BofA SECURITIES

Global Software

Agentic AI: software ramifications deep dive

Industry Overview

Agentic AI deep dive: is software commoditised?

Advances in frontier models and agentic AI have reinforced the narrative that software creation is approaching commoditisation. We believe this view overstates the impact. This note provides an in-depth analysis of how agentic AI is disrupting the software sector. We examine **1/** the current state of agentic coding, **2/** the implications of software cloning and increasingly capable frontier models, and **3/** how software vendors are adapting their strategies and pricing models. **4/** We also review AI revenue disclosures and monetisation strategies.

Agentic coding: separating the hype from reality

Despite rapid progress in frontier models and agentic coding, we believe software vendors remain a durable and critical part of any technology stack. Our view is that the "everyone can now build software" narrative is well ahead of reality. AI has materially lowered the cost of prototyping, but many demonstrations fail to capture the challenges associated with scalability, reliability, security, integration and ongoing maintenance.

Software moat: evolving, but not disappearing

We believe agentic AI will weaken some software moats while reinforcing others. UI/UX and standardised implementation work may become easier to automate, but trusted data, industry expertise, AI control, partner relationships and bespoke deployment capabilities should either remain relevant or gain importance. Network effects could also strengthen as more customers and partners add data and insights. In our view, the deeper AI moves into enterprise workflows, the more it depends on proven software, business rules and systems of record.

Pricing model & AI spend cannibalisation

Pricing remains a key concern, as the shift from seat-based to consumption or outcome-based models remains uncertain, while inference costs leave gross-margin economics an open question. Overall, Enterprise software spending is still expected to grow c. 13% 2025-30 CAGR, although AI introduces uncertainty over agentic cannibalisation. In this note, we examine how value capture and pricing may evolve.

AI revenue disclosures across EU / US

AI revenue disclosure remains inconsistent across the software sector and is generally more developed among US vendors than European peers. Current disclosures suggest AI generally accounts for low to mid-single digit of group revenues, within a broad c.1–11% range. However, we believe reported figures understate the true contribution, given that AI monetisation is often embedded within bundles or captured through broader price increases. Medium-term targets remain scarce, although companies providing them expect AI to represent c.5–30% of group revenues. See Exhibit 19 for details.

20 October 2026

Equity
Europe
Software

Victor Cheng, CFA >>
Research Analyst
MLI (UK)
+44 20 7996 0272
victor.cheng@bofa.com

Frederic Boulan, CFA >>
Research Analyst
MLI (UK)
+44 20 7995 7101
frederic.boulan@bofa.com

Tal Liani
Research Analyst
BofAS
+1 646 855 5107
tal.liani@bofa.com

Koji Ikeda, CFA
Research Analyst
BofAS
+1 415 913 2106
koji.ikeda@bofa.com

Tomer Zilberman
Research Analyst
BofAS
+1 646 855 3203
tomer.zilberman@bofa.com

Carla Courthial >>
Research Analyst
MLI (UK)
44 77 6120 0059
carla.courthial@bofa.com

Aditya Buddhavarapu, CFA >>
Research Analyst
MLI (UK)
+44 20 7995 7318
aditya.buddhavarapu@bofa.com

UI / UX = User Interface / Experience

See Glossary on pages 37 - 38

>> Employed by a non-US affiliate of BofAS and is not registered/qualified as a research analyst under the FINRA rules.
Refer to "Other Important Disclosures" for information on certain BofA Securities entities that take responsibility for the information herein in particular jurisdictions.
BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.
Refer to important disclosures on page 47 to 49. Analyst Certification on page 44. Price Objective Basis/Risk on page 39.
13030043

Executive Summary

The rise of agentic coding and its limitations

Vibe coding has quickly moved beyond experimentation and is evolving to become an important enterprise development model. In the last few months, better frontier models and advances in reverse engineering are lowering the barriers to software creation, but we believe enthusiasm remains well ahead of reality. AI agents can already generate code and accelerate development, yet they remain dependent on existing software tools, experienced engineers and clear product direction. Without these foundations, they often produce lower-quality clones or replicate established solutions rather than create genuinely differentiated products. Coding, viewed in this context, is therefore far from "solved". As capabilities improve, we expect software companies to adapt and remain critical in turning raw model output into enterprise-grade products that are reliable, auditable, secure and cost-efficient.

Software moat is evolving, not disappearing

A central question for investors is whether incumbent SaaS vendors can defend their positions as AI model providers move up the technology stack and lower barriers to software creation. In our view, moat erosion is highly uneven.

> Agentic AI can commoditise much of the standardised workflow, but the last mile remains difficult and expensive. Deep integration, process complexity, governance and exception handling are where established enterprise software vendors continue to differentiate.

Overall, we believe software moats persist, although their relative importance changes as agentic AI develops. Interface quality and user familiarity should become less differentiating as agents navigate applications on users' behalf. Trusted business data, embedded industry knowledge and control over AI actions could become stronger defences. Established reseller networks should remain valuable through customer relationships and implementation expertise, although routine configuration work may face pressure. Genuine network effects, where additional users or partners improve the product for others, also remain defensible and could strengthen through shared learning. We see competitive advantage shifting towards assets that are difficult to replicate beyond the software code itself.

Pricing: AI challenges seat-based monetisation

Another key concern from investors with whom we've spoken is whether seat-based pricing can survive the shift towards AI agents, which may perform more work without increasing user numbers. In our analysis, we see software providers generally responding with consumption and outcome-based pricing models, charging for measures such as tokens, actions or results. However, fast evolving capabilities, token efficiencies and customer expectations mean pricing remains experimental, with customer resistance potentially slowing the transition from subscription models.

> In our view, without clearer evidence that software providers can successfully transition pricing and monetize AI usage, this concern is unlikely to fade in the near term and will continue to weigh on software sentiment.

Based on our analysis, adoption within customers currently remains uneven and is likely to be a key constraint on near-term value creation for software providers. Enterprises are achieving productivity gains in discrete, measurable tasks, but scaling these benefits across end-to-end workflows requires significant organisational change. Consequently, many deployments remain at the pilot or assistive stage rather than becoming autonomous, enterprise-wide systems. This points to a more gradual monetisation curve for software providers.

Buy vs Build debate

Investors also often question whether AI will enable companies to build software internally rather than buy it from established vendors. We think this places too much focus on the initial coding process. The speed of initial development should not be conflated with the full cost of owning and operating software in production. In enterprise, much of the cost sits outside the first build, e.g. integration with existing systems, workflow management, access controls, exception handling, auditability, compliance and ongoing maintenance.

> Our industry checks suggest some companies have struggled to replace incumbent software in-house, citing cost and quality challenges. Many have renewed existing contracts for another cycle while continuing to develop internal alternatives.

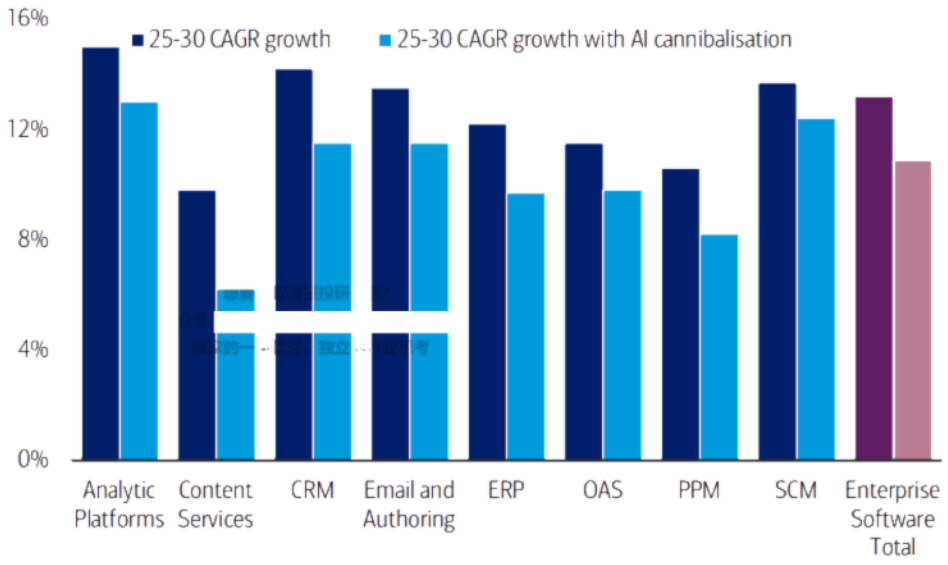
Looking from this perspective, AI may compress the initial software development without removing the complexity required for reliable deployment and maintenance at scale. In some cases, AI may even expand the governance burden by introducing new requirements around model oversight, traceability, security and policy enforcement. Faster software creation does not necessarily result in proportionally lower total cost of ownership.

As such, we believe software providers can remain relevant even as AI improves coding productivity. Their value may increasingly lie less in code generation itself and more in providing the surrounding operational framework needed to run software safely and consistently across the enterprise.

Enterprise software growth: a 2.3pp AI cannibalisation potential

Enterprise software spending appears set for healthy double-digit growth over the next few years, although AI could materially redistribute that growth. Gartner forecasts application software spending growing at 13.2% annually over 2025–2030. However, it also estimates that $116bn of cumulative spending will shift from traditional functional categories towards cross-functional agentic applications, reducing the growth of the remaining siloed markets to 10.9%. This should not necessarily represent an equivalent headwind for large vendors able to provide both core applications and cross-enterprise orchestration, as they can recapture spending moving between these categories. We therefore see the greatest risk for narrower providers confined to standalone functions, while broader platforms could potentially sustain growth closer to pre-cannibalisation market rates.

Exhibit 1: Enterprise software growth 2.3pp lower if cannibalization happens
Impact of cross functional cannibalization on Enterprise Software



Source: Gartner, BofA Global Research
CRM = Customer Experience and Relationship Management, ERP = Enterprise Resource Planning, OAS = Other Applications Software, PPM = Project and Portfolio Management, SCM = Supply Chain Management

BofA GLOBAL RESEARCH

AI revenue disclosure across companies remains uneven

Currently, AI revenue disclosure remains inconsistent across the software sector and is generally more developed among US vendors than European peers. European software providers currently focus more on usage metrics rather than revenue contribution. Revenue disclosures across the two groups suggest AI generally accounts for low to mid-single digit of group revenues, within a broad c.1–11% range. However, we believe reported figures understate the true contribution, given that AI monetisation is often embedded within bundles or captured through broader price increases. Medium-term targets remain scarce, although companies providing them expect AI to represent c.5–30% of group revenues. See Exhibit 19 for details.

We expect AI pricing models to remain fluid in the near term, as rapid improvements in AI capabilities continue to reshape both customer value and the underlying unit economics.

What might happen long term when agentic AI gets much better?

Our view is that the "everyone can now build software" narrative is well ahead of reality and is likely to remain so, even as frontier models become more capable. AI has materially lowered the cost of experimentation and prototyping, but many demonstrations fail to capture the challenges associated with scalability, reliability, security, integration and ongoing maintenance.

If AI eventually enables everyone to develop software more efficiently, it is reasonable to assume that leading software vendors will be among the greatest beneficiaries, leveraging the same technologies at scale and often with greater speed and effectiveness. Admittedly, business models will likely evolve, with traditional seat-based pricing increasingly supplemented by hybrid subscription and consumption-based models. Our bottom line is that the enterprise software industry is unlikely to disappear. The more important question is how economic value will ultimately be distributed: whether software vendors can retain the benefits of AI-driven productivity gains, or whether a growing share of customer spend migrates towards AI platforms and model providers. We believe the likely outcome lies between these extremes: existing software vendors should capture part, but not all, of the value created. However, if AI expands the overall market, this could still represent a net benefit for the industry.

Stock Implications

SAP (Buy, €226): defensive position, but upside to be demonstrated

We see SAP as less exposed to AI disintermediation vs peers, however it has yet to demonstrate meaningful AI traction and commercial momentum. SAP is a mission critical ERP across finance, HR, supply chain, procurement and CX. Deep data models, embedded controls and regulatory workflows create high switching costs, making generic agents poor substitutes. AI is being embedded inside transactional processes rather than layered at the edges. Joule Copilot and Joule Studio underpin c.40 domain agents and allow customers to compose/customise workflows; monetisation blends seat based and consumption based pricing. Commercial traction is already visible: >2/3 of cloud orders include AI, 90% of the 50 largest deals included BDC or AI, c.60% of cloud customers already use AI, and €2bn of BDC contracts are signed. The main risk is GenAI substitution at the periphery. Mitigating factors, regulatory requirements, mission critical coverage and material switching costs tilt the balance in SAP's favour while it builds an AI led upsell runway.

Industry software: strong moat, but seat-based pricing a concern

We believe AI will materially reshape CAD software and user interaction, but for **Dassault Systemes (Neutral, €25.6)**, **Autodesk (Buy, $300)** and **Nemetschek (Buy, €93)**, the key moat is not simply creating 3D models, but maintaining an accurate digital record across design, testing, manufacturing and product life. AI-generated geometry must still satisfy engineering constraints, undergo simulation and connect with downstream processes. However, these solutions are particularly exposed to seat-based pricing model disruption by AI as their pricing model remains mostly seat-based.

Specifically for AEC software (**Autodesk** and **Nemetschek**), AI adoption in AEC (whether a generic AI model or a first-party tool from AEC providers) will likely be slow and gradual, as key considerations of trust and data governance remain obstacles. Discussions with industry participants suggest headless interfaces and Generic AI models could be used for bounded, repetitive workflows; but first party tools are likely to continue to exist for deeper workflows.

Travel IT: distribution disruption ahead

In Travel IT, we believe personal agents could durably reshape how travellers discover and compare options at the top of the distribution funnel. We expect a bigger impact to airline distribution for **Amadeus (Underperform, €53)** and **Sabre (Buy, $2.5)** compared to hotel distribution **HBX Group (Buy, €12)**.

Personal AI agents may struggle to become mainstream booking channels, given suppliers' control over inventory, limited user trust and privacy constraints. However, we believe they can still disrupt the top of the distribution funnel without executing transactions.

Communication-platform-as-a-service (CPaaS)

For vendors like **Twilio (Buy, $320), Sinch (Underperform, 29SEK)** and **Bandwidth (Neutral, $60)**, we think agentic-driven communication (SMS, phone calls, etc.) will drive higher volume growth over the long-term. We would track consumer-facing agent adoption, and messaging and voice volume trends to validate our thesis.

AI data platforms and databases: AI reinforces value

We believe AI reinforces the strategic importance of **Snowflake (Buy, $470)** and **MongoDB (Buy, $540)** as agents increase the value of governed and secured enterprise data. Also, Snowflake's AI products provide an additional monetization opportunity to unlock more business intelligence.

Observability software
We believe AI is generating higher volumes and more complex workloads that need to be monitored with AI products from vendors like **Datadog (Buy, $305)**, **Dynatrace (Buy, $65)**, and **Elastic (Neutral, $108)**.

Software development tools
We think that AI is changing the software development lifecycle and introducing higher complexity and increased need for security, which could make software development software tools like **JFrog (Buy $110)** and **GitLab (Neutral, $54)** more strategic.

Deterministic robotic process automation (RPA) software
For vendors like **UiPath (Underperform, $15)**, we think probabilistic AI agents introduce questions around which gains higher share of future workloads. Customers may adopt UiPath for high-reliability and lower cost per task versus other AI options today, but advances in AI capabilities or lower AI costs could eat into the value proposition of RPA vendors.

Other individual software names

- **Temenos (Buy, CHF 97):** We see minimal near-term displacement from general purpose agents: banks' high complexity, high risk environments favour co-innovation around AI rather than core replacement. Where AI may appear first is in surrounding workflows (e.g., drafting, case notes, query resolution), not in the transactional heart of the platform.

- **Planisware (Buy, €31):** We see a particular risk linked to its commitment to seat based pricing, as buyers test consumption aligned models and agentic usage grows. Offsetting this, the unified platform accelerates AI delivery and product velocity; complex integrations and workflows are harder for generic tools to replace.

- **TeamViewer (Buy, €9.5):** We see a particular risk linked to its commitment to seat based pricing, as agentic AI capable of resolving basic IT issues increasingly erodes into the lower-end segment. That said, the strategy to reposition around a broader digital workplace platform is credible, anchored by expansion into Digital Employee Experience (DEX) and automation.

Viewpoints from industry participants

Our discussions with various industry experts and participants suggest enterprise AI adoption remains early and pragmatic. Most expect hybrid deployment, persistent implementation complexity, durable incumbent advantages in data and workflows, growing orchestration importance, evolving pricing models and only gradual disruption.

- **Enterprise AI** remains early stage: customers are progressing from broad experimentation and proof of concepts towards narrower, workflow -specific deployments with clearer ROI, ownership and production feasibility.
- The debate extends beyond **buy versus build**: most enterprises are likely to combine incumbent software, cloud infrastructure, external implementation support and customised internal development.
- **Production deployment** remains difficult: security, governance, integration, reliability, maintenance and long tail edge cases create a substantial gap between compelling prototypes and scaled production applications.
- **Enterprise SaaS**: some participants expect Enterprise software to face longer term pressure, not immediate displacement: failed or costly internal builds are encouraging many customers to renew incumbent platforms while developing alternatives over several years.
- **Systems of record**: almost all participants expect systems of record to retain structural advantages: proprietary data, workflow semantics, embedded integrations, enterprise controls and switching costs continue to anchor AI agents to existing software platforms.
- **Forward deployed engineering** is becoming mainstream: software vendors, infrastructure providers and consultancies increasingly embed technical teams with customers to customise, integrate and operationalise AI deployments.
- **AI pricing** remains unsettled: vendors are experimenting with subscriptions, usage, milestones and outcomes, but unpredictable consumption, baseline measurement and attribution continue to constrain outcome-based adoption.
- **Costs** remain a material constraint: scaled agent deployments can be expensive to operate and maintain, increasing customer demand for cost caps, measurable returns and tightly defined use cases.
- **Model** level differentiation may erode: declining inference costs, improving open-weight models and model agnostic tooling could commoditise underlying LLMs and shift value towards application and distribution layers.
- **The orchestration layer** is becoming more important: durable differentiation may increasingly sit in agent harnesses, governance, workflow context, cost controls and coordination across multiple applications.

Agentic coding: rapidly maturing

Vibe coding has rapidly evolved from an experimental approach into a potential enterprise development model. As frontier models improve and advances in reverse engineering lower barriers to software creation, more development is likely to become agentic. However, we believe the hype remains well ahead of reality. Currently, AI-built clones are often lower quality, models still depend on underlying software tools, and without senior engineers, AI agents tend to mimic rather than innovate.

What is agentic coding?

Agentic coding refers to the use of AI tools that do more than simply suggest lines of code. These systems can interpret instructions in natural language, generate software, run tests, fix errors and in some cases complete multi-step tasks with limited human input. "Vibe coding" was the early, informal version of this idea: describing what you wanted in plain English and letting AI build something quickly, often for experimentation rather than production use. What has changed over the past 12–18 months is that this approach has moved well beyond hobbyist projects. It is now evolving into a broader enterprise software development model, with more structure, more oversight and clearer use cases across the development lifecycle.

So, is coding "solved"?

Our view: Agentic coding may remove the need to write code manually, but frontier models still depend on existing software tools and engineering expertise to provide meaningful output. Without clear technical and product direction, they often default to replicating established solutions rather than developing genuinely novel and differentiated products.

As it stands, while AI agents can already accelerate software development significantly, building a successful product still requires experienced engineering talent, product judgment and deep domain expertise. Today's AI models excel at replicating, adapting and iterating on established designs and known workflows, but often struggle to generate truly novel concepts beyond what they have previously learned or can reference. AI capabilities will undoubtedly continue to improve, as the pace of progress over the past few years has demonstrated. However, we also expect software companies to adapt and evolve alongside these advancements.

The rise of software cloning and reverse engineering

Software cloning gained visibility in early October with a popular tool REA (Reverse Engineer Anything), which connects AI coding agents to established software-analysis tools, including Ghidra, Hopper and IDA. These tools can examine the packaged files used to install and run software, then work backwards to understand its underlying logic without seeing the original source code. Rather than recovering that code, REA gives AI

agents a simplified blueprint and supporting information that they can use to recreate and test similar functionality. It can investigate traditional applications, websites and individual features, rather than requiring an entire product to be rebuilt.

> In our view, this could reduce the cost and effort of replicating existing features, although it does not yet prove that complete commercial software products can be cloned reliably.

Developers do not always need to examine a program's underlying machine code to recreate parts of it. AI coding tools can learn from what users see, how the product responds, and information contained in manuals, public file formats and other documentation. Developers can then use screenshots, sample files and test cases to check whether the new product produces similar results.

This means a competitor may not need to copy an entire application to disrupt it. Recreating a limited number of widely used features could be enough to serve less demanding customers or put pressure on pricing. However, reproducing isolated functions is very different from matching the reliability, performance and breadth of a mature commercial product.

There are three main ways to recreate software:

- **Copying what the product does:** Developers study the software's interface, documentation and outputs, then use AI to build an alternative that behaves in a similar way. The new product does not need to work in the same way internally, provided that it produces comparable results for users.
- **Working backwards from the installed software:** Most commercial software is distributed in a form designed for computers rather than people to read. Reverse-engineering tools translate parts of it into a more understandable form, helping AI identify how certain features may work. This provides more insight than simply observing the interface, but connecting thousands of components into a stable product remains difficult.
- **Studying connections to cloud services:** Some features are processed on remote servers rather than on the user's computer. Developers can study the information exchanged between the application and those servers to understand how a feature is accessed. This may help them build a compatible connection, but it does not recreate the proprietary technology running in the cloud.

Adobe-style clones attract significant attention

The most visible case of software cloning emerged earlier this month in October, when ArtCraft released seven open-source applications modelled on Adobe products including Photoshop, Illustrator, Premiere Pro and Lightroom. The applications closely reproduce the layout, tools and workflows familiar to Adobe users, making them look surprisingly similar at first glance. Their rapid development and strong public interest illustrate how AI coding tools can lower the cost of building a recognizable software interface. However, the applications remain early-stage products and have significant missing or incomplete features. ArtCraft itself continues to describe several applications as being in early alpha or under development.

In our view, replicating menus, panels and basic functions is materially easier than matching the depth, reliability, performance and file compatibility accumulated over decades in a professional software suite. A similar pattern is emerging in CAD, where open-source SolidCraft offers familiar parametric modelling workflows but describes

itself as suitable primarily for simpler parts, while major capabilities such as simulation, advanced surfacing, CAM and 2D drawings remain outstanding.

We acknowledge that these gaps could narrow as AI improves and active open-source communities contribute code, testing and frequent releases. However, building a dependable replacement for complex production software is likely to remain a much harder task than reproducing its interface.

Implications for software development

> We believe vibe-coded clones may struggle to achieve lasting commercial adoption. Customers need confidence that software will remain supported, secure and compatible over time. Community-led projects can become abandonware if contributors lose interest, with unclear accountability for failures, data loss or security breaches creates liability concerns, particularly for enterprise customers.

While AI tools can now support multiple stages of development autonomously, this does not mean developers are being replaced. Instead, their role is shifting from writing every line of code towards reviewing AI-generated work, checking how the system has been designed, testing security and ensuring the final product meets the required standard.

Tools that turn a simple prompt into an application can make basic software easier to create. However, the more durable enterprise opportunity is likely to come from AI being built into existing professional development processes, where organisations can still review changes, track decisions and connect new code with their current systems. We explore more in depth in the evolution of software moat in a later section of the note.

Types of agentic coding

Software cloning is only one application of agentic coding. The broader market spans several development approaches, differing mainly in how much control is delegated to AI and how closely developers remain involved. We can classify agentic coding broadly into 4 categories:

- **Pure vibe coding (prompt to app):** The user provides only natural language and never touches the underlying code. The AI generates a complete, runnable application. Representative tools include Replit Agent, Bolt.new, Lovable, and v0.
- **Conversational development:** A developer works within an IDE but drives development primarily through natural language dialogue with an AI agent. The developer reviews and edits AI-produced code, but the primary interaction mode is language rather than direct authorship. Representative tools include Cursor, Codex and Claude Code. Increasingly, these tools operate not just as conversational partners but as autonomous agents capable of executing multistep tasks, running tests, and iterating on their own output with limited human intervention.
- **AI-accelerated coding:** The developer writes code directly, while the AI suggests completions, generates functions, and fills in boilerplate. GitHub Copilot has operated in this mode since 2022. Although frequently grouped with vibe coding, it diverges from Karpathy's premise: The developer actively reads, reviews, and shapes the code throughout.
- **Spec-driven AI development:** The developer provides natural language requirements, and the AI generates formal specifications and task plans before producing code. The human reviews the specification before implementation begins.

Representative tools include AWS Kiro and IBM watsonx Code Assistant with Project Bob in Enterprise solutions. However, Claude Code and Codex remain the more popular option as their capabilities continue to expand.

These four modes show that agentic coding is not one single product category, but a spectrum. At one end, AI acts almost like a no-code builder, turning prompts into working applications with minimal user involvement. At the other, AI becomes a more controlled enterprise development layer, where natural language is used to define requirements, create plans and support implementation inside governed workflows. The common thread is that natural language is becoming a real interface for software creation, not just a novelty feature.

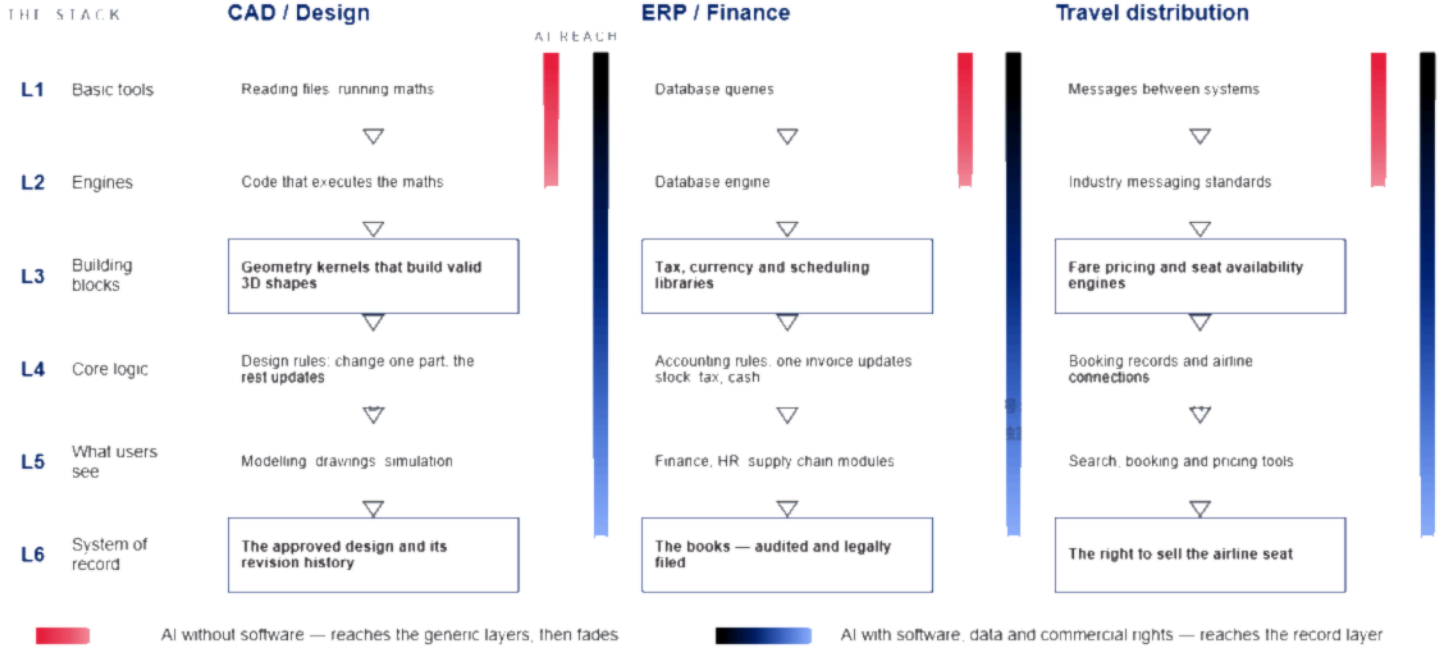
Software moat in the age of agentic AI

Overall, we believe software moats persist, although their relative importance changes as agentic AI develops. UI / UX becomes less important as agents navigate applications. Trusted business data, embedded industry knowledge and control over AI actions could become stronger defences. Established reseller networks should remain valuable through customer relationships and implementation expertise, although routine configuration work may face pressure. Genuine network effects, where additional users or partners improve the product for others, also remain defensible and could strengthen through shared learning. However, neither a large installed base nor simply holding customer data guarantees a moat. We therefore see competitive advantage shifting towards assets that are difficult to replicate beyond the software code itself.

The software stack remains a core moat

Frontier models can handle generic tasks, but useful real-world outputs still depend heavily on existing software, data and tool calls. The deeper AI moves into a workflow, the more it relies on proven engines, business rules and systems of record. Rebuilding these capabilities from scratch would be costly, unnecessary and often unproven. Without the underlying software stack, AI can generate answers and recommendations, but remains limited in its ability to execute reliable, accurate and legally recognised business processes.

Exhibit 2: Frontier models use, instead of bypass, software to produce quality outputs
Ability of AI frontier models alone vs using software



Source: BofA Global Research

BofA GLOBAL RESEARCH

The Exhibit above illustrates this across three areas. In CAD, AI can assist with design, but still relies on geometry engines, design rules and approved records sitting within a CAD software. In ERP, it needs existing databases, tax logic and accounting systems to update legally recognised books. In travel, AI can help users search and book, but still depends on pricing engines, live seat availability and distribution intermediaries such as Amadeus and Sabre. AI may change the interface, but the underlying software, data and commercial agreements remain essential.

Is owning the AI control surface the new moat?

We believe owning the AI control surface is emerging as a new moat for large software providers. This is the layer that decides what AI agents can access, which actions they can take and how they are monitored. As agentic AI evolves, we expect agents to increasingly move beyond single tasks in one application and run complete business processes across departments and systems. A sales order, for example, may pass through CRM, finance and supply-chain systems without a human stepping in. In this world, the most valuable position belongs to whoever coordinates and governs that work.

Leading vendors are already competing for this position. SAP requires AI agents, including those built on third-party models, to access its systems through its own governed platform, and it restricts outside AI tools that operate around it. ServiceNow's AI Control Tower aims to discover, monitor and govern AI agents across the whole enterprise, including agents running on other vendors' systems, and can shut down agents that act outside their permissions. We see this as clear evidence that control over AI agents, not the interface, is where software vendors now expect value to sit.

However, we do not expect one company to control all AI agents. Instead, we see a layered set-up. At the top, an enterprise-wide "control tower" is likely to oversee all agents and set the overall rules. Below it, specialist platforms co-exist to run the agents in their own areas: Microsoft for productivity tools such as Office and Teams, SAP for ERP, ServiceNow for IT and workflows, and CPaaS players (e.g. Twilio and Sinch) for customer communications. Each would manage its own domain while still answering to the control tower above.

UI / UX becoming a weaker moat

We believe UI/UX is becoming a weaker moat for software. As AI agents take on more agentic tasks, they interact with applications mainly through APIs rather than user interfaces, so a familiar or well-designed interface matters less. Agents can also work through clunky or legacy interfaces on users' behalf, making a superior user experience less relevant. We note that software providers are already responding by basing their defensibility on proprietary data, business context and governed control over what agents can access and execute.

SI / VAR network: still one of the stronger moats

> We expect reseller ecosystems to remain a meaningful and durable software moat. Established networks of systems integrators and value-added resellers provide trusted customer relationships, implementation expertise and local support, which remain important as enterprises adopt AI. Software vendors with strong partner ecosystems should therefore retain advantages in distribution, implementation and customer retention.

We see SI/VAR ecosystems as an underappreciated software moat. Enterprise deployments rarely end with the software sale. They typically require integration, migration, process redesign, localisation, training and ongoing optimisation. Partner-influenced revenue should therefore become more important, not less.

Deep partner networks reduce adoption risk. Customers prefer platforms supported by credible implementers, established deployment methods and certified talent. This improves market reach and time-to-value, while making the platform harder to displace after deployment.

Ecosystem scale also creates a talent moat. Certifications and specialist expertise encourage consultants, developers and customers to standardise around a platform. Customers can recruit skilled staff, partners can generate revenue and practitioners can invest in relevant training. This reinforces retention, particularly in complex areas such as CRM, ITSM, ERP, HCM and vertical software.

Agentic AI may alter partner economics, but is unlikely to remove this moat. It can automate routine configuration and coding, while creating new demand around data readiness, governance, controls, auditability and coordination across systems.

Data moat: powerful in the right architecture

> We view data moats as selectively strong, not universal. They are strongest where vendors capture proprietary workflow data and control the action loop, as SAP and ServiceNow do; weakest where vendors mainly host portable customer data and expose functionality that agents can access elsewhere.

In the agentic AI era, the strongest data moat is less about owning the data and more about understanding its context, controlling who can use it and linking actions to outcomes. Data moat is therefore one of the most misused ideas in enterprise software. Unlike consumer internet companies, enterprise software vendors often do not own the underlying data. It belongs to the customer, comes from its operations and is simply stored, combined or analysed within the application. Salesforce, for example, brings customer information from different systems into a single view, but the underlying data largely remains the customer's.

We therefore distinguish between imported data and data generated through workflows. Imported data can be valuable, but is often portable. Data created as software runs real business processes is harder to replicate. This includes approval histories, exceptions, errors, escalation paths, resolution times, manual overrides and final outcomes. **ServiceNow** highlights more than 85bn production workflows and 7tn transactions as the business context underpinning its AI capabilities. **SAP** similarly generates valuable context as customers run finance, supply-chain and other core processes. This is very different from simply loading customer documents into an AI model.

This distinction becomes more important with agentic AI because agents work better when data is clearly linked to actions and outcomes. The moat is therefore not the amount of raw data, but the quality of the business context around it: what happened, why it happened, which action was taken and whether it worked. This creates a feedback loop through which software and agents can improve over time.

We therefore see the strongest data moats in systems of record and systems of action. Vendors that sit inside core workflows and record the resulting changes can build operational histories that competitors cannot easily buy or copy. By contrast, software that mainly displays or analyses information from other systems may face greater pressure, because customers or competing agent platforms could access the same underlying data through APIs.

Network effects: real but category dependent

Network effects can be a powerful moat, but only in selected software categories. They matter most in platforms, marketplaces and communication tools, where each additional customer, developer or partner increases the product's value to others. They matter much less in standalone applications used largely within one company.

We think genuine network effects arise when greater participation makes the product more valuable to others. This can happen directly, through more users or transactions, or indirectly, as a larger installed base attracts developers, plug-ins, integrations, partners and trained professionals.

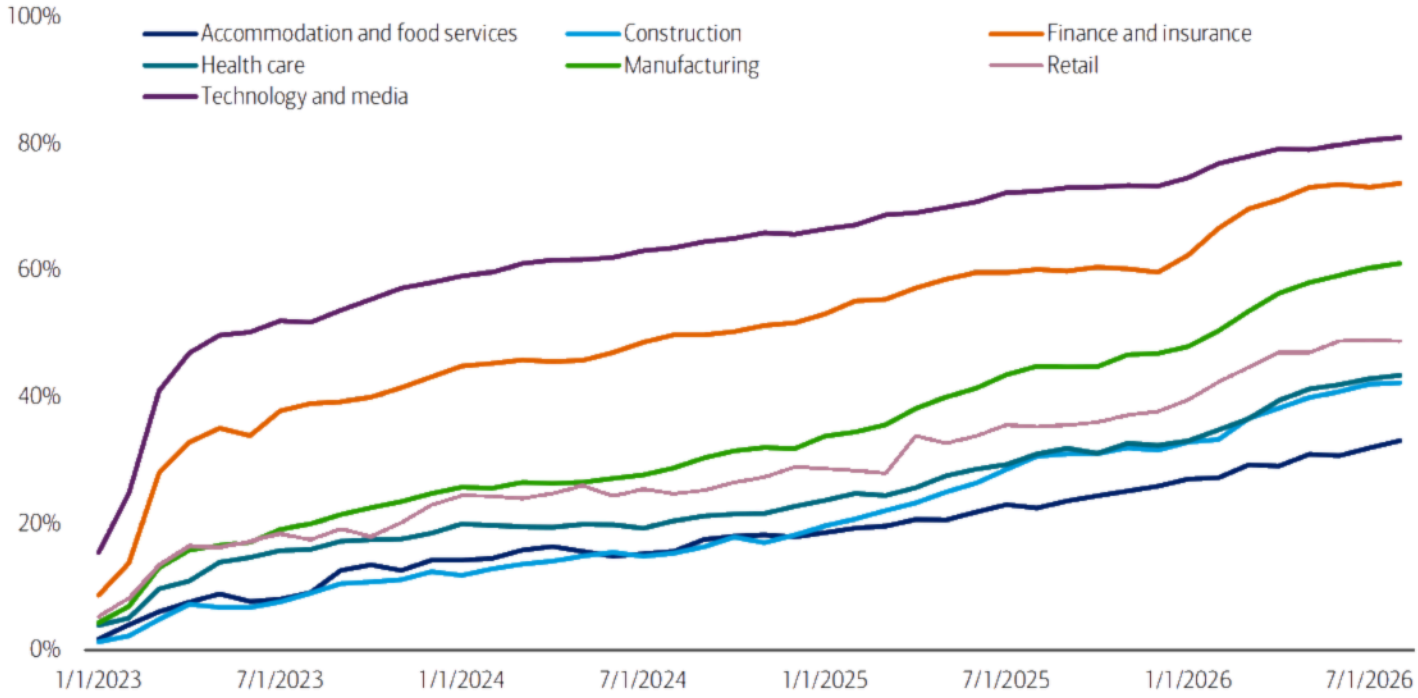
The clearest enterprise examples are app marketplaces, payment and procurement networks, industry collaboration platforms, and communication products. Salesforce illustrates the platform model, where more customers attract developers and implementation partners; more applications and expertise make the platform more useful; and this can attract further customers. Salesforce's marketplace now extends beyond conventional applications to third-party agents, tools and capabilities.

However, a standalone ERP, HR or analytics product does not automatically become more useful to one customer simply because another customer adopts it. Scale can still strengthen the moat indirectly: widespread adoption creates a larger pool of users familiar with the product, attracts more systems integrators and resellers, and supports more integrations and implementation expertise. These benefits can lower deployment risk and reinforce customer retention. However, they are better described as ecosystem, distribution and talent advantages than direct network effects, because each new customer does not necessarily improve the product for existing users.

Where agentic AI shows up in enterprise

Most companies now adopt some form of AI in 2026, but few run it deep, with top 1% companies (by AI spending per employee) spending over 500x the median. Agentic AI is arguably already widely available and distributed for large horizontal software players, though productivity and pricing remain unresolved. Easy wins and use cases are more widespread, but advanced and complex workflows remain challenging. For vertical software, agentic AI adoption varies but is generally slower.

Exhibit 3: AI adoption continues to grow across sectors
Share of U.S. businesses with paid subscriptions to AI models, platforms, and tools

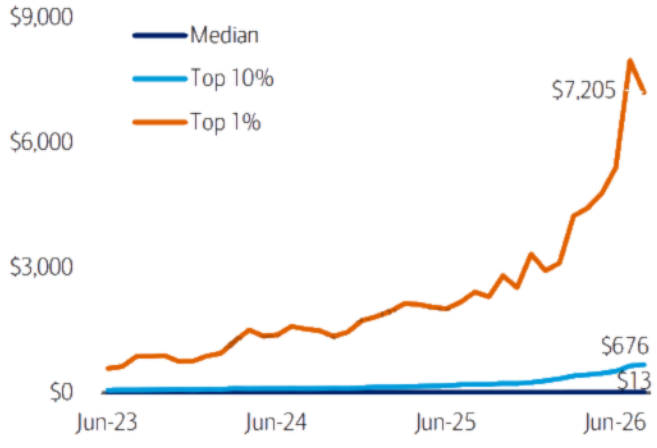


Source: Ramp AI Index, business spend data from Ramp

BofA GLOBAL RESEARCH

While many companies have adopted some form of AI, the depth of adoption remains limited. According to the Ramp AI Index, which covers c.70,000 small and medium-sized businesses, median AI spend is just $13 per employee per month, while the top 1% of spenders allocate more than 500x that level. Although technology companies are disproportionately represented among the highest spenders, their median spend is still only $80 per employee per month, more than 100x below the top 1%. In our view, this underscores how uneven AI adoption remains, not only across sectors but also in terms of intensity of usage among companies.

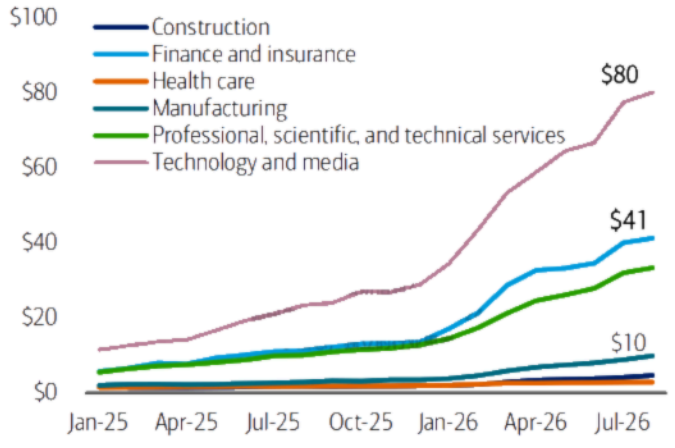
Exhibit 4: Top 1% spend 500x more on AI per employee
Monthly AI spend per employee by percentile



Source: Ramp AI Index, business spend data from Ramp

BofA GLOBAL RESEARCH

Exhibit 5: AI spending uneven across sectors
Median monthly AI spend per employee by sector



Source: Ramp AI Index, business spend data from Ramp

BofA GLOBAL RESEARCH

Current state of agentic AI across business functions

Across business functions, agentic AI is moving beyond small pilots and into real production use, although progress remains uneven. Adoption is strongest in areas such as cybersecurity, software development and HR, where tasks tend to be more structured, repeatable and easier to measure, making both the business case and the deployment process more straightforward. At the same time, companies are increasingly viewing agentic AI as a broad, horizontal capability that can support many different parts of the organisation, rather than something limited to one or two isolated domains. However, in many functions, the areas seen as most promising have not yet reached the same level of production maturity. That gap creates a clear opportunity for vendors: focus on use cases where value is visible, complexity is manageable, and implementation barriers can be reduced to move demand from interest to live deployment.

Exhibit 6: Agentic AI is reaching production across a broad set of horizontal use cases
Top 3 use-cases where AI agents are widely used and in production

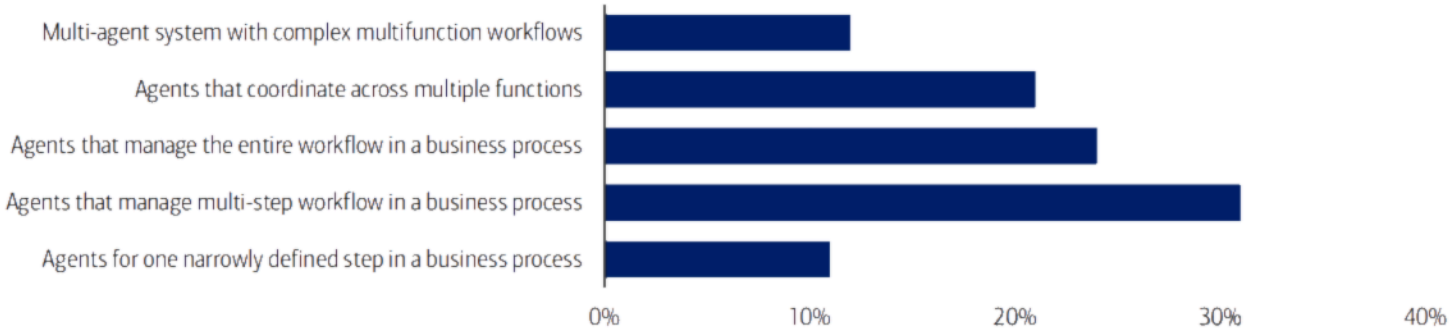
	Top Use Case #1		Top Use Case #2		Top Use Case #3	
Cybersecurity	Security assessment	61%	Automating response playbooks	60%	Auto-remediation of decision-making	58%
Development	Documentation (e.g., code, solutions, models)	61%	Vulnerability prioritization	50%	Quality and testing	50%
HR	Employee self-service	51%	AI-assisted goal setting	47%	Interview scheduling automation	47%
IT operations	AI-driven root cause analysis (RCA)	50%	GenAI post-incident review and knowledge capture	43%	Runbook and automation generation from natural language	41%
Engineering/R&D	Develop training/simulation environments	49%	Supply-driven design	45%	ChatPLM	45%
Customer service	Voice of the employee and voice of the customer feedback loops	48%	Contextualized Responses to customer queries	48%	Automated 24 x 7 customer service	45%
Marketing	Pipeline optimization	45%	Autonomous campaign	42%	Automated creative	40%
Facilities	Automated equipment health response	43%	Augmented employee tracking	39%	Connected property management	39%
Sales	Automated knowledge management	43%	AI-driven chatbots for lead qualification	42%	Proposal document generation	39%
Supply chain	Agile production planning	42%	Automated inventory replenishment	41%	Workforce planning/labor forecasting	41%
Procurement	Risk identification	41%	Supplier performance management	34%	Drafting contracts	33%
Finance	Cash flow forecasting	40%	Preparing tax filings	38%	Virtual booking assistant	37%
Legal	Predictive coding	39%	Repository management	37%	Language translation and localization	36%

Source: IDC Survey (March 2026, n = 1967), BofA Global Research

BofA GLOBAL RESEARCH

Enterprise AI agents are advancing beyond narrowly defined task automation toward broader process execution and more complex cross-functional workflows. Process-level agents remain the most common starting point, but adoption is broadening into cross-functional and multi-agent deployments, which should materially raise the strategic importance of integration, oversight and governance.

Exhibit 7: AI agents are expanding beyond narrow tasks into cross-functional workflows
Question: How would you define the scope of the most complex AI agents currently deployed within your function?



Source: IDC Survey (March 2026), BofA Global Research (n=1967)

BofA GLOBAL RESEARCH

In this context, we see large existing horizontal vendors increasingly position their offerings around orchestration, integration, and business process transformation rather than standalone automation.

Strong adoption, pricing & productivity lags

Enterprise AI is now showing up in large horizontal players. SAP got Business AI in over 90% of top 50 deals, ServiceNow has over $1bn ACV in AI, Microsoft hit 30m paid co-pilot seats. However, strong AI attachment rate is different from fixing the seat-based pricing problem. In vertically focused software, the landscape is more nuanced.

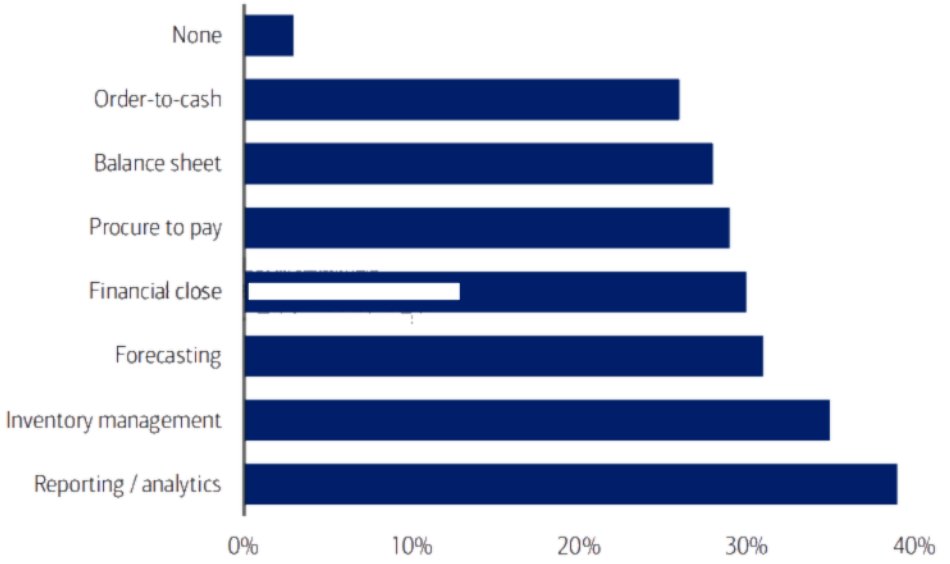
Which workflows are most "agentifiable" now

The most agentifiable enterprise workflows today are those that are high-volume, text-heavy, rules-guided, and spread across multiple systems rather than physically executed. In practice, that points to customer service, IT service desk, finance operations (AP, reconciliations, close support), procurement, HR onboarding/offboarding, and parts of sales operations such as lead qualification, proposal drafting, and CRM updates. These workflows are increasingly agentifiable because modern agents can reason, plan, call tools, update records, and orchestrate multi-step actions with limited human intervention.

Enterprise Resource Planning

AI adoption in ERP appears to be moving from selective pilots to broader operational use across finance, planning and transactional workflows. Early ERP AI spend has been concentrated on insight generation, reporting efficiency and decision support, rather than full workflow delegation. Importantly, no single process category dominates current usage, which indicates adoption is diffusing in parallel across core enterprise functions rather than remaining confined to isolated use cases.

Exhibit 8: Current AI-powered ERP usage
Q: Which end-to-end business processes has your organization invested in AI-powered ERP? And in the next year, where will your organization invest in for AI-powered ERP?



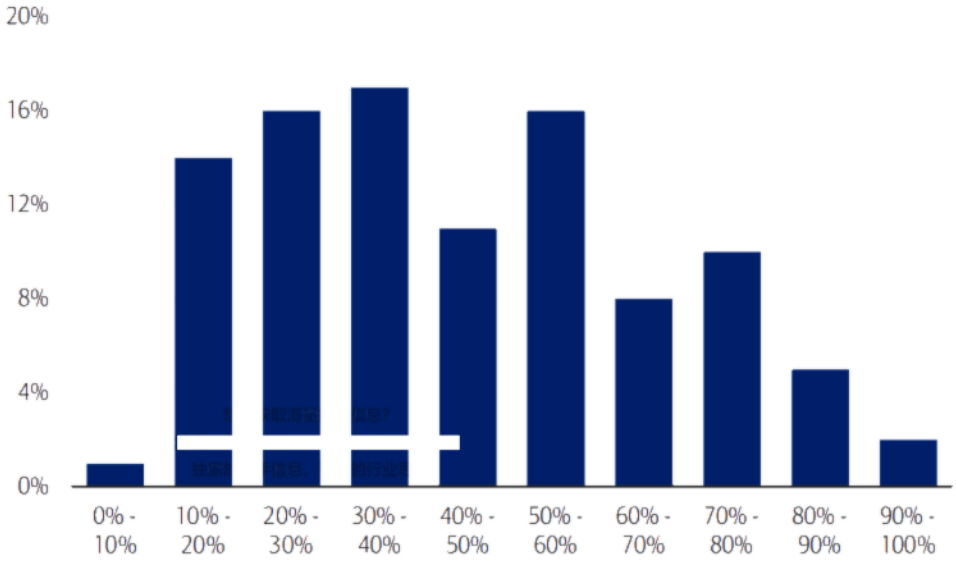
Source: IDC Survey (June 2025, n=150), BofA Global Research

BofA GLOBAL RESEARCH

Customer Relationship Management / Sales

Sales has moved beyond experimentation, with AI agents already in production across multiple implementation approaches. The shift is now toward orchestrating broader sales workflows, including lead qualification, knowledge management, proposal generation, coaching, and deal progression, rather than automating isolated tasks.

Exhibit 9: Agentic AI adoption in Sales appears relative mature relative to other corporate functions
Q: For what proportion of processes across all of marketing have agentic AI initiatives been started?

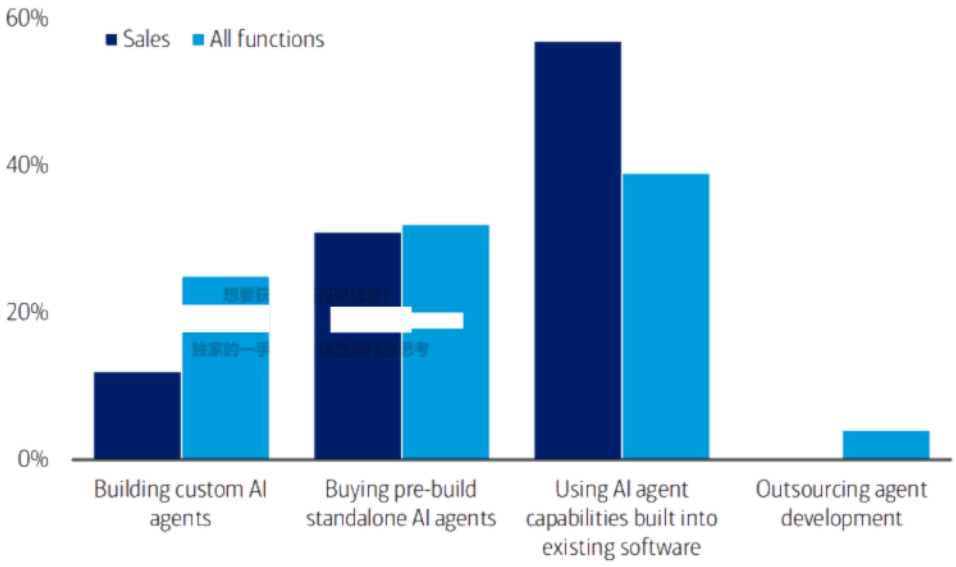


Source: IDC Survey (March 2026, n=145), BofA Global Research

BofA GLOBAL RESEARCH

While sales sees broad applicability for agentic AI, scaling is constrained by a number of factors, including data quality and governance, security and compliance, skills gaps and reliability, among other factors.

Exhibit 10: Sales is favoring embedded and prebuilt agentic AI approaches
Q: What is or will be your organization's usual approach to implement agentic AI within your sales function?



Source: IDC Survey (March 2026, n = 145 for Sales, n=1967 for all functions), BofA Global Research

BofA GLOBAL RESEARCH

According to a recent IDC survey, Sales favor embedded AI capabilities within existing software applications and prebuilt agents over custom development. This should favor existing incumbents, though we note that vendor selection increasingly tied to integration, responsible AI support, road map maturity, and measurable commercial impact.

Vertical players: regulation, system of records

Across industries, meaningful data on agentic AI adoption remains limited, but broader GenAI usage appears uneven by vertical. Adoption is highest in computing, business and finance, and office-based workflows, while more labour-intensive functions continue to show lower penetration.

Travel IT – Amadeus, Sabre, Expedia, HBX Group

Travellers have increasingly adopted AI for trip planning over the past few years, although its use for bookings remains limited. Early personal agents, including OpenClaw and Hermes, primarily served tech-savvy users. More recent launches, including Meta's Muse, ChatGPT's Dots and Instinct, have broadened the appeal of booking travel through AI. Travel is a key use-case for personal AI, with Instinct reporting half the transaction volume it processes relates to travel. Our analysis highlights several inventory-access models, ranging from browser-based website access to contractual API integrations. While the model is still evolving, we believe personal agents could durably reshape how travellers discover and compare options at the top of the distribution funnel. (see note here for further details: Personal AI agent reshapes travel distribution)

For airline IT, Recent case studies from airlines and IT service providers highlighted faster development alongside lower costs from AI adoption. Most airlines remain reluctant to commit to Offer/Order migration: only c. 6% of passengers boarded have committed to moving from existing PSS (Passenger Service System), flat yoy. We also saw more demonstrations of airlines working with systems integrators and IT service providers to access AI capabilities and specialist talent. However, we see few signs that the airline IT transition will prompt airlines to outsource existing in-house solutions. Improved development productivity may instead enable them to bring more technology work in-house over time. (see note here for further details: Airline tech conference feedback: AI dominates the debate)

CAD software: frontier model does not threaten CAD software yet

In early Sep, OpenAI's launch of GPT-6 Astra showcased its enhanced 3D modelling capabilities. Viral online demonstrations from early-access users included transforming a single house photograph into a detailed photorealistic 3D model and navigable environment, generating exploded views of complex objects, and producing circuit-board designs through electronic design software. Following increasing AI investment in the industrial software space, with the announcements by Mistral AI (partnership with a number of larger DSY customers including Airbus, BMW) and Jeff Bezos' Prometheus $12bn fund raising to build an "artificial general engineer", this raises further concerns on increasing disintermediation in a segment initially seen as more insulated from AI risk. Overall, we believe AI will materially reshape CAD software and user interaction; and incumbents that fail to evolve can become obsolete quickly. Further iterations by OpenAI and other AI native players will create credible alternative for at least part of industrial design workflows. However, software providers that successfully embed AI will likely strengthen their value proposition and can capture meaningful medium-term upside.

While the progress is significant, we believe the outputs can appear more disruptive than the underlying reality, with the market overlooking cost inefficiencies, reliance on existing software and the relatively standardised nature of some tasks demonstrated. Many advanced AI outputs still depend heavily on specialist software to complete the underlying work. AI can interpret intent and propose a design, but CAD tools create precise geometry, simulation engines validate performance, and circuit-design software applies engineering rules. Purpose-built software remains better suited to deterministic tasks requiring accuracy, consistency and verification, while refinement, compliance, production validation and engineering judgement remain difficult to automate. Over time, we expect vendors to embed AI optimised for specific workflows, using efficient specialist models or software layers that guide frontier models through complex processes. This should deliver more reliable and cost-efficient results than asking general-purpose models to discover each workflow independently, while reducing the technical burden on users.

Construction software: trust and data readiness remains with incumbents

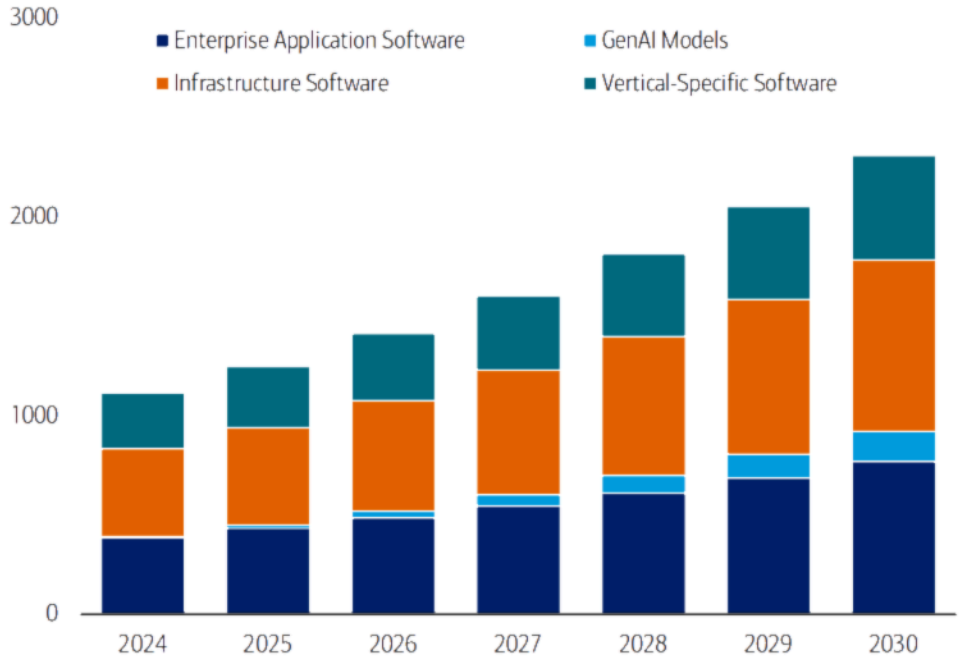
We expect existing AEC software players are unlikely to be disrupted by AI because the industry's workflows are exceptionally complex, regulation-heavy, and deeply embedded in long-standing BIM ecosystems. BIM remains the most widely adopted digital backbone in AEC, establishing a high switching-cost environment that naturally protects incumbents such as Autodesk and Nemetschek. Furthermore, practitioners consistently note that AI still needs to evolve before it can match the maturity and specificity of BIM-centric workflows, reinforcing that AI is currently an enhancer—not a replacement—of incumbent platforms. These technical and structural realities make it difficult for new players, especially AI-only entrants, to gain traction in mission-critical design and engineering environments.

Moreover, the biggest barriers to adopting new digital tools are training, integration complexity, and the risk of disrupting billable project delivery—factors that favour incremental upgrades from trusted vendors rather than wholesale adoption of new systems. AI adoption in AEC remains focused on augmenting tasks such as predictive analytics, project management, and design optimization—not core engineering judgments, which remain highly regulated and liability-sensitive. As key incumbents, Autodesk and Nemetschek are already integrating AI into their existing software stacks, and is well positioned to capture the benefits of AI while preserving customer lock-in. This dynamic ensures that on balance AI strengthens, rather than threatens, the competitive position of established AEC software vendors. (see our note: Autodesk University 2026).

IT budget: zero sum or expansion?

Enterprise application software spending is expected to grow at a 13.2% CAGR over 2025–30. Even after factoring in potential cannibalisation as AI platforms capture a share of spend, traditional application markets are still expected to grow at a 10.9% CAGR, suggesting AI reshapes how spending is distributed rather than shrinking the overall opportunity.

Exhibit 11: Global software expected to grow 13% CAGR FY25-30E
Global software spend forecast (2024 to 2030)



Source: Gartner

BofA GLOBAL RESEARCH

The Jevons' paradox in AI

Jevons' paradox is the idea that when a technology makes something cheaper, faster or easier to use, total demand for that capability often rises rather than falls. In technology, this has happened repeatedly. As internet bandwidth improved over time, usage did not plateau; global internet traffic in 2014 was >2bn times late 1980s levels because better, cheaper connectivity unlocked entirely new applications and behaviours. Similarly, as cloud lowered infrastructure costs and improved scalability, software did not become less important; software's share of tech budgets still rose during the period.

We think AI is likely to follow the same pattern in enterprise software. As intelligence becomes cheaper and more accessible, companies should automate more workflows, generate more content, analyse more data and embed software into more day-to-day decisions. That does not mean enterprises will bypass vendors and build everything themselves. In practice, the difficult part is rarely the model alone, but integrating AI into the organisation. Our view, therefore, is that AI should increase software consumption overall, but value should accrue disproportionately to platforms that already own workflow, data and trust, rather than to standalone model access alone.

Exhibit 12: Software spend expect to grow double digits while GenAI models lead the way
Software CAGR growth 2025 to 2030 by type across verticals

Industry	Application Software	GenAI Models	Infrastructure Software	Vertical-Specific Software
Banking and Investment Services	15.0%	59.5%	12.9%	13.8%
Communications, Media and Services	12.1%	57.2%	10.8%	12.2%
Education	10.7%	48.1%	12.3%	12.0%
Government	10.6%	59.3%	11.0%	9.8%
Healthcare	11.7%	58.8%	11.0%	5.8%
Insurance	13.6%	58.7%	14.6%	8.9%
Manufacturing and Natural Resources	10.7%	59.2%	11.1%	11.1%
Oil and Gas	11.6%	60.2%	13.7%	10.3%
Other Industries	11.6%	58.8%	11.4%	12.0%
Power and Utilities	14.5%	58.4%	14.6%	16.5%
Retail	9.7%	58.4%	10.0%	9.7%
Transportation	10.1%	59.3%	9.8%	10.4%
Total	**13.2%**	**58.2%**	**11.9%**	**11.3%**

Source: Gartner, BofA Global Research

BofA GLOBAL RESEARCH

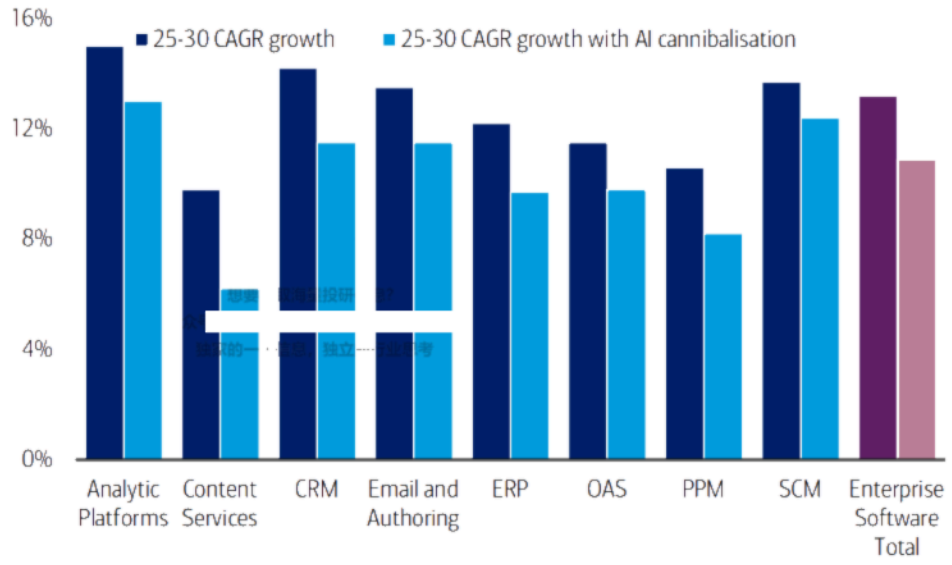
Enterprise software growth remains resilient despite cannibalisation risk

According to Gartner's latest forecast (September 2026), it now expects the Enterprise software market to grow 13.2% CAGR over 2025–2030. Growth spans major categories, with baseline CAGRs of 14.2% for CRM, 12.2% for ERP and 13.7% for supply chain management. Importantly, its separate scenario incorporating cross-functional agentic cannibalisation still shows the traditional application markets growing 10.9% CAGR. Our interpretation is that AI changes the distribution and pace of application spending more than it eliminates growth: the debate concerns which software categories or products retain enterprise budgets, rather than whether enterprise software demand disappears.

AI adoption is moving towards cross-enterprise execution

Currently, AI deployments concentrate on discrete activities such as invoice reconciliation, customer service and inventory management, where outcomes are measurable and governance is relatively contained. Gartner expects the next phase to involve coordinated agents executing workflows across CRM, ERP, analytics and supply chain systems, rather than operating within individual applications. Buyers would increasingly prioritise end-to-end business outcomes over improvements within a single department, redirecting spending towards cross-functional orchestration. Our interpretation: we believe AI can substitute for some application spending while expanding demand for the software foundations required to operate autonomous workflows reliably.

Exhibit 13: Enterprise software growth 2.3pp lower if cannibalization happens
Impact of cross functional cannibalization on Enterprise Software



Source: Gartner, BofA Global Research
CRM = Customer Experience and Relationship Management, ERP = Enterprise Resource Planning, OAS = Other Applications Software, PPM = Project and Portfolio Management, SCM = Supply Chain Management

BofA GLOBAL RESEARCH

The application-growth headwind is back-end loaded

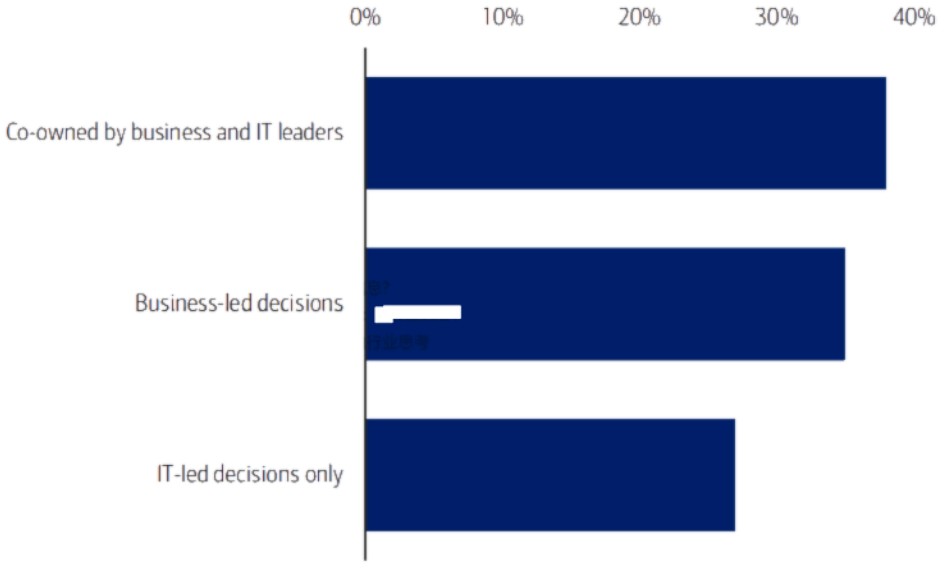
Gartner estimates $116bn of cumulative cannibalisation from standalone application markets between 2026 and 2030. Annual cannibalised-spending rises from $0.4bn in 2026 to $2bn in 2027, $11bn in 2028, $25bn in 2029 and $79bn in 2030. Gartner also anticipates revenue pressure from 2028 for agentic providers confined to a single function or market. This represents spending redirected towards cross-functional offerings, rather than an equivalent disappearance of total software expenditure. Our interpretation is that exposure depends on whether vendors participate in cross-functional execution or remain confined to increasingly substitutable functional capabilities.

Furthermore, cross-functional agentic applications are expected to account for more than 10% of application-software spending, with value concentrating among vendors controlling orchestration, enterprise data foundations and workflow execution. Adoption remains dependent on governance, security, compliance and organisational change. Our interpretation is that large incumbents with integration, trusted data access and orchestration capabilities can participate in the transition, while narrower vendors face greater pressure. The key distinction is between losing functional spending and capturing the emerging execution layer.

Agentic AI investment is more business-owned (vs traditional IT spend)

Agentic AI appears more business-owned than traditional software procurement, although this does not necessarily prove line of business (LoB) budget ownership. According to the IDC survey below, 73% of agentic AI investment decisions are either co-owned by business and IT leaders or led by business stakeholders. This suggests agentic AI spending may increasingly be justified against LoB productivity pools (e.g. sales, marketing, engineering, customer service or finance) rather than being viewed solely as centralised IT spend. This could reduce the degree of budget displacement from incumbent software spend, as some AI investment may be funded or justified against LoB productivity outcomes rather than competing solely within centralised IT software budgets.

Exhibit 14: Majority of agentic AI investment decisions are co-owned or led by line of business
Q. Who is making the investment decisions for agentic AI initiatives within your function?



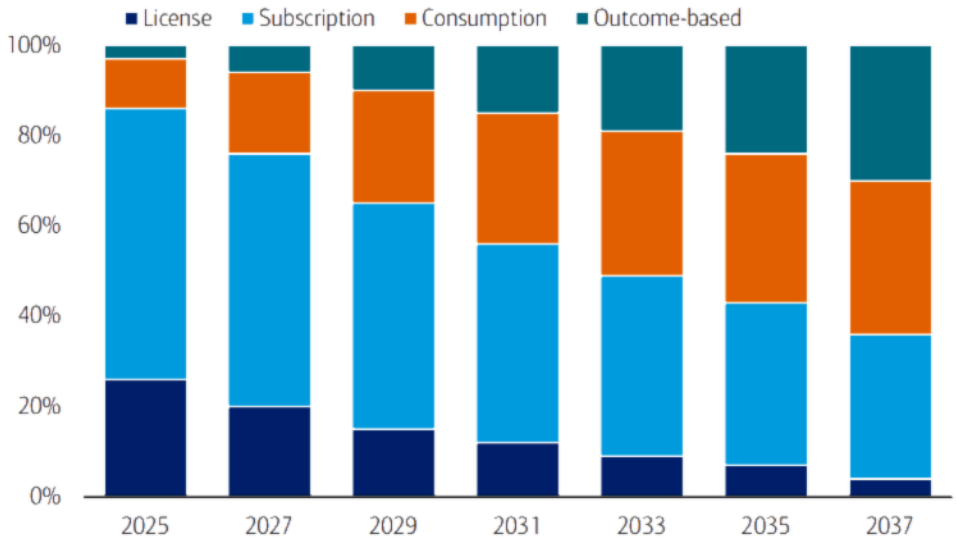
Source: BofA Global Research, IDC's Early Adopter Agentic AI Functional Use Case Study, March 2026 (n = 1,967)

AI economics for enterprise software

A key concern from investors is whether seat-based pricing can survive the shift towards AI agents, which may perform more work without increasing user numbers. We see software providers generally responding with consumption and outcome-based models. However, changing capabilities, costs and customer expectations mean pricing will remain experimental, with customer resistance potentially slowing the transition from subscriptions.

Pricing model: from subscription to consumption

Exhibit 15: Enterprise software by pricing model
We expect consumption and outcome-based pricing to grow in the next 10 years

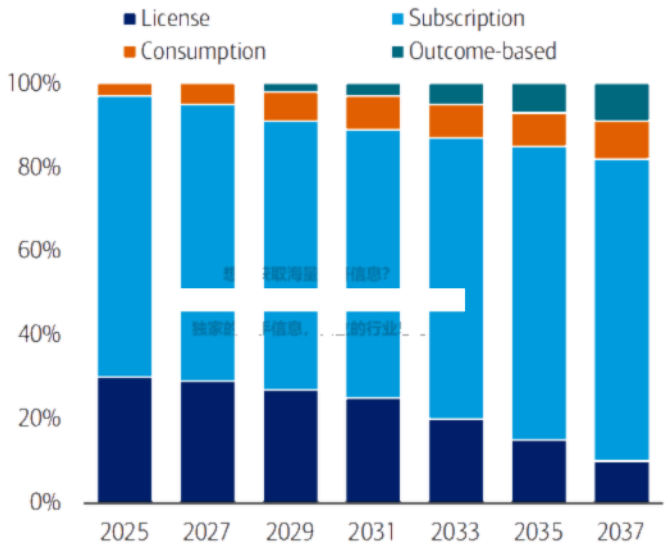


Source: BofA Global Research estimates, IDC

BofA GLOBAL RESEARCH

IDC expects uneven disruption by application category. Marketing, HCM, Financials, and SCM are highlighted as among the fastest markets to introduce agentic capabilities, with Marketing and SCM advancing towards fully agent-led workflows faster than HCM and Financials. In contrast, PLM, Digital Commerce, EAM/ALM, and Salesforce performance & productivity are characterised as laggards in embedding agentic capabilities to date—although sales force performance & productivity and EAM/ALM should see a significant increase in embedded agentic functionality over the next two years. In lay terms, the early movers tend to be domains with high-volume workflows and repeatable decisions where agents can take clear actions (e.g., optimise campaigns or manage supply chain exceptions). Later movers are often areas where workflows are more specialised, heavily governed, or tightly coupled to complex operational processes, slowing the path from "assist" to "act".

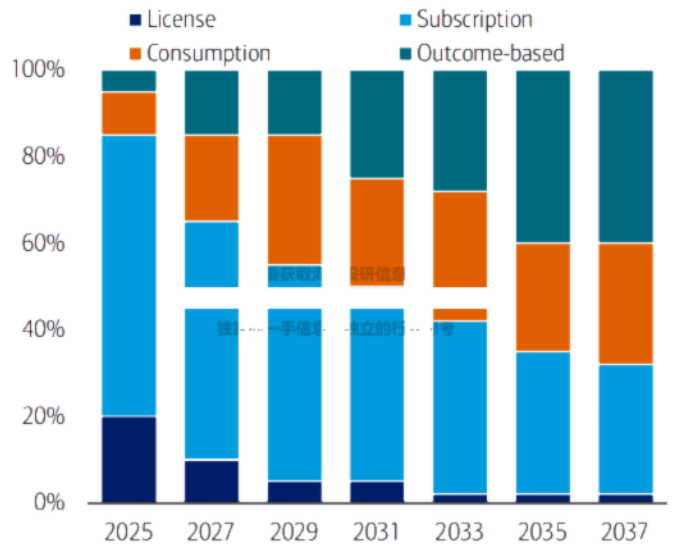
Exhibit 16: ERP software by pricing model (2025 to 2037)
Consumption & outcome-based pricing to account for 18% by 2037



Source: BofA Global Research estimates, IDC

BofA GLOBAL RESEARCH

Exhibit 17: CRM software by pricing model (2025 to 2037)
Consumption & outcome-based pricing to account for 68% by 2037



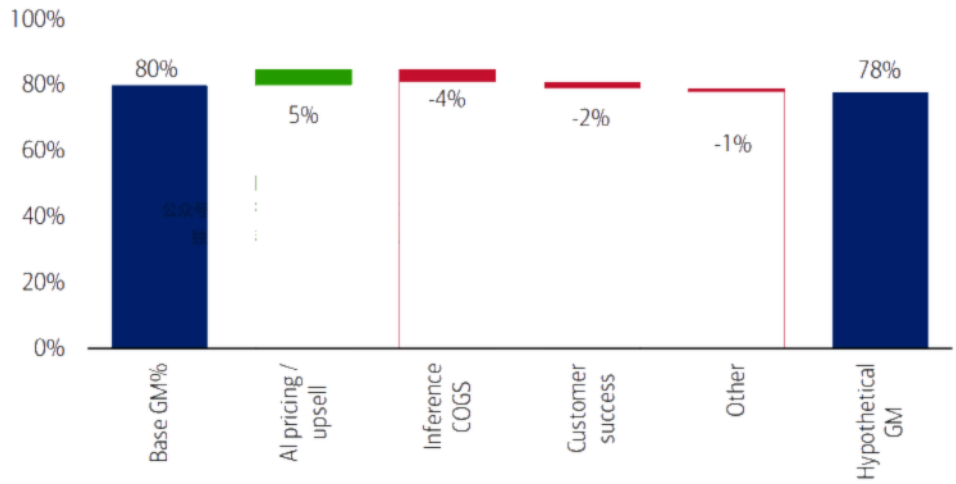
Source: BofA Global Research estimates, IDC

BofA GLOBAL RESEARCH

AI Margins, Platform Defensibility and MCP

We believe AI does not necessarily imply a structural gross margin reset for enterprise software. The near-term impact is more nuanced. Vendors face incremental costs from inference, retrieval, third-party models and infrastructure during early customer ramp-up when usage is less predictable and systems are still being tuned. That can create temporary pressure on gross margin, especially where pricing is not yet fully aligned to consumption. Over time, however, the economic model should improve as vendors optimise model routing, reduce unnecessary token usage, introduce more deterministic workflows, and shift more tasks away from expensive general-purpose reasoning towards narrower, lower-cost execution paths. We walk through how we think gross margins can evolve for a software provider below.

Exhibit 18: Hypothetical gross margin waterfall with AI rolled out widely
Base gross margin to hypothetical gross margin with the impact of AI



Source: BofA Global Research estimates

BofA GLOBAL RESEARCH

Key factors include:

- **AI monetisation:** In the medium term, the main upside should come from AI being embedded into higher-value plans, premium workflows and usage-based modules. In a scaled state, the key assumption is that AI is no longer given away to drive adoption, but supports sustainable ARPU uplift and a structurally better revenue mix. Please refer to our AI revenue disclosure section for details.

- **Model efficiency and infrastructure optimisation**: At scale, gross margin should benefit from lower unit costs through better model routing, caching, prompt optimisation, improved utilisation and lower third-party model pricing. We expect these benefits to only become material as AI matures, and when usage patterns are predictable and engineering teams have had time to optimise.

- **Customer success**: As AI deployments become more complex, software providers will need stronger customer-success teams to redesign workflows, integrate data, train users, monitor outcomes and turn promising pilots into reliable, scaled production deployments.

- **Inference and model-serving COGS**: Even in a mature state, inference remains the largest structural headwind because AI features consume compute on every interaction. The steady-state question is whether pricing and usage monetisation more than offset this recurring delivery cost, particularly as customers expect frequent, always-on AI functionality across the product suite.

- **Data, retrieval and platform delivery costs:** Scaled AI products often require more than model inference alone, including vector databases, retrieval, embeddings, storage, logging and data transfer. These costs may become more efficient over time, but they remain a persistent gross margin drag because they are integral to delivering enterprise-grade AI reliably and at low latency.

In most enterprise settings, the commercial value sits not in raw AI model output, but rather higher in the stack: workflow integration, permissions, governance, testing, monitoring, reliability and business context. If that remains true, AI monetisation should extend beyond simple cost pass-through. If software vendors can price against business outcomes, productivity and embedded workflow value, AI can potentially remain margin-neutral over time, even if underlying model usage is meaningful.

The central debate for investors is therefore where value accrues as the stack matures. The bear case is that model capability improves quickly, costs fall, and competition compresses pricing such that application vendors retain little economic rent. In that scenario, AI features become table stakes and gross margin comes under pressure as infrastructure costs remain visible while differentiation narrows. The bull case is that enterprise software vendors preserve pricing power because the hard problem is not intelligence in isolation, but production-grade implementation across existing systems of record and systems of work. In that world, models commoditise faster than workflow, context and distribution.

MCP and similar open connectivity standards should be viewed through that lens. MCP can reduce integration friction and make it easier for agents, models and enterprise systems to exchange context. That is strategically important, but interoperability is a building block rather than a complete architecture. The harder layer remains orchestration inside real business processes: deciding which agent acts, with what permissions, against which data, under what controls, and with what latency, accuracy and auditability. Open standards may broaden the ecosystem, but they do not by themselves remove the need for a trusted application layer. If anything, they may increase the value of platforms that already sit at the intersection of data, workflow and user activity.

AI revenue disclosure across US / EU

AI revenue disclosure remains inconsistent across the software sector and is generally more developed among US vendors than European peers. European software providers currently focus more on usage metrics rather than revenue contribution. Revenue disclosures across the two groups suggest AI generally accounts for low to mid-single digit of group revenues, within a broad c.1–11% range. However, we believe reported figures understate the true contribution, given that AI monetisation is often embedded within bundles or captured through broader price increases. Medium-term targets remain scarce, although companies providing them expect AI to represent c.5–30% of group revenues.

Exhibit 19: AI revenue disclosure across US and EU software companies
AI revenue disclosure remains inconsistent across the software sector

Company	Current AI revenues	AI medium term target	Pricing Model			Leading AI Products
			Embedded	Seat add-on	Consumption	
US						
Adobe	AI-first ARR >$650m, up >150% yoy	see it as next billion dollar business		✓	✓	Firefly, Acrobat AI Assistant, GenStudio
Atlassian	N/A (Rovo in >80% of F500 companies)	N/A	✓		✓	Rovo, Jira Coding Agent
Autodesk	N/A	N/A	✓		✓	Autodesk Assistant, NeuralCAD, Forma
Cadence	N/A	33% EDA share of R&D budgets long term		✓	✓	ChipStack, ViraStack, AgentStack
Datadog	AI-native customers: 12% of revenue (3Q25; metric since discontinued)	N/A	✓		✓	Bits AI, Agent Observability, LLM Observability, MCP server
DocuSign	10.8% of ARR	18% to 19% of ARR by FY27 exit		✓	✓	IAM, Iris, Agent Studio
Dropbox	N/A	N/A	✓			Dash, Dropbox Protect
Dynatrace	AI adoption (not revenue): 1,000 customers monitoring AI applications in production; 800 using agentic capabilities (F1Q27)	N/A	✓		✓	Dynatrace Intelligence, MCP server
Elastic	N/A	AI to add 5% yoy growth contribution	✓		✓	Agent Builder, Jina embedding
Gitlab	Duo Agent Platform: ~$22.5mn ARR (F2Q27), ~2.1% of our ARR estimate for that quarter	Modest contribution in FY27	✓	✓	✓	Duo Agent Platform, Orbit
Hubspot	N/A (19% of Pro+ customers use agents)	N/A	✓		✓	Breeze Assistant, Customer Agent, Prospecting Agent
Intuit	$146m IES run rate	N/A	✓	✓		Intuit Enterprise Suite, Agent Studio
Microsoft	>$37bn AI run rate	N/A		✓	✓	M365 copilot, Github Copilot, Foundary, Agent365
MongoDB	AI-native customers: 3%-5% of Atlas revenue (Investor Day); 165% y/y growth	N/A	✓		✓	Vector database (Voyage AI), MCP server, Atlas Agent Engine, Atlas Vector Search
Salesforce	3.6% of group revs	N/A		✓	✓	Agentforce, Aiforce, Data 360
ServiceNow	(Now Assist ACV >$1bn	30% of ACV by 2030		✓	✓	Now Assist, AI Control Tower, Autonomous Workforce
Shopify	N/A	N/A	✓			Sidekick, Agentic Storefronts, UCP
Snowflake	$100m in AI ARR				✓	Snowflake CoCo, CoWork, Cortex AI
Synopsys	N/A (>50 agentic engagements)	N/A		✓	✓	AgentEngineer, Autopilot Platform, GPT-Synopsys
Twilio	N/A (Voice AI rev up 60% yoy)	N/A	✓		✓	ConversationRelay, Conversational Intelligence, Agent Connect, Stytch
UIPath	AI ARR >$250m, >50% yoy)	N/A	✓		✓	Agent Builder, Maestro, IDP
Veeva	N/A	Not material in FY27	✓		✓	Vault AI agents, Falcon, Falcon MLR
Workday	6.3% of group revs (AI ARR $600m)	N/A		✓	✓	Illuminate agents, Sana, Extend Pro

Exhibit 19: AI revenue disclosure across US and EU software companies
AI revenue disclosure remains inconsistent across the software sector

Company	Current AI revenues	AI medium term target	Pricing Model			Leading AI Products
			Embedded	Seat add-on	Consumption	
Zoom	N/A	N/A	✓	✓	✓	AI Companion, Zoom Virtual Agent, ZoomMate
Europe						
Amadeus	N/A	N/A	✓	✓		Nevio, Amadeus Max
ATOSS	N/A	N/A	✓	✓	✓	ATC AI Agent
Dassault Systems	50m to 100m in FY26	5% of software revs by 2029			✓	Vertual Companions, Generative Experiences, 3D UNIV+RSES
IONOS	N/A	80% of incremental WPP revs by FY28	✓	✓		AI Phone Receptionist, AI App, Site Builder, Momentum
Nemetschek	mid-single-digit EUR m ARR in 2026	N/A		✓		Bluebeam Max, Firmus AI
Planisware	N/A	N/A	✓	✓	✓	Prisma, Oscar
Sage	N/A	N/A	✓	✓	✓	Sage Copilot, AP automation, Finance Intelligence Agent
SAP	AI in >90% of top 50 deals	N/A	✓	✓	✓	Joule, Business AI Platform
Sabre	N/A	N/A	✓	✓		Agentic APIs and MCP Server, SabreMosaic
TeamViewer	N/A	N/A	✓		✓	Tiia agent, AI Session Insights
Temenos	N/A	N/A				FCM AI Agent, Copilot for Core, Intelligent Core

Source: BofA Global Research estimates, Company

BofA GLOBAL RESEARCH

Software providers currently monetise AI in three main ways:

- **Included in the existing plan:** Customers pay no separate AI fee. Vendors use inclusion to encourage adoption and may seek to recover the cost through higher list prices over time.
- **Per-user add-on or premium tier:** Some or all AI functionality is reserved for an additional fee or higher-priced tier. This creates a direct revenue uplift, but monetisation does not increase with usage.
- **Consumption-based pricing:** Customers purchase credits or usage units, which are consumed by AI actions. This links vendor revenue more directly to usage, but makes spending less predictable and harder for customers to budget.

Consumption pricing is constantly evolving

Software providers continue to experiment with what they measure and monetise, ranging from model tokens to actions, conversations and completed outcomes. Pricing structures are also shifting, with vendors adjusting included allowances, overage terms and discounts, sometimes within existing contracts. Early moves towards outcome-based pricing are emerging, including charges per resolved customer conversation.

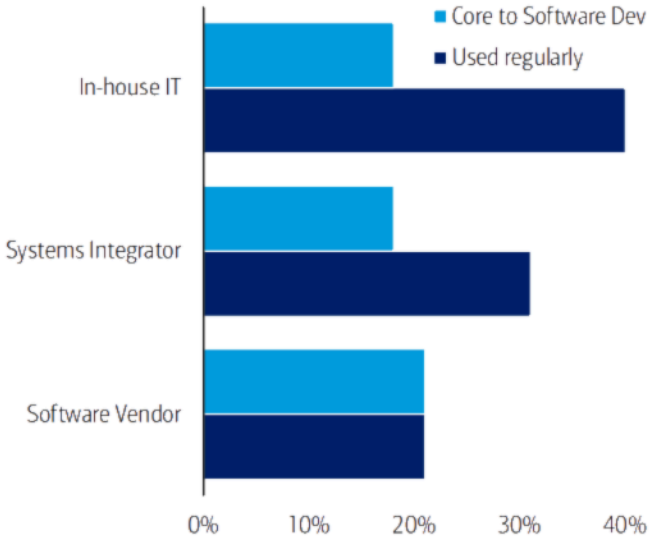
Overall, we expect AI pricing models to remain fluid in the near term, as rapid improvements in AI capabilities continue to reshape both customer value and the underlying unit economics.

Productivity gains & bottlenecks

Agentic AI is delivering measurable gains in targeted tasks, but company-wide adoption remains difficult. Many customers are still moving from pilots towards broader deployment, creating a gap between technical capability and commercial value. For software providers, monetisation will therefore depend not just on better agents, but on helping customers embed them into everyday work and demonstrate clear, repeatable returns.

Productivity gains for IT players...

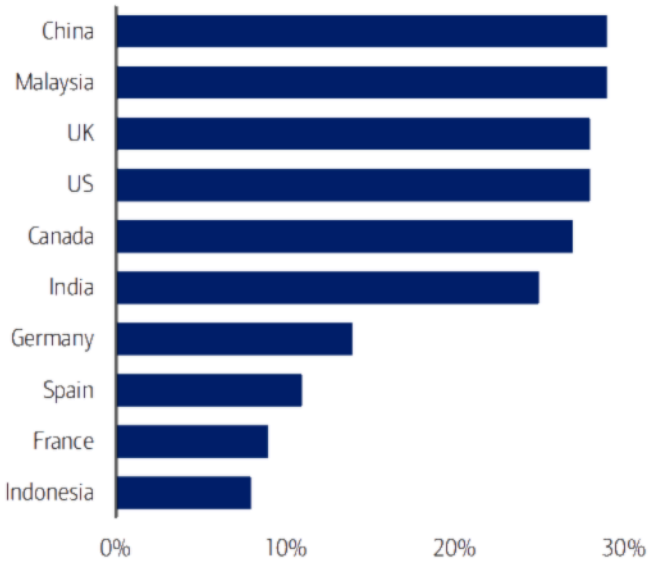
Exhibit 20: Agentic AI adoption for software dev varies by org type...
Q: To what extent have you or your team used agentic tools in your work?



Source: BofA Global Research, IDC DevOps Survey (Dec 2025, n=508)

BofA GLOBAL RESEARCH

Exhibit 21: ...and regions, with many EU markets showing low adoption
Q: To what extent have you or your team used agentic tools in your work?



Source: BofA Global Research, IDC DevOps Survey (Dec 2025, n=508)

BofA GLOBAL RESEARCH

The pace of adoption is meaningful

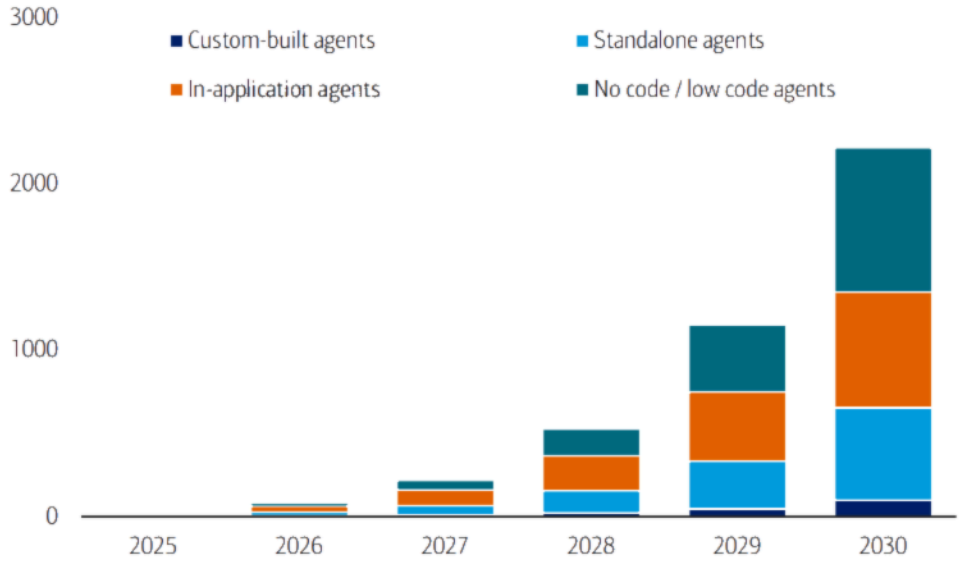
According to a recent IDC Survey, evidence suggests developers are already broadly engaging with these tools, particularly for ideation, experimentation and productivity gains. At the same time, the data suggests organisations are not treating this as a frictionless rollout. Security and compliance concerns remain significant, as do workflow integration, maintainability, version control and uncertainty around returns on investment.

AI can help generate code, but vendors remain critical in making agentic coding enterprise-ready, auditable and cost-efficient.

From an investment perspective, that creates one useful way to think about winners and losers. The strongest positioning is unlikely to come only from having the most impressive demo, but from combining productivity gains with governance. Tools that can show clear time savings, integrate into existing development environments and provide

visibility over what the AI changed, why it changed it and how it can be rolled back are more likely to win enterprise budgets. More broadly, the category should benefit vendors that can embed AI deeply into the software development lifecycle rather than treat it as a standalone assistant.

Exhibit 22: We are still in the very early days of scaling AI agents
Worldwide agents count (millions) by type, 2025-2030



Source: IDC

BofA GLOBAL RESEARCH

- **In-application agents** are seeing the fastest early uptake because they are easier to add into existing SaaS software and fit naturally into how companies already buy technology. They are also likely to be the main starting point for smaller businesses. Over time, however, their usefulness may be constrained as larger companies look to automate more complex workflows that span multiple functions. As a result, custom-configured and standalone agents are likely to become more important over the next few years as companies build broader agent systems of their own.

- **Custom-configured agents** built using low-code and no-code tools are likely to become the most widely deployed agent type by 2029. They allow business users to create agents tailored to their own needs, while pulling together information and actions from several different software applications. This should benefit software vendors that offer embedded agent-building tools, although competition is also likely to increase from independent platforms that specialise in coordinating agents across systems.

- **Standalone agent** deployment is still relatively limited today because the category is at an early stage, but adoption could increase meaningfully as agent-to-agent (A2A) standards improve. Over time, these agents are likely to be used mainly for industry-specific or process-specific tasks that need to work across multiple software ecosystems. In many cases, they could offer a cheaper and more scalable option than building highly customised agents from scratch.

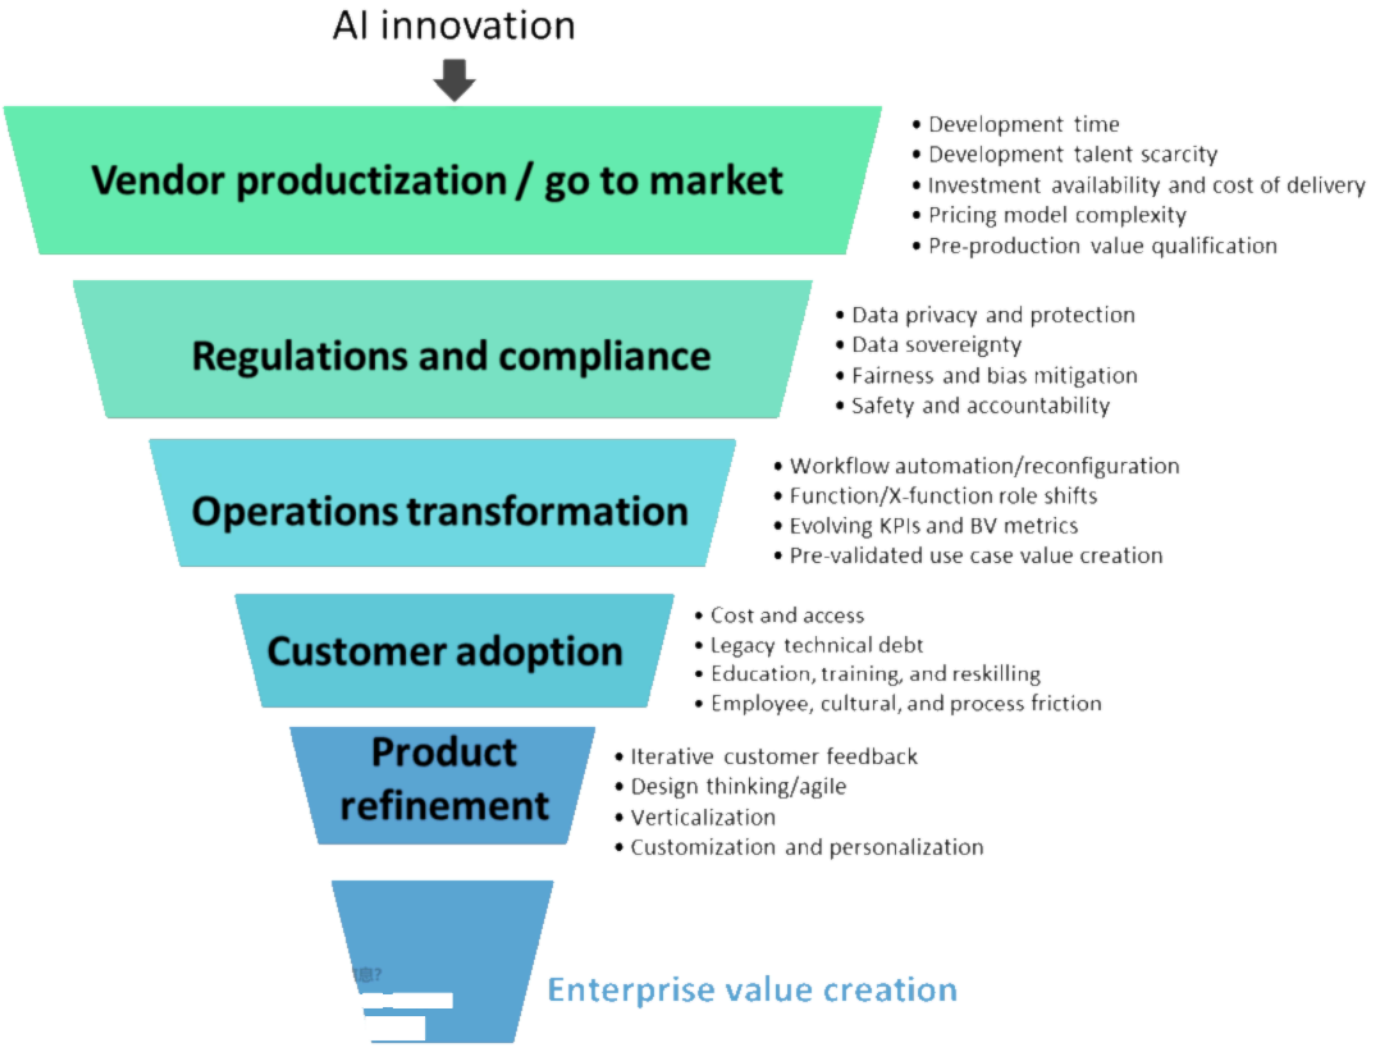
- **Bespoke agents** are likely to remain a smaller part of the market and will mainly be used for proprietary, highly secure or sovereign workloads. They are usually the most expensive to build and run, partly because they require dedicated infrastructure and stricter compliance controls. Even so, they should remain strategically important in situations where organisations place a high value on control, exclusivity and data protection.

...but bottlenecks for most non-IT organizations

Assistance is easy, workflow redesign is harder

AI value creation in enterprise software is progressing more slowly than the pace of innovation would imply because the bottleneck has shifted from model capability to enterprise execution. In practice, value is delayed by a chain of frictions spanning vendor productization, governance, customer adoption, workflow redesign and ongoing refinement. The result is that many organisations capture early productivity benefits, but struggle to scale those gains into broader, measurable improvements in throughput, decision quality and operating performance.

Exhibit 23: Enterprise value creation lags behind AI innovation
The speed of AI value creation



Source: BofA Global Research, IDC

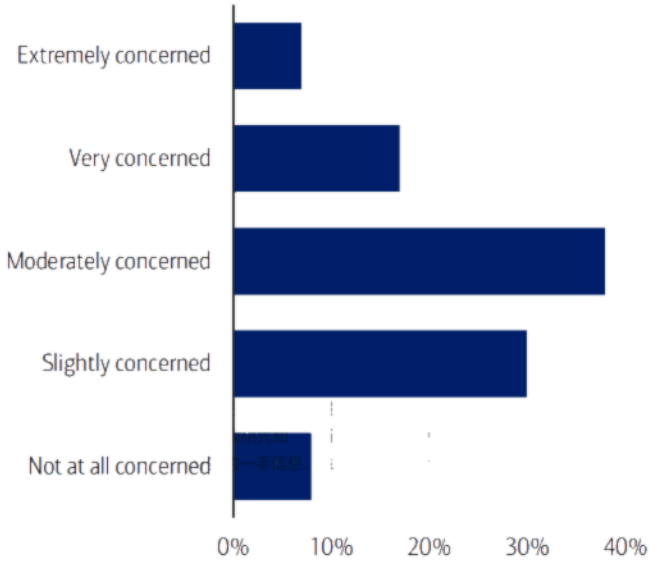
BofA GLOBAL RESEARCH

Five points on the funnel:

- **Vendor productization and monetisation** — AI capability has become a leading software buying criterion, and a meaningful share of customers are willing to pay a premium for the right functionality. Yet supply remains constrained by development complexity, specialist talent scarcity, rising delivery costs and evolving pricing models, while premature launches can widen the gap between product ambition and realised customer value.

- **Regulation, compliance and trust** — As AI moves closer to core decision-making, privacy, protection, sovereignty, explainability and accountability become gating factors rather than side issues. These barriers are manageable, but they can slow the transition from assistive use cases to more autonomous deployment, particularly in regulated sectors and where governance remains fragmented.
- **Operational transformation** — The sharpest loss of value typically occurs when AI is deployed without redesigning the underlying workflow. In that setting, AI improves individual tasks but leaves end-to-end processes, decision rights and cross-functional operating models largely intact, limiting impact to incremental gains. Legacy KPIs can compound the issue by under-measuring improvements in speed, throughput, execution quality and decision effectiveness.
- **Customer adoption** — Slow uptake is usually more structural than cultural. Technical debt, data-access constraints, scaling costs, training gaps, leadership misalignment, process friction, underfunded change management and weak user experience can all prevent AI from becoming habitual and mission-critical. That helps explain why many enterprises still see pockets of experimentation rather than broad operational deployment, with agentic AI remaining far from enterprise wide in most organisations.
- **Product refinement** — Value creation does not end at general availability; it depends on continuous iteration after launch. Customer feedback, agile development, verticalization and configurable workflows are central to improving relevance, reducing friction and shortening time to value. Vendors that refine around real workflows and operating requirements, rather than feature volume alone, are better positioned to sustain adoption and convert innovation into measurable outcomes.

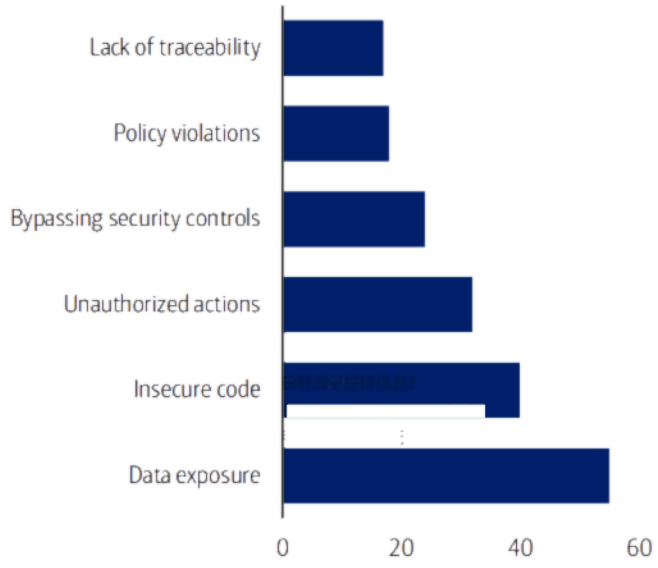
Exhibit 24: Concerns on agentic AI security risks remains...
Q: How concerned is your organization about security risks with agentic AI in software development workflows?



Source: IDC Survey, BofA Global Research

BofA GLOBAL RESEARCH

Exhibit 25: ...with data exposure the #1 concern
Q: Which of the following security risks are of greatest concern to your organization when it comes to agentic AI tools in software development workflows?



Source: IDC Survey, BofA Global Research

BofA GLOBAL RESEARCH

The next phase of AI differentiation in enterprise software is unlikely to be defined simply by who launches the most features, or even by who has the strongest access to models. More durable advantage should accrue to vendors that can validate use cases before launch, align pricing with customer value, simplify deployment, verticalize effectively and refine products continuously once they are in market. On the customer side, faster value realisation should favour organisations willing to redesign workflows, refresh KPI frameworks and embed clearer governance and accountability around AI-enabled execution. In short, the real separator is not innovation at the top of the funnel, but the ability to operationalise that innovation deeply enough for value to compound at scale.

The rise and fall of "tokenmaxxing"

The rise of "tokenmaxxing" reflected a growing assumption that more token usage meant greater productivity, particularly as agentic AI moved into everyday coding workflows. However, tokens measure model activity, not useful business output: usage is difficult to compare across vendors and can reward verbose or inefficient behaviour. For enterprises, adoption therefore depends less on raw usage than on confidence in costs, accuracy and measurable returns. This matters because heavy token consumption may signal experimentation, but without clear outcomes it risks becoming another expanding cost centre.

Software providers are responding with measures focused on work completed rather than compute consumed, such as Salesforce's Agentic Work Units. This is directionally more useful, but completed tasks must still be assessed against quality and business impact. Sustainable adoption will require software providers to demonstrate lower cost per useful task, faster processes, fewer errors and more predictable returns. Ultimately, enterprises need evidence that agents are improving outcomes, not simply generating more activity.

Glossary

A2A = Agent-to-Agent

ACV = Annual Contract Value

AEC = Architecture, Engineering and Construction

AI = Artificial Intelligence

ALM = Application Lifecycle Management

AP = Accounts Payable

API = Application Programming Interface

ARR = Annual Recurring Revenue

ARPU = Average Revenue per User

AWS = Amazon Web Services

AWU = Agentic Work Unit

B2B = Business-to-Business

BDC = Business Data Cloud

BIM = Building Information Modelling

BV = Business Value

CAD = Computer-Aided Design

CAM = Computer-Aided Manufacturing

CAGR = Compound Annual Growth Rate

COGS = Cost of Goods Sold

CPaaS = Communications Platform as a Service

CRM = Customer Relationship Management

CX = Customer Experience

EAM = Enterprise Asset Management

EDA = Electronic Design Automation

ERP = Enterprise Resource Planning

EU = European Union

FAA = Federal Aviation Administration

F500 = Fortune 500

FCM = Financial Crime Mitigation

GenAI = Generative Artificial Intelligence

GM = Gross Margin

GPT = Generative Pre-trained Transformer

HCM = Human Capital Management



HR = Human Resources

IAM = Intelligent Agreement Management

IATA = International Air Transport Association

IDC = International Data Corporation

IDE = Integrated Development Environment

IDP = Intelligent Document Processing

IES = Intuit Enterprise Suite

IT = Information Technology

ITSM = Information Technology Service Management

KPI = Key Performance Indicator

LLM = Large Language Model

LoB = Line of Business

M365 = Microsoft 365

MCP = Model Context Protocol

MLR = Medical, Legal and Regulatory

PCI = Payment Card Industry

PLM = Product Lifecycle Management

PSS = Passenger Service System

R&D = Research and Development

RCA = Root Cause Analysis

REA = Reverse Engineer Anything

RPA = Robotic Process Automation

SaaS = Software as a Service

SCM = Supply Chain Management

SDK = Software Development Kit

SI = Systems Integrator

SMS = Short Message Service

UI / UX = User Interface / User Experience

UCP = Universal Commerce Protocol

VAR = Value-Added Reseller

WPP = Web Presence and Productivity

Exhibit 26: Stocks mentioned
Prices and ratings for stocks mentioned in this report

BofA Ticker	Bloomberg ticker	Company name	Price	Rating
AMADF	AMS SQ	Amadeus	EUR 54.3	B-3-7
AMADY	AMADY US	Amadeus	US$ 60.67	B-3-7
ADSK	ADSK US	Autodesk	US$ 240.7	B-1-9
BAND	BAND US	Bandwidth	US$ 43.5	C-2-9
DASTY	DASTY US	Dassault	US$ 26.11	B-2-7
DASTF	DSY FP	Dassault Systemes	EUR 23.29	B-2-7
DDOG	DDOG US	Datadog Inc	US$ 260.36	C-1-9
DT	DT US	Dynatrace	US$ 61.54	B-1-9
ESTC	ESTC US	Elastic	US$ 94.6	C-2-9
GTLB	GTLB US	GitLab	US$ 53.29	C-2-9
XHVXF	HBX SM	HBX	EUR 7.725	C-1-9
FROG	FROG US	JFrog	US$ 91	C-1-9
MDB	MDB US	MongoDB	US$ 348.69	C-1-9
NEMTF	NEM GY	Nemetschek	EUR 76.25	B-1-7
XPZAF	PLNW FP	Planisware	EUR 26.85	C-1-7
SABR	SABR US	Sabre	US$ 1.64	C-1-9
SAP	SAP US	SAP	US$ 226.03	B-1-7
SAPGF	SAP GY	SAP	EUR 199.84	B-1-7
CLCMF	SINCH SS	Sinch	SEK 46.16	C-3-9
SNOW	SNOW US	Snowflake	US$ 315.35	C-1-9
TMVWF	TMV GY	TeamViewer	EUR 7.01	C-1-9
TMNSF	TEMN SW	Temenos AG	CHF 72.4	B-1-7
TMSNY	TMSNY US	Temenos AG	US$ 87.12	C-1-7
TWLO	TWLO US	Twilio	US$ 257.29	C-1-9
PATH	PATH US	UiPath	US$ 13.12	C-3-9

Source: BofA Global Research

BofA GLOBAL RESEARCH

Price objective basis & risk

Amadeus (AMADF / AMADY)
We value Amadeus at a price objective (PO) of €53 (ADR: US$62) per share using a discounted cash flow (DCF) approach. Key assumptions in our DCF model are 1) weighted-average cost of capital (WACC) of 7.5%, 2) long-term revenue growth of 1% and 3) 38% in terminal EBITDA margin.

Downside risks to our price objective are: 1) further penetration of the Direct Connect model between travel agencies and airlines, driving more disintermediation, 2) share loss in the company's home European market, 3) mis-execution on a large airline reservations project, and 4) prolonged travel weakness due to the COVID-19 pandemic.
Upside risks to our price objective are: 1) faster air bookings growth due to share gain vs GDS, slower disintermediation or reintermediation, 2) stronger pricing upside and new wins in air IT, 3) revenue per booking continues to grow due to pricing effects and booking mix.

Autodesk (ADSK)
Our $300 PO is based on an EV/FCF multiple of about 21x our CY27 FCF estimate. The multiple is at a slight discount to the design software peer group average of 23x. We believe that a discount multiple is warranted by the company's near-term growth outlook, durable business model, and favorable margin profile, offset by the overhang associated with elevated go-to-market risks from the recent reduction in force.

Downside risks to our PO are increased competition leading to higher churn, greater-than-expected headwinds to FCF in association with Autodesk's billing transition, spending reductions in the company's end markets arising from an economic slowdown, and lower-than-expected net subscriber additions.

Upside risks to our PO are better-than-expected economic activity, greater-than-projected net subscriber additions and net revenue retention, lower-than-expected impact of billing transition on FCF, and better-than-expected execution leading to strong operating margin expansion.

Bandwidth Inc (BAND)
Our $60 PO is based on 22.5x EV/CY27E FCF. The multiple is a discount to infrastructure software peers at 29.2x EV/CY27E FCF, reflecting our forecast for slower CY27E FCF growth (10% vs. 20% for peers) and lower CY27E FCF margins (11% vs. 19% for peers). We believe this discount appropriately reflects execution risks around AI Voice and Maestro becoming material growth drivers and the FCF multiple is the best valuation metric given Bandwidth's revenue growth and FCF margin profile (4% and 11% in CY27E) versus the infrastructure software peer averages of 12% and 19%.

Downside risks to our PO are 1) AI voice adoption could take longer to scale, 2) CPaaS pricing pressure could persist, 3) Messaging remains exposed to political cycles, 4) Large enterprise deployments lengthen, 5) Competition intensifies from CPaaS providers, hyperscalers, and AI-native infrastructure vendors.

Dassault Systemes (DASTF / DASTY)
Our valuation and price objective (PO) of €25.6/ US$30 per share is based on a discounted cash flow (DCF) model, assuming 2% long-term growth, 26.3% long-term EBIT margins (from 25%), and 7.2% weighted-average cost of capital (WACC)

Upside risks to our price objective are (1) increasing traction of 3DX driving faster growth, (2) successful expansion into new verticals, and (3) successful AI monetization

Downside risks are (1) slowing spending in the automotive, aerospace sectors, life science, dampening growth, (2) competitive pressure esp in the life science segment, (3) disruption from GenAI

Datadog Inc (DDOG)
Our $305 PO is based on an EV/revenue multiple of 20.6x EV/CY27E revenue. This is a premium compared to its infrastructure software peers at 5.0x, which we think is justified given its revenue growth and free cash flow upside potential and category leadership.

Downside risks to our PO: 1) valuation risk, 2) competition, 3) deteriorating net revenue retention, 4) execution, and 5) enterprise software spend can be cyclical.

Dynatrace (DT)
Our $65 PO is based on 27.0x EV/C27E FCF, comparable to infrastructure software peers at 26.9x, which we think is warranted given category leadership balanced against execution risk.

Downside risks: macroeconomic downturn, increase in competition, ability to hire direct sales reps at a healthy clip, ability to expand into adjacent markets, and ability to sustain high R&D innovation.

Elastic NV (ESTC)
Our $108 PO is based on 4.8x EV/C27E revenue multiple, lower than infrastructure software peers at 7.1x, warranted for its category leadership, differentiation, and positioning in the attractive AI opportunity, balanced against execution risk.

Risks to our PO are 1) competition, 2) relatively unproven profitability, 3) M&A, which could weigh on margins and execution, 4) application spending being highly cyclical, 5) increasing Software as a Service (SaaS) demand, which could weigh on gross margins, and 6) conversion ratio to paid subscriptions, which could decrease.

GitLab Inc. (GTLB)
Our $54 PO is based on an EV/revenue multiple of 6.2x to our FY28E (C27E) revenue estimate, lower versus infrastructure peers at 7.0x. We believe the discount multiple to peers is warranted given potential for increasing competition from other AI code generation alternatives and execution risk from the revenue model transition, balanced against its revenue and free cash flow upside potential.

Downside risks to our PO: 1) valuation risk, 2) competition, 3) deteriorating NRR, 4) extended timeline to profitability, 5) execution, and 6) enterprise software spend can be cyclical. Upside risks to our PO: 1) faster than expected adoption of Duo Agent Platform, 2) accelerating developer seat growth, 3) faster adoption of new consumption-based products like Secrets Manager, 4) higher than expected incremental margins, and 5) Flex driving higher platform usage.

HBX Group (XHVXF)
We value HBX Group at a price objective (PO) of €12 per share using a discounted cash flow (DCF) approach. Key assumptions in our DCF model are 1) weighted-average cost of capital (WACC) of 8.9%, 2) long-term revenue growth of 1% and 3) 57% in terminal EBITDA margin. We are also applying a 10% discount to our DCF derived valuation to account for incremental uncertainty on long term growth prospects amid potential AI disruption and intensifying competitive dynamics.

Downside risks to our price objective are: 1) Competition: While HBX is the largest bed bank globally, it also competes with larger companies in the B2B distribution space including Expedia and Booking. 2) Disruptions: Like other travel tech companies, HBX is inherently susceptible to travel disruptions. 3) AI: Generative AI could enable faster innovation/disruption from new entrants or decrease the need for intermediaries.

JFrog Ltd (FROG)
Our $110 PO is based on 17.2x EV/C27E revenue, which compares to JFrog's infrastructure peers at 5.6x revenue, warranted in our view given JFrog's better growth profile vs. the peer average (i.e., c.20% vs. low double-digit).

Downside risks to our PO are macroeconomic downturn, increase in competition, ability to upsell installed base, ability to convert free users to paid subscriptions, ability to expand into adjacent markets, and ability to sustain high R&D innovation.

MongoDB Inc (MDB)
Our PO of $540 is based on 12.0x CY27E revenue, above the infrastructure software peer group average at 7.7x to reflect continued healthy opportunity in the core market as the company scales, and AI demand tailwind optionality.

Downside risks to our PO are: 1) MongoDB trades at a significant premium to its software large-cap peer group on an EV/Sales basis. In the event of changes to market sentiment stemming from global macro uncertainty, or potentially disappointing quarterly revenue results, there could be risk of a potential pullback. 2) The company recognizes a large majority of its revenue based on actual consumption which implies more revenue volatility than a traditional subscription model. 3) MongoDB serves in a highly competitive market, which consists of deep-pocketed next-gen NoSQL database, public cloud vendors and legacy database vendors. An inability to execute on a product roadmap for added capabilities could result in slowing share gains or even share losses.

Nemetschek (NEMTF)
Our PO of €93 per share is based on a DCF model, assuming 2.3% long-term growth, 32% long-term EBIT margins, and 7.8% WACC.

Upside risks to our price objective are (1) increasing traction of MAXON driving more adoption across the whole ecosystem and (2) successful transition to subscription and cloud.



Downside risks are (1) slowing capex spending in the AEC sector, dampening growth, and (2) product delays meaning slower cloud adoption and consequently slower subscription transition.

Planisware SA (XPZAF)
Our PO of €31 is based on a DCF assuming 2.5% long-term growth and 8.1% WACC. We see substantial long-term upside potential, with EV/EBITDA and FCFE metrics in line with the sector despite a higher growth potential.

Downside risks are competition from larger rivals, generative AI lowering barriers to entry, execution on guidance, FX exposure and macro risks.

Sabre (SABR)
Our PO of $2.5 per share is based on a DCF approach. Assumptions in our DCF model are 1) WACC of 9.7%, 2) long-term revenue growth of 0.5% and 3) 22% in terminal EBITDA margin.

Upside risks are (1) better macro/faster recovery from COVID-19 that drives up bookings volumes, and (2) further potential new wins with hospitality groups and/or airlines for the reservations IT platform.

Downside risks are (1) weak demand for corporate travel due to a secular shift to video calls, (2) market share loss in distribution and IT solutions.

SAP SE (SAPGF / SAP)
Our price objective of €226 (US$264) is based on a Sum of the Parts model, valuing Cloud, on-prem and services separately. We use a 20x 2027E EV/EBITDA for the Cloud business, in line with US peers adjusted for growth. We value the on-prem business using a DCF (-8% revenue growth to 2031, terminal growth 0%, 40% terminal EBIT margins, WACC 6.7%), while we value the services business at a 10x EV/EBITDA multiple, in line with IT Services peers.

Upside risks to our price objective are (1) a faster multiple re-rating as the market prices in a successful cloud transition and (2) slower erosion in the core licence business.

Downside risks are increased competition, especially from Cloud native peers, and execution challenges leading to a weaker-than-expected growth or margin profile.

Sinch (CLCMF)
Our PO of SEK29 is based on a DCF model, assuming 1% long term growth, 12% long term EBIT margins, and 10.7% WACC.

Upside risks to our price objective are: i) reacceleration in the pace of digital transformations, ii) stronger US execution and cross-sell under new management iii) rebound in core Messaging margins.

Downside risks are: i) a weak demand environment in 2026, ii) lengthening sales cycles as enterprises put digital transformation strategies on hold iii) concentrated efforts by Microsoft/Amazon to strengthen CPaaS offerings iv) sharp USD depreciation.

Snowflake (SNOW)
Our PO of $470 is based on 22.4x CY27E revenue, a premium to the infrastructure software group at 5.5x for durability of growth stemming from Snowflake's traction in the AI software market, which has a large, growing TAM of $155bn.

Upside risks to our PO: Snowflake recognizes revenue based on actual consumption (versus subscription ratably in arrears). In the event of changes to market sentiment on

macro abating and IT spend returning to normal, consumption and revenue growth could quickly reaccelerate which could drive multiple expansion.

Downside risks to our PO: 1) Snowflake trades at a premium to its software large-cap peer group. In the event of changes to market sentiment stemming from global macro uncertainty, or potentially disappointing quarterly revenue results, there could be risk of a potential pullback, 2) The company's consumption model implies more revenue volatility than a traditional subscription model, 3) Snowflake serves in a highly competitive market, which consists of deep-pocketed next-gen DBMS vendors, public cloud vendors and legacy DBMS. An inability to execute on a product roadmap for added capabilities could result in slowing share gains or even share losses.

TeamViewer (TMVWF)
We value TeamViewer using a DCF approach. Key assumptions in our DCF model are 1) WACC of 7.7%, 2) long-term revenue growth of 0.5% and 3) 25% in terminal EBIT margin. We then applied a 10% discount to our DCF-derived equity value to capture heightened long-term uncertainty linked to AI, resulting in our €9.5 PO.

Upside risks to our PO are 1) SMB segment to stabilise and remain more resilient than expected, 2) material expansion in Enterprise use-case, particularly on DEX, IoT and AR opportunities.

Downside risks to our PO are: 1) The long-term trend towards cloud will eventually diminish the need for remote desktop solutions.2) Big platform competitors with deep pockets may enter the market with "out of the box" solutions. 3) High profile security breaches may impact brand image and divert business users away.

Temenos AG (TMNSF / TMSNY)
Our price objective is CHF97 / US$125. This is driven by a DCF valuation with the following assumptions: 1/ terminal growth rate of 2%, 2/ WACC of 6.7%.

Downside risks are banks' growth slowdown and/or slowdown in third-party software outsourcing, lack of product traction, implementation issues. Upside risks are large deal signings, better traction in the US, and a faster economic recovery.

Twilio (TWLO)
Our $320 PO is based on 35.5x EV/CY27E FCF.

Our target multiple is above infrastructure software peer group trading at 32.5x EVCY27E FCF. We believe a premium to the peer multiple is justified based on an accelerating growth profile, expanding profitability profile, balanced against execution risks.

Upside risks to our PO are 1) faster-than-expected reacceleration of organic revenue growth, 2) macro recovery driving better-than-expected usage rates, and 3) higher-than-expected revenue mix of software versus core messaging, voice and email.

Downside risks to our PO are 1) increased competition from larger platform vendors, 2) possible price pressure, 3) deterioration in the macro economy, 4) decrease in ability to monetize the platform via developers, 5) deterioration in gross margins, 6) missteps causing growth to dip below management's targets, 7) major network and product issues.

UiPath (PATH)
Our $15 PO is based on 3.3x CY27E rev, a discount to the infrastructure group trading at 7.1x, given questionable durability of growth and margin expansion and execution risk as it transforms into an agentic automation company.

Downside risks are 1) Application spending is highly cyclical, thus during an economic slowdown, when firms are faced with shrinking IT budgets, projects involving application upgrades, migrations or new installations are often deferred which could present a higher degree of risk for a bookings deceleration for UiPath and other application vendors, and 2) UiPath faces competition from point solutions and platform vendors - the absence of ongoing feature enhancements and new module development could present the risk of slowing share gains or even share losses to application platform vendors investing in machine learning and automated workflow offerings and/or other RPA vendors.

Upside risks are 1) Discounted valuation presents risk of reversion in the event of changes to multiple expansion across the software group, or potentially positive quarterly revenue and 2) the company sees material revenue from its AI agent offerings, which could likely drive revenue acceleration.

Analyst Certification

We, Victor Cheng, CFA, Frederic Boulan, CFA, Koji Ikeda, CFA and Tomer Zilberman, hereby certify that the views each of us has expressed in this research report accurately reflect each of our respective personal views about the subject securities and issuers. We also certify that no part of our respective compensation was, is, or will be, directly or indirectly, related to the specific recommendations or view expressed in this research report.

Special Disclosures

BofA Securities is currently acting as financial advisor and providing a fairness opinion to Amadeus IT Group SA in connection with its proposed acquisition of IDEMIA's public security unit, which was announced on April 29, 2026.

BofA Securities is currently acting as financial advisor to Workday Inc in connection with its proposed acquisition of ConductorOne Inc (doing business as C1.ai), which was announced on October 13, 2026.