Futurum®

2026 Research Agenda

Key Topics & Coverage Areas

Futurum Intelligence provides critical insights into digital transformation, focusing on adoption, innovation, and disruption. Backed by a team of industry experts, we deliver research through personalized analyst-client interaction and client portals with visualization dashboards and qualitative and quantitative data reports.

Our research is organized into practice areas aligned with key digital transformation topics, addressing critical business questions. Each area includes analyst coverage and planned deliverables for the year. Through collaboration and strong industry relationships, we identify emerging trends early, helping clients make informed business decisions.
Please select an image to view the agenda of a specific practice area.

Please select an image to view the agenda of a specific practice area.

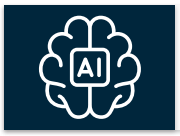

AI Platforms

CIO & Technology Buyers

Cybersecurity & Resilience

Data Intelligence, Analytics & Infrastructure

Enterprise Software & Digital Workflows

Ecosystems, Channels & Marketplaces

Hybrid Cloud & Infrastructure

Intelligent Devices

Networking

Observability

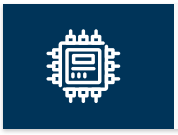

Semiconductors, Supply Chain & Emerging Tech

Software Lifecycle Engineering

Last updated: February 23, 2026

AI Platforms

The Futurum AI Platforms research agenda for 2026 focuses on the critical shift from experimental AI to the deployment of industrial-scale, resilient, and autonomous systems. This transition, which is moving the market's focus from general-purpose chatbots to agentic applications, represents a significant inflection point for enterprises, vendors, and investors. For enterprises, navigating the confluence of infrastructure constraints, diverse global regulations, and model optimization is essential to achieving competitive advantage and operational safety. Vendors must align their offerings to address these complex enterprise demands, while investors must understand the technological and regulatory hurdles to identify the next wave of market leaders. Our coverage will focus on these seven areas, all aimed at understanding the path from isolated pilots to the realization of a unified autonomous enterprise, whether that comes in 2026 for the enterprise or beyond.

Key Issues for 2026

The Realities of Agentic AI

While Agentic AI will be a dominant theme, 2026 is less about universal scale and more about the foundational struggle for reliability. We are tracking the shift from rigid, human-led interfaces to agents that can navigate multi-step workflows for use cases such as customer service and complex data analysis. This transition may trigger a broad industry shift from per-seat licensing to agent-usage-based pricing, potentially consolidating parts of the software market.

However, significant hurdles remain:

- Legacy systems often lack the real-time connectivity and identity management necessary for these agents to act autonomously without regular failure.
- Enterprises must decide whether to adopt siloed agent platforms or a unified agentic mesh that abstracts complexity across their entire technology stack.

Inference-time Compute

Even if an enterprise never trains a single foundational model, inference-time compute (the 'thinking' phase of AI, where value is realized) will become a significant factor in technology budgets in 2026. Still, there are ways to mitigate potential runaway costs. As the value of AI shifts from static knowledge ingrained during the training phase to compute applied at the moment of query, this shift directly impacts the model consumer, who, rather than the model trainer, incurs the recurring operational costs of running these systems in production. Because inference is a metered, utility-like bill that scales with every interaction, it represents a recurring operational expense that can quickly spiral if not managed. To avoid spiraling costs when a proof-of-concept hits production, enterprises are moving beyond simple API calls. They are increasingly opting for dedicated cloud- and on-premises inference services that provide more predictable throughput and better cost management for large-scale autonomous operations. Plus, the strategic use of inference-time scaling enables non-model builders to make smaller, cheaper models perform at elite levels without the need for expensive, proprietary fine-tuning. This is a critical lever for organizations to gain a competitive edge while maintaining a disciplined bottom line. Inference-time scaling enables

Go back to the first page

organizations to buy exactly as much intelligence as they need at any given moment, ensuring their AI spend remains aligned with the actual business value of the output. But it's not necessarily easy to achieve, so expect more tools to help enterprises do so in 2026.

Energy & Cooling as Primary Bottleneck

The scarcity of power and cooling will become the primary constraint on AI expansion in 2026, leading to delays in data center deployment:

- Modern AI workloads generate heat levels that make liquid-cooling technologies mandatory, forcing global infrastructure redesigns.

- The power demand of advanced facilities is reaching scales that can outstrip local grid capacity, so a new discipline is emerging: carbon-aware AI scheduling. This involves automatically shifting non-critical model processing to geographic zones or times of day where renewable energy is most available, ensuring sustainability goals don't conflict with operational requirements.

Model Fragmentation

The physical pressures mentioned above are contributing to the fragmentation of models, moving the industry away from one-size-fits-all architectures. While the largest models continue to grab headlines for their broad capabilities, enterprises are increasingly deploying specialized Small Language Models (SLMs) at the edge for latency-critical tasks such as local voice assistants, IoT device control, and privacy-sensitive data processing. This fragmentation is driven by a move toward domain-specific intelligence where models are trained on narrow, high-quality enterprise data rather than the entire web. The enterprise IT is evolving to manage this multi-modal complexity, ensuring that data - including images, video, and audio - is processed by the appropriately sized architecture.

Sovereign AI & Corporate Control

Digital sovereignty, once a niche geopolitical idea, is now a global business imperative for companies aiming to safeguard their proprietary intellectual property.

- In 2026, Sovereign AI will emphasize an organization's capacity to govern its entire intelligence stack, encompassing local data centers, bespoke models, and specialized security measures.

- A significant number of organizations are moving essential AI operations back from public clouds to private environments to maintain corporate autonomy. In contrast, cloud providers will counter with their own sovereign offerings.

The expectation is that enterprise sovereign AI will insulate businesses from external, often volatile policy changes, minimize third-party risk, and meet global data sovereignty mandates.

Looking for the Exit: M&A & IPOs

The market landscape in 2026 will likely be defined by a thinning of the herd, as the AI sector moves from speculative growth to a high-stakes capital reckoning. We are entering an initial consolidation phase as sub-scale model labs and mid-tier startups struggle with rising compute costs and the challenge of turning research into profitable products. This consolidation will be largely driven by an M&A wave, as sub-scale model labs and mid-tier startups find themselves hitting a money wall - the point where the immense capital required for next-generation compute and infrastructure outstrips their

Go back to the first page

ability to generate immediate profit. To bridge this gap, large cloud providers and incumbents are moving to absorb these independent innovators to rapidly integrate specialized engineering talent and proprietary data into their own platforms. Simultaneously, 2026 is emerging as the potential breakout year for AI IPOs.

Three primary candidates stand out as bellwethers for the industry's long-term sustainability, and likely IPO candidates, should the market support it:

- OpenAI: The current AI poster child will look for a huge valuation
- Anthropic: The enterprise safety hedge will have lesser but still significant ambitions
- Databricks: Which is working its way through the alphabet of investment rounds

Global AI Regulation & Compliance

Finally, all these developments are taking place under the shadow of global AI regulation, which has transitioned from voluntary frameworks to aggressive enforcement:

- The EU AI Act will reach a significant milestone in August 2026, bringing stringent requirements for high-risk systems across sectors such as healthcare and finance into full force.
- In the United States, a fragmented landscape of state-level laws, particularly in California, Colorado, and Texas, is effectively shaping domestic governance, even as it faces potential preemption by federal policy.

This creates a highest-common-denominator compliance challenge, forcing global enterprises to implement automated data lineage, risk assessments, and transparency measures to maintain market access across all jurisdictions.

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast
- Decision-Maker Survey - bi-annual IT Decision Maker survey
- Analyst Insight Report - a report on critical issues in the industry
- State of the Market Report- a report on technology, markets, products, and vendors
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

CIO & Technology Buyers

Futurum's Enterprise Technology Buyers practice provides comprehensive coverage of how enterprise technology decisions are made across the organization, from the CIO and IT leadership to the growing universe of business executives who now directly influence technology strategy, budgets, and outcomes. As digital capabilities become embedded in every function, enterprise technology buying has evolved from a centralized IT process into a distributed, multi-stakeholder model spanning marketing, data, security, revenue, operations, and customer experience.

While the CIO remains a central orchestrator of architecture, governance, and enterprise platforms, buying authority increasingly resides across roles such as the CMO, CDO, CISO, and CRO. AI, cloud platforms, automation, and data systems are no longer implemented solely as infrastructure investments but as business capability engines tied directly to growth, efficiency, risk management, and customer engagement. Futurum's research reflects this reality by examining technology demand through the lens of business outcomes, buyer intent, and real-world purchasing dynamics, not solely IT strategy.

A core focus of the practice is understanding how enterprises operationalize advanced technologies such as AI, analytics, and agentic systems at scale. As organizations move from experimentation to execution, buyers across multiple functions must align on governance models, data foundations, security frameworks, and integration strategies. Our research provides insight into how enterprise buyers evaluate emerging capabilities, where market offerings fall short of expectations, and how organizations balance innovation velocity with trust, compliance, and cost discipline.

The Enterprise Technology Buyers practice also delivers deep analysis of technology ecosystems and vendor landscapes, helping organizations understand how platforms, hyperscalers, best-of-breed vendors, and service partners intersect with evolving buyer needs. By connecting buyer demand signals with market supply realities, Futurum equips technology providers and enterprise leaders alike with a clear view of where investment is accelerating, where friction persists, and where future opportunity lies.

This practice is grounded in continuous engagement with senior enterprise decision-makers and powered by Futurum Intelligence, combining quantitative survey research with qualitative insight to track shifting priorities, spending patterns, and buying behavior throughout the year.

Key Issues for Buyers in 2026

CIOs as Enterprise Orchestrators in a Decentralized Buying Model

In 2026, CIOs remain foundational to enterprise technology success, but their role increasingly centers on orchestration rather than ownership:

- CIOs are responsible for enabling scalable platforms, integration architectures, data governance, and AI operating models that support dozens of semi-autonomous buyers across the business.
- Success depends on balancing speed and flexibility with enterprise-wide controls, ensuring that distributed technology adoption does not result in fragmentation, security exposure, or unsustainable cost structures.

Go back to the first page

CMOs as Primary Drivers of Digital Spend and Customer-Facing AI

The CMO has emerged as one of the most influential technology buyers outside IT:

- Marketing leaders now control significant budgets for MarTech, customer data platforms, personalization engines, and AI-driven engagement tools.
- CMOs in 2026 are under pressure to demonstrate measurable revenue impact from technology investments, while navigating growing complexity across data privacy, AI governance, and platform integration with core enterprise systems.

Data, AI, and Growth Leaders Redefining Competitive Advantage

Chief Digital, Data, AI, and Revenue leaders increasingly shape enterprise technology direction:

- These roles focus on monetizing data, deploying AI agents into revenue and operations workflows, and accelerating decision cycles.
- Their influence is reshaping buying criteria toward outcome-driven platforms that integrate analytics, automation, and AI into business processes rather than standalone tools.

Operating in a Multi-Cloud, Sovereign, and Cost-Constrained Environment

Enterprise buyers are rethinking where workloads run and why:

- Decisions around cloud placement, data residency, inference economics, and platform consolidation now involve both IT and business leaders.
- The challenge is aligning architectural choices with regulatory requirements, performance expectations, and long-term financial sustainability.

Scaling AI Safely Across the Enterprise

AI adoption in 2026 is defined by execution, not experimentation:

- Buyers must operationalize agentic systems, define accountability for automated actions, and ensure transparency across data usage and decision models.
- Governance frameworks must enable innovation while protecting enterprises from reputational, regulatory, and security risk.

Go back to the first page

Planned Deliverables

- Enterprise Technology Buyers Survey: Bi-annual global buyer study spanning CIOs, CDOs, CMOs, and other key senior business technology decision-makers
- Analyst Insight Report - Report on critical issues in the industry, including:
  - Enterprise Technology Buyers Report: Bi-annual flagship outlook
  - Specialty Buyer Reports: Rotating role- and segment-based deep dives (e.g., CMO, Data & AI, Security, Revenue & Growth)
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Cybersecurity & Resilience

Futurum covers a broad spectrum of cybersecurity-related technologies, including application, cloud, data, endpoint, network security, identity and access management (IAM), and integrated risk management and Security Operations Center (SOC) markets. With this in mind, our vantage point spans key use cases, including threat hunting and intelligence, incident response, attack detection, infusion of AI, and cyber-recovery. Key themes of our coverage include modernizing security operations, security infrastructure, security management, and related areas.

Key Issues for 2026

In 2026, the cybersecurity landscape shifts from the initial rush of AI adoption to a more complex phase of "industrialization" and autonomy. While short-term threat vectors remain familiar, the long-term architecture of security is undergoing a meaningful transformation. Our coverage for 2026 focuses on the rise of Agentic AI, the explosion of Non-Human Identities (NHI), and the "Data Fusion" required to protect unstructured data. We will also track the critical strategic pivots organizations must make to address Shadow AI, API complexity, and the looming requirements of quantum readiness and cybersecurity risk quantification.

The Era of Agentic AI and Non-Human Identities (NHI)

The AI narrative evolves from simple acceleration to autonomy. We are entering the era of Agentic AI, where autonomous agents (e.g., Agent365) execute complex workflows without human intervention. This drives an explosion in Non-Human Identities (NHI), creating a massive, under-protected attack surface.

- **The "High-Wire" Act:** While organizations deploy fine-tuned and neurosymbolic models to scale operations, adversaries are "industrializing" their toolkits, aiming for a larger share of wallet through efficiency rather than just "doomsday" events.
- **The Identity Crisis:** Security teams must pivot from user-centric IAM to managing machine-to-machine interactions and securing the credentials of autonomous agents.

Governing the "Shadow AI" Explosion and Expanding Supply Chains

The attack surface is no longer just "sprawling"; it is deepening. The rapid adoption of unmanaged AI tools has birthed a Shadow AI situation, necessitating a new era of CASB (Cloud Access Security Broker) and SSPM (SaaS Security Posture Management) capabilities.

- **Supply Chain Convergence:** Application Security (AppSec) and Third-Party Cyber Risk Management (TPCRM) are converging as supply chain risks move closer to the code level.
- **Browser & Compute Security:** We see a resurgence of browser-based security concerns and a rise in Confidential Computing to protect data-in-use, driven by the need to secure high-value AI computations.

Go back to the first page

Platform Dynamics: API Complexity and Complexity at the Edge

The platform vs. point-solution debate continues, but is nuanced by the extreme complexity at the edge. Modern applications are becoming increasingly complex at the edge - a tangled web of APIs, edge dependencies, and content delivery networks, now increasingly adding AI elements as well.

- **Integration over Consolidation:** While vendors expand platforms to include Identity and Data, the unique demands of AI are forcing a re-evaluation of point solutions. The market focus shifts from pure consolidation to deep **integration**, enabling specialized AI defense tools to operate within broader ecosystems.

- **External influences:** The complexity at the edge will also be increasingly affected by geopolitical changes, as organizations grapple with sovereignty considerations arising from regulatory mandates. The choices organizations make will increasingly include considerations from compliance, legal, and policy stakeholders.

Cyber Resilience: Data Fusion and the Unstructured Challenge

Data protection matures into Data Fusion—the convergence of DLP, DSPM, and traditional backup into a unified data resilience strategy.

- **The Unstructured Frontier:** The battleground shifts to unstructured data, which AI models voraciously consume. Securing this data against ransomware and theft requires distinct strategies from structured data protection.

- **Recovery Assurance:** As ransomware persists, the line between primary storage and backup blurs. Resilience is no longer just about recovery; it is about ensuring data purity and availability for AI training and business continuity.

Strategic Imperatives: Quantum Readiness and Risk Quantification

Beyond immediate threats, strategic drivers are reshaping the C-level agenda.

- **The Quantum Scramble:** With the 2030 horizon approaching, the "quantum scramble" begins in earnest. Organizations must start updating infrastructure—from fleet management to printers—to support Post-Quantum Cryptography (PQC).

- **Cyber Risk Quantification (CRQ):** As government roles in cyber regulation increase, boards are demanding more rigorous CRQ to justify spend and measure exposure in financial terms.

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast
- Decision-Maker Survey - bi-annual IT Decision Maker survey
- Analyst Insight Report - a report on critical issues in the industry
- State of the Market Report - a report on technology, markets, products, and vendors
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Data Intelligence, Analytics & Infrastructure

Enterprise data is the lifeblood of business. No endeavor, whether a simple order-to-cash process or a complex, agentic AI solution, can survive without timely access to accurate, high-quality, secure, and governed data. As we move into 2026, the market is pivoting from experimentation to engineering, demanding a stack that is not just "AI-ready" but explicitly architected to accelerate AI.

This shift is reshaping the four pillars of Futurum's market coverage:

- **Data Foundation & Storage:** The market is witnessing the end of blind, deaf, and dumb storage. The convergence of database systems and the demand for high-throughput object storage are turning this layer into the high-performance memory tier required to feed hungry GPU clusters and agentic systems.
- **Data Processing & Orchestration:** The era of the monolithic platform is fracturing into a composable, open lakehouse architecture. Most notably, the semantic layer has graduated from a BI feature to a critical, standalone infrastructure component that translates business context for AI agents.
- **Data Analysis & Intelligence:** The data professional is evolving into an "AI Shepherd." Consequently, Futurum is tracking the reinvention of Business Intelligence into "Generative BI," where natural language replaces code, and the focus moves from building dashboards to managing the lifecycle of AI models and agents.
- **Data Management & Trust:** Governance is no longer a passive activity. It is now an active defense system. Futurum tracks how data observability and FinOps are together becoming a "radar" for the modern stack, enforcing accountability through data contracts to prevent model hallucinations and spiraling compute costs.

Futurum monitors these evolving dynamics across the entire lifecycle, analyzing everything from the raw physical storage of digital assets to the polished, agent-delivered insights that drive decision-making.

Key Issues for 2026

- General-purpose databases and data storage layers are commoditizing vector search capabilities, placing immense pressure on specialized vector database vendors.
- Metadata catalogs are evolving from static documentation repositories into active control planes that automate security and lifecycle policies.
- The semantic layer is decoupling from visualization tools to become a "headless" standard accessible by both BI dashboards and AI agents.

Go back to the first page

- Data Contracts are shifting quality accountability upstream, blocking schema changes in CI/CD pipelines before they break downstream products.
- Federated governance catalogs are finally turning the theoretical Data Mesh concept into a practical, operational reality for the enterprise.
- Data FinOps is emerging as a critical discipline for attributing and controlling the spiraling compute costs associated with agentic AI workloads.
- Zero-ETL architectures are becoming the preferred method for high-volume data sharing, reducing reliance on brittle replication pipelines.
- Knowledge graphs are resurging as an essential grounding truth infrastructure needed to prevent hallucinations in RAG architectures.
- Streaming data architectures are moving from niche use cases to the default standard for ingestion to support real-time AI context.
- Data clean rooms are proliferating as the primary mechanism for privacy-safe external collaboration and first-party data analysis.
- Text-to-SQL engines are reaching the maturity level required to make natural language the primary interface for complex data analysis.
- Sovereign AI clouds and localized infrastructure are rising to address strict data residency and privacy regulations globally.

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast
- Decision-Maker Survey - bi-annual IT Decision Maker survey
- Analyst Insight Report - a report on critical issues in the industry
- State of the Market Report- a report on technology, markets, products, and vendors
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Enterprise Software & Digital Workflows

Enterprise Applications are the lifeblood and framework for accomplishing work in the modern organization. We examine 12 categories of applications used in the enterprise, including Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), Workplace Collaboration, Human Resources, Supply Chain & Logistics, Analytics & Business Intelligence (BI), Project & Portfolio Management (PPM), Industry/Vertical-Specific Applications, and Communication Services, and delve into how they shape the broader enterprise information architecture. We also focus on the underlying technologies and systems that power these applications, including artificial intelligence and automation, and assess how trends in employee engagement and experience impact the market.

Key Issues for 2026

The Increasing Use of Agentic AI to Manage More Complex Workflows and Processes

In just over a year, agentic AI has evolved from a nascent technology with limited use cases and capabilities to a core technology embedded in a wide range of enterprise applications and platforms. And while agentic AI delivers results in relatively simple scenarios or tasks, the substantial ROI promised by vendors is unlikely to materialize until agentic technology can be applied across more complex workflows that incorporate near-real-time data, multi-step reasoning, and self-optimization capabilities. It is these more complex processes that consume significant time, effort, and resources to address and often have the most significant impact on customer, employee, and partner metrics, including experience, effort, and satisfaction, which directly affect a business's overall health and success:

- Vendors that can help their customers deploy agentic AI to address these complex scenarios likely will see the most success in monetizing agentic AI.
- Ensuring the accuracy and efficiency of these agents, as well as building trust among decision-makers, workers, and customers, will be top of mind throughout the year as the technology continues to mature.

Shifting Pricing and Business Models

As the types and complexities of AI workflows and use cases continue to expand, vendors are still struggling to effectively monetize AI. While the traditional, seat-license-based approach appeared to be on the way out in 2025, the challenges of generating a solid ROI from AI led to increased vendor flexibility, with some offering a choice of pricing models, ranging from seat-based to consumption-based to outcome-based. As technology improves and the diversity of use cases continues to expand, customers may adopt a variety of pricing approaches tailored to specific usage patterns and risk tolerances. Vendors that can provide this flexibility will be best positioned to attract new business across a broader range of scenarios. Notably, vendors that are able to successfully automate agentic workflows and tasks that are highly repeatable and scalable will likely shift to an outcome-based pricing model.

Go back to the first page

The Integration of Data and Applications Across the Technology Stack

Vendors will continue to support the integration of data from disparate apps and systems into the front end of customers' choosing, while also highlighting the benefits of a unified platform approach:

- Instead of being forced to work with data in a specific application, vendors are increasingly making it easy to pull in and manipulate data and initiate workflows from within the application of their choosing, a technique designed to make enterprise software adapt to the user, rather than the other way around, thereby driving more efficiency and productivity while reducing friction.

- The challenge will lie in strengthening and delivering the right message (the power of using a unified platform vs. application and data flexibility) to the right customer at the right time, in a way that does not dilute other messaging.

This will also lead to a growing convergence of disparate functional areas, such as contact center operations, customer service and support, marketing, sales, and fulfillment, into a more unified customer experience that delivers the right messaging, actions, and process flows across the entire customer journey, fed by a unified and real-time data-driven strategy.

The Era of SaaS Platforms as Orchestrators

As the era of agentic AI continues to evolve and mature, SaaS players have realized the value of not only managing their own data, AI agents, and workflows but also serving as an enterprise-wide orchestration layer capable of monitoring and managing third-party workflows and AI agents. These market participants are realizing that control over enterprise data and workflows – both AI- and human-augmented – drives platform utilization, revenue, and, perhaps most importantly, stickiness, which increases the likelihood of contract renewals and value expansion.

In 2026, expect to see both major SaaS platform vendors, as well as third-party integration vendors, consultants, and application-management platforms, enter the market and fight to control this important function. In fact, it is also likely that a new class of vendors will emerge as agnostic (or mostly agnostic) arbiters that serve as a master control plane for managing agents, humans, and workflows, regardless of the company or organization that built or provided the agents. These vendors may become increasingly important as agentic workflows span disparate systems, organizations, and jurisdictions.

A Renewed Focus on Contextual Assistance and Training

While agentic AI may become the front door for basic horizontal functionality, the predictions of the demise of SaaS applications are premature, due the large backlog of software implementations already on the books, the complexity of domain-specific workflows and processes, and the desire of many organizations to extract value from their existing technology investments.

As humans and AI increasingly work together in the enterprise application space, there will be a growing need for contextual assistance and training to ensure customers derive the maximum value from their software investments. The increasing use of AI and agentic AI technology portends the deployment of relevant, easily digestible, and context-based assistance and training features on top of or within enterprise applications.

The new and emerging generation of workers, who have little or no patience for reading documentation, combined with the rapid pace of innovation, has rendered obsolete most traditional learning and training resources. Organizations will fail to quickly realize value from their software investments if their workers and customers are unable to adopt and utilize these new learning tools and capabilities within the flow of normal work.

Go back to the first page

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast
- Decision-Maker Survey - bi-annual IT Decision Maker survey
- Analyst Insight Report - a report on critical issues in the industry
- State of the Market Report - a report on technology, markets, products, and vendors
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

GTM Ecosystems, Channels & Marketplaces

Across the technology landscape, vendors are increasingly leaning into their indirect go-to-market strategies and fostering more ecosystem partnerships. There are several drivers for this changing mindset. Economically, vendors are scrutinizing their cost structures more than ever as we have moved away from near-zero interest environments. Companies can no longer justify 'rampant hiring' in their sales & marketing divisions as a GTM tactic. Meanwhile, from a technology standpoint, customer IT environments are becoming more complex, spanning multiple clouds and architectures, while the application landscape is becoming increasingly customized. In short, no technology company can meet a customer's entire IT needs through its own portfolio; partnering is the only way to close the gaps.

Futurum will explore the discipline of partnering in the technology space amidst these trends and future disruptors such as AI. We will explore how GTM is evolving horizontally (e.g., partnering with other technology stacks) and vertically (e.g., embracing an ecosystem of partners engaged in product deployment and service offerings).

Key Issues for 2026

The Birth of AI Frontier Partners

- As enterprise demand for applied AI surges, partners with services capabilities will be expected to move beyond traditional advisory roles and deliver tangible, measurable value at the bleeding edge of technology adoption. The cutting-edge companies in this space will become AI Frontier Partners.
- AI Frontier Partners are those that demonstrate deep agentic capabilities, strong advisory services, data management and automation expertise, and ISV-like innovation.
- These partners aren't just integrating AI; they're inventing the next wave of business possibilities. From large language models to custom-built AI agents, they shape intelligent systems that propel organizations into the future. With a keen understanding of business challenges and strategic foresight, they steer organizations beyond theory to real-world transformation; bridging vision with actionable, AI-powered roadmaps.
- These partners don't settle for off-the-shelf. They design, build, and launch original solutions that raise the bar, continually pushing the boundaries with proprietary technology and fresh thinking. Their expertise extends to the very foundations of intelligent business, transforming messy data into a launch pad for seamless process automation and smarter decision-making.
- Partners that invest aggressively in AI talent, industry-specific solution accelerators, and proprietary methodologies will be able to co-create new business models with their clients. Those that remain anchored in legacy advisory or undifferentiated outsourcing will risk disintermediation as technology vendors, cloud marketplaces, and next-gen integrators move up the value chain.

Go back to the first page

The Rise of the Adaptive GTM Stack in Partner Functions Will Collapse the Distance Between Strategy and Execution

- Traditional operating rhythms of decide, plan, deploy, measure, adjust will be replaced by a fluid, continuous loop, powered by AI that dynamically interprets market signals and autonomously tests new approaches.
- Adaptive GTM stacks will leverage AI agents to realign pricing, ICP definitions, sales territories, channel incentives, and content sequencing in near real time, removing the lag between insight and action.
- The role of GTM leaders will shift from designing "perfect" annual plans to governing, coaching, and fine-tuning the systems and rules that intelligent agents use to make thousands of micro-decisions every day.
- Organizations that cling to static structures, annual planning cycles, and rigid partner tiers will fall behind as competitors iterate and respond to market changes instantly.
- The winners will be those who invest in AI-enabled feedback systems, dynamic data sources, and governance frameworks that empower continual learning and adaptation across the GTM organization.

Marketplaces Will Help Unlock Agentic Commerce

- Cloud and software marketplaces will move beyond being transactional hubs and play host to a new generation of AI agents that buy, sell, negotiate, and personalize on behalf of both vendors and customers.
- The emerging "agentic commerce" model will enable autonomous procurement, dynamic bundling, and personalized offers, making the buying experience both frictionless and highly adaptive.
- ISVs and partners will need to design offerings and operational processes with these digital agents in mind, ensuring compatibility with API-driven processes for negotiation, fulfillment, and support.
- Marketplace platforms will become critical arenas for experimentation, allowing companies to quickly test new product configurations, pricing models, and partnership combinations using real-time agent-driven feedback.
- Those able to build robust agent-to-agent commerce capabilities will access new revenue streams and unlock GTM opportunities not possible with manual, human-mediated sales alone.

Planned Deliverables

- Analyst Insight Report - an analyst report on critical issues facing partner leaders
- Partner Survey - bi-annual survey of IT partners assessing key issues they are facing and opportunities they are focused on
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Hybrid Cloud & Infrastructure

Data is the lifeblood of a modern business, and the underlying storage technology plays a vital role in delivering it. The storage industry is continually adding more data and application-aware capabilities and services to what was historically a box of storage devices with a network attached. The requirements for cloud applications, and now for AI applications, have blurred the line between storage and data, with more software-defined capabilities bringing innovations as fast as new hardware platforms and capabilities.

Key Issues for 2026

Storage Architectures for AI

Customers realize that AI is not a single workload with a well-defined set of storage requirements. The AI training pipeline has a set of requirements, and an AI-enabled business application will have different requirements:

- Scalability and data governance within data lakes and object stores will continue to be a focus; model training and fine-tuning will remain challenges for enterprise organizations.
- 2026 will see far more focus on the long-term viability of AI-enabled applications, with inference built into applications and delivered as agents.
- Vendors must clearly demonstrate their ability to scale KV-cache capacity and persistence for inference, either through local storage on GPU-enabled hosts or through centralized KV-cache stores.
- Customers also need to see how their storage can natively support RAG, generating vector embeddings autonomously as data is ingested.

Purchasing and Consumption Models

The model of outright storage purchase no longer dominates; more storage-as-a-service models are being embraced. In many cases, the adoption of STaaS was driven by customers' inability to predict capacity requirements. These same customers often lack the FinOps rigor to ensure business value from the as-a-Service model. Customers are also recognizing that these newer models do not have the forced decision point of array retirement, which traditionally triggered re-evaluation of suppliers. Vendors must demonstrate that the STaaS model is not a "golden handcuff" and that their TCO is lower than that of an outright purchase. The demands of AI systems are also driving requirements for hybrid- and multi-cloud data mobility from STaaS products.

Go back to the first page

Storage Cyber Resiliency

Unfortunately, ransomware has not died out; it has evolved to enable attackers to earn large sums of money. The ability to guarantee rapid and complete recovery from persistent ransomware tools is vital to customers. While data protection software vendors typically handle recovery orchestration, ransomware often targets backup copies before encrypting primary data. It falls to storage hardware vendors to provide truly immutable storage or offline media to guarantee recoverability. Customers who have been impacted by ransomware are usually very aware of the long restore time for offline media and the challenges of recovering data that may have been compromised gradually over time, requiring offline media from multiple points in time. Auditable recovery capabilities will become mandatory as business insurance, due diligence, and regulatory compliance drive technology adoption.

Hybrid and Multi-Cloud Data Management

Mature enterprise cloud adoption is characterised by the use of multiple clouds, both public and private, to achieve business outcomes. The resulting fragmentation of corporate data across multiple clouds impedes the creation of further value from the data estate:

- Customers need ways to access distributed data holistically, either through data copy mechanisms or remote data access acceleration.
- Capabilities to normalize data across platforms will be in demand, especially to enable AI adoption.
- Traditional data lake methodologies may not provide timely access to data and may limit where data is accessible while maintaining acceptable performance.
- European businesses have been demanding cloud and data sovereignty for some time; these demands will only increase as geopolitical tensions increase.

Emerging Disruptive Technologies

Disruptive technologies often seem to appear from nowhere when they are released, yet in any hardware design, the lead time from concept to production is usually measured in years. CXL has been on the brink of transforming hardware system design for a few years. Recent market moves suggest 2026 will be a big year for CXL products. Computational storage has failed to deliver on its promise as a general-purpose solution in the data center, but may find its place in edge devices, where hardware solutions are often customized to specific applications and deployed at large, distributed scale.

Planned Deliverables

- Analyst Insight Report - a report on critical issues in the industry
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Intelligent Devices

Futurum covers a broad range of AI-enabled consumer and commercial devices. This includes PCs and peripherals, tablets, mobile handsets, XR, hearables and wearables, as well as IoT, IIoT, automotive, and robotics segments. The expansion of AI training and inference from a cloud-centric model to a more hybrid edge-to-cloud model is driving a rapid transformation across the devices segment of the tech stack. With next-gen AI-capable PCs, mobile handsets, and wearables now capable of handling increasingly large AI models, AI and agentic workloads are beginning to expand from thermally expensive cloud-based silicon to the more thermally efficient silicon powering AI-capable consumer and commercial devices.

The next 6 to 12 months will see a significant acceleration in AI capabilities in devices, and will see the start of an expansion towards physical AI, which includes robotics. This transition will disrupt not only core device segments but the entire technology ecosystem around them as silicon vendors, cloud service providers (CSPs), independent software vendors (ISVs), and their partners adapt to new use cases, form factors, and hybrid, interwoven AI services models.

Additionally, the Intelligent Devices practice will also work with adjacent practices to 1) clearly map how AI orchestration will work across platforms and form factors up and down the entire technology stack, 2) validate the roadmaps of the most critical vendors in the accelerating robotics segment, and 3) quantify potential impacts of memory and storage supply chain constraints on key AI device segments.

Key Issues for 2026

PC Segment

- New generations of NPU-equipped AI PCs powered by entirely new AI-capable processors will continue to transition the PC segment towards a broader mix of cloud-based and on-device AI-enabled workloads.
- X86 architecture (Intel, AMD) will remain extremely competitive against Arm (Apple, Qualcomm) processors in the enterprise, with Intel's Core Ultra Series 3 processors applying most of the performance pressure for at least the first half of the year.
- While current AI PC processor vendors Qualcomm, Apple, AMD, and Intel will continue to compete for market share in the AI PC segment, NVIDIA will likely enter the market later this year. This could bring additional disruption and performance targets to the AI-accelerated PC segment reset.
- The proliferation of high-end, professional-grade AI desktop and deskside systems - dubbed desktop AI servers - will continue to drive innovation and push AI performance targets in the PC segment.

Go back to the first page

Mobile Segment

- Mobile chipset vendors, including Qualcomm, MediaTek, Google, and Apple, will continue to focus on introducing AI-enabled, multimodal, contextual, and agentic capabilities to mobile platforms.
- An additional focus for handset OEMs will be to leverage further hybrid multimodal AI capabilities (on-device and in the cloud) to deliver greater differentiation and more personalized AI-powered experiences for users.

XR Segment

- The next wave of AI-enabled capabilities in XR headsets will integrate natural language, computer vision, and location to bring more utility to AI glasses (formerly known as smart glasses) and increasingly turn them into "agentic glasses."

Other Device Segments

- Automotive: Multimodal and agentic AI capabilities will continue to advance the cockpit, infotainment, and ADAS capabilities of next-gen smart vehicles, making them smarter, more capable, and more secure.
- Robotics: As on-device AI capabilities continue to improve, 2026 will see a wave of investments and innovation towards disruptive new "physical AI" form factors and systems. While humanoid robots will be part of that focus, mobile, static, geofenced, robotic arms, and other robot form factors will begin to scale across a number of key segments like manufacturing, healthcare, logistics, and retail.
- Printers and workspace peripherals: Highly secure agentic AI is also beginning to enter the printer and workspace peripherals segments, driving both utility and value for users.
- Hearables and wearables: We are also seeing AI's impact on smartwatches, health wearables. and audio solutions. The next 6 to 12 months will bring significantly improved UI, predictive, and agentic user experiences to those product categories.

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast
- Decision-Maker Survey - bi-annual IT Decision Maker survey
- Analyst Insight Report - a report on critical issues in the industry
- State of the Market Report- a report on technology, markets, products, and vendors
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Networking

Enterprise networking is the transport system for advanced technologies on-premises and in the cloud. The landscape of data center networking has shifted in the past few years to focus less on cloud computing and direct Internet access, and now is primarily focused on providing high-speed interconnections for AI workloads. Bandwidth is increasing rapidly as AI models evolve and require more and more resources to execute in a reasonable amount of time. Innovation in the market must also embrace sustainability to ensure that development doesn't outpace the ability of modern data center infrastructure to provide power and cooling.

Key Issues for 2026

East-West Traffic Fabrics

Networking architecture has historically been optimized to serve applications or direct users out of the network toward the Internet. This North-South traffic flow has been disrupted by the needs of AI clusters. Traditional Clos architecture with leaf/spine connectivity cannot keep up with GPUs that frequently exchange large data sets.

- Traditional leaf/spine design is being supplanted by purpose-built fabrics. This could include standards like Ultra Ethernet or manufacturer-specific protocols.
- Operations teams must now manage two distinct network architectures to deliver data to the AI cluster and also to exchange data between cluster nodes.

Beyond 800G Ethernet

The need for increased performance has forced manufacturers to bring faster connectivity to the market. 2026 will see 800G Ethernet become the standard for new deployments, and early adopters will begin to look past the terabit Ethernet mark toward 1.6T Ethernet options.

- Data center rack designs will need to accommodate faster modules while also ensuring signal integrity. Any imperfection at these speeds will lead to data loss and reduced performance across the network.
- Current optical technology is capable of reaching necessary speeds, but the power consumption and heat generation are also increasing rapidly. Manufacturers will need to investigate technologies like co-packaged optics or linear pluggable optics (LPO) to reduce environmental impacts, while newer technologies like silicon photonics are still in the development stages.

Go back to the first page

Rethinking Cooling and Power Needs

800G Ethernet optics and switches with high port counts are increasing the amount of power that each device draws from the available power budget for the rack. Additionally, the mass of cabling needed to interconnect the various networks for the rack is impeding air flow and creating issues for traditional air cooling of these already-hot devices.

- Optical technology currently draws a significant amount of power to produce high performance. Manufacturers are researching methods to reduce power consumption, but will need time to implement them into new device architectures.

- Cooling for networking devices will follow a similar path to high-density server technology. Liquid or immersion cooling, as well as cold plates for hot ASICs, are options that must be considered to avoid stressing existing cooling infrastructure.

Planned Deliverables

- Analyst Insight Report - a report on critical issues in the industry

- State of the Market Report- a report on technology, markets, products, and vendors

- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Observability

In 2026, observability shifts from providing operational visibility into systems toward narrowing the trust gap introduced by non-deterministic AI and autonomous agents. As agents plan, decide, and execute work across the SDLC, observability becomes the mechanism that makes behavior understandable, governable, and safe at scale. This moves observability upstream from post-incident analysis into execution, control, and management, where trust is established through evidence, not assumption. The following key issues define how observability platforms must evolve to support agent-driven systems without sacrificing accountability, control, or enterprise confidence.

Key Issues for 2026

Embedded Observability-native Execution

AI-driven execution is continuous and non-deterministic, making post-execution explanation insufficient for governance or control. Observability must exist at the moment work is performed, embedded directly into AI and agent execution across the SDLC rather than added as a downstream analysis layer.

- Treat observability as a downstream analysis layer, or embed it directly into AI and agent execution so behavior is visible as work occurs across the SDLC.
- Over the next year, differentiation concentrates around platforms that expose execution-time visibility into agent intent, decisions, and outcomes and carry that visibility across development, pipelines, runtime, and operations. Embedded observability is required to operate systems that execute continuously and non-deterministically.
- Platforms that rely on post hoc reconstruction from infrastructure telemetry lose control as AI-driven execution accelerates and are bypassed as buyers consolidate around execution-native platforms.

Observability as the Control Surface for Non-deterministic AI

Non-deterministic AI challenges assumptions that systems can be trusted through predictability or replay alone. Observability becomes the control surface that makes probabilistic execution explainable, bounded, and governable in production systems.

- Accept non-determinism as opaque risk, or use observability as the control surface that makes probabilistic execution explainable and governable.
- Over the next year, differentiation concentrates around platforms that capture decision context, constraints, and intent as part of execution. Observability turns non-determinism into an observable system property rather than an unmanaged liability.

Go back to the first page

- Platforms that cannot close this gap force customers to restrict AI autonomy to preserve trust and are excluded from regulated and production-critical environments.

Observability as the Agent OS Management Surface

Agentic systems shift operational control away from managing system health toward managing behavior at scale. Observability becomes the management surface through which agent state, coordination, and impact are directed in real time, forming a core capability of an emerging agent OS.

- Agent behavior is managed, not merely observed. Observability provides real-time visibility into agent intent, state, coordination, and impact, and uses that visibility to direct execution as it occurs. This shifts observability from reporting outcomes to actively managing autonomous behavior.

- Control operates through behavioral signals. Operational control is exercised through signals that govern how agents plan, escalate, and interact with systems and each other. Observability supplies the inputs required to throttle execution, enforce boundaries, and intervene before actions compound into systemic risk.

- Management moves upstream from operations to execution. Observability functions at the same layer as planning and decision-making rather than after deployment. Platforms that confine observability to post-execution analysis lack a viable way to manage agentic systems at scale and force customers to limit autonomy to retain control.

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast*
- Decision-Maker Survey - bi-annual IT Decision Maker survey*
- Analyst Insight Reports - a report on critical issues in the industry
- State of the Market Report - a report on technology, markets, products, and vendors*
- Futurum Signal Report on Observability

* Timing is expected to be aligned with Software Lifecycle Engineering

Go back to the first page

Semiconductors, Supply Chain & Emerging Tech

Semiconductors have created a technology super-cycle powering the AI revolution. The industry is on pace to approach $1 trillion in revenue in 2026, marking a third consecutive year of elevated growth driven by AI training, inference, and new classes of intelligent systems. The semiconductor industry now spans a deeply interdependent global supply chain where constraints at any layer shape overall performance and economics. Beyond traditional data center compute, emerging technologies are expanding the market through new computing form factors that depend on breakthrough semiconductor innovation and frontier AI models, including intelligent robotics, domain-specific XPUs, and early hybrid classical-quantum platforms. Together, these forces are shifting the industry from a focus on standalone chips toward tightly integrated, system-level platforms that define the next phase of AI-driven growth.

Key Issues for 2026

Supply Chain Bottlenecks

- As transistor scaling slows, the key bottlenecks in semiconductors have shifted from the wafer to the system.
- Performance and economics are now constrained by advanced packaging capacity, memory supply, power delivery, and thermal management, as rising chip- and rack-level power density turns efficiency and cooling into first-order design constraints.

Responding to Energy Scarcity with Token Efficiency

- Energy availability has replaced silicon supply as the primary constraint on AI expansion, forcing a shift in the industry's North Star metric.
- As global data center demand doubles and the power gap between grid capacity and cluster requirements widens, the success of an AI deployment is no longer measured by peak FLOPS, but by its tokens per dollar per watt. This is driving a move toward power-capped computing, where software orchestration and hardware efficiency are tuned to ensure that fixed power envelopes produce the maximum possible revenue-generating intelligence.

Storage Hierarchy for Long-Context Reasoning

As AI shifts from chatbots to agents capable of multi-step reasoning, the storage of inference context has become a first-order design challenge. Long-context windows generate massive Key-Value (KV) caches that quickly exhaust expensive HBM on GPUs. In 2026, leading AI clusters will utilize a tightly integrated hierarchy of HBM, SSD, and HDD to optimize the cost and power profile of every token generated, creating a tiered storage strategy where the speed of context retrieval defines the practical utility of frontier models.

Go back to the first page

Foundry Diversification

- Competition for semiconductor fabrication will intensify in 2026 across memory, advanced logic, and packaging as AI-driven demand pushes the industry to its physical and organizational limits.

- Demand for logic, memory, and packaging is forcing tighter coordination between foundries, memory vendors, OSATs, and system integrators. It is elevating fabrication access itself into a strategic differentiator, where early commitments and long-term capital partnerships increasingly determine who can scale AI systems in 2026 and beyond.

Innovating beyond GPUs

XPUs, robotics, and quantum innovation increasingly drive progress in 2026:

- XPUs are evolving into workload-specific AI platforms that combine heterogeneous compute engines, tightly coupled memory hierarchies, and chiplet-based designs to improve price-performance.

- Robotics processors push this model further by integrating AI inference, real-time control, sensor fusion, and functional safety onto single packages for edge inference.

- In parallel, quantum processors are advancing through hybrid systems, where control electronics, cryogenic interfaces, error mitigation, and software orchestration determine practical utility.

Government-spurred Acceleration

In 2026, governments are no longer passive observers of the semiconductor industry; they are active participants shaping its long-term direction. Rising geopolitical tensions and recent supply-chain shocks have reframed semiconductors as core national security infrastructure, on par with energy and defense.

As a result, policy is moving beyond broad trade restrictions toward direct intervention across the value chain. Governments are actively redesigning domestic chip industries to reduce exposure to globalized risks. This intervention shows up in partnership formation, trade policy, and infrastructure investment. The net effect is that governments will accelerate semiconductor innovation while expanding the total addressable market.

Planned Deliverables

- Market Data - bi-annual market sizing & five-year forecast
- Decision-Maker Survey - bi-annual IT Decision Maker survey
- Analyst Insight Report - a report on critical issues in the industry
- State of the Market Report - a report on technology, markets, products, and vendors
- Futurum Signal Report Access - any Signal Report published in the relevant practice area

Go back to the first page

Software Lifecycle Engineering

The Software Lifecycle Engineering market is moving decisively from AI experimentation to AI accountability across the SDLC. In 2026, enterprises will be required to demonstrate AI-driven business value, operational impact, and measurable risk reduction in development, not just incremental developer productivity gains. Vendors that cannot connect AI investment to durable outcomes will face growing scrutiny from customers, buyers, and boards.

At the same time, the industry is racing to industrialize AI systems capable of meeting enterprise expectations. Vendors are assembling a new agent software stack for AI, agents, workflows, management, and infrastructure, but most stacks remain incomplete. Prompts, LLM modes, and agent builders alone do not produce production-ready systems. The hard work now lies in designing AI-native lifecycle platforms that embed agent identity, control planes, behavioral governance, security guardrails, testing, operational management, observability, and end-to-end lifecycle control. Decisions made here will either enable enterprise-scale agent adoption or quietly constrain it.

These are not abstract platform choices. They are commitments that shape how vendors earn trust, scale deployments, and remain relevant as buyers consolidate around fewer, AI-native lifecycle platforms.

Key Issues for 2026

Emergence of a New AI-centric Development Model

The next 6–12 months will determine whether software development fully transitions from code-centric execution to intent-directed systems built around AI and agents:

- Vendors must decide whether to preserve familiar developer patterns and workflows or introduce new AI-centered primitives for planning, delegation, and execution that enable AI to perform end-to-end collaborative work.
- The winners will be those who successfully support human-AI collaboration across code, specs, security, and runtime behavior, and who provide customers with a clear strategy and roadmap for moving to new AI-centric development models, patterns, and architectures.

Vendors that fail to make this transition explicit risk locking customers into architectures that cannot scale agent-centered systems, capping the impact of AI, and undermining long-term value.

AI Shifting from Feature to Core Capability Across the SDLC

Vendors across application development, testing, security, operations, and platform engineering face a structural choice:

- Layer AI into existing tools or design solutions so AI becomes a first-class execution capability across their portions of the lifecycle.

Go back to the first page

- Over the next year, differentiation will move toward shared context, persistent AI state, and cross-stage decision-making that allows AI systems to plan, act, and adapt across development, platform engineering, systems reliability engineering, and operations.

- Platforms that fail to embed AI into how work flows, policies are enforced, and systems are operated across the SDLC will lose relevance as buyers consolidate around fewer, AI-native lifecycle platforms.

Once buyers standardize on AI-native platforms, late rearchitecture will be exponentially more complicated, with less market impact.

The New AI Stack Emerges

In 2026, AI-centered development will consolidate into a defined software stack that underpins both software engineering and AI applications. This stack establishes how intent is captured, how work is delegated to agents, and how execution is governed across the SDLC, infrastructure, DevOps, platform engineering, and operations.

We will see the emergence of agent-native infrastructure, the functional equivalent of what Kubernetes brought to the microservices era, built around agent control planes, orchestration, and scalable agent operations. At its core, an agent control plane manages identity, permissions, memory, lifecycle, policy, and observability. Orchestration layers coordinate specialized agents across planning, building, testing, security, deployment, and operations, enabling parallel, asynchronous, interdependent, and long-running execution with human oversight. As this AI stack emerges and matures, existing CI/CD, workflow automation, and policy enforcement layers that cannot operate at agent speed or scale will be bypassed or absorbed.

Establishing Agent Control Planes and Harnesses

The next 6–12 months will set the foundation for how agent environments are controlled, observed, and trusted in production. Vendors are racing to define control planes that provide agent identity, scoped authority, behavioral constraints, and real-time observability integrated directly into development and operational platforms. Those that succeed will make agent-based software viable at enterprise scale, while those that treat governance, security, and observability as add-ons risk being locked out of production deployments.

Agent control models adopted now will determine whether agents are viewed as manageable systems or as ungovernable risks.

AI Open Standards: Control, Interoperability, and Trust

The next 12 months will determine where agent-based ecosystems fragment or converge. While existing open standards such as MCP and A2A continue to mature, new open standards and open-source efforts will emerge to address gaps and interoperability challenges in agent harnesses, control planes, policy enforcement, memory management, security, behavior, and more. Vendors must make key decisions about where they lead, where they contribute, and where they align with existing standards.

Misjudging this balance risks either ecosystem isolation or loss of strategic control over core platform value.

Go back to the first page

Hybrid Development: Parallel AI and Traditional Software Pipelines

The near-term challenge is not only deciding how much work to give AI, but how to run parallel development models without breaking delivery. Over the next 6–12 months, the key innovation will be:

- Platforms that allow AI-centric workflows and traditional software development to operate side by side, sharing artifacts, context, and governance while moving at different cadences.
- Deployment pipelines and patterns will evolve to support inference-heavy workloads, AI models, and traditional software stacks. We will see both specialization in operational infrastructure and offerings designed to support a wide range of workloads. Removing friction in the deployment and infrastructure is of high importance for getting these workloads into production.
- Vendors that reconcile AI-generated plans, code, tests, and operational changes with existing pipelines, review processes, and ownership models will enable adoption without disruption.

Platforms that force customers to choose between AI-native and traditional workflows risk slower adoption and increased organizational friction.

Planned Deliverables

- Analyst Insight Reports - a report on critical issues in the industry
- State of the Market Report - a report on technology, markets, products, and vendors
- Futurum Signal Report on Software Development Platforms and Agent Orchestration

Go back to the first page