

Futurum

Futurum Research 2025

Key Issues & Predictions

Last updated on March 07, 2025

The Futurum Group All rights reserved 2024

Welcome to Futurum's 2025 Key Issues & Predictions Report

As we stand on the edge of 2025, I'm reminded of one fundamental truth: disruption waits for no one. Organizations and leaders who can anticipate, adapt, and act quickly will thrive. At The Futurum Group, we are focused on helping you decode the complexities of today's digital-first world so you can stay ahead of the curve.

This year's predictions dive deep into the trends reshaping industries—from AI's pervasive influence across enterprise applications to seismic shifts in hardware, cloud, and customer experience. Our team of analysts—some of the brightest minds in research and strategy—has dissected the forces driving change and outlined actionable insights for what's next.

Whether it's the rise of agentic AI disrupting software consumption, cloud marketplaces revolutionizing GTM strategies, or the accelerating impact of AI PCs on productivity, the message is clear: we are entering a new era where business agility, intelligence, and innovation are the ultimate differentiators.

The insights shared here are not just about spotting trends but about preparing for transformation. As customer expectations soar and competition intensifies, companies that embrace change as an opportunity rather than a challenge will set the pace for the next decade.

I invite you to explore the predictions in this report and reflect on how your organization can harness these shifts to drive growth, improve outcomes, and elevate experiences for customers and employees alike.

Here's to meeting the future head-on—together.



Tiffani Bova
Chief Strategy and Research Officer
The Futurum Group





AI Software & Tools: Agentic AI Disrupt the Business Application Universe

Prediction: By the end of 2025, we will see a significant shift in how enterprise software is consumed. At least 30% of routine business software interactions will be mediated through AI agents rather than direct user interfaces, leading to a fundamental restructuring of software licensing models.

Why This is Trending:

The convergence of three key factors accelerates this trend.

- First, recent breakthroughs in large language models have dramatically improved **agents' ability to understand context and execute complex instructions** across multiple systems.
- Second, the widespread adoption of APIs and standardized integration protocols has made it **technically feasible for agents to interact with diverse software systems**.
- Third, the **increasing complexity of enterprise software stacks** - most large companies have multiple hundreds of SaaS applications - has created an urgent demand for solutions that can **abstract away this complexity** for end users while maintaining operational efficiency.

Use cases:

- **Cross-platform business process automation:** An AI agent could manage an entire employee onboarding process by automatically coordinating across HR systems, IT provisioning platforms, and training modules. It could handle everything from creating accounts and setting up permissions to scheduling orientation sessions and ensuring compliance requirements are met, all without manual intervention across multiple systems.
- **Intelligent resource optimization:** An agent could continuously monitor and manage enterprise resource planning (ERP) systems, procurement platforms, and inventory management software simultaneously, making real-time decisions about stock levels, supplier orders, and logistics planning while reconciling data across all systems to maintain optimal operations.
- **Integrated customer experience management:** Rather than requiring separate teams to monitor different customer touchpoints, an AI agent could simultaneously interact with CRM systems, support ticketing platforms, social media management tools, and email marketing software to provide unified, context-aware customer service and relationship management, automatically escalating issues and coordinating responses across all channels.

"The rise of agentic AI represents a significant shift in enterprise software. Instead of employees juggling dozens of different applications and interfaces, they'll simply tell AI agents what they need to do - onboarding a new hire or reconciling financial data across systems - and the agents will handle the complex coordination behind the scenes. This isn't just about automation; it's about fundamentally changing how businesses - and the people within them - interact with their software systems, potentially saving billions in training costs and dramatically reducing the cognitive load on workers to free them up for more productive and creative tasks."



Nick Patience

Vice President & Practice Lead AI Software & Tools

[Request Analyst Session](#)



NPU-Equipped AI PCs Will Upend the PC Market

Prediction: AI-capable PCs (PCs equipped with an NPU and capable of running some AI training and inference workloads locally) will come to represent at least 40% of new PC shipments by the end of 2025.

Why This is Trending:

Three primary reasons are driving this change.

- **NPU for PCs:** The introduction of NPUs into device system architectures, which includes PCs, is enabling devices to perform previously energy-intensive tasks far more efficiently than they could with traditional CPUs and GPUs. This new capability unlocks next-gen AI training and inference capabilities directly on the device, which in turn creates entirely new horizons of added utility for users and their organizations. NPU-equipped PCs also happen to deliver vastly superior performance per watt to their predecessors, translating into all-day (and even multi-day) battery life to users.
- **OEM Commitment to the Transition.** Every major PC OEM is fully committed to this market transition, with aggressive competition between silicon vendors Qualcomm, AMD, and Intel accelerating performance improvements at both the processor and system levels. NVIDIA is also rumored to enter the market within 6-12 months. The PC ecosystem is moving forward, not backwards. AI PCs are already beginning to replace soon-to-be-obsolete pre-AI PCs.
- **PC Refresh Cycle.** The end of support for Windows 10 (slated for October 2025) will also help drive the PC refresh cycle towards AI PCs and accelerate the adoption of AI PCs in the commercial segment.

Use cases:

As AI-capable PCs are an evolution of pre-AI PCs, all previous use cases for PCs still apply. However, new use cases have already begun and will continue to emerge.

Moving AI Processing from the Cloud to Devices. As large language models and large mixed models (multimodal AI) become more efficient, they are able to move from the cloud to devices. For instance, many of the large language models trained in the cloud a year ago can now be trained directly on-device today. This means that, as that trend continues, organizations will be able to train and test many of their models securely, onsite and at a fraction of the cost they would have otherwise incurred, directly on PCs.

Agentic AI in the PC. As agentic AI begins to transform the way users interface with apps and software, AI-capable PCs will be uniquely positioned to deliver agentic-AI forward experiences for users, most of which are expected to save them time and significantly increase their productivity. Examples of this range from drafting emails responses and managing calendars in seconds to reducing the time it takes to design a presentation or a report from hours to minutes.

All Day & Multi-Day Battery Life. PCs capable of delivering all-day and multi-day battery life even in thin-lightweight form factors will also transform the way users work and play with their PCs, not only in hybrid and remote work scenarios but at the office as well, with notebook PCs becoming far easier to carry around between meetings.

"The AI PC is, first and foremost, a radically better PC than pre-AI PCs. It is tangibly faster, more powerful, more capable and more useful. The all-day battery life alone is such a radical system improvement that even without its AI capabilities, it would be worth the upgrade. But perhaps more importantly in the long term, the AI PC also lays the necessary foundation for the next generation of software experience, which will be dominated by agentic AI. As agentic AI begins to insert itself into every application, from search, system management and security to productivity and creativity software, users in both the consumer and the commercial segments will need PCs designed to handle agentic AI workloads locally in order to take full advantage of the coming disruption/opportunity."



Olivier Blanchard

Research Director & Practice Lead AI Devices

[Request Analyst Session](#)



Channel & GTM: The Use of Cloud Marketplaces Surges Forward

Prediction: Cloud marketplaces will become as big a Go-to-Market (GTM) for Independent Software Vendors (ISVs) as traditional distribution is for commercial hardware. Nearly \$400 billion of committed cloud spending will continue to help fuel this engine.

Why This is Trending:

- **Marketplace fees have gradually been coming down.** When marketplaces first came on the scene, fees were north of 20%, making them a very expensive proposition. Now, they are at about 3% as standard, and in some cases, as low as 1.5%. At this price point, marketplaces are as cost competitive as traditional distribution channels, and leave more room in the margin stack for ecosystem partners to take part. This fee reduction comes as a result of increased volume in marketplace activity, as well as the underlying goal that hyperscalers have: driving more infrastructure consumption.
- **Cloud commits across the major hyperscalers continue to surge.** As the cloud becomes increasingly pivotal to enterprises the world over, companies are increasingly entering into long-term contracts with hyperscalers that ensures they have the best pricing, and a guarantee of resource availability. As of Q3 2024, cloud committed spending across the leading three hyperscalers surged to \$393 billion (representing nearly 30% growth year-on-year). Certain portions of this commitment can be utilized on third-party products on the cloud marketplace, leading to a ready-made marketplace economy for ISVs to tap into.
- **The hyperscalers have all launched programs that allow their partners to participate in the cloud marketplace.** These 'Private Offer' programs enable partners to create custom offers for their customers via the marketplace. This could include pricing, bundling as well as their own value-add services. These programs ensure that ISVs that want to participate in the cloud marketplace can still leverage their partner ecosystem, and crucially, reward them for that activity via their own partner programs. Partners will play an integral role in cloud marketplaces. AWS, which has the most mature program, has indicated that north of 30% of marketplace transactions already feature a partner as the selling agent. This number will continue to grow.

Use cases:

- **CrowdStrike** is one of a handful of companies that has surpassed \$1 billion in total sales in the AWS Marketplace. Since launching on AWS in 2017, it has been its fastest growing route-to-market, and is also responsible for delivering a higher-than-average deal size (compared to its other sales channels). CrowdStrike has over 20 integrations with AWS products including AWS Control Tower and AWS GuardDuty.
- **NetApp** recently launched NetApp Data Infrastructure Insights on the Azure Marketplace to help customers planning an Azure migration with streamlined observability and real-time telemetry data. NetApp has leaned heavily into its hyperscaler GTM strategy. In addition to offering products on the marketplaces, it is the first vendor to offer first-party services with all three leading hyperscalers.
- **Salesforce** and AWS announced a wide-reaching strategic partnership agreement in 2023. One aspect of the agreement was the availability of select Salesforce products in the AWS marketplace for the first time, including Data Cloud, Service Cloud, Sales Cloud, Industry Clouds, Tableau, MuleSoft, Platform, and Heroku. In its Q3 2024 earnings, Salesforce highlighted AWS as a key growth driver with transactions doubling quarter-over-quarter and 10 deals exceeding \$1 million in sales.



Alex Smith

Vice President & Practice Lead Channels & GTM

[Request Analyst Session](#)

"Every vendor is trying to figure out their marketplace strategy, which ones to prioritize, how to operationalize it, and how to bring traditional partners on that journey. The most successful vendors in the marketplace will be the ones that understand how to include service delivery partners as part of their marketplace strategy."



CIO: Confronted with the Realities of Strategically Operationalizing AI, IT Leaders Will Rethink How They Wield the Cloud

Prediction: Enterprise IT will undergo a dramatic transformation in 2025, as CIOs strategically re-architect their cloud infrastructure to meet the demands of AI-driven workloads. According to Futurum's latest CIO Insights survey, 89% of CIOs report leveraging AI for strategic improvements, with 71% reevaluating the optimal environments for running cloud workloads.

Why This is Trending:

Three primary reasons are driving this change.

- **Legacy Architectures Fall Behind.** The rise of generative AI and large-scale machine learning models has introduced unprecedented compute and storage requirements that legacy architectures cannot support.
- **Balancing Cloud Deployments.** Organizations are seeking to balance the flexibility of public clouds with the control and cost predictability of private or hybrid cloud environments.
- **Compliance & Data Sovereignty.** Increased awareness of data sovereignty and compliance needs is driving CIOs to redesign their cloud strategies with AI in mind.

Use cases:

- **AI-Optimized Data Centers:** Enterprises deploying on-premises GPU-based architectures to support cost-effective training and inferencing workloads while also maintaining data control.
- **Making AI Affordable:** Making AI workloads inexpensive enough to operate, especially for complex knowledge work in price-sensitive industries like healthcare and insurance, to achieve ROI requirements.
- **Dynamic Cloud Bursting:** Leveraging hybrid cloud environments to seamlessly scale AI workloads to public clouds during peak demands.
- **AI-Enhanced Business Resiliency and Disaster Recovery:** Implementing AI-driven predictive analytics to optimize failure-resistance, operational failover and recovery processes across multi-cloud architectures.

"As AI becomes integral to business strategy, CIOs are being forced to reconsider how and where it's optimal to deploy compute resources. The need for low latency, cost efficiency, and compliance in AI applications is driving a rapid shift toward hybrid and multi-cloud strategies. For IT leaders, this means 2025 will be a pivotal year for a comprehensive realignment of their infrastructure with the realities of the AI era."



Dion Hinchcliffe

Vice President & Practice Lead CIO Insights

[Request Analyst Session](#)



Cybersecurity: Artificial Intelligence (AI) Accelerates the Race between Attackers and Defenders

Prediction: Generative AI will become an increasingly important differentiator across the cybersecurity toolchain. In fact, more than 40% of cybersecurity decisionmakers noted the integration of new technologies such as generative AI as a top spending priority relating to cybersecurity, and 39% note new AI-driven security platforms, in The Futurum Group's Cybersecurity Decision Maker IQ data.

Why This is Trending:

Three primary reasons are driving this change.

- **Speed & Efficiency of Attacks.** Malicious actors are using AI to develop more evasive and sophisticated threats with greater speed and efficiency.
- **New Response Tools.** In response, defenders must also adopt AI-powered tools to reduce the likelihood of a successful breach and mitigate the resulting damage in the event of a successful breach.
- **AI-Powered Cyber Solutions.** Vendors are responding in kind by baking AI-based capabilities into their solutions to support security and infrastructure teams alike from the standpoint of their cyber-resiliency.

Use cases:

- **Automated Threat Detection:** AI can be used to analyze vast amounts of data, such as logs, and to draw correlations between multiple disparate data stores. As a result, it can uncover patterns and anomalies, such as unusual network traffic or suspicious user behavior, that may indicate a potential threat - and that would be difficult or impossible to uncover without AI.
- **Rapid Incident Response:** Generative AI can automate incident response tasks, such as identifying affected systems and recommending recovery points and remediation steps, as a result significantly reducing response time.
- **Enhanced Security Awareness Training:** AI can create realistic phishing simulations, tailored to specific roles or even individual users, that can be used from a training perspective to improve security awareness and diligence amongst employees.

"As attackers use AI to elevate their game, organizations must also start evaluating AI as an important tool to fight back. Cybersecurity vendors are responding by steadily integrating generative AI into their solutions for a myriad of outcomes. These include enhancing vulnerability and threat detection and automating routine tasks, in order to reduce the likelihood of a successful attack, and the resulting impact of successful breaches. This approach will only increase in its criticality to helping organizations to protect their most important data assets and optimizing the resilience of their most critical business services."



Krista Case

Research Director & Practice Lead Cybersecurity

[Request Analyst Session](#)



Data Management & Analytics:

Prediction: 2025 will be the year of metadata for enterprises seeking faster and easier access to the data necessary to fuel impactful business and AI outcomes. As the year progresses, companies will increasingly set aside the complexities of data fabrics and data meshes to instead focus on metadata to deliver data- and AI-borne business value. Employing the philosophy of a meta grid (a concept originally defined by Ole Olesen-Bagneux), companies will seek to unify disparate databases, warehouses, lakes, lakehouses, and even knowledge management systems. At its core, a meta grid functions as a decentralized architecture comprising metadata-driven data catalogs. When used together, these islands of knowledge can be used to sidestep a centralized, company-wide master view of data and instead surface data and data-driven insights where and when needed with little to no IT friction.

Why This is Trending:

Pragmatically, this trend recognizes that a centralized, company-wide master view of data is no longer a realistic or necessary goal. In today's decentralized, hybrid/multi-cloud data landscape, data is often scattered across multiple systems, formats, platforms, and departments, making it nearly impossible to create a single, unified view.

A meta grid approach acknowledges this reality and instead focuses on providing a more nuanced and contextual understanding of that data, surfacing insights where and when needed, with minimal IT involvement. This approach enables businesses to tap into the power of their data without getting bogged down in complex integration and governance efforts.

More strategically, the idea of a meta grid also has the potential to truly democratize access to data and AI-driven insights across the organization. By providing a more flexible and agile method by which business users and data professionals can access and utilize contextual, timely and accurate data, a meta grid can help companies make more informed decisions faster and with greater confidence.

Use cases:

- **Data quality and validation.** A meta grid can automatically validate data quality, comparing data against predefined standards and recommending improvements to data pipelines while also identifying data inconsistencies and anomalies in near real-time.
- **Observability and the examination of the entire data landscape.** A meta grid can quickly and continuously examine an entire portfolio of data assets to identify redundancies, quality issues, security vulnerabilities, non-compliant data practices, etc.
- **Breaking down data silos.** At a high level, a meta grid can bring together the disconnected domains of knowledge spanning information management, data management, IT management, and knowledge management. Tactically, this translates into fewer routine data migrations and consolidations, thereby minimizing repeated and wasteful data movements while also enabling more effective data sharing and collaboration.
- **Refactoring and modernizing data architectures.** A meta grid can perform basic data improvements, such as optimizing data storage and retrieval, to more complex tasks, such as examining an organization's data architecture and developing multiple approaches to refactoring legacy data systems into modern, cloud-based data platforms.



Brad Shimmin

Vice President & Practice Lead, Data Management & Analytics

[Request Analyst Session](#)

"In the race to deliver timely, accurate, and contextual insights, many organizations have gotten bogged down in the act of building out the perfect data stack and data landscape, one that enforces consistency (e.g., a data fabric) while also encouraging collaboration and self-reliance (e.g., a data mesh). As a result, in 2025, we expect to see companies seek out a third, middle way, with metadata catalogs serving as a unifying principle across disparate data and AI assets."



DevOps and Application Development: GenAI's Biggest Gains Are In Old & New Code

Prediction: By the end of 2025, generative AI will be valued just as much, if not more, for its ability to reduce developer toil and tackle technical debt compared to its ability to generate new code.

Why This is Trending:

Three primary reasons are driving this change.

- While generating software code from a natural language prompt garners great interest, generative AI is proving to be particularly well suited to software development tasks that require an understanding of software's inner workings.
- Generative AI's strengths in natural language processing and inferencing are highly applicable to software, which is written in symbolic language such as programming languages, scripts, configuration files and markup languages.
- AI agents are conducive to tasks including writing unit and functional tests, performing code and security reviews, and performing routine updates and bug fixes. More complex tasks, such as refactoring and modernizing existing code, require a fuller and more comprehensive examination of the inner workings of software.

Use cases:

- **Streamlining Development Tasks.** AI agents simplify the creation of software and technical documentation, including API documentation, architectural diagrams, readme and configuration files, and software unit and function software test cases.
- **Code Reviews and Improvement.** AI agents can automatically perform code reviews, comparing uses of software patterns and adherence to modern programming standards while also recommending improvements to code.
- **Examining Full Codebases.** AI agents can quickly and continuously examine entire codebases to identify reusable code, reoccurring software bugs, security vulnerabilities, non-compliant coding practices and code that is no longer used.
- **Tackling Technical Debt.** AI agents trained in upgrading and integration testing software packages and libraries can increasingly perform routine software upgrades.
- **Refactoring and Modernizing Code.** AI agents can perform basic code improvements, such as removing inefficiencies in code, to more complex tasks, such as examining an application codebase and developing multiple approaches to refactoring legacy code into microservices.

“Generative AI is not only revolutionizing new software development, it is also fundamentally changing how software developers work with existing codebases. Software development relies upon the human ability to understand the interdependencies and inner workings of software codebases, a natural use case for generative AI. Agentic AI will have a profound impact at tackling the mountain of technical debt and reducing the toil developers deal with as part of their everyday work.”



Mitch Ashley

Vice President & Practice Lead DevOps
& Application Development

[Request Analyst Session](#)



Enterprise Applications: Pricing Model Shifts

Prediction: Generative AI-powered features will enter widespread use in 2025, thereby requiring significant shifts in pricing models, with seat-license models being supplanted by consumption-based and outcomes-based approaches. A 2024 Futurum Intelligence survey of 895 decision makers and influencers found that 40% of respondents were paying for software on a consumption-based pricing model, and 15% were using an outcomes-based model.

Why This is Trending:

Three primary reasons are driving this change.

- Both vendors and enterprise customers are realizing that AI is enabling work to be completed more quickly and efficiently than ever before, and paying for a full-seat license is both inefficient and also lacks a direct connection to business results.
- Consumption-based models more closely track usage of AI services to cost, whereas outcomes-based models ensure that customers do not pay for software that is not delivering promised results.
- As AI agents proliferate, we expect a strong shift to outcomes-based pricing models in 2025.

Use cases:

There are several examples of these new pricing models being used today:

- **Adobe** offers generative credits or tokens, which are correlated with consumption of compute used when images or videos are generated, as a way for customers to understand the value and costs associated with its powerful text-to-image and text-to video services.
- **Zendesk** has announced the use of outcomes-based pricing with its AI agents, under which customers will only pay for successful interactions, based on agreed-upon interaction metrics.
- **Workhuman** has fully shifted to an outcomes-based model, under which larger overarching metrics, such as employee engagement or retention, are used to assess whether Workhuman has delivered on its promises and will be compensated for the use of its platform.

“The combination of generative AI and automation technologies has rapidly transformed the ability of vendors to deliver far more functionality within its software offerings, resulting in significant advances in productivity, efficiency, accuracy, and other metrics. However, these advanced features are the direct result of compute power, which can be expensive for vendors to provide, either because their volume of utilization is low and the unit cost is high or because the volume of customer usage is so high that they’re simply burning through a high volume of compute resources. New pricing models are required to align the economic model for both customers and vendors more closely and fairly while providing additional transparency.”



Keith Kirkpatrick

Research Director & Practice
Lead Enterprise Applications

[Request Analyst Session](#)



Semiconductors: Chiplets, Heterogeneous Integration and Advanced Packaging Will Be Critical Enabling Technologies for High-Performance Applications

Prediction: Chiplets will account for an increasing share of semiconductor foundry services with some leading advanced packaging equipment suppliers expecting chiplets to account for 25% of foundry revenue by 2030

(source: BESI Investor Presentation, November 2024)

Why This is Trending:

Industry Economics: For decades, Moore's Law delivered performance and cost improvements through shrinking transistor geometries, increased die sizes and larger diameter wafers. As overcoming the technical challenges associated with this monolithic scaling has become too expensive to justify for all but a few semiconductor device types, industry investment has shifted more towards advances in packaging technology.

Cost Optimization: By partitioning the chip into separate functional elements (chiplets), only the advanced logic functions need to be fabricated on leading-edge process nodes. Less-critical functions can be processed using legacy nodes or a different process technology altogether. Chiplets also have the advantage of the higher yields typically associated with smaller die.

Performance Benefits: In addition to facilitating higher-speed data processing and transfer, heterogeneous integration also enables lower power consumption and better heat dissipation, which are increasingly important system performance metrics.

Use cases:

- **Stacked Die:** Heterogeneous integration utilizes a variety of 2.5D and 3D advanced packaging process technologies, including Hybrid Bonding and Thermo Compression Bonding (TCB)
- **High-bandwidth memory (HBM):** For use in applications such as high-performance computing (HPC), HBM achieves much higher bandwidth than standard DRAM by stacking multiple vertically interconnected DRAM dies.
- **Co-packaged Optics:** Co-Packaged Optics (CPO) is the advanced heterogeneous integration of optical components and semiconductor devices on a single package, aimed at addressing performance and power challenges in high-bandwidth data-center applications.

"As the monolithic scaling associated with Moore's Law has become increasingly economically unviable as a driver of semiconductor performance and cost improvements, so the demands of future applications, such as AI processing, have become even more onerous. We will see an increasing share of industry investment (both R&D and CAPEX) directed towards the advanced packaging technology required to deliver heterogeneous integration of chiplets."



Richard Gordon

Vice President & Practice Lead Semiconductors

[Request Analyst Session](#)

Futurum

Become a Client

