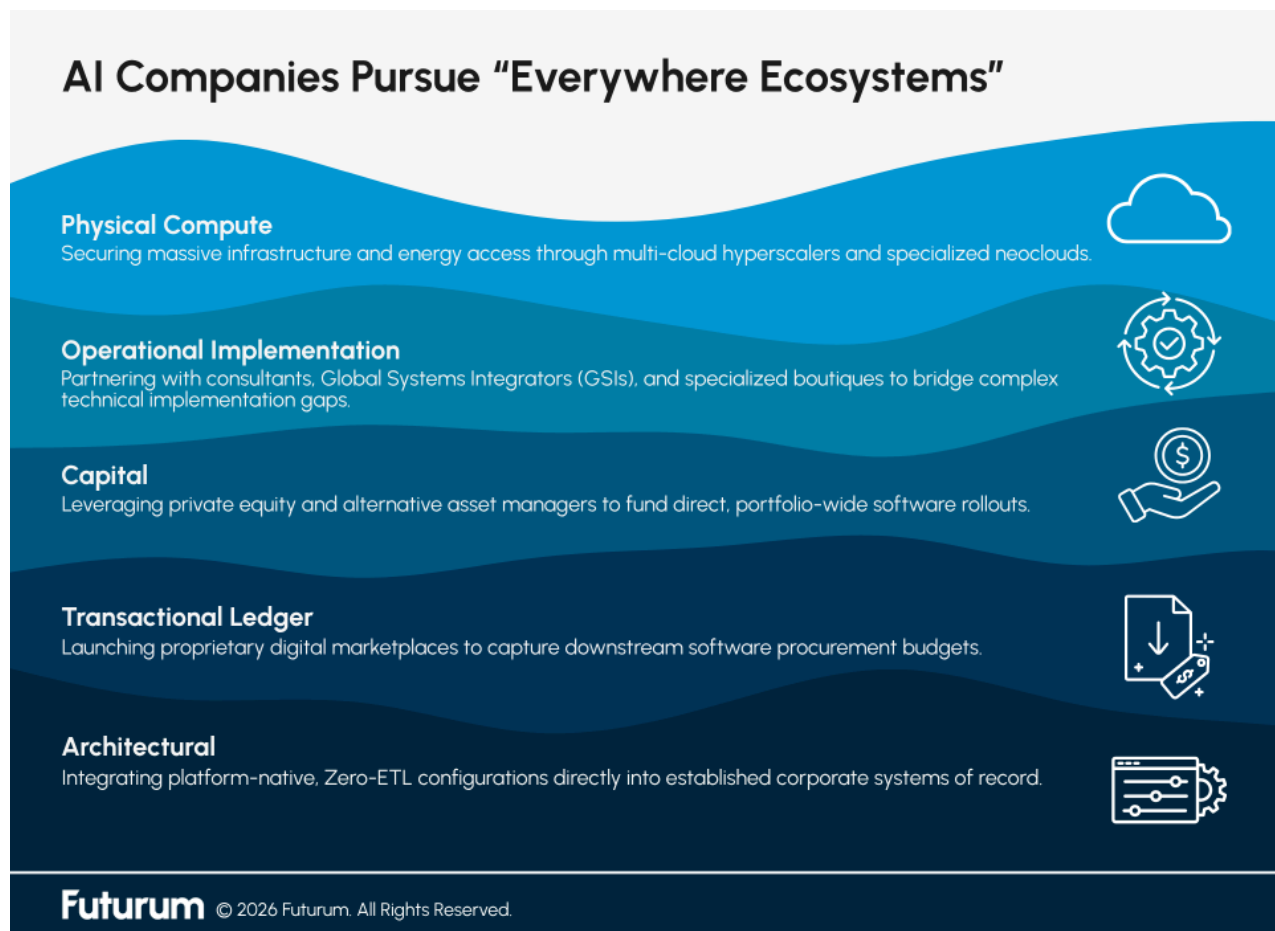
The commercial growth and operational scaling of OpenAI and Anthropic have completely broken the historical benchmarks set by the SaaS era. They represent the two most rapidly expanding organizations in modern corporate history, rewriting the growth trajectories once championed by giants such as Salesforce and Slack.

Their shared strategic objective is to integrate foundational artificial intelligence directly into the operational fabric of global enterprises. Achieving this has outpaced conventional go-to-market strategies and traditional alliance frameworks. Consequently, both Anthropic and OpenAI are breaking open every traditional partnership mechanism to capture surging corporate demand. By doing so, these organizations are pioneering a highly aggressive, multi-dimensional distribution framework that moves beyond traditional SaaS dynamics. We define this new paradigm as the "Everywhere Ecosystem."

## Defining the Everywhere Ecosystem

Traditional enterprise technology adoption typically relies on a linear distribution model, where a vendor sells a standalone software layer directly to a customer or through a single primary cloud host. The Everywhere Ecosystem represents a structural break from this model. While industrial frameworks, such as NVIDIA’s five-layer AI cake, describe the physical stack required to manufacture intelligence from energy to chips, the Everywhere Ecosystem describes the commercial stack required to capture mass enterprise penetration. It is characterized by the simultaneous injection of a single cognitive reasoning engine into every layer of the enterprise technology environment, occurring across five distinct vectors:

Figure 1: AI Companies Pursue “Everywhere Ecosystem



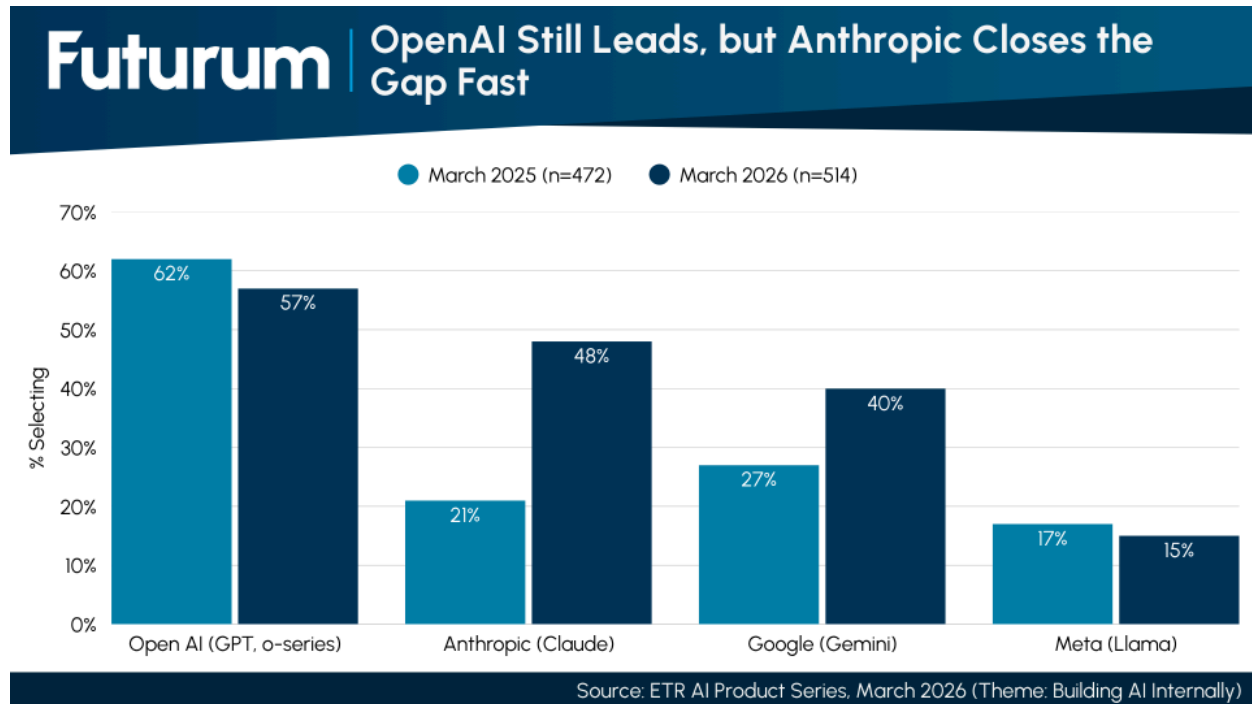
Source: Futurum Research, June 2026

Ultimate market velocity is no longer determined by raw model benchmarks or standard SaaS sales cycles. Instead, it belongs to the laboratory that can embed its intelligence so deeply into these five vectors that it becomes structurally impossible for an enterprise to remove.

This systemic shift is already reflected in current enterprise buying behavior. In ETR’s AI Product Series March 2026 report, OpenAI GPT “o-series” models were shown to remain the most

widely used foundation model in enterprises at 57%, while Anthropic Claude saw the most significant gains, rising from 21% to 48% in a year. While the well-documented model test performances will undoubtedly play an essential role in determining which LLMs get adopted, the technology industry has long shown that ecosystem entrenchment can trump product capabilities; this will be a key battleground between the two labs.

Figure 2: OpenAI Still Leads, but Anthropic Closes the Gap Fast



Source: ETR AI Product Series, March 2026 (Theme: Building AI Internally)


## 1. Infrastructure and Compute Alliances: Hyperscalers, Neoclouds, and Frontier Infrastructure

In the current AI race, compute capacity is no longer just a backend utility. Instead, it is the ultimate determinant of market distribution and enterprise adoption. Frontier models demand so much raw power that physical infrastructure has become a critical strategic weapon in go-to-market execution. For OpenAI and Anthropic, commercial survival hinges on securing specialized, resilient, and immediately available compute through a complex matrix of partnerships. This relationship is entirely symbiotic: AI labs need the data centers and energy grid access to train and serve their models, while cloud providers need access to top-tier LLMs to drive consumption across their broader cloud ecosystems.

As a result, alliances with major hyperscalers have become the primary distribution engine for both labs. Microsoft, AWS, Google Cloud, and Oracle are no longer just infrastructure hosts because they have evolved into critical channel partners. The data backs this up: according to ETR's March 2026 report, cloud providers remain the dominant procurement route for foundation models, with 60% of enterprise respondents buying directly through them.

Futurum’s AI Platforms Decision Maker Survey independently validated this, showing that 64% of customers were deploying AI models on provider-managed clouds.

Figure 3: Where Models Get Deployed



Source: Futurum Research, March 2026

Both labs have abandoned single-vendor exclusivity in favor of multi-cloud architectures, securing capacity through hyperscalers (Microsoft, AWS, Google Cloud, Oracle) and specialized "Neocloud" providers (CoreWeave, Nebius) to meet extreme compute demands.

Yet despite these alliances, the severe constraint on global compute has driven AI companies to establish unconventional partnerships at the absolute frontier of physical infrastructure. This shift is illustrated by diverse and unconventional collaborations. These range from Anthropic’s landmark alliance with SpaceX to secure immediate capacity from the 300-megawatt Colossus 1 supercluster, to the rise of sovereign-backed infrastructure entities such as Humain in the Middle East.

## 2. The Implementation Battleground: Management Consultants, GSIs, Frontier Partners, and Direct-to-Enterprise Motions

Embedding AI into actual enterprise workflows requires deep, localized implementation. This "Implementation Gap" is the single biggest bottleneck to enterprise adoption. Getting access to models is one thing, but integrating these models into enterprise software stacks and proprietary workflows presents a highly complex technical challenge. This dynamic has initiated a strategic contest for ownership of the next critical phase of enterprise integration, featuring distinct categories of service and technology partners.

GSI such as Accenture, KPMG, and TCS have moved beyond traditional advisory roles to become active implementation partners. By establishing dedicated AI organizations and natively integrating models such as Claude and GPT into their own global service gateways, these firms now provide the operational scale and technical labor required to embed cognitive architectures directly into core corporate systems.

Simultaneously, management consultancies such as BCG, McKinsey, and Bain exert significant influence over early-stage strategy. Because generative AI is a CEO-level priority, these firms design strategic roadmaps that frequently dictate model selection long before technical procurement begins, exemplified by initiatives such as OpenAI's "Frontier Alliance" with McKinsey.

In addition to traditional service providers, a new category of "[Frontier Partners](#)" has emerged. These highly specialized, artificial intelligence-native technical entities prioritize advanced agentic execution and proprietary cognitive system engineering over high-level advisory or commoditized software deployment. Despite competing against entrenched GSIs and elite management consultancies, these agile partners maintain a distinct operational advantage derived from their singular focus on model behavior, customization, and deployment dynamics. Unburdened by legacy enterprise software architectures or massive labor-force optimization mandates, Frontier Partners operate as pure-play model orchestrators capable of designing, refining, and scaling highly integrated, autonomous workflows that legacy firms often struggle to execute.

Despite these partner networks, direct-to-enterprise implementation has become an increasingly prominent strategy for foundational laboratories. Driven by the necessity to accelerate adoption and bypass traditional partner delivery timelines, both OpenAI and Anthropic are deploying internal engineering cohorts directly into key accounts. The establishment of the OpenAI Deployment Company and Anthropic's expansion of its internal applied engineering divisions represent direct-to-market strategies designed to ensure the success of high-value deployments. By embedding internal engineers within client organizations, the laboratories frequently find themselves in dual roles. They act as technology enablers for their GSI partners while simultaneously competing with them for high-margin implementation and digital transformation services.

### 3. Capital-Backed Direct Deployment: The Intersection of Private Equity and Applied Engineering

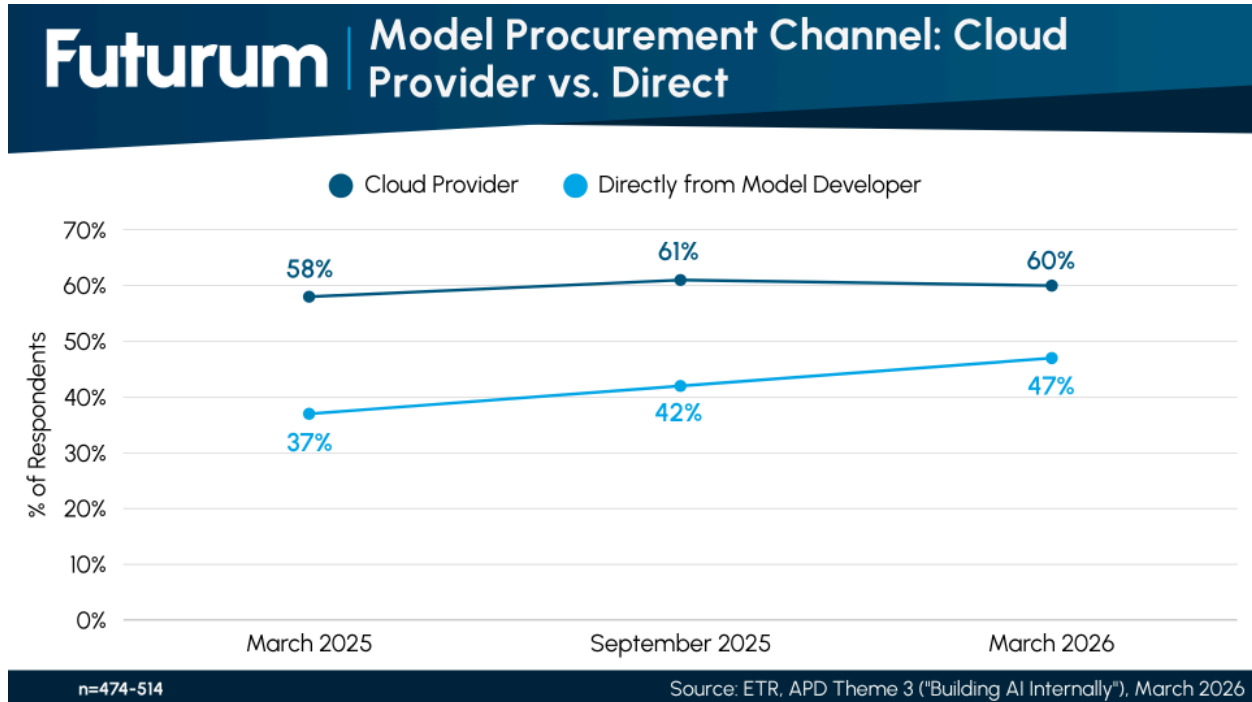
Despite the extensive network of indirect alliances, direct enterprise engagement remains a proven, historically validated methodology for technology deployment. Foundational model companies have recognized that enterprise adoption cannot be left entirely to traditional third-party capabilities. Direct deployment allows companies to bypass the prolonged enablement cycles that any ecosystem building requires. This ensures that highly complex,

agentic cognitive systems are deployed with maximum velocity, priority, and operational fidelity directly within the customer's technical environment.

However, a more aggressive playbook has emerged at the center of this direct-to-market push. Rather than establishing standard, high-overhead corporate direct sales forces, OpenAI and Anthropic are executing this direct-to-market motion through strategic alliances with major players in the global capital markets, who are equally hungry to capitalize on the AI opportunity. Private equity firms, alternative asset managers, and investment banks are no longer acting as passive financial sponsors; instead, they have emerged as active operational orchestrators seeking to capitalize on the artificial intelligence transition by injecting massive funding directly into deployment vehicles. This structure creates a powerful convergence of interests. Capital markets provide the necessary multi-billion-dollar funding and immediate, captive customer portfolios, while the foundational laboratories deliver the specialized engineering personnel required to operationalize those investments across entire corporate portfolios.

This capital-backed direct deployment strategy has manifested in two fundamentally distinct institutional models: OpenAI's direct corporate outsourcing approach and Anthropic's portfolio-wide standardization framework. OpenAI's strategy is executed through "The OpenAI Deployment Company" (DeployCo), a multi-billion-dollar direct services business backed by TPG and Bain Capital that embeds cohorts of forward-deployed engineers to build bespoke integrations for large enterprises. This model is already demonstrated by scaled deployments at BBVA and John Deere. Conversely, Anthropic has established a \$1.5 billion standalone enterprise services firm in partnership with Blackstone and Goldman Sachs Alternatives to deploy the Claude model family systematically across a highly diversified mid-market network. While OpenAI's model utilizes direct technical acquisitions such as Tomoro (as well as those to come) to solve localized integration challenges on an account-by-account basis, Anthropic's model leverages the captive, aggregate footprint of institutional portfolios. This includes Blackstone's network of 250 companies representing \$300 billion in collective revenue, creating a uniform, platform-native artificial intelligence operating system. Both examples, however, represent an investment path designed to inject AI into the enterprise.

Figure 4: Model Procurement Channel: Cloud Provider vs. Direct



Source: ETR, APD Theme 3 ("Building AI Internally"), March 2026

## 4. Digital Marketplaces as Consolidated Procurement Engines and Strategic Ecosystem Moats

Foundational AI companies are launching proprietary marketplaces to replicate the integration moats historically built by major software and cloud platforms. These centralized marketplaces allow enterprises to identify, evaluate, and deploy certified third-party integrations directly into existing operational stacks, positioning the model labs as central orchestrators of the corporate technology environment.

OpenAI leverages the GPT Store to cultivate an expansive, community-driven ecosystem. By democratizing the creation and searchability of custom assistants, OpenAI secures its position within managed Enterprise and Team workspaces. The framework enables secure, governed deployment of proprietary agentic workflows. By encouraging employees to develop and distribute specialized tools, OpenAI embeds ChatGPT as the default execution layer, creating high operational friction for any potential migration to competitive model architectures.

Anthropic's Claude Marketplace replicates the hyperscaler playbook to secure deeper customer lock-in. Through a "Commitment Burn-down" model, enterprises apply pre-negotiated spending commitments toward third-party tools built on Claude, such as GitLab Duo or Harvey. This creates a powerful commercial incentive: once customers commit millions in annual spend to Anthropic, they are financially motivated to procure all downstream cognitive software through the marketplace to fulfill contractual quotas, effectively consolidating procurement.

These marketplaces build formidable ecosystem moats, further positioning model companies at the center of the technology landscape. Controlling both the cognitive reasoning engine and the financial ledger of software procurement creates profound lock-in. This centralization threatens to disrupt traditional SaaS distribution models, establishing foundational labs as the definitive gateways for future enterprise computing.

## 5. Platform-Native Architecture and Zero-ETL Integration: Unlocking Systems of Record

Early market narratives predicted that generative AI would trigger a "SaaS apocalypse," systematically wiping out legacy software architectures. In operational reality, established enterprise applications function as deeply embedded, mission-critical systems of record that are highly insulated from direct displacement due to substantial data complexity and process lock-in. Rather than replacing systems such as customer relationship management (CRM) or enterprise resource planning (ERP) suites, frontier Large Language Models are functioning as analytical accelerators, unlocking latent economic value, automation, and predictive insights directly from within these governed data repositories. Consequently, the strategic imperative has shifted from technology displacement to native, platform-level integration. This establishes a prevalent model-embedding paradigm that operates via platform-native, often "Zero-ETL" (Extract, Transform, Load) interfaces.

This architectural shift is demonstrated across a diverse, highly competitive grid of enterprise software alliances, where models operate as localized, native intelligence engines. For example:

- **Adobe:** Integrates Claude and GPT architectures to execute design projects within its Firefly platform.
- **CrowdStrike:** Utilizes frontier model intelligence for automated vulnerability discovery and risk mitigation via Project QuiltWorks.
- **Databricks:** Partners with Anthropic and OpenAI to power context-aware data workflows and its Lakewatch security platform.
- **HubSpot:** Connects with both labs to power the Breeze suite for automated content, sales outreach, and database management.
- **Microsoft:** Features diversified model access in Microsoft 365 Copilot and licenses Anthropic's Copilot Cowork framework.
- **Salesforce:** Uses Claude as the primary engine for Slack and the preferred model for Agentforce deployments.
- **SAP:** Collaborates with Anthropic to embed Claude's reasoning capabilities into its Joule assistant.

- **ServiceNow:** Integrates frontier models to drive agentic operations and low-latency, speech-to-speech technologies.
- **Snowflake:** Partners with OpenAI and Anthropic to deliver frontier intelligence natively within the Cortex platform for customized applications.
- **Workday:** Collaborates with Anthropic on its Solopreneurship Accelerator Program to provide grants and training to entrepreneurs.

## Conclusion: Institutional Omnipresence and the Untapped Transactional Channel

The strategic trajectories of OpenAI and Anthropic demonstrate a fundamental shift in how disruptive technologies achieve enterprise scale. By weaving cognitive architectures into physical compute assets, systems-integrator pipelines, private equity portfolios, proprietary marketplaces, and native workflows, both companies are successfully building the foundations of an "Everywhere Ecosystem." This multi-layered approach has effectively institutionalized generative artificial intelligence within core enterprise architectures.

While OpenAI and Anthropic have turned their attention to high-touch enterprise deployment, the next phase of market maturity depends on scaling downstream transactional channels. Monetizing the long-tail of global technology spend requires transitioning from bespoke, high-touch engineering to standardized, automated distribution.

This absence represents a significant barrier to long-term market optimization. Legacy channel entities possess the exclusive, deeply entrenched operational relationships and earned trust required to efficiently unlock the vast small-to-medium enterprise (SME) and mid-market sectors. Because this segment of the market is highly fragmented, it cannot be economically addressed via customized direct-sales cohorts or bespoke, account-by-account joint ventures.

Consequently, the next phase of the market maturity curve will depend on how successfully these companies can scale their downstream transactional channels. Early operational adjustments, such as inaugural partner events, strategic hires of veteran channel executives, and initial investments into programmatic partner certifications, indicate a growing institutional recognition of the need to invest in partner motions. Establishing transactional channels akin to those that historically helped scale past technology cycles is key to full market penetration. Only then will the "Everywhere Ecosystem" fully manifest.

## Other Insights from Futurum:

[The Hidden Moat: Why Operational Depth Defeats the 'Build It Yourself' Narrative](#)

[From Storage to Action: Why Autonomous AI is Forcing a Database Revolution](#)

[The Hard\(er\) Challenge in Agent Governance Is Authorization](#)

## About Us

### About the Authors

Alex Smith is GM, Futurum Research, and the VP & Practice Lead, Ecosystems, Channels & Marketplaces at Futurum. He is responsible for establishing and maintaining the Channels Research program within the overall Futurum GTM and Channels Practice. This includes overseeing the channel data rollout in the Futurum Intelligence Platform, conducting primary research activities such as research boards and surveys, delivering thought-leading research reports, and advising clients on their indirect go-to-market strategies. He can be reached at [asmith@futurumgroup.com](mailto:asmith@futurumgroup.com).

### About The Futurum Group

Every day, The Futurum Group's analysts, researchers, and advisors help business leaders worldwide anticipate tectonic shifts in their industries and leverage disruptive innovation. Unlike traditional analysts, The Futurum Group works not only in analysis and research but also takes that insight and knowledge even further, engaging all the way through the go-to-market process.

Futurum Research provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing. It also creates, distributes, and amplifies rich media content that all stakeholders read, watch, and listen to.

See more details on The Futurum Group at [futurumgroup.com](http://futurumgroup.com).

## Copyright & Use License

### Copyright Notice

Copyright ©2026 by The Futurum Group, LLC. All rights, including that of translation into other languages, are specifically reserved. No part of this publication may be reproduced in any form, stored in a retrieval system, or transmitted by any method or means, electrical, mechanical, photographic, or otherwise, without the express written permission of The Futurum Group [futurumgroup.com](http://futurumgroup.com). United States copyright laws and international treaties protect this publication. Unauthorized distribution or reproduction of this publication, or any portion of it, may result in severe civil and criminal penalties and will be prosecuted to the maximum extent necessary to protect the publisher's rights.

### License Notice

This document may be distributed within the licensed organization only.

The following acts are prohibited:

Transmittal to others outside your immediate organization including partners, resellers, external consultants, etc. in any media format

Posting on a website which is accessible to others outside your immediate organization

The possession or use within an unlicensed organization

### Limitation of Liability Notice

The information contained herein has been obtained from sources believed to be reliable. The Futurum Group shall have no liability for errors, omissions, or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve their intended results. The opinions expressed herein are subject to change without notice.