

June 6, 2025 04:58 AM GMT

Software | North America

AI Agents Knocking at the Door

The growing power of deep reasoning models pave the way for a new wave of Agentic AI applications able to automate broadening business functions, unlocking tremendous value. The customer journey to Agentic AI will take time and software players will have to evolve business + pricing models to thrive

Key Takeaways

- Over decades software was developed based on the assumption that humans play the central role in executing tasks, workflows and business processes.
- Agentic computing removes this assumption, replacing it with the idea that Agents increasingly become the entity executing tasks and making decisions
- The value creation opportunity is enormous as Agentic applications penetrate the \$630+ billion application software market over time
- However, the transition to agentic will come with change – redefining how apps are architected and requiring new business and pricing models
- Software companies well positioned to benefit include: AMZN, CRM, GOOG, HUBS, MSFT, NET, NOW, OKTA, OS, PLTR, SAIL, SNOW, & TEAM

For Deep Dive in the Agentic Computing Opportunity Please see our [Agents Knocking on the Door Presentation](#) attached to this report

Bringing Agency to Software. Over the last several decades, software has been developed based on the assumption that humans play the central role in executing tasks, workflows and processes. The expanding capabilities of Agentic AI removes this long-held assumption, replacing it with a framework where AI Agents increasingly become the entity executing tasks and making decisions. The coalescing forces driving this shift – the increasing power of deep reasoning models, the ability to have a semantic understanding of data and the declining price of intelligence – create a substantial opportunity for enterprises to automate entire business functions and processes, unlocking enormous value via digital labor replacement. However, investors should view the march towards agentic computing in the enterprise as an evolutionary journey rather than a revolutionary sprint. The inherently indeterminate nature of foundational models, a proliferation of application and data silos built up over decades, persistent challenges with data quality, an ever-growing list of security concerns and a lack of standardization with respect to governance models and application frameworks present significant obstacles to deploying Agentic AI applications – challenges the industry will have to solve over time.

Bracing for Change. Taking in proper historical context, Agentic computing represents the next progression of a long-standing evolution towards greater abstraction in software. At a basic level, an AI Agent is a software program defined

MORGAN STANLEY & CO. LLC		
Sanjit K Singh		
Equity Analyst		
Sanjit.Singh@morganstanley.com		+1 415 576-2060
Keith Weiss, CFA		
Equity Analyst		
Keith.Weiss@morganstanley.com		+1 212 761-4149
Brian Nowak, CFA		
Equity Analyst		
Brian.Nowak@morganstanley.com		+1 212 761-3365
Josh Baer, CFA		
Equity Analyst		
Josh.Baer@morganstanley.com		+1 212 761-4223
Elizabeth Porter, CFA		
Equity Analyst		
Elizabeth.E.Porter@morganstanley.com		+1 212 761-3632
Chris Quintero		
Equity Analyst		
Chris.Quintero@morganstanley.com		+1 212 761-1686
Oscar R Saavedra		
Research Associate		
Oscar.Saavedra@morganstanley.com		+1 212 761-0827



SOFTWARE	
North America	Attractive
Industry View	

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

by key attributes including the ability to: 1) understand and interact with its environment (receive inputs), 2) collect and connect to data, 3) take action autonomously without using pre-determined rules and 4) learn from its experience over time. Like past waves, the transition towards AI Agents will unleash greater automation, productivity and efficiency, but also force the software industry to redefine how applications are built and architected, as well as evolve business models including instituting new approaches to pricing.

System of Record vs. System of Engagement – A Framework for Organizing the Agent Landscape. The emerging Agentic AI landscape can be organized through a bifurcated framework distinguishing between System of Record (SoR) Agents and System of Engagement (SoE) Agents, each addressing distinct automation paradigms and market opportunities.

- **SoR Agents function as domain-specific, functional automation solutions** attempting to directly replace human roles in well-defined processes such as coding, customer service, sales lead generation, business intelligence, etc. Examples include GitHub Copilot, Salesforce Agentforce, and HubSpot Prospecting Agent. These Agents target a narrower market opportunity, which we closely align with IDC's labor-centric automation market segment, representing a \$16 billion opportunity today projected to reach \$33 billion by 2028.
- **SoE Agents operate as supervisory orchestrators focused on workflow and process automation across multiple systems**, delivering value through improved system utilization and indirect time saving rather than direct labor replacement. Examples include ServiceNow Now Assist, GitLab Workflows, and Atlassian Rovo. These Agents target a broader market opportunity, which we closely align with IDC's system-centric and decision-centric automation, and business value engineering market segments, representing a \$36 billion opportunity today projected to reach \$69 billion by 2028.

Exhibit 1: Framework for Organizing the Agent Landscape – System of Record style Agents vs System of Engagement style Agents.

	SoR Agent	VS	SoE Agent
Role:	Functional		Supervisor
Automation Focus:	Domain or Role Specific		Workflow / Process-Oriented
Value Proposition:	Labor Displacement		System Utilization
Quantification of ROI:	More Direct - Automation of Role		Indirect - Time Savings on Broader Workflows
Difficulty of Implementation:	Hard - New Tech Learning Curve		Harder - New Tech Plus Broader Integrations
Market Opportunity:	Narrow - Defined by Role of Domain		Broad - Defined by Scope of Workflows

Source: Morgan Stanley Research

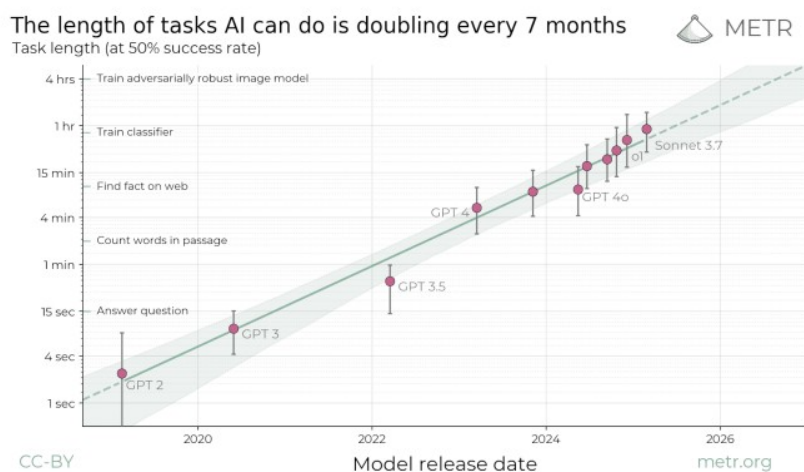
Where Are We on the Path to Agentic AI? Our industry conversations on

production deployments of AI Agents have been mixed – customers report that, when "working", results are impressive and easily exceed ROI mandates to warrant further investment. In a significant portion of cases, however, reliability, accuracy and precision have been disappointing. While this may suggest that the AI reasoning models are limited in their ability to handle complex tasks, we stress that the situation is fluid and major advancements in reasoning capabilities appear to be unfolding in a time frame measured in weeks. Underscoring this point, according to AI research firm METR, **AI agents have doubled the length of tasks they can automate every 7 months for the last 6 years.**

To contextualize where the market stands in terms of agent-based automation capabilities, AI agent provider Sema4.ai, uses a 5 level framework. Currently, the market is at level 3 on the path to fully autonomous Agentic AI. This represents the first level of *constrained autonomy* – moving beyond rule-based automation and task-specific co-pilots. At this stage, AI Agents can create, execute, and adapt plans based on feedback while handling complex tasks through multiple reasoning cycles.

For example, Salesforce's Agentforce agents can reconcile a 100-page invoice against internal systems with human-like reasoning. This represents a substantial leap from the lower levels of rule-based automation/robotic process automation (RPA) (level 0), basic AI-enhanced process automation (level 1), and task-specific AI assistants/copilots (level 2). However, the state of the market remains below level 4 capabilities – systems that can self-improve and modify their own instructions, and well below level 5 capabilities – fully autonomous AGI-type systems that are capable of independent reasoning and creativity without oversight. This progression indicates that while we have made significant progress in autonomous planning and execution capabilities, human oversight is still necessary in most cases.

Exhibit 2: Length of Tasks That Generalist Frontier Model Agents Can Complete Autonomously Has Been Doubling Approx. Every 7 Months for the Last 6 Years

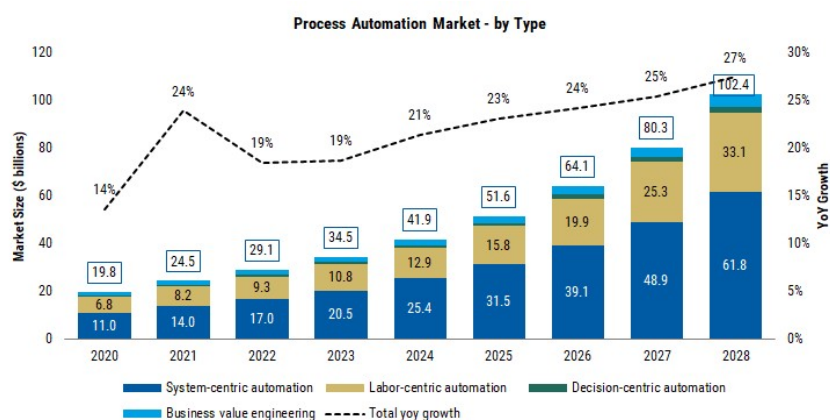


Source: Model Evaluation & Threat Research (METR) "Measuring AI Ability to Complete Long Tasks"

How Big is the Opportunity? We evaluate the addressable opportunity for agentic computing across three distinct lenses that reflect varying degrees of market penetration and technological maturity. A narrow view, focused on AI-powered automation, represents a \$6 billion opportunity today with projected growth to \$20

billion by 2028. A wider lens encompassing the broader process automation market represents a \$52 billion opportunity today with projected growth to \$102 billion by 2028. An expansive view of the opportunity sees Agentic AI applications penetrating the entirety of the application software market over time, which represented a \$637 billion opportunity today.

Exhibit 3: Expansive View on the Agentic AI Market Points to a \$52B Opportunity Today, Expected to Reach \$102B by 2028 (+26% CAGR)



Source: IDC Research

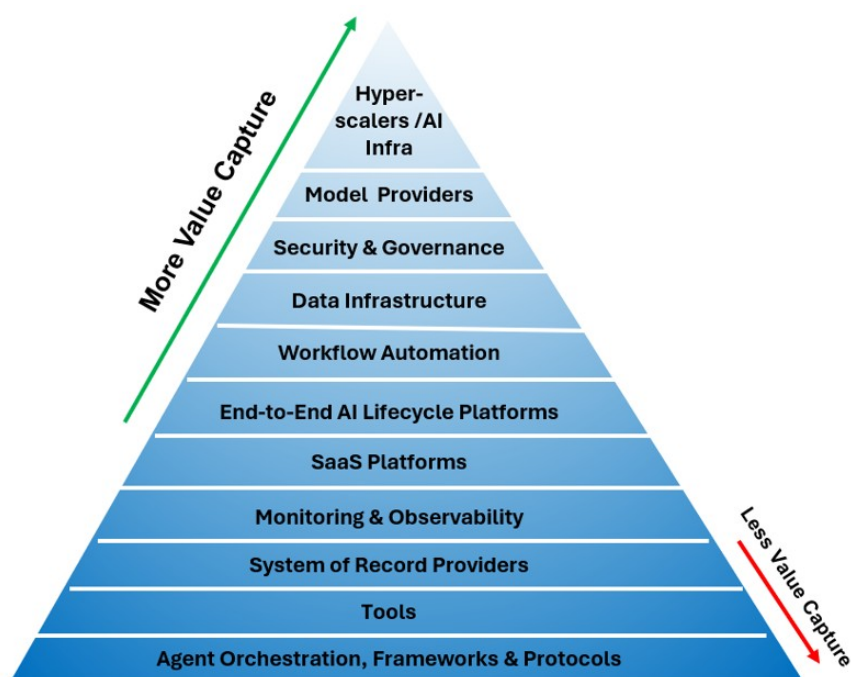
Where Will Value Accrue Across the Stack? Looking across the different categories of the Agentic AI ecosystem, we see the highest value capture potential in the following areas.

1. **Hyperscalers/AI infrastructure:** AI Infrastructure providers have a significant opportunity in providing the compute to enable the real-time inference that deep reasoning agents require. This gives leaders in this category license to not only provide the infrastructure but also the data retrieval, security, governance and policy enforcement capabilities required to operate, manage and secure agentic applications.
2. **AI Model providers:** AI model providers extend the frontier around reasoning and intelligence that will make AI Agents more capable over time. This positioning provides the opportunity to extend into areas such as Agentic frameworks, data processing and integration so as to evolve into modern application platforms.
3. **Security & Governance:** We see a significant monetization opportunity in making Agentic applications more trustworthy, reliable, accurate and secure. Agent architectures create a large surface area for cyber attacks including data exfiltration, supply chain attacks, prompt injection and open source contamination.
4. **Data Infrastructure:** For AI Agents to be useful, they need governed access to the most important data sources in real-time. We see two key opportunities for data infrastructure providers. First, in helping make Agentic architectures possible by modernizing a customer's data estate and improving data quality. Second, by enabling fast, reliable and accurate information retrieval at runtime.
5. **Workflow Automation:** With AI Agents embedded in workflow automation

platforms, the unstructured components of common business processes can be more fully automated. As workflow automation platforms span across key enterprise systems, they become well-positioned to serve as an orchestration/management layer for both 1st and 3rd party agents.

6. **AI Lifecycle Platforms:** Agents within AI lifecycle platforms automate the workflows and tasks associated with developing and deploying AI applications including building data pipelines, integrating enterprise data with LLMs, building and evaluating models and visualizing data. Furthermore, Agents can be embedded into external applications to enhance customer engagement and to enable highly interactive experiences.

Exhibit 4: Most Value Likely Accrues to the Hyperscalers and AI Infrastructure Layer as AI Agent Deployments Likely Result in Significant Utilization of Significant Compute Resources



Source: Morgan Stanley Research

Which Companies are Poised to Benefit? In evaluating how software companies are positioned for the agentic opportunity, we organize the players into three categories – 1) Agent Beneficiaries, 2) Agent Contenders and 3) Agent Wildcards.

- **Agent Beneficiaries** are vendors where we have conviction that they can either a) directly monetize AI Agents/Agentic architectures today or in the near future such that it becomes a meaningful contributor to growth or b) see their core business benefit materially from agentic adoption
- **Agent Contenders** are vendors that are in the right categories to see benefit but require a high degree of execution and market maturity that is not yet evident
- **Wildcards** are vendors where there is debate on whether agentic adoption results in a tailwind to growth or proves disruptive to the core business

Based on our assessment of strategic positioning, product roadmaps and the ability to execute to the opportunity, we highlight the following public companies as Agent Beneficiaries.

AI infrastructure

- Microsoft (OW) – Covered by Keith Weiss
- Amazon (OW) – Covered by Brian Nowak
- Google (OW) – Covered by Brian Nowak

Cybersecurity

- Okta (OW) – Covered by Keith Weiss
- Sailpoint (EW) – Covered by Keith Weiss
- Cloudflare (OW) – Covered by Keith Weiss

SaaS Platforms

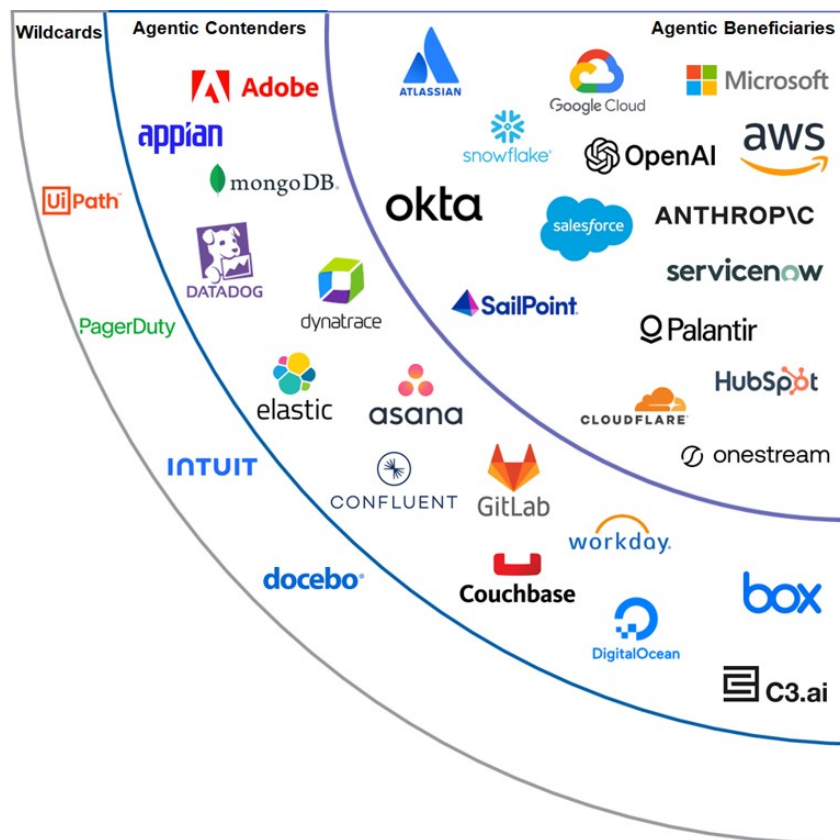
- Salesforce (OW) – Covered by Keith Weiss
- HubSpot (OW) – Covered by Elizabeth Porter
- OneStream (OW) - Covered by Chris Quintero

Data Infrastructure

- Snowflake (EW) – Covered by Keith Weiss
- Palantir (EW) – Covered by Sanjit Singh

Workflow automation

- ServiceNow (EW) – Covered by Keith Weiss
- Atlassian (OW) – Covered by Keith Weiss

Exhibit 5: Vendors Best Positioned to Benefit from Agentic AI Opportunity

Source: Morgan Stanley Research

Agent Pricing: Multiple Approaches But No Market Consensus Yet. AI Agent pricing today is highly fragmented, as companies explore multiple approaches to best capture their fair share of the value being created by their Agentic offerings. Current approaches can be categorized along two dimensions – activity-based vs outcome-based models, and fixed vs variable pricing structures.

In the activity-based segment, companies like Cognition Labs and Harvey AI have adopted an "employee replacement" type model charging a fixed price per agent. Microsoft, OpenAI, and Anthropic utilize a "consumption model" with variable pricing based on agent actions/tokens. In the outcome-based segment, Salesforce and Artisan implement a "process automation model" charging on a per completed workflow basis (e.g., an Agentforce conversation regardless of successful outcome). Zendesk and Sierra employ a "results-based model" where customers only pay if the desired outcome was achieved. Meanwhile, vendors like Atlassian, ServiceNow and HubSpot are taking a hybrid approach by embedding agent capabilities within higher subscription tiers – with usage limits and ability to overconsume with additional pay.

Today most organizations still rely on traditional seat-based models, though increasingly with usage-based components, while workflow and outcome-based approaches continue to gain traction. This diversity in pricing strategies reflects the nascent stage of the enterprise opportunity and underscores that the market is still in a discovery phase in terms of settling on a pricing model for agentic solutions.

Exhibit 6: Most Companies Still Use Traditional Seat-Based Models, Though Increasingly Combined with Usage-Based Component. Workflow and Outcome-Based Models are Gaining Traction



Source: <https://www.growthhunched.com/p/ai-agent-pricing-framework>, Morgan Stanley Research

The Obstacles to Widespread Adoption Are Significant. While most enterprises today are looking for ways to leverage GenAI, widespread adoption of Agentic AI across organizations will inevitably take time due to several unresolved obstacles. First, the indeterminate nature of models today means that accuracy and hallucinations remain an issue, making it difficult for businesses to fully trust AI-driven decision making. Second, the lack of adequate security guardrails raises serious concerns about the potential exfiltration of proprietary data or IP by bad actors. Third, the absence of agreed-upon standards, protocols, and frameworks for building multi-agent applications hinders interoperability and slows development. Fourth, data quality issues and a fragmented data and application silos limit the ability to effectively deal with complex business processes. Lastly, agentic systems lack sufficient semantic understanding of proprietary business data, especially given the unique processes and data nomenclature of each customer organization. Collectively, these obstacles mean that enterprises will likely proceed in a deliberate fashion as they work towards broader adoption of Agentic AI.

Agents Knocking on the Door Presentation

Morgan Stanley

Morgan Stanley Research



AI Agents Knocking at the Door

June 6, 2025

Sanjit Singh

sanjit.singh@morganstanley.com
(415) 576-2060

Keith Weiss

keith.weiss@morganstanley.com
(212) 761-4149

Oscar Saavedra

oscar.saavedra@morganstanley.com
(212) 761-0827

Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

Morgan Stanley

MORGAN STANLEY RESEARCH

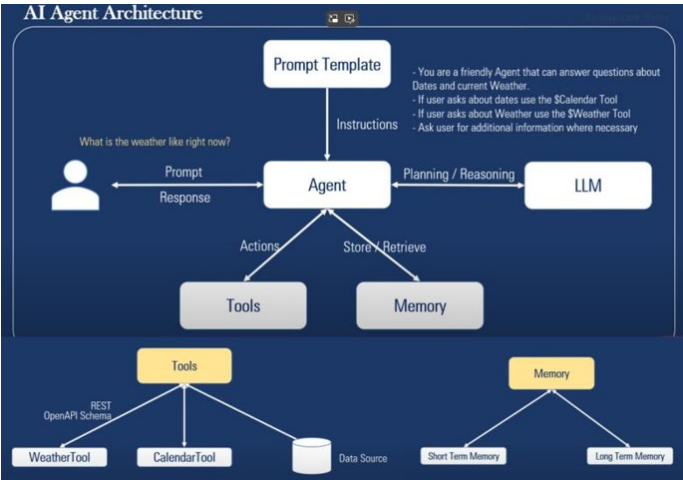
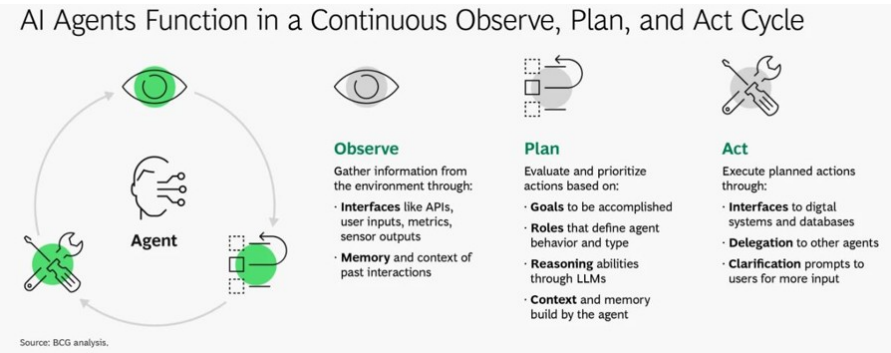
Executive Summary

Executive Summary

What's the Big Idea?	<ul style="list-style-type: none">• Bringing Agency to Software: Over the last several decades software has been developed based on the presumption that humans play the central role in executing tasks, workflows and processes. The adoption of AI Agents and Agentic Architectures means more and more of the software development effort will utilize a framework where AI agents will be the entity executing tasks and making decisions.
Implications	<ul style="list-style-type: none">• Bracing for Change: AI agents represent the next phase of a long-standing evolution in software towards greater abstraction of the tools used to automate business and consumer processes. Like past waves, the transition towards AI Agents will result in greater automation, productivity and efficiency, but also force the software industry to evolve business models and institute new approaches to pricing.
Opportunity	<ul style="list-style-type: none">• A narrow view on the Agentic AI market points to a \$6 billion opportunity today, expected to reach \$20 billion by 2028• An expansive view on the Agentic AI market points to a \$52 billion opportunity today, expected to reach \$102 billion by 2028.
Beneficiaries	<ul style="list-style-type: none">• Hyperscalers and AI infrastructure providers• AI Model providers extending into the application layer• Security vendors enabling authorization, threat detection, vulnerability/posture management and policy enforcement• Data management providers• Workflow and process automation providers• SaaS Platforms
Risks	<ul style="list-style-type: none">• Indeterminate Nature of the Models: Accuracy and hallucinations are still a problem• Lack of adequate security guardrails to ensure proprietary data or IP is not exfiltrated by bad actors• Lack of agreement on standard, protocols and frameworks to build multi-agent applications• Lack semantic understanding of enterprise's business data from one organization to the next• Limited reasoning capabilities to navigate complex business processes

What is an AI Agent?

What is an AI Agent	An AI Agent is a system (software program) that understands and interact with its environment (receive inputs), collect data, make decisions, executes tasks and meet specified goals without using pre-determined rules.
Autonomous AI Agent	Autonomous AI Agents are systems that leverage large language models (LLMs) to plan and execute a task or process without a high degree of human interaction. These agents distinguish themselves from standard GenAI chatbots by combining tools (i.e., databases, websites, APIs, and even other agents) with memory of past interactions to chain multiple thoughts and actions together, enabling them to autonomously work toward defined objectives while drawing on both external information and learned experiences to produce relevant outputs.
Agentic Architecture	Agentic Architecture generally refers to the idea of bringing together multiple engines for understanding queries/requests, evaluating strategies for solving problems, accessing necessary data and executing the planned actions. Under this architecture, a large language model (LLM) typically serves as a reasoning engine, connected to tools and memory. Tools help connect the LLM to other sources of data or computation and can be used to take actions (run code, modify files, etc.). Memory (short or long term) helps the agent remember previous interactions with either humans, other agents, or tools.



Source: Company Data, Morgan Stanley Research, www.simform.com

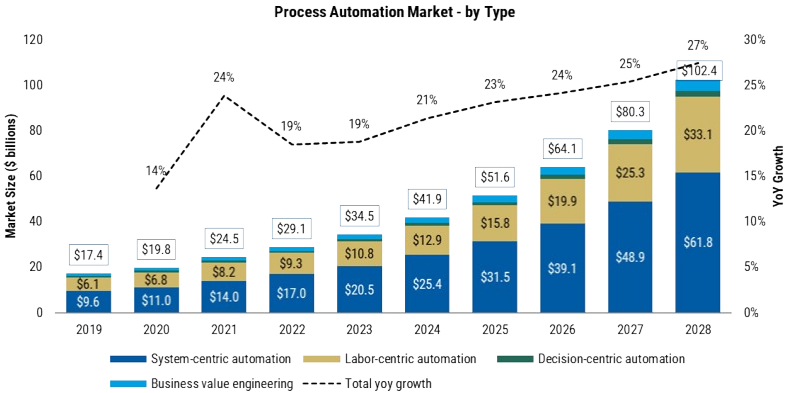
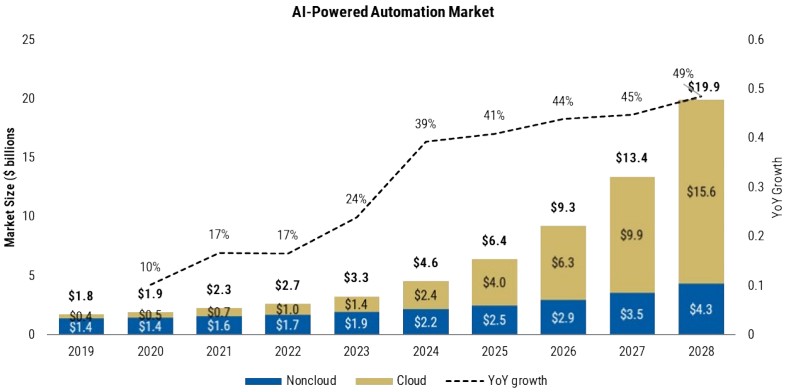
How Big is the Market Opportunity?

Narrow View

Evolution of Process Automation: A narrow view on the Agentic AI market points to a \$6 billion opportunity today, expected to reach \$20 billion by 2028

Expansive View

Broad Expansion of Application Functionality: An expansive view on the Agentic AI market points to a \$52 billion opportunity today, expected to reach \$102 billion by 2028.



Key Drivers

- Increased demand from enterprises to further automate processes.
- Ongoing Improvements in foundational models results in more business value being realized, driving further demand.
- Continued declines in the price of intelligence which raises ROI for AI initiatives driving more demand for agents.

Source: IDC Worldwide Intelligent Process Automation Software Forecast

System of Record (SoR) vs System of Engagement (SoE) AI Agents

SoR vs. SoE
Ai Agents

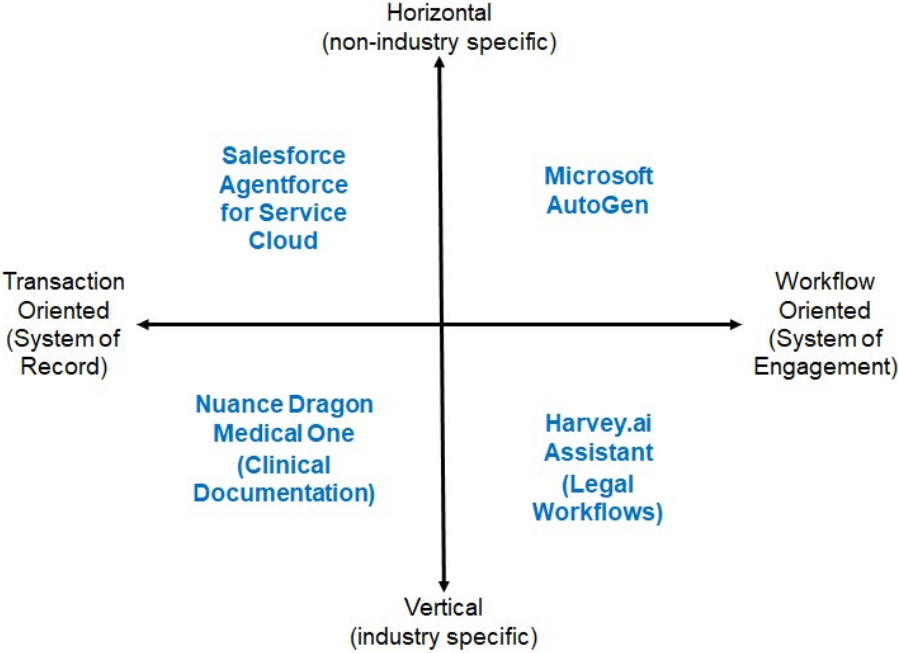
High level framework for organizing the agent landscape is the distinction between System of Record style agents from System of Engagement style agents.

System of Record (SoR) Agent:

- Automates transactions within one application or domain
- For example, Salesforce’s Agentforce automated the interactions of a Call Center agent

System of Engagement (SoE) Agent:

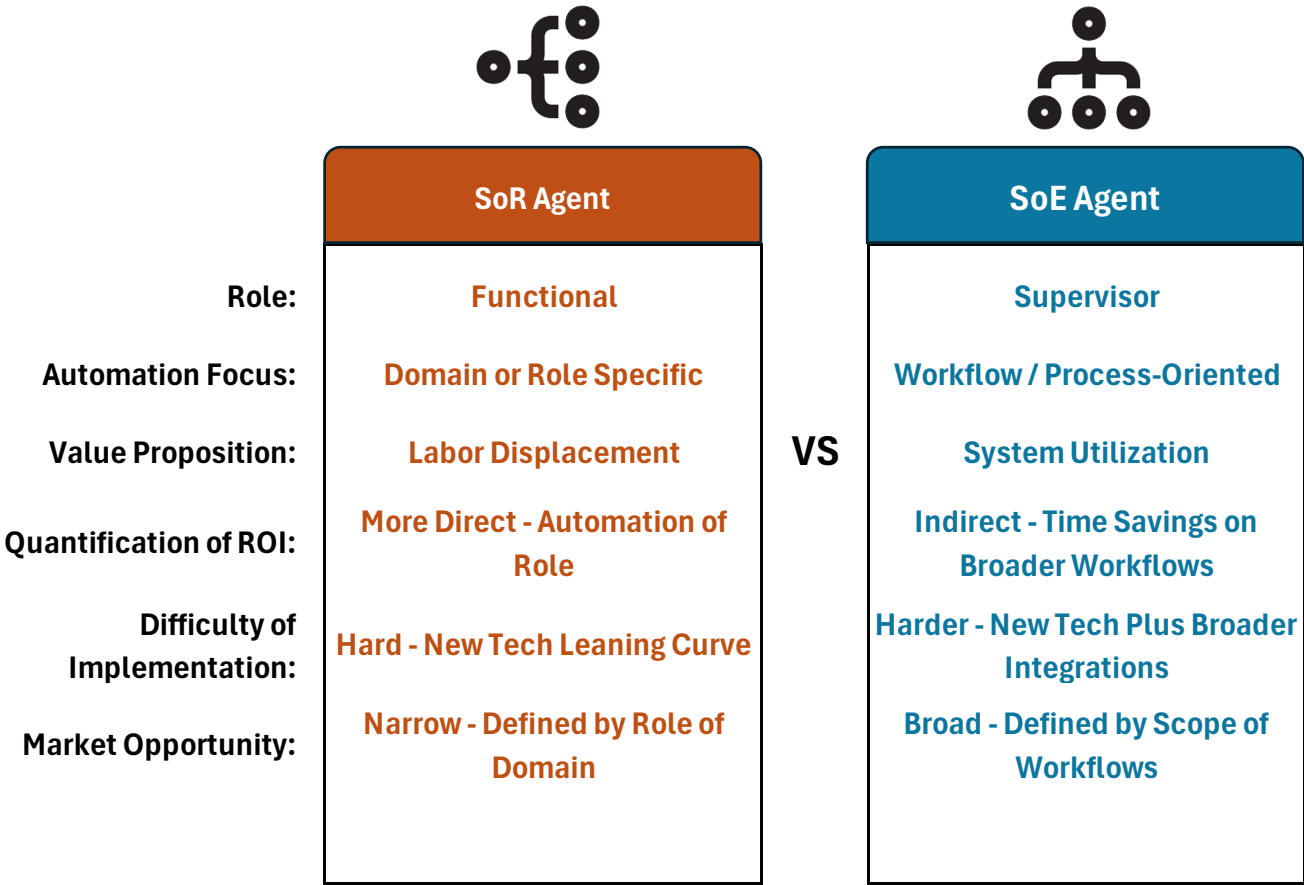
- Automates business processes spanning multiple applications or systems
- For example, an HR Onboarding workflow spanning Payroll, Benefits, IT Procurement systems to complete a workflow



Source: Morgan Stanley Research

System of Record (SoR) vs System of Engagement (SoE) AI Agents

SoR vs. SoE Ai Agents One framework for organizing the agent landscape is the distinction between System of Record style agents from System of Engagement style agents.



Source: Morgan Stanley Research

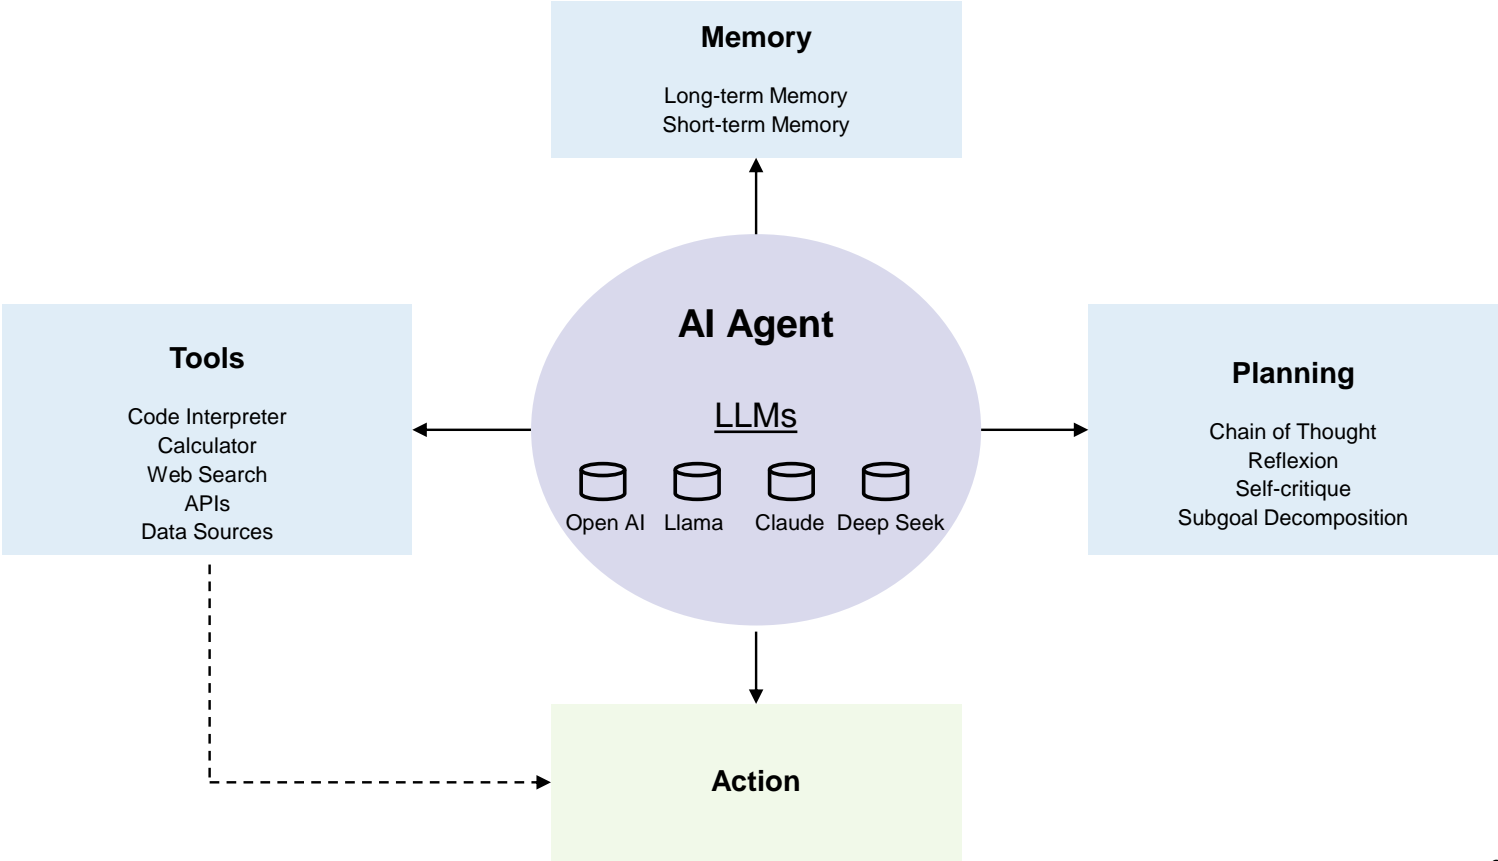
Morgan Stanley

MORGAN STANLEY RESEARCH

The Emerging AI Agent Operating System

Key Elements of the Agent OS

The emerging AI Agent Operating System (OS) consists of: 1) the LLM, 2) Tools, 3) Access to Memory and 4) the ability to plan

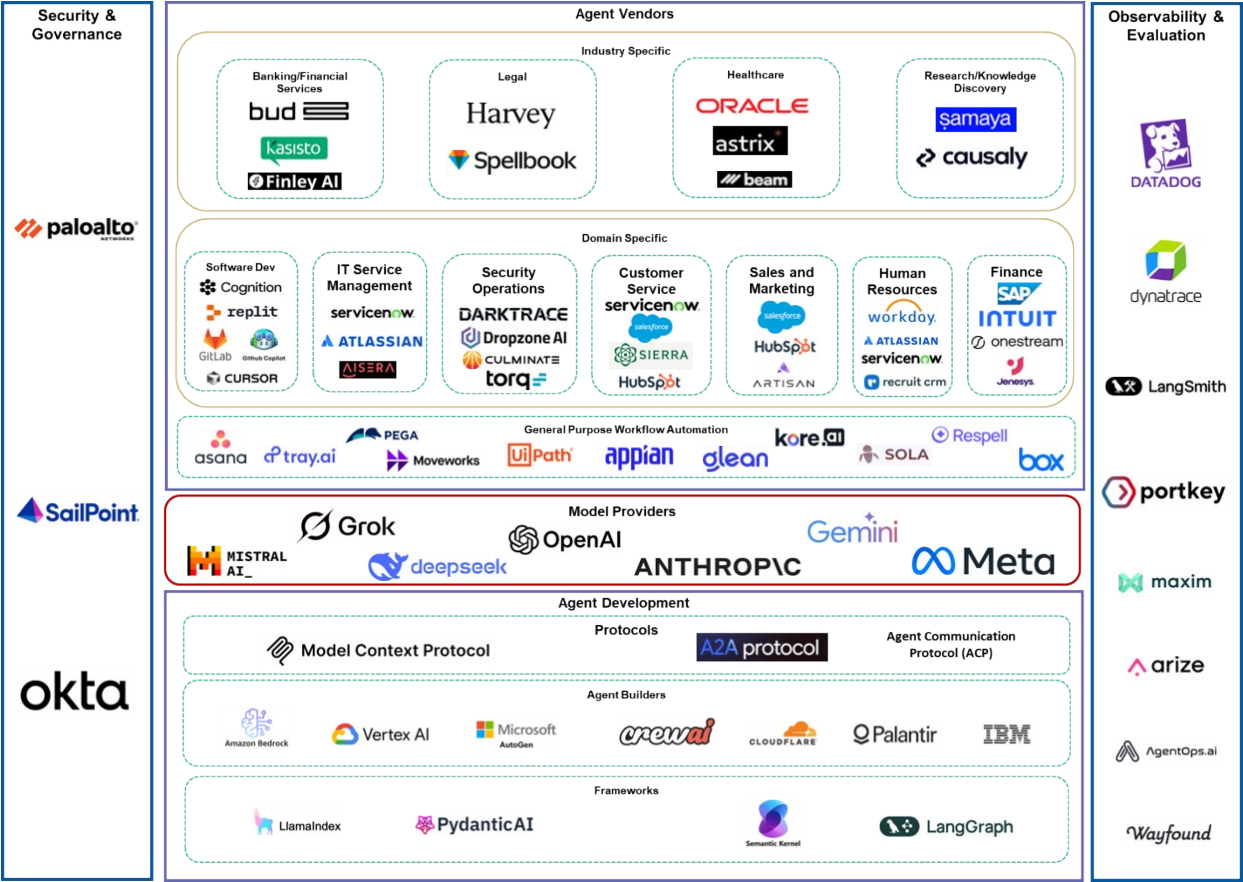


Source: Lilian Weng's Lil'Log (<https://lilianweng.github.io/>)

Current Market Landscape

Bottom Line

Agentic AI market landscape extends from vendors enabling Agent development to vendors providing purpose-built and general-purpose Agents



Source: Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

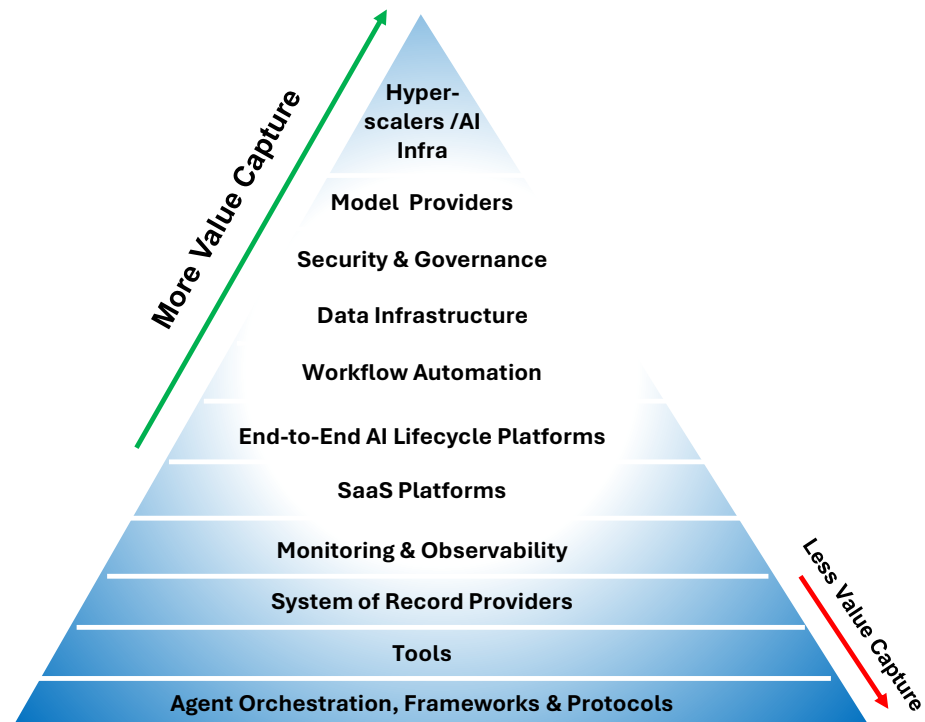
Where in the Stack Will Value Accrue?

Bottom Line

The most value likely accrues to the hyperscalers and AI infrastructure layer as AI Agent deployments likely result in significant utilization of significant compute resources

• Other Parts of the Stack Likely to Benefit:

- **Model Providers:** Serving as the “brain” of any agentic solution gives license for model providers to own more of the AI app stack
- **Security, Governance & Guardrails:** The ability to infuse trust, reliability and accuracy while enforcing policy present significant opportunities for monetization
- **Data Infrastructure:** A necessary requirement to build useful, effective and accurate agentic systems require a modern data infra that make proprietary unstructured and structured data easily governed and accessible
- **Workflow Automaters:** Software providers that are ingrained in customers’ key workflows and business processes can create substantial value by further automating those workflows or re-engineering altogether with agentic AI



Source: Morgan Stanley Research

10

Morgan Stanley

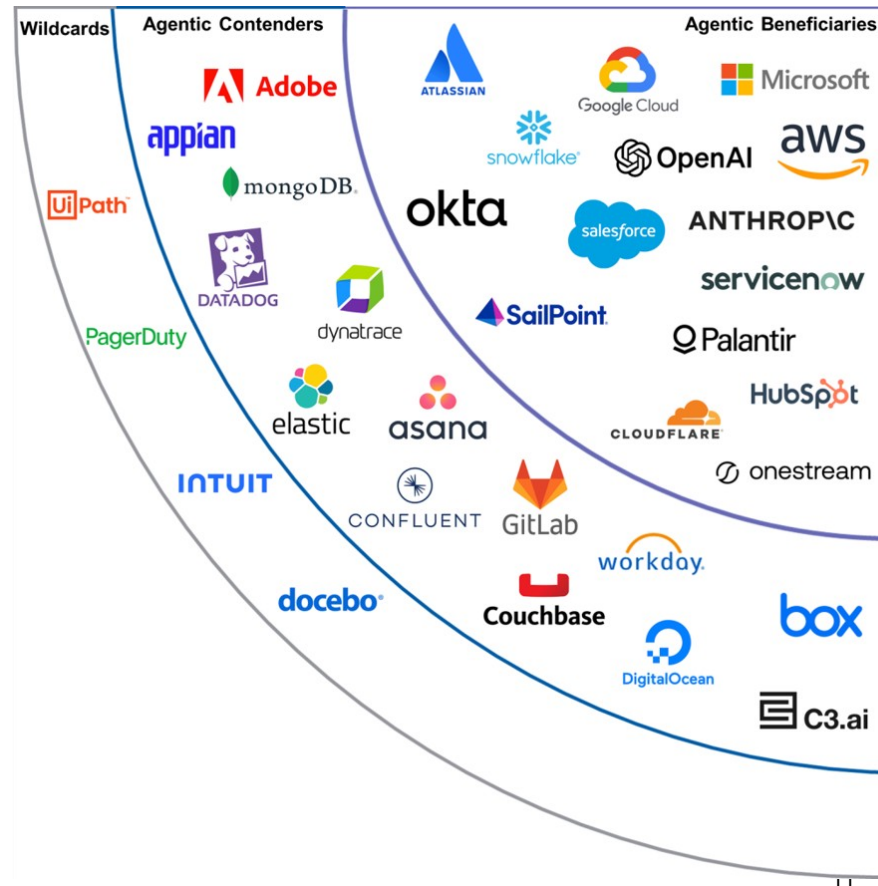
MORGAN STANLEY RESEARCH

Who is Poised to Capture the Agentic Opportunity?

Bottom Line

Software vendors best positioned to benefit based on a) ability to monetize AI Agents/Agentic architecture or b) core business benefitting from enterprise adoption

- **Agent Beneficiaries** are vendors where we have conviction that can either a) directly monetize AI Agents/Agentic architectures today or in the near future such that it becomes a meaningful contributor to growth or b) see their core business benefit materially from agentic adoption
- **Agent Contenders** are vendors that are in the right categories to see benefit but require a high degree of execution and market maturity that is not yet evident
- **Wildcards** are vendors where there is debate on whether agentic adoption results in a tailwind to growth or proves disruptive to the core business

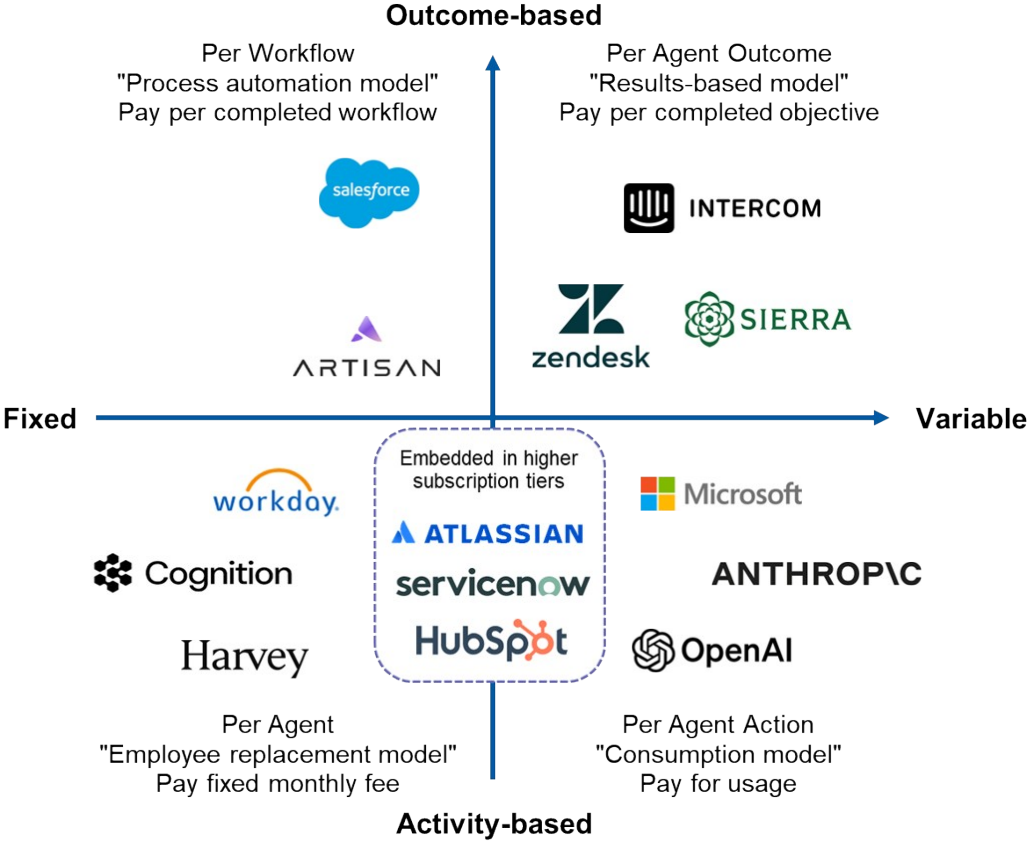


Source: Morgan Stanley Research

How Will AI Agents Be Priced?

Bottom Line

Coming from an era of seat-based pricing, agentic pricing is taking on several forms with no clear consensus at this early stage



Source: <https://www.growthunhinged.com/p/ai-agent-pricing-framework>, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Part I

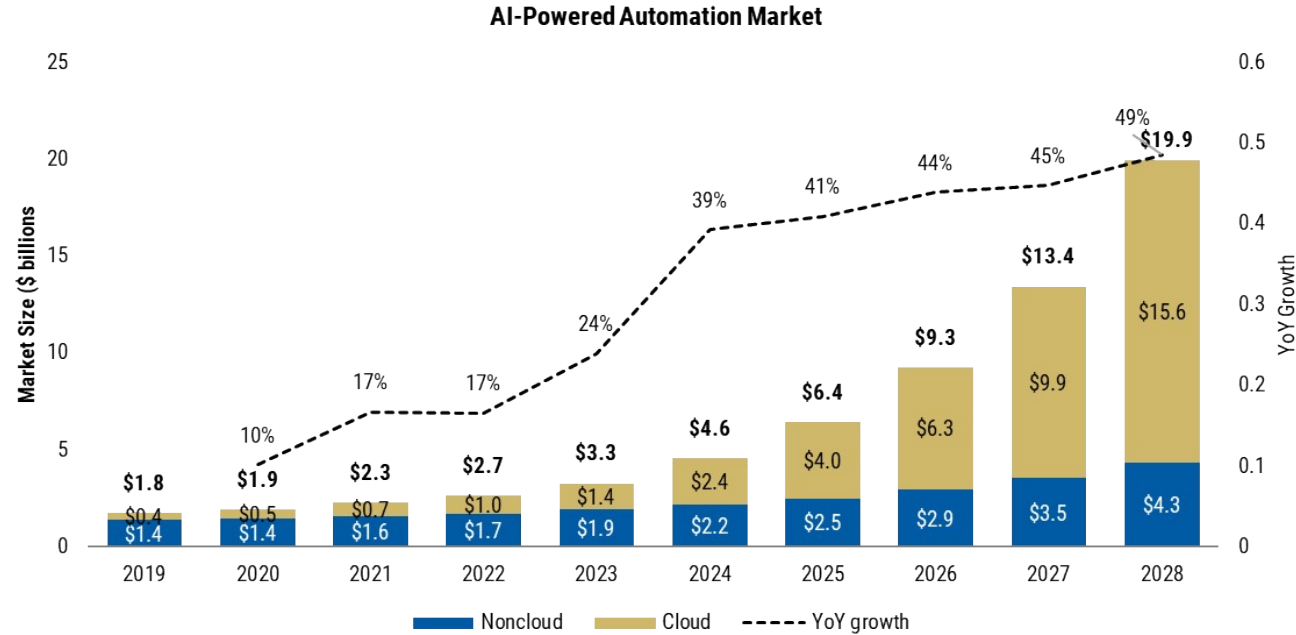
Key Debates

Debate #1: How Big is the Opportunity?

Bottom Line

A conservative view on the Agentic AI market points to a \$6 billion opportunity today, expected to reach \$20 billion by 2028

- IDC definition of AI-powered automation includes agentic automation, intelligent document processing (IDO), business process-specific specialized models, and predictive analytics



Source: IDC Worldwide Intelligent Process Automation Software Forecast

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #1: How Big is the Opportunity?

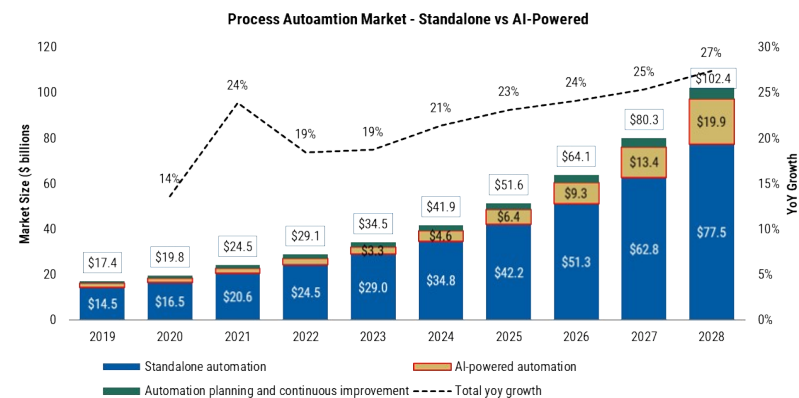
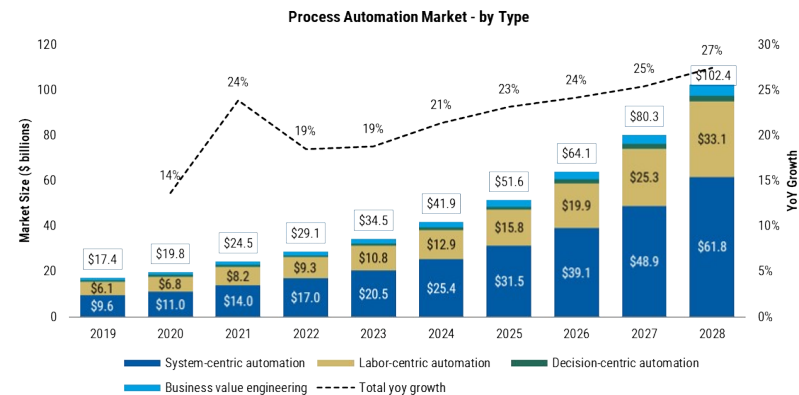
Bottom Line

An expansive view on the Agentic AI market points to a \$52 billion opportunity today, expected to reach \$102 billion by 2028

- **System of Record (SoR) AI Agent:** Aligning the SoR market opportunity with IDC's labor-centric automation market segment
 - \$16B market opportunity today → \$33B by 2028
 - Labor-centric automation includes business process automation and robotic process automation
 - Functionalities typically support people-oriented activities that require manual steps or human decision-making (i.e., loan approvals, new employee onboarding)
- **System of Engagement (SoE) AI Agent:** Aligning the SoE market opportunity with IDC's system-centric, decision-centric, and business value engineering market segments
 - \$36B market opportunity today → \$69B by 2028
 - Functions typically involve back-end automation of two or more systems to support interoperability, exchange info, and process transactions, along with rules-based decision automation

While our definitions of SoR/SoE AI Agents do not exactly match IDC's market segment definition, we believe these serve as close proxies

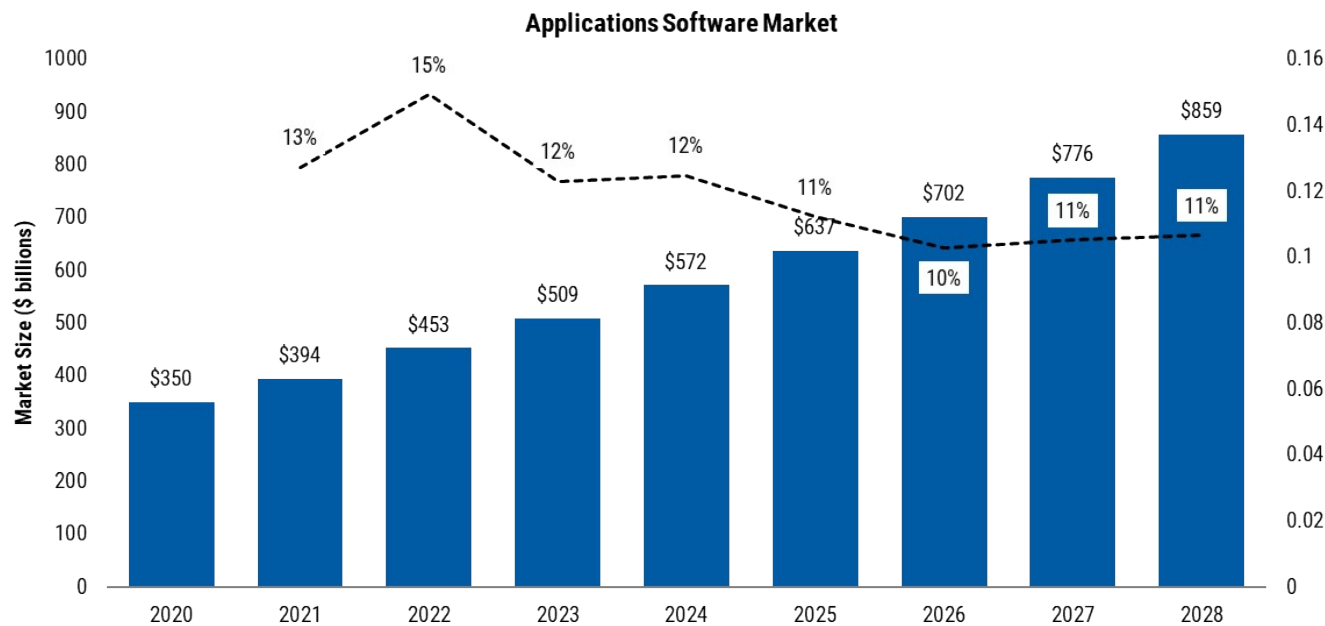
Source: IDC Worldwide Intelligent Process Automation Software Forecast



Debate #1: How Big is the Opportunity?

Bottom Line

As Agentic AI solutions mature, more of the Application Software market should open up as potential opportunity



Source: IDC Software Tracker Forecast

Morgan Stanley

MORGAN STANLEY RESEARCH

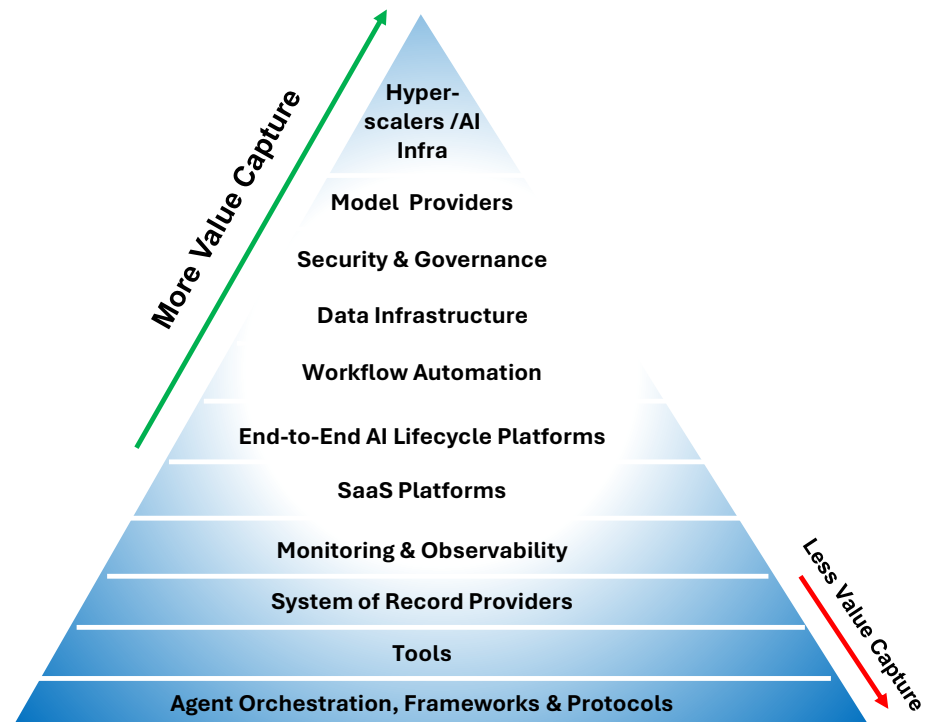
Debate #2: Where Will Value Accrue?

Bottom Line

The most value likely accrues to the hyperscalers and AI infrastructure layer as AI Agent deployments likely result in significant utilization of significant compute resources

• Other Parts of the Stack Likely to Benefit:

- **Model Providers:** Serving as the “brain” of any agentic solution gives license for model providers to own more of the AI app stack
- **Security, Governance & Guardrails:** The ability to infuse trust, reliability and accuracy while enforcing policy present significant opportunities for monetization
- **Data Infrastructure:** A necessary requirement to build useful, effective and accurate agentic systems require a modern data infra that make proprietary unstructured and structured data easily governed and accessible
- **Workflow Automaters:** Software providers that are ingrained in customers’ key workflows and business processes can create substantial value by further automating those workflows or re-engineering altogether with agentic AI



更多请关注公众号：洞研纪要更新
添加微信：jiyaosx

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #2: Assessing Where Value Accrues in the Stack

Layer	Definition	Rationale	Examples	Monetization
Hypescaler / AI Infra	Includes the compute, network, storage and model serving capabilities for the AI models that power agentic architecture	AI Infrastructure providers have a significant opportunity in providing the compute to enable the real-time inference that deep reasoning agents require. This gives leaders in this category license to not only provide the infrastructure but also the data retrieval, security, governance and policy enforcement capabilities required to operate, manage and secure agentic applications.	Amazon AWS, Google Cloud, Microsoft Azure, Coreweave, Lambda	Very High
AI Model Providers	Includes the state-of-the-art language, image, video and audio AI models both commercial and open source	AI model providers extend the frontier around reasoning and intelligence that will make AI agents more capable over time. This positioning provides the opportunity to extend into areas such as agentic frameworks, data processing and integration so as to evolve into modern application platforms.	Open AI, Anthropic, Llama, Deep Seek, Mistral, Grok	Very High
Security & Governance	Includes the state-of-the-art language, image, video and audio AI models both commercial and open source	We see a significant monetization opportunity in making agentic applications more trustworthy, reliable, accurate and secure. Agent architectures create a large surface area for cyber attacks including data exfiltration, supply chain attacks, prompt injection and open source contamination.	Palo Alto Networks, CrowdStrike, Okta, Sailpoint, Rubrik	Very High
Data Infrastructure	Includes the data management software that enables and govern access to an organization's structured, semi-structured unstructured data	For AI agents to be useful, they need governed access to the most important data sources in real-time. We see two key opportunities for data infrastructure providers. First, in helping make agentic architectures possible by modernizing a customer's data estate and improving data quality. Second, by enabling fast, reliable and accurate information retrieval at runtime.	Oracle, Snowflake, MongoDB, Confluent	High
Workflow	Workflow automation reduces manual work by automating the flow of tasks, documents and data involved in common work related activities and business processes	With AI agents embedded in workflow automation platforms, the unstructured components of common business processes can be more fully automated. As workflow automation platforms span across key enterprise systems, they become well-positioned to serve as an orchestration/management layer for both 1st and 3rd party agents.	ServiceNow, Atlassian, Monday Appian, UiPath	High
AI Lifecycle Platforms	Platforms that automate the lifecycle for building + running AI applications including data integration, data processing, ontology mapping, model building/evaluation and app deployment	Agents within AI lifecycle platforms automate the workflows and tasks associated with developing and deploying AI applications including building data pipelines, integrating enterprise data with LLMs, building and evaluating models and visualizing data. Furthermore, agents can be embedded into external applications to enhance customer engagement and to enable highly interactive experiences.	Palantir	High

Source: Company Data, Morgan Stanley Research

18

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #2: Assessing Where Value Accrues in the Stack

Layer	Definition	Rationale	Examples	Monetization
SaaS Platforms	Includes incumbent SaaS platforms who serve as system of record across multiple domains resulting in these platforms operating multiple data silos for their end customers	With agents, SaaS platforms will be able drive much deeper levels of automation and efficiency for each business function or domain they operate. AI Agents add value by unlocking the full capabilities of SaaS platforms. However, the challenge will be pricing as SaaS providers have often priced per seat or per user heading into an era where seats (at least for some functions) may stop growing or decline	Salesforce, Adobe, HubSpot	Medium
Observability	software that monitors the health and performance of the digital business including the underlying infrastructure, applications and end user activity	The adoption of agentic architectures increases the complexity of the app environment and expands the surface area for monitoring. Furthermore, observability agents can automate the investigation process and kick off remediation workflows typically executed by engineering staff. The debate will be whether incumbents capture the opportunity or whether there will be a need for a new breed of AI-native observability solutions.	Datadog, Dynatrace	Medium
System of Record	Systems of record software are information retrieval systems that serve as the source of truth and central repository for data of specific domain ensuring consistency and accuracy across the entire organization	AI agents, simplify access to systems of record and make it easier for customers to build operational and analytical applications on top of these core business systems. In addition, agents can replace human labor for common tasks such as cash collection associated with key business functions. However, to fully automate business processes that span multiple systems, requires agents that can coordinate across multiple systems	Workday, ADP, Intuit	Medium
Tools	Encompass the set of actions that an AI agent can take by calling specific services or functions it has access to.	Tools make AI agents useful by giving Agents the ability to take action in the face of query or trying to solve a problem or task. Tools can include calendars, calculators, code interpreters, web search, enterprise search, APIs, databases and much more.	Calculators, web search, code interpreters, databases	Low
Frameworks & Protocols	Includes frameworks and protocols to build multi-agent applications, foster agent to agent communication, enable multi-agent orchestration and connect LLMs to external tools and data sources	The abstraction and standardization benefits that application frameworks and management/orchestration tools provide play a crucial role in accelerating the pace of development of modern applications. However, the monetization in this area of software has been mixed looking back over the past few decades, particularly as open source technologies have played a more dominant role within application infrastructure	Autogen, MCP, A2A, ACP	Low

Morgan Stanley

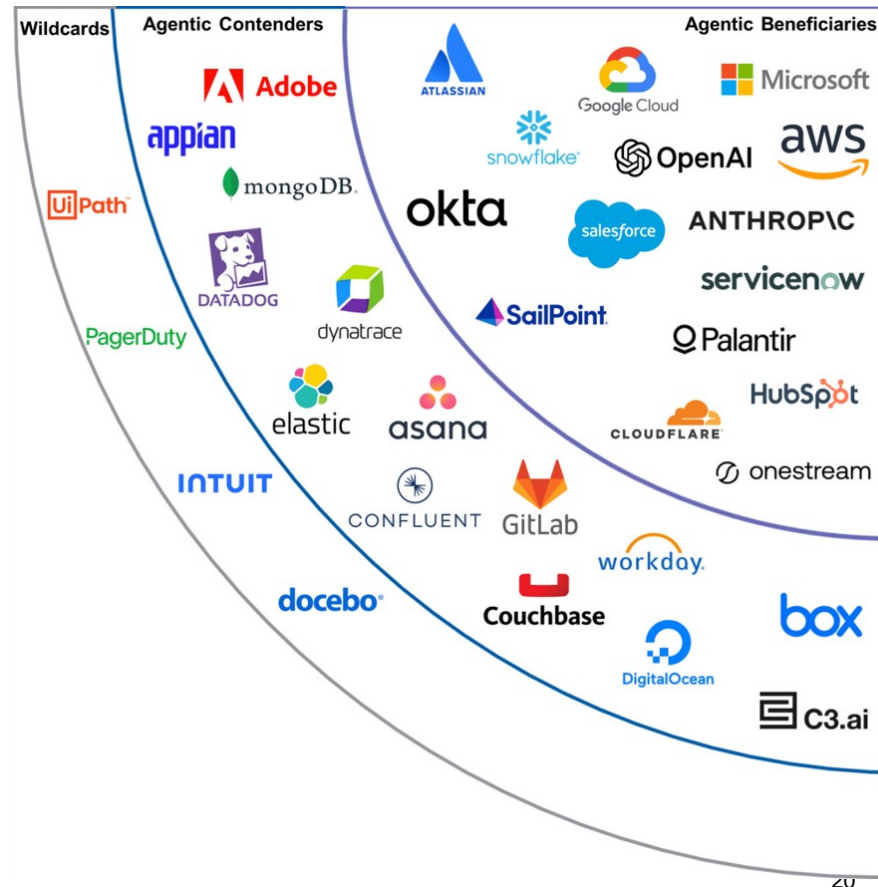
MORGAN STANLEY RESEARCH

Debate #3: Which Public Companies Are Positioned to Capture the Opportunity?

Bottom Line

Software vendors best positioned to benefit based on a) ability to monetize AI Agents/Agentic architecture or b) core business benefitting from enterprise adoption

- **Agent Beneficiaries** are vendors where we have conviction that can either a) directly monetize AI Agents/Agentic architectures today or in the near future such that it becomes a meaningful contributor to growth or b) see their core business benefit materially from agentic adoption
- **Agent Contenders** are vendors that are in the right categories to see benefit but require a high degree of execution and market maturity that is not yet evident
- **Wildcards** are vendors where there is debate on whether agentic adoption results in a tailwind to growth or proves disruptive to the core business



Source: Morgan Stanley Research

Debate #3: Which Vendors Are Best Positioned to Win?

Agent Beneficiaries – Hyperscalers	
Microsoft	<ul style="list-style-type: none">• In our 1Q25 CIO Survey, CIO’s indicated Microsoft as the top vendor expected to be used for Agentic AI strategies• Microsoft’s focus and expertise on security and data privacy create critical trust within vast existing customer installed base to develop and deploy agents responsibly• Microsoft’s ubiquitous productivity & business apps allow for deep integrations w/ Enterprise workflows and data; its strong developer ecosystem and platform offerings provide a natural adoption path of Microsoft’s Agentic AI products• Microsoft benefits from its comprehensive AI portfolio and its breadth of Agent specific offerings, including: GitHub Copilot’s asynchronous coding agent; Azure AI Foundry Agent Service, Microsoft 365 Copilot Studio / Agent Build, Agents for Microsoft Teams, M365 Agents toolkit for Visual Studio, M365 Agents SDK, Data Agents in Microsoft Fabric, SRE agents (among others)
Amazon AWS	<ul style="list-style-type: none">• AMZN continues to be the share leader within the public cloud market, with the largest base of customers to cross sell GenAI capabilities into. AWS has already generated multiple billions of dollars in AI revenue.• In our 1Q25 CIO Survey, Amazon ranked as the second largest gainer of incremental share of both overall IT budgets and GenAI workloads.• AMZN offers Amazon Q to businesses, integrating with a user’s AWS account, business intelligence tools, contact center and business apps to provide comprehensive and personalized answers to user questions as well as perform tasks on behalf of the user.
Google Cloud	<ul style="list-style-type: none">• In our 1Q25 CIO Survey, Google Cloud was the third largest gainer of incremental share of IT budgets as enterprises shift to the cloud in the next three years, trailing Microsoft and Amazon• Google has announced various agentic offerings and shipping products at an accelerated pace, including Project Mariner (which is coming to Vertex AI), Gemini Code Assist, Vertex AI Agent Builder, and Gemini for Google Workspace• Gemini 2.5 Pro is currently the leading model across a variety of benchmarks and Google’s integration of Gemini across GCP could improve capabilities and enterprise adoption

Source: Company Data, Morgan Stanley Research

Debate #3: Which Vendors Are Best Positioned to Win?

Agent Beneficiaries – Model Providers	
OpenAI	<ul style="list-style-type: none">• OpenAI has established itself as a frontrunner in Agentic AI development, with several foundational models serving as the intelligence backbone for agentic systems. The company has released a complete development ecosystem for building agentic applications including 1) Agents SDK, 2) Responses API, and 3) built-in tools integration• The company has moved beyond theoretical frameworks and delivered practical agentic applications that demonstrate its potential including 1) Operator, an autonomous AI agent capable of performing complex tasks through web browser interactions, 2) Codex, an agentic web-based coding tool, and 3) Deep Research, an agent that conducts in-depth web searches and analyses to generate comprehensive reports• OpenAI has significant financial backing enabling it to innovate, having raised a total of \$48 billion in private funding with the latest record-setting \$40 billion funding round at a \$300 billion valuation – would make it the fourth largest publicly traded software company after MSFT (\$3.4 trillion), ORCL (\$560 billion), and PLTR (\$318 billion)
Anthropic	<ul style="list-style-type: none">• Anthropic has emerged as a formidable contender to OpenAI, with differentiated technological capabilities, strategic partnerships, and a governance framework uniquely suited for enterprise adoption• The company's Claude models provide the cognitive backbone for advanced agentic systems, including Claude 3.7 Sonnet which introduced a breakthrough “hybrid reasoning” framework that enables users to control the computational resources used, optimizing performance based on task complexity. The model family (Haiku, Sonnet, Opus) offers right-sized solutions across use cases, from real-time customer service agents to longer term tasks with many steps• Strategic partnerships – AWS and Accenture – and constitutional AI framework is helping drive enterprise adoption, particularly in highly regulated industries. Partnership with AWS and Palantir to provide U.S. intelligence and defense agencies access to Claude, validating security credentials for sensitive applications

Source: Company Data, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #3: Which Vendors Are Best Positioned to Win?

Agent Beneficiaries – Cybersecurity	
SailPoint	<ul style="list-style-type: none">IGA solutions are focused on helping organizations set and enforce policies around who has access to what applications and data in an organization and as AI agents are deployed at scale across an enterprise, they will need an increasing level of access to both data and applications to complete more complex tasks. With SailPoint's foundation in IGA, existing human identity security capabilities should translate well to non-human'sAnnounced its "Agent Identity Security" solution in March with capabilities including 1) managing lifecycle of AI agents with automated governance, and 2) enforcing access certifications to prevent unauthorized access and security risks.
Okta	<ul style="list-style-type: none">A leader in human authentication, a growing presence within governance, and the company's "Auth for GenAI" solution highlight the continued expansion of Okta's capabilities across multiple identity security growth vectors.While Okta is taking a different approach to positioning itself within the agentic marketplace, its leadership in authentication and success with newer product cycles highlight the company as another potential winner as the use of AI agents across the enterprise continue to expand.Okta's "Auth for GenAI" solution includes 1) user authentication, 2) calling APIs on the user's behalf, 3) async authentication, and 4) fine grained authorization for retrieval augmented generation (RAG).
Cloudflare	<ul style="list-style-type: none">In early innings of monetizing its multi-product platform and has emerged as a potential GenAI winner as orgs. look to adopt and experiment with LLMs. By leveraging its "Edge Network" with 330+ PoPs globally, it's well-positioned to gain share given increased AI inference at the edge, especially as it relates to cost, latency and data security/residency.Cloudflare's Developer Services Platform already has >3 million developers and continues to see growing adoption across R2 Object Storage, Workers and AI gateway for training, inference and monitoring AI workloads.Cloudflare's Workers AI, a serverless GPU-powered platform on Cloudflare's global network, provides all the foundational elements of running AI inference on GPUs in 190+ cities.

Source: Company Data, Morgan Stanley Research

Debate #3: Which Vendors Are Best Positioned to Win?

Agent Beneficiaries – Data Infrastructure	
Snowflake	<ul style="list-style-type: none">• Snowflake is poised to benefit from the deployment of AI agents and agentic architectures in two fundamental ways. First, as one of the popular stores of business data and increasingly semi-structured and unstructured data, Snowflake will be a key source of context and information for a variety of 3rd party and custom developed AI agents. Second, Snowflake is also well positioned to monetize agentic architectures with its Cortex Agents.• Cortex agent leverage two key AI “tools” from Snowflake – 1) Cortex Analyst and 2) Cortex Search. Cortex agents call Cortex Analyst in scenarios when the agent needs to respond to queries requiring structured data leveraging Cortex Analyst’s text-to-sql capabilities. In cases, where the agent needs to respond to an unstructured data request (such as document or an invoice) Cortex Search is called given its hybrid search (semantic + keyword) capabilities.• Snowflake’s agents can be embedded in any application and the core value proposition is the ability to respond to structured and unstructured queries from a single platform.
Palantir	<ul style="list-style-type: none">• As part of its Foundry platform for commercial customers, Palantir AIP allows customers to securely deploy LLMs within organizations and provides a set of integrated capabilities including data integration, ontology mapping, management services, workflow automation, application deployment and many others to build AI applications.• AIP supports the building and deployment of AI agents that can then be embedded within AIP applications via the AIP Agent Studio module within AIP• With Agent Studio, customer can build agents of varying complexity including simple ad-hoc agents, task-specific agents, agentic applications (that can manage and update state) and fully automated agents that can be delegated tasks to fully execute complex workflows.

Source: Company Data, Morgan Stanley Research

Debate #3: Which Vendors Are Best Positioned to Win?

Agent Beneficiaries – Workflow Automation	
ServiceNow	<ul style="list-style-type: none">• ServiceNow’s expansion in the scope of workflows automated and ability to combine AI, data, and workflow automation position it to become the de facto “AI operating system for enterprises”.• AI agents and their ability to solve more complex use cases will be based upon the data they’re built upon and ServiceNow’s expansive workflow automation capabilities (IT, HR, customer service, etc.) combined with its CMDB serving as the system of record for enterprises’ IT infrastructure (>85% of Fortune 500 are customers) provide the company with a robust data asset that it can leverage to offer differentiated AI agents.• Building upon this data advantage, ServiceNow’s Workflow Data Fabric offering provides customers with the ability to connect their data residing outside of ServiceNow into the platform, unifying customer data and expanding the number of use cases AI agents can address.
Atlassian	<ul style="list-style-type: none">• Atlassian’s core products are deeply embedded in the daily workflows of enterprises and teams, giving the company a unique vantage point on how modern organizations collaborate, manage projects, and share knowledge. This foundation, along with a highly efficient innovation engine, positions it well to integrate agentic AI across core workflows and benefit from the emerging opportunity• Atlassian has embedded Rovo and Atlassian Intelligence – its AI/Agentic AI functionalities – into its core products (Jira, Confluence, and Jira Service Management). We view this approach as the right strategic move for driving near-term adoption and usage to secure a better positioning for the long-term agentic AI opportunity• Today there are >1.5 million MAUs on Rovo and Atlassian Intelligence, up from >1 million a quarter ago

Source: Company Data, Morgan Stanley Research

Debate #3: Which Vendors Are Best Positioned to Win?

Agent Beneficiaries – SaaS Platforms	
Salesforce	<ul style="list-style-type: none">Salesforce remains a key play on the incumbency advantage to operationalizing AI agents, with competitive differentiation across: 1) Defined Role, where Salesforce can extrapolate clearly outlined jobs for the agents that exist inside an organization, de-risking the process & time to defining the experience; 2) Data, where Data Cloud ingests structured, unstructured, metadata, & federated data both internally from Salesforce and from external sources (Snowflake, Databricks, etc), to drive knowledge behind the agents; 3) Actions, or the workflows needed to perform a job, where Salesforce can capitalize on years of experience building a broad swath of workflows, all while leveraging the Trust Layer to place ethical and accuracy guardrails around each step of the agent processSalesforce’s solid Agentic positioning is growing increasingly apparent with building Agentforce momentum: as of F1Q26, the company sustained 120% YoY growth in Data + AI ARR (~\$100M of which is specific to Agentforce) & reported ~8k Agentforce deals, up from ~5k in 4Q25
HubSpot	<ul style="list-style-type: none">HubSpot is well positioned to address the agent opportunity given: 1) deep data sets combining rich structured and unstructured data; 2) a unified platform approach that provides context across data, engagement (myriad different hubs), and action (AI-powered execution); and 3) a vibrant agent ecosystem developed through the company’s Agent.ai network. Further, HubSpot’s value proposition as an easy-to-use, out of the box solution with AI embedded across the platform allows for quick time-to-deployment and potential to scale over timeStrategically, HubSpot remains relentlessly focused on driving strong adoption and utilization with new products before introducing direct monetization, a framework we believe should foster strong and durable monetization tailwinds aheadWith usage and adoption statistics starting to materially inflect, HubSpot’s move to integrate credit-based pricing strategies for the Customer Agent product bridges an intuitive hybrid Subscription/Consumption model which should well align price to value for the solution
OneStream	<ul style="list-style-type: none">AI Agents have a higher likelihood of initial adoption and are monetizable within domains that are: 1) Primarily cost centers, where organizations can realize tangible ROI from reduced headcount; 2) Experiencing talent shortages, where organizations are increasingly looking to technology as the solution for the labor gap; and 3) High levels of manual data processing. Furthermore, given enterprise CFO buyers are inherently more risk-averse, we believe this will allow for existing vendors in this space to benefit versus new vendors.OneStream has been developing AI/ML tools since 2017 (SensibleML) and recently launched multiple agents (Finance Analyst, Operations Analyst, Search Analyst, Deep Analysis) that have seen strong initial customer feedback as the agents can handle a lot of the manual and data intensive parts of their jobs.

26

Source: Company Data, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #3: Which Vendors Are Best Positioned to Win?

Gartner defines this segment as “tools that augment and extend IT Service Management workflows using AI.”

- They analyze ITSM data and metadata to provide advice and actions on IT service desk and support activities.
- They can either be stand-alone solutions, capabilities within an ITSM platform, or an add-on to an ITSM platform.

While this definition and solutions are not strictly “Agentic”, this gives us a rough sense of which vendors are best positioned in the ITSM space as this technology progresses and garners adoption.

- **ServiceNow** offers “thousands” of pre-built AI agents for “every” workflow and an AI Agent Studio to build fully customized agents. The company also recently introduced its AI Agent Orchestrator, a “control tower” that conducts planning and leads collaboration between ServiceNow and third-party AI Agents.
- **Aisera’s** Agent Assist is an Agentic AI-driven tool designed to boost agent productivity by providing real-time, context-aware answers and case summaries, recommend next best actions, and draft knowledge-base articles. Embeds directly within SaaS apps like ServiceNow, Salesforce, Zendesk, Jira, Freshworks, etc.
- **Moveworks’** platform, powered by its Agentic Automation Engine, enables users to create AI Agents for automating workflows across the organization.

Magic Quadrant for AI Applications in IT Service Management



Debate #3: Which Vendors Are Best Positioned to Win?

Gartner defines this segment as “tools that assist in generating and analyzing software code and configuration.”

- They use LLMs fine-tuned for code and/or program-understanding technology, and integrate into developer tools like code editors, command-line terminals and chat interfaces.
- Developers can prompt it to generate, analyze, debug, fix, and refactor code, to create documentation, and to translate code between languages.

While this definition and solutions are not strictly “Agentic”, this gives us a rough sense of which vendors are best positioned in the Software Developer space as this technology progresses and garners adoption.

- **GitLab Duo** provides AI capabilities across the software delivery lifecycle, including 1) code generation, completion, explanation, refactoring, 2) security vulnerability and root cause analysis, and 3) value stream analytics. In addition, GitLab is working on **GitLab Duo Workflows**, moving beyond code assistant and into Agentic AI. It'll behave like an autonomous always-on agent monitoring projects to identify and resolve
- **GitHub Copilot** developed by GitHub in collaboration with OpenAI, provides real-time code suggestion, completions, and assistance throughout the software delivery lifecycle. It was first to market (ahead of GitLab) and is the most widely used code assistant today.
- **Google’s Gemini Code Assist**
- **Amazon CodeWhisperer**

Magic Quadrant for AI Code Assistant



Source: Gartner Magic Quadrant for AI Code Assistant

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #3: Which Vendors Are Best Positioned to Win?

Gartner defines this segment as “*software platforms used to build, orchestrate and maintain multiple use cases and modalities of conversational automation.*”

- Most notably addresses customer service and support, IT service desk, HR, and call center automation.

While this definition and solutions are not strictly “Agentic”, this gives us a rough sense of which vendors are best positioned in the conversational agent space (i.e., customer service, sales and marketing, etc.) as this technology progresses and garners adoption.

- **Kore.ai** offer an Agent Platform with pre-built and custom AI Agent for internal support (i.e., HR, IT), external (i.e., customer service), and process automation.
- **Cognigy** offers customer service specific AI agents that can be easily deployed in contact centers. It offers voice and chat agents, as well as agent copilots for human service agents.

Magic Quadrant for Enterprise Conversational AI Platforms



Source: Gartner Magic Quadrant for Enterprise Conversational AI Platforms

Morgan Stanley

MORGAN STANLEY RESEARCH

Debate #4: How Will Agents be Priced?

Bottom Line	Many vendors focused on driving adoption and usage rather than generating meaningful revenue from their AI Agents. Despite having a pricing model in place, most are experimenting with different forms, including usage-based, workflow-based, and outcome-based.
Per Agent	<ul style="list-style-type: none"> • Fixed Monthly Fee: Customers pay a recurring fee for access to the AI agent, regardless of usage (i.e., Devin) • Tiered Pricing: Different pricing tiers offer varying levels of access and features • Per-User Pricing: Customers pay a fee per user who has access to the AI agent (i.e., NOW)
Per Action / Consumption	<ul style="list-style-type: none"> • API Calls: Vendors charge based on the number of requests made to the AI agent's API (CRM's Flex Credits) • Compute Time/Data Processed: Pricing is tied to the amount of processing power or data the AI agent uses • Token-Based: Based on the number of tokens used for input and output by the AI agent (i.e., Anthropic, OpenAI)
Per Workflow	<ul style="list-style-type: none"> • Per Conversation: Customers pay a fee for each AI agent interaction or conversation (i.e., CRM)
Per Outcome	<ul style="list-style-type: none"> • Latest evolution of pricing, whereby customers pay based on the successful outcomes achieved by the AI agent, such as resolved conversations, saved cancellations, or increased sales. • While not publicly available, Sierra emphasizes an "outcome-based pricing" approach for their customer service AI agent.
Embedded	<ul style="list-style-type: none"> • Access to AI Agents/features included in the higher subscription tier. • Current state of many vendors who are focusing on driving customer adoption and usage (i.e., TEAM and HUBS).

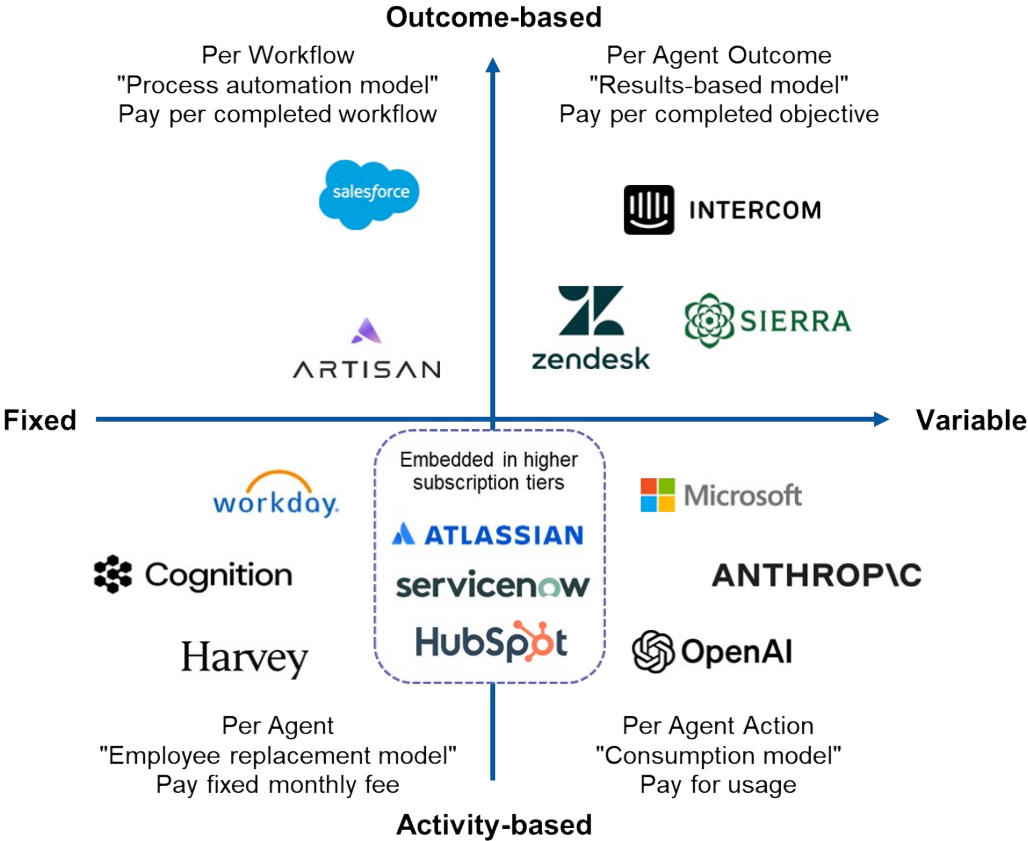
Source: Company Data, Morgan Stanley Research

30

Debate #4: How Will Agents be Priced?

Bottom Line

Most companies still use the traditional seat-based models, though increasingly combined with a usage-based component. Workflow and outcome-based models are gaining traction



Source: <https://www.growthhunched.com/p/ai-agent-pricing-framework>, Morgan Stanley Research

Debate #4: How Will Agents be Priced?

Case Study: Workflow-based Pricing – Salesforce (Agentforce)

Initially launched with straightforward **\$2 per conversation** model

- Conversation types include:
 - **ASA Messaging:** Conversation window starts when the user first sends a message and ends 24 hours later
 - **Sales Coach:** Conversation is consumed when a user clicks “Get Feedback” following a role play or a stand-and-deliver coaching session
 - **Sales Development Rep (SDR):** Conversation is consumed when an SDR agent sends out the initial email to a lead. Includes all subsequent activities carried out by the agent for that lead
- Each conversation includes 200 Einstein Requests, with additional requests drawing from a shared pool or incurring overage charges

Introduced three additional pricing models on May 15, 2025, including:

- **Flex Credits:** Offers granular consumption-based pricing at \$0.10 per action (20 Flex Credits), allowing customers to pay only for specific agent activities. Credits are sold in packs of 100K for \$500
- **Flex Agreement:** Convert user licenses into Flex Credits or vice-versa, allowing you to move your budget between people and AI agents as needed
- **Agentforce User Licenses,** and add-ons with included Agentforce usage: Unlimited employee-facing agent usage through traditional per-user, per-month pricing. *Pricing coming Summer 2025*

Source: Company website

Get started with Agentforce today.

Currency converter ⓘ
US Dollar (USD) ▼

Flex Credits
Pay per action as you scale
Agentforce.

\$500
USD/Per 100k Credits

- ✓ Out-of-the-box agent skills
- ✓ Agent Wizard
- ✓ Agent Builder
- ✓ Prompt Builder

Contact us

Conversations
Drive efficiency with digital labor
24/7.

\$2
USD/Per conversation

- ✓ Out-of-the-box agent skills
- ✓ Agent Wizard
- ✓ Agent Builder
- ✓ Prompt Builder

Contact us

Standard Success Plan
Set your company up for success
with self-guided, always-on
resources.

**Included in all
licenses**

- ✓ Access to Knowledge Articles,
documentation, and
Trailhead
- ✓ Global Trailblazer Community

Contact us

Agentforce Rate Card

Actions are the individual tasks your agents perform to complete customer requests. Examples of actions include updating account details, scheduling appointments, summarizing case information, or providing answers using our knowledge base.
This page is provided for information purposes only and is subject to change. Contact a sales representative for detailed pricing information.

Debate #4: How Will Agents be Priced?

Case Study: Outcome-based Pricing – Zendesk (AI Agents)

Pricing strategy exemplifies the evolving landscape of enterprise AI monetization, blending traditional subscription with outcome-based pricing

- Existing subscription plans embed AI capabilities, with incremental features on higher tiers and add-ons
- Outcome-based pricing: Introduced in Sept 2024 as a model directly dependent on the AI agent performance
 - Pay \$1.50-\$2 per successfully resolved AI interaction – with volume discounts available
 - Exempts escalations to human agents from charges, directly tying costs to business results

Outcome-based model addresses key customer pain points:

- Cost predictability: Avoid overpaying for underperforming AI
- Scalability: Unlimited AI capacity during peak periods w/o fixed staffing costs
- ROI transparency: Compare cost per AI-resolved ticket to human cost

We note some customer confusion over what constitutes an “automated resolution”, with abandoned chats sometimes triggering charges

Source: Company website

Traditional Subscription

MOST POPULAR			
<div>Support Team</div> <div>Get up and running fast with streamlined email support.</div> <div>\$19</div> <div>per agent/month billed annually</div> <div>Buy now</div> <div>Starting with the basics<ul style="list-style-type: none">✓ Simplify email support with ticketing✓ Automate with routing and triggers✓ Action insights with prebuilt reporting✓ Unify data with 1,000+ integrations</div>	<div>Suite Team</div> <div>Automate with AI across email, messaging and live chat, phone, and social media.</div> <div>\$55</div> <div>per agent/month billed annually</div> <div>Buy now</div> <div>Support Team, plus<ul style="list-style-type: none">✓ Provide service across any channel✓ Offer 24/7 support with a help center✦ Automate resolutions with AI agents FEATURED✦ Shrink queues with generative replies FEATURED✦ Improve CSAT with quality assurance FEATURED</div>	<div>Suite Professional</div> <div>Optimize operations with AI, data, and customizations</div> <div>\$115</div> <div>per agent/month billed annually</div> <div>Buy now</div> <div>Suite Team, plus<ul style="list-style-type: none">✓ Collaborate to close tickets✓ Refine operations with CSAT insights✓ Build trust with working hours and SLAs✓ Up security with HIPAA and local storage✦ Reduce wait times with smart routing FEATURED✦ Assist and optimize with agent Copilot FEATURED</div>	<div>Suite Enterprise</div> <div>Adapt service safely with change management tools.</div> <div></div> <div>Talk to Sales</div> <div>Suite Professional, plus<ul style="list-style-type: none">✓ Test updates safely in sandbox✓ Control access with custom permissions✓ Streamline approvals and track changes✓ Handle spikes with extra capacity✓ Tailor by brand with 300 help centers✓ Save time with tailored workspaces</div>

Outcome-based Pricing

	Support Team	Suite Team	Suite Professional	Suite Enterprise
AI agents				
Essential plan	✓*	✓	✓	✓
Advanced plan	Add-on	Add-on	Add-on	Add-on
Included automated resolutions (up to 10K)	5 ARs per agent/month	5 ARs per agent/month	10 ARs per agent/month	15 ARs per agent/month
Committed automated resolutions (volume discounts available)	\$1.50	\$1.50	\$1.50	\$1.50
Pay-as-you-go automated resolutions	\$2	\$2	\$2	\$2

Debate #4: How Will Agents be Priced?

Case Study: Combined Model – HubSpot (Breeze)

Two-tiered model combining traditional subscription with consumption-based credits

Understanding the Philosophy

- Given HubSpot's strategic approach to driving healthy and consistent adoption/utilization on new products (including Agents) before moving to direct monetization, the company initially rolled out Breeze Agents in Beta across all Platforms at September 2024's INBOUND Conference


Direct Monetization Phase

- HubSpot Credits:** As of June 2025, HubSpot is rolling out HubSpot Credits, a pay-as-you-scale model that allows customers to pay for usage-based features. Credits are utilized when an action is performed that consumes them (ex, Service tickets resolved)

Key Considerations

- How Many Credits are Allocated?*
- Free: N/A
- Starter: N/A
- Pro: 3k credits/month
- Enterprise: 5k credits/month
- What Subscriptions/Seats are Eligible?
 - All Hubs (Marketing, Sales, Service, Operations, Content) across all tier levels (Starter, Pro, Enterprise)
 - Only Paid & Partner seats (not free/view-only users)
- How do Users Purchase Incremental Credits Above What is Allocated?
 - Small: 5k incremental credits/month, priced at \$45/month
 - Medium: 30k incremental credits/month, priced at \$270/month
 - Large: 100k incremental credits/month, priced at \$900/month

Highest level of HubSpot subscription	Included HubSpot Credits in subscription per month
Starter	500
Professional	3,000
Enterprise	5,000

 **Breeze Customer Agent**
Your AI concierge for seamless, instant customer connections.

Small
5,000 credits/mo
Starts at
\$45/mo

Unlock Breeze Customer Agent with a Professional or Enterprise plan

Talk to Sales

How you can use 5,000 credits:
✓ Handle up to 50 customer conversations per month with Breeze Customer Agent

Medium
30,000 credits/mo
Starts at
\$270/mo

Unlock Breeze Customer Agent with a Professional or Enterprise plan

Talk to Sales

How you can use 30,000 credits:
✓ Handle up to 300 customer conversations per month with Breeze Customer Agent

Large
100,000 credits/mo
Starts at
\$900/mo

Unlock Breeze Customer Agent with a Professional or Enterprise plan

Talk to Sales

How you can use 100,000 credits:
✓ Handle up to 1,000 customer conversations per month with Breeze Customer Agent

Source: Company website

Morgan Stanley

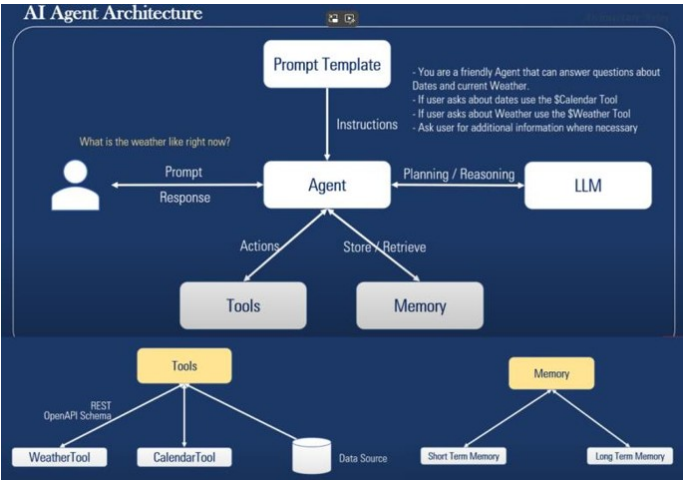
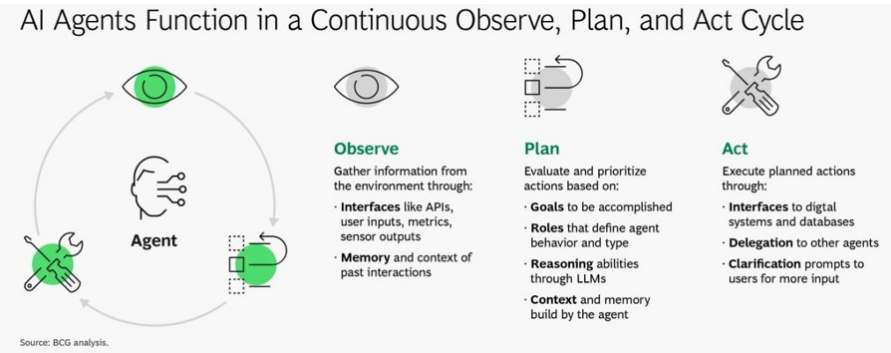
MORGAN STANLEY RESEARCH

Part II

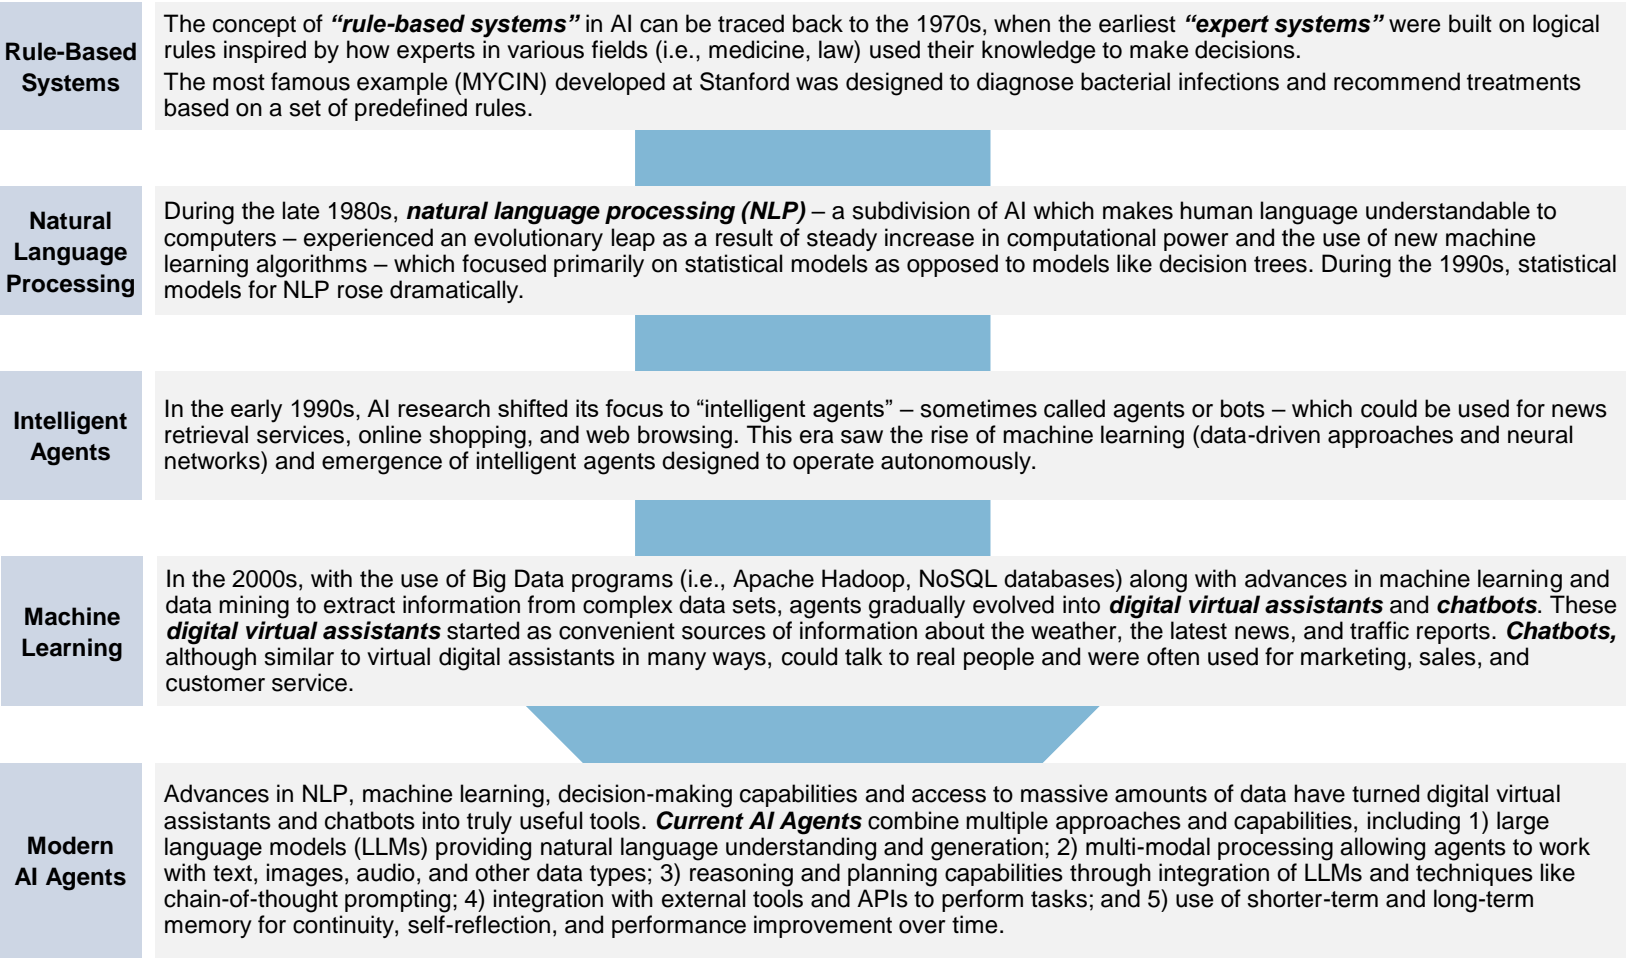
How Did We Get Here?

What is an AI Agent?

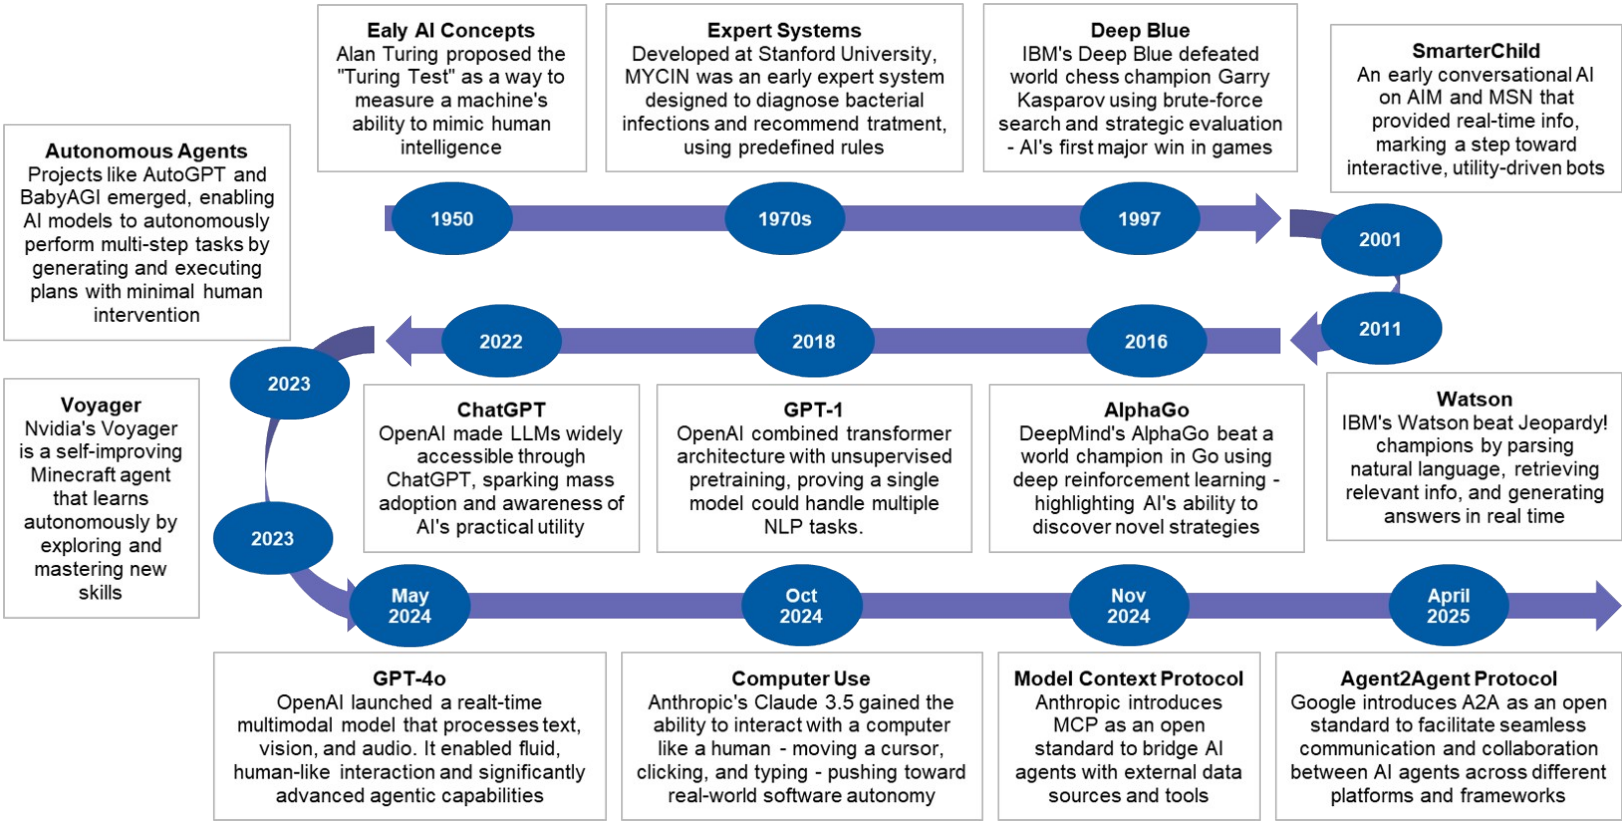
What is an AI Agent	An AI Agent is a system (software program) that understands and interact with its environment (receive inputs), collect data, make decisions, executes tasks and meet specified goals without using pre-determined rules.
Autonomous AI Agent	Autonomous AI Agents are systems that leverage large language models (LLMs) to plan and execute a task or process without a high degree of human interaction. These agents distinguish themselves from standard GenAI chatbots by combining tools (i.e., databases, websites, APIs, and even other agents) with memory of past interactions to chain multiple thoughts and actions together, enabling them to autonomously work toward defined objectives while drawing on both external information and learned experiences to produce relevant outputs.
Agentic Architecture	Agentic Architecture generally refers to the idea of bringing together multiple engines for understanding queries/requests, evaluating strategies for solving problems, accessing necessary data and executing the planned actions. Under this architecture, a large language model (LLM) typically serves as a reasoning engine, connected to tools and memory. Tools help connect the LLM to other sources of data or computation and can be used to take actions (run code, modify files, etc.). Memory (short or long term) helps the agent remember previous interactions with either humans, other agents, or tools.



Evolution Over Time



Key Milestones and Recent Breakthroughs



Source: Morgan Stanley Research

RPA vs AI Agents?

Key Distinction

Robotic Process Automation uses pre-defined rules to automate tasks that involve structured data. In contrast AI agents are able to automate tasks and processes using indeterministic, reasoning and context while showing the ability to adapt to change

Robotic Process Automation (RPA)

RPA uses software robots or “bots” to automate repetitive, rule-based tasks by mimicking human interactions with digital systems. It operates based on predefined workflows and structured data.

- **Decision Making:** Follows pre-defined rules and decision trees
- **Data Handling:** Structured, predefined inputs
- **Orchestration:** Centralized control, follows fixed workflows
- **Interaction:** UI-based, mimics human clicks and keystrokes
- **Fault Tolerance:** Low; fails when encountering unexpected inputs
- **Learning:** No learning; runs static rules to execute routine tasks

VS

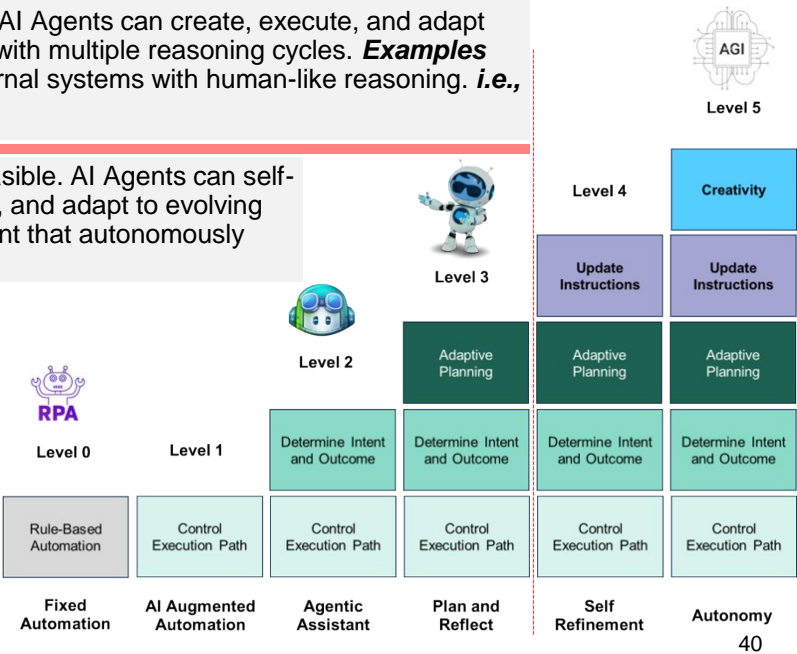
AI Agents

AI Agents are autonomous systems that perceive their environment, analyze data, and take actions to achieve specific goals. They use machine learning, natural language processing, and reasoning to adapt, learn from experiences, and make decisions without explicit human intervention.

- **Decision Making:** Autonomous, undeterministic based on reasoning and context
- **Data Handling:** Structured and unstructured, dynamic inputs
- **Orchestration:** Decentralized, can self-organize and collaborate
- **Interaction:** Conversational, API-driven, and reasoning-based
- **Fault Tolerance:** High; can adjust to errors and uncertainties
- **Learning:** Continuous learning allows for reasoning, problem-solving and decision making

Levels of Agentic Automation

Level 0	No Agentic behavior – pure rule-based automation. Tasks follow a fixed, deterministic process with no planning or execution control. Examples include Data entry and web scrapping. <i>i.e., RPA</i>
Level 1	Basic AI-enhanced automation where LLMs assist in decision-making but within predefined steps. Limited benefits over traditional automation. Examples include routing customer support emails.
Level 2	Task-specific AI assistant capable of using tools, interpret intent, execute actions like summarization or content generation, but follow static, short-term plans. Examples include AI co-pilots for search, summarization, and email drafting. <i>i.e., GitHub Copilot</i>
Level 3	Current Level – First level constrained autonomy. AI Agents can create, execute, and adapt plans based on feedback, handling complex tasks with multiple reasoning cycles. Examples include reconciling a 100-page invoice against internal systems with human-like reasoning. <i>i.e., Agentforce</i>
Level 4	While still theoretical today, they are technically feasible. AI Agents can self-improve, modify their instructions, create new tools, and adapt to evolving tasks. Examples include invoice reconciliation agent that autonomously adds vendors and improves accuracy.
Level 5	Full AGI (artificial general intelligence) – capable of independent reasoning, creativity, and solving novel problems beyond training. Examples include digital knowledge worker handling complex tasks with no oversight.



Source: sema4.ai "The Five Levels of Agentic Automation"

Morgan Stanley

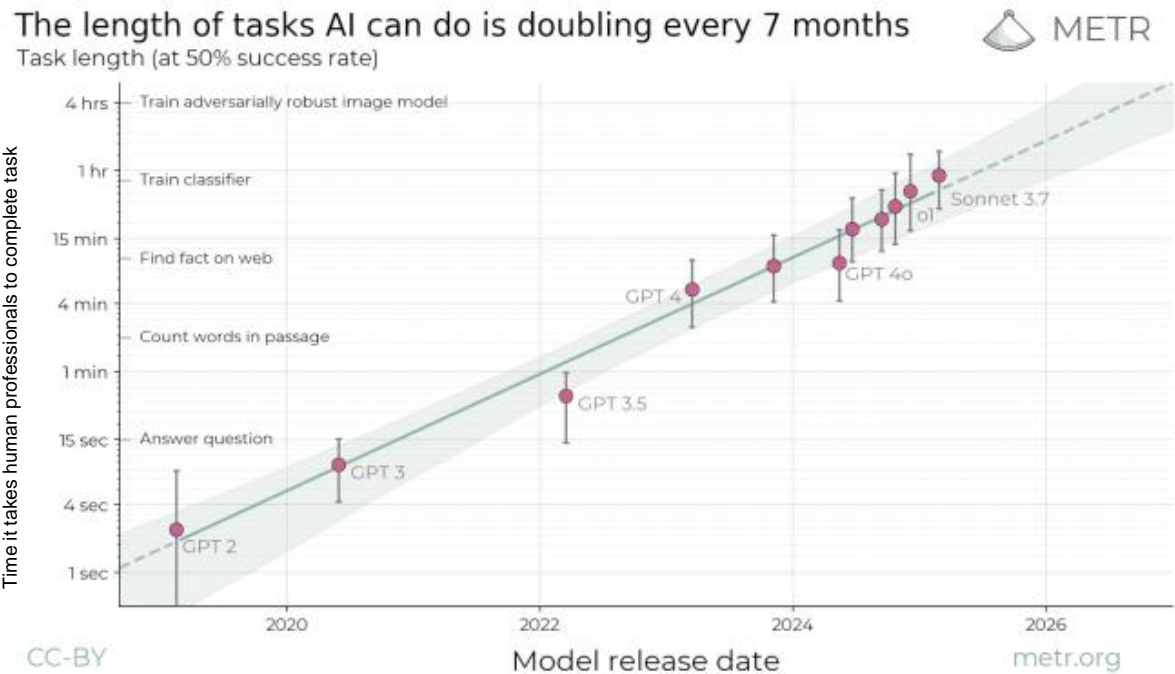
MORGAN STANLEY RESEARCH

AI Agents Are Rapidly Improving the Ability to Complete Long Tasks

Bottom Line

Length of tasks that generalist frontier model agents can complete autonomously has been doubling approximately every 7 months for the last 6 years.

In under a decade, AI Agents are expected to independently complete a large fraction of software tasks that currently take humans days or weeks



Source: METR “Measuring AI Ability to Complete Long Tasks”
Note: Shaded region represents 95% confidence interval

Morgan Stanley

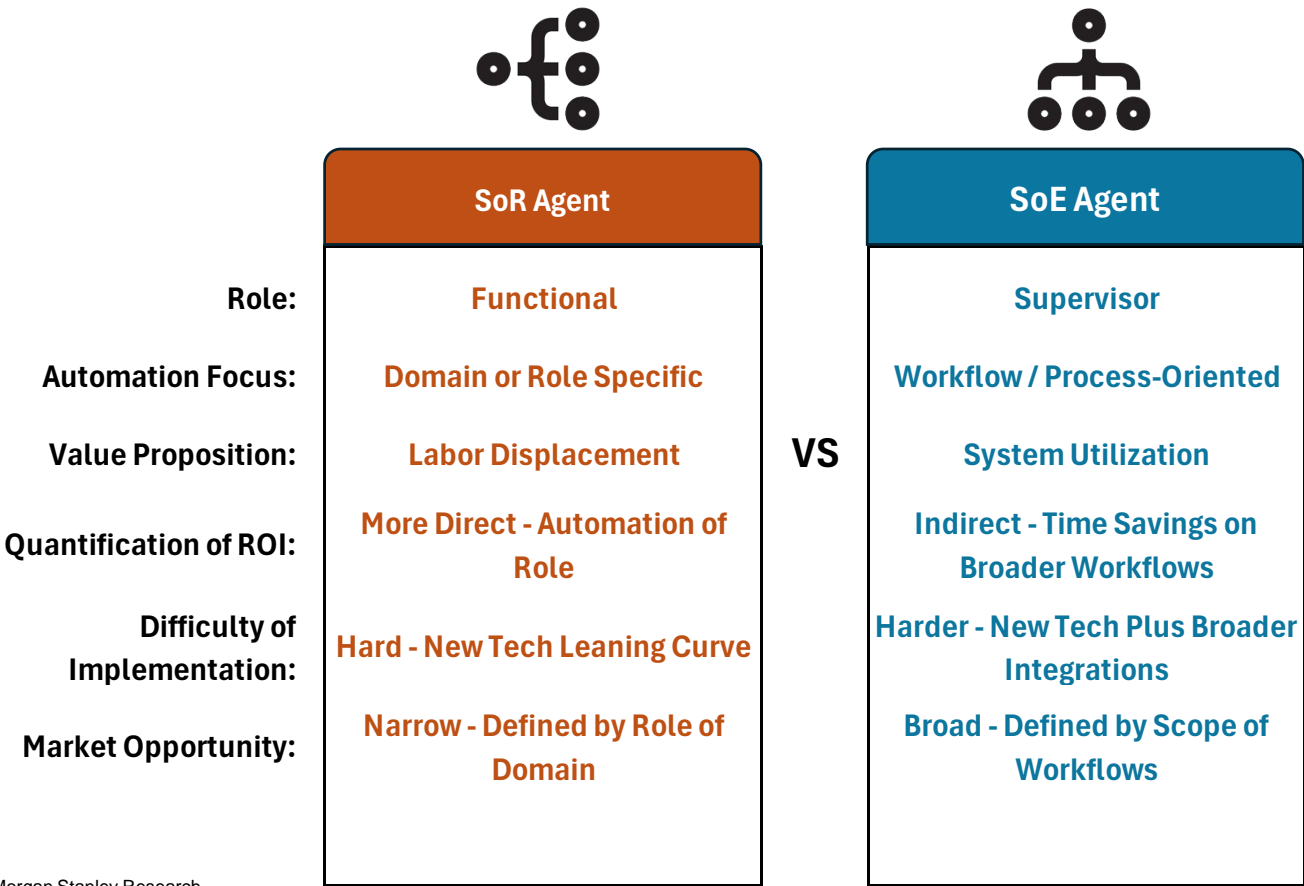
MORGAN STANLEY RESEARCH

Part III

Organizing the Agent Landscape

System of Record (SoR) vs System of Engagement (SoE) AI Agents

**SoR vs. SoE
Ai Agents** One framework for organizing the agent landscape is the distinction between System of Record style agents from System of Engagement style agents.



Source: Morgan Stanley Research

System of Record (SoR) Agents – Role Specific Automation

SoR Agents	System of Record style agents focus on automating functions related to specific domains such as coding, customer service, sales lead generation, business intelligence, etc. Examples include: GitHub Co-pilot, Workday, Salesforce Agentforce and HubSpot Prospecting agent
Agent Role	SoR agents take a functional role in terms of their automation objectives. This means that the agents operate within a particular business function or application-specific domain
Agent Focus	Focus of SoR agents is on specific role(s) within a particular domain such as customer service agents, software developers, business analysts, research assistants, sales development representatives etc.
Value Proposition	The value proposition for SoR agents centers on labor displacement or labor avoidance
Quantifiable ROI?	With a value proposition often tied to labor displacement/avoidance, directly ascribing cost savings should be easier with SoR style agents
Time to Market Impact	Given the ability to deliver hard, measurable ROI, the monetization timeline for SoR agents should be sooner and generally ahead of SoE agents

Source: Morgan Stanley Research

System of Engagement (SoE) Agents – Workflow / Process-Oriented Automation

SoE Agents	Unlike System of Record agents, System of Engagement agents focus on managing and automating workflow that span across different systems and applications. Examples include: ServiceNow Now Assist, GitLab Workflows and Atlassian Rovo
Agent Role	SoE agents take on a more supervisory / managerial role in that they manage a workflow across multiple applications and functional systems
Agent Focus	Compared to SoR agents that focus on automating particular roles within a given domain, SoE agents focus on automating workflows or an entire business process
Value Proposition	The value proposition for SoE agents is about enabling better utilization of enterprise resources as well as better utilization of the software vendor’s solution
Quantifiable ROI?	With a value proposition more associated with utilization than labor displacement, measuring ROI is generally more difficult for SoE style agents than SoR agents
Time to Market Impact	Given less quantifiable ROI, the monetization timeline for SoE agents should be later and generally after SoR style agents

Source: Morgan Stanley Research

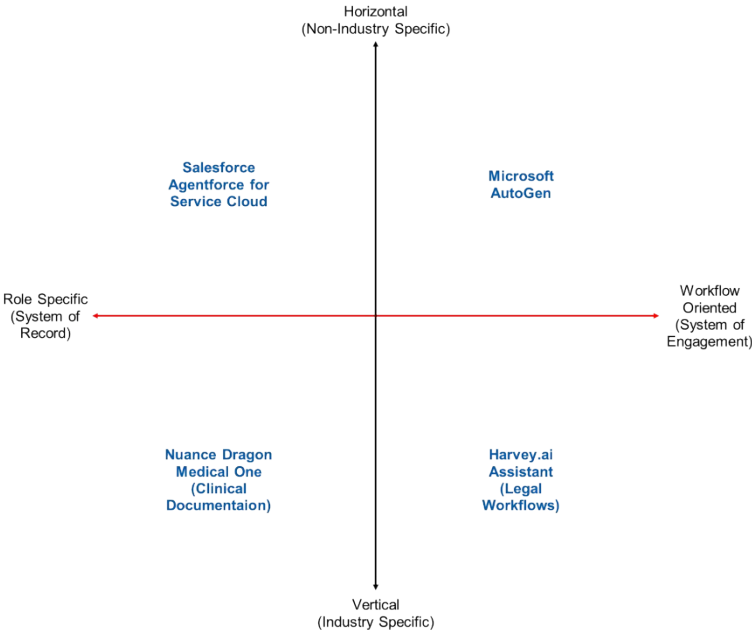
Organizing Across Two Vectors

Bottom Line

As investors look to understand the competitive landscape, we suggest organizing across two vectors, the first being Workflow versus Role Specific, inline with the industry discussion around 'Systems of Record' versus 'Systems of Engagement'.

System of Record vs. System of Engagement // Horizontal vs. Vertical

- **System of Record** agents take a functional role in terms of their automation objectives. This means that the agents operate within a particular business function or application-specific domain.
 - **Examples include:** GitHub Co-pilot, Workday, Salesforce Agentforce and HubSpot Prospecting agent
- **System of Engagement** agents take on a more supervisory / managerial role in that they manage a workflow across multiple applications and functional systems We expect an analogous dynamic within agents
 - **Examples include:** ServiceNow Now Assist, GitLab Workflows and Atlassian Rovo



Source: Company Data, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Organizing Across Two Vectors

Bottom Line

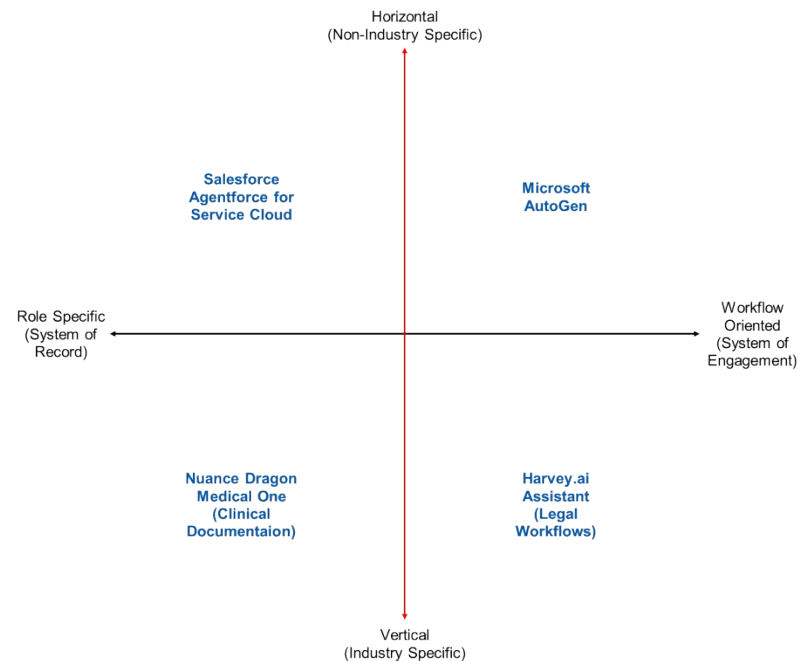
As investors look to understand the competitive landscape, we suggest organizing across two vectors, the second vector would be horizontal versus vertical oriented solutions

System of Record vs. Workflow Process // Horizontal vs. Vertical

- **Vertical-oriented** solutions look to automate process within one specific industry (i.e., Harvey.ai looks to automate legal workflows)
- **Horizontal-oriented** agents look to automate processes common across all industries (i.e., HR and software development).

These organizing principles likely presents **trade-offs** between time to market and overall size of the market opportunity.

- **Well defined, role-specific solutions from a transaction-oriented agent**
 - Should be easier to get up and running
 - Specifically targets cost savings/revenue opportunity for clear ROI
 - Strong vertical orientation likely furthers both these benefits.
- **Workflow-oriented agents, with more horizontal use case**
 - Target broader and bigger opportunities over time
 - Take more time to get built out and deployed.
- The Microsoft 365 Copilot solution (inc. agentic capabilities) targets the >400m current commercial Office users,
- The Agentforce solution for ServiceCloud serves 3 million Call Center representatives in the US (based on Bureau of Labor Statistics data) and perhaps 12-15 million globally.

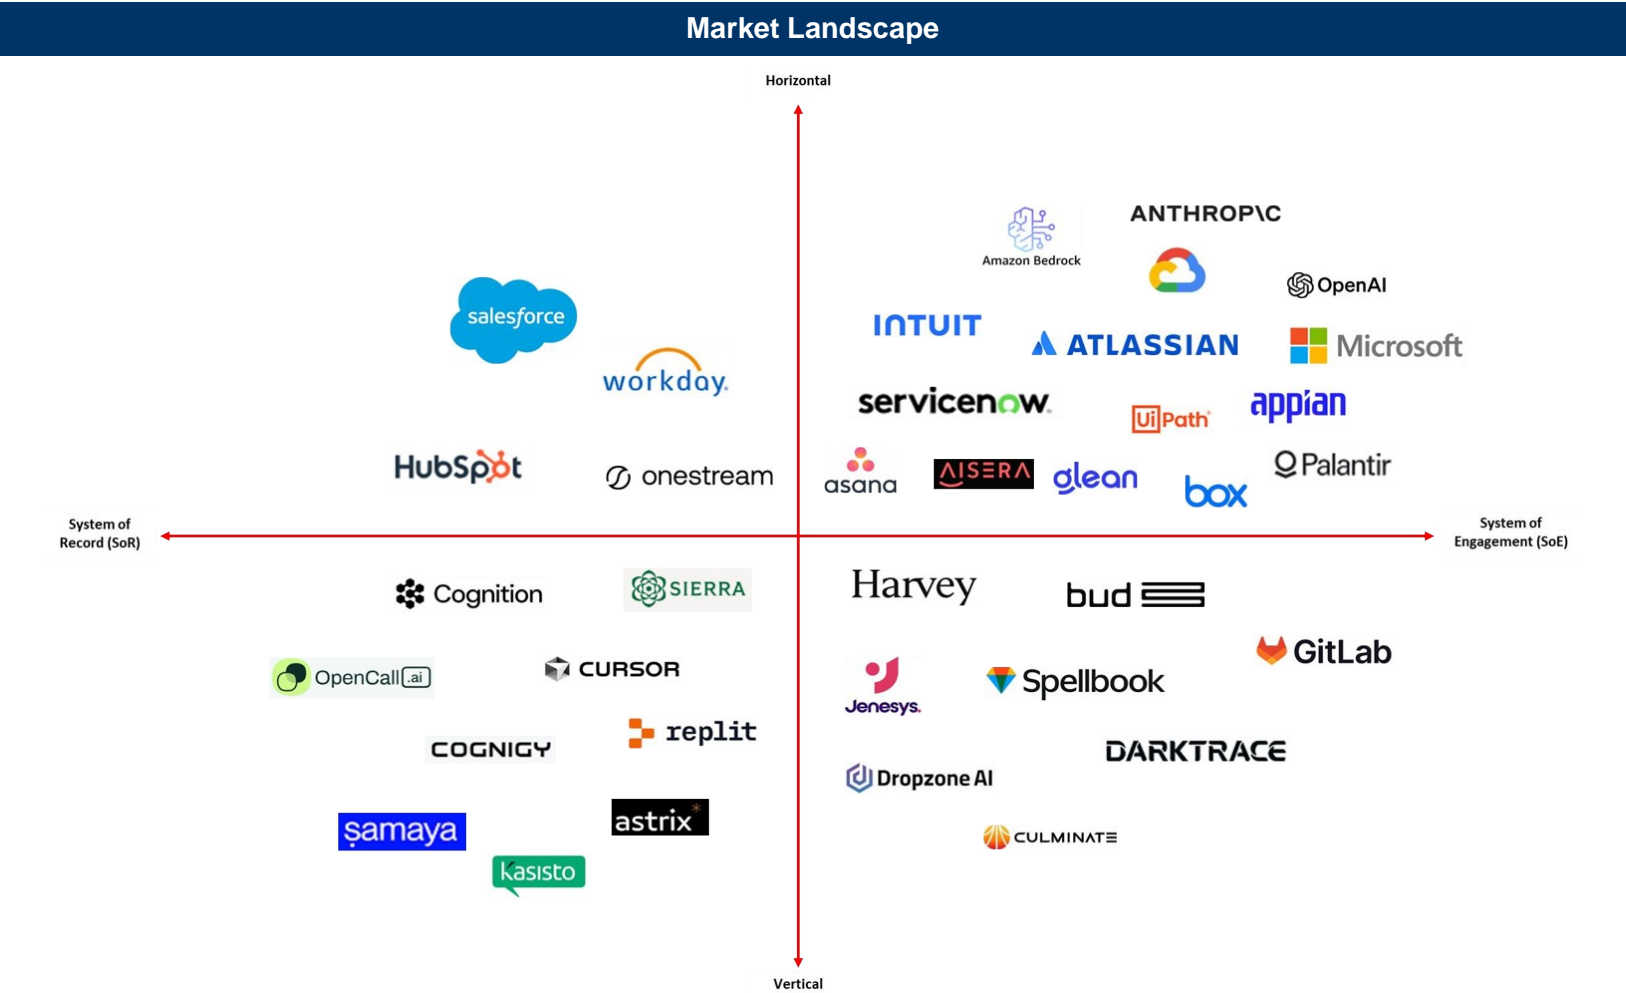


Source: Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Organizing Across Two Vectors

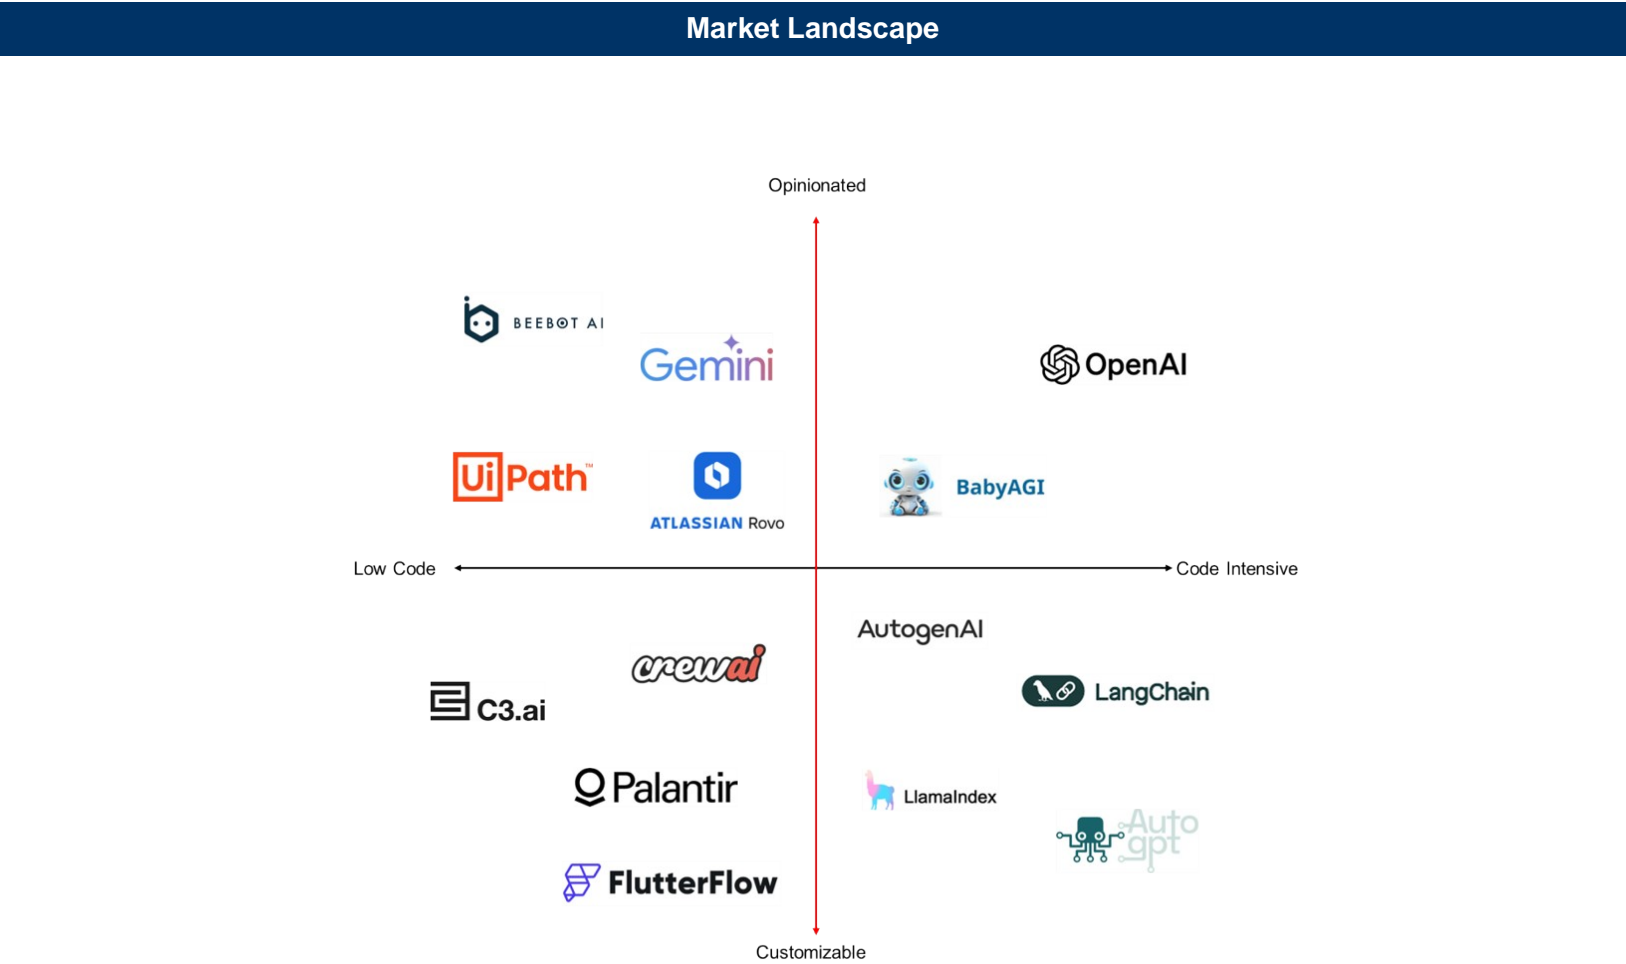


Source: Company Data, Morgan Stanley Research
Note: Positioning within quadrants not indicative of relative performance or competitive positioning

Morgan Stanley

MORGAN STANLEY RESEARCH

Organizing Across Two Vectors



Source: Company Data, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Classifying AI Agents By Reasoning Paradigms

Reasoning Paradigms

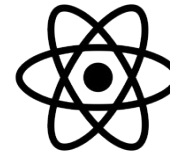
Another framework to organize AI Agents is by the different classes of reasoning paradigms they can employ. A few of the popular reasoning paradigms include: Chain of Thought, ReAct, Reflexion and Tree of Thought

Chain of Thought (CoT)



Prompts the agent to think in a linear, step by step by fashion by generating an internal reasoning chain before producing the final answer. Chain of Thought reasoning does not use any external tools

ReAct



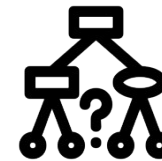
This style of prompting brings together reasoning in the form of thoughts with actions. Unlike CoT, ReAct agents are able to interact with tools and its environment to generate an answer

Reflexion



Extends ReAct style reasoning by adding two key elements: self-reflection and long-term memory. By adding these capabilities these agents can improve and learn over time with each successive run.

Tree of Thought (ToT)



Unlike the linear reasoning employed by ReAct agents, Tree of Thought agents search through a tree of possible reasoning steps and evaluate each path before choosing the optimal one

Chain of Thought Agents

What are CoT Agents?	A type of AI agent that employs linear, step-by-step reasoning (thoughts) before producing an output
Logic Structure	Thought 1 → Thought 2 → Thought 3 → Final Answer
Advantages	Simple and easy and easy to implement and good for math, logic puzzles and general reasoning tasks
Limitations	Does not interact with its environment, cannot access tools and does not make any observations during reasoning

Standard Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought Prompting

Model Input


Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅

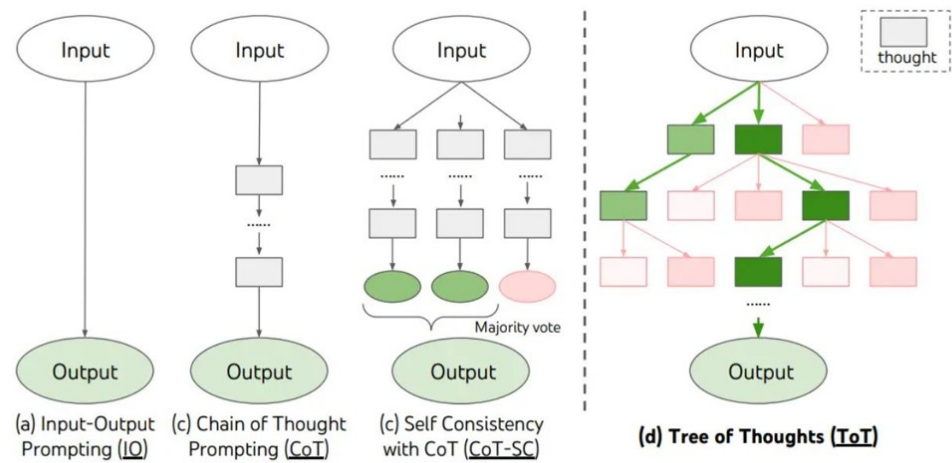


更多调研关注公众号: 调研纪要更新
添加微信: jiyaox

Source: Prompt Engineering Guide (<https://www.promptingguide.ai/>)

Tree of Thought Agents

What are ToT Agents?	AI agents that decompose tasks into a hierarchical structure of potential actions or routes, enabling the agent to assess possible solutions and enhance decision-making by discarding less viable paths at each stage before reaching the ultimate solution.
Logic Structure	Agent generates multiple solution paths → Forms decision tree → Evaluates paths → Selects the Optimal Path at each step
Advantages	Allows for planning and a structured evaluation of different potential solutions while eliminating unpromising solutions at each step. Useful in tasks that require complex decision making and deliberate planning
Limitations	Do not interact with the environment, complex to implement, require capable evaluators and is computationally expensive



Source: Tree of Thoughts: Deliberate Problem Solving with Large Language Models; Shunyu Yao et al (2023)

ReAct Agents

What are React Agents?	AI agents that integrate reasoning with action are defined by their capacity to engage with the environment through tools, their capability to make observations, and their ability to further reason based on those observations.
Logic Structure	Thought → Take an Action → Observe → Repeat
Advantages	Can interact with the environment and access various tools, adapts based on feedback and good at use cases involving interactive tasks such as coding, Q&A augmented by retrieval
Limitations	Does not have persistent memory across interactions and does not learn from past experience

Question: Aside from the Apple Remote, what other device can control the program Apple Remote was originally designed to interact with?

Answer: iPod

Thought: Let's think step by step. Apple Remote was originally designed to interact with Apple TV. Apple TV can be controlled by iPhone, iPad, and iPod Touch. So the answer is iPhone, iPad, and iPod Touch.

Answer: iPhone, iPad, iPod Touch

Standard

Reason only

Act only

ReAct

Thought 1: I need to search Apple Remote and find the program it was originally designed to interact with.

Act 1: Search[Apple Remote]

Obs 1: The Apple Remote is a remote control introduced in October 2005 by Apple ... originally designed to control the Front Row media center program ...

Thought 2: Apple Remote was originally designed to control the Front Row media center program. I need to search Front Row next and find what other device can control it.

Act 2: Search[Front Row]

Obs 2: Could not find [Front Row]. Similar: ['Front Row Seat to Earth', 'Front Row Motorsports', 'Front Row (software)', ...]

Thought 3: Front Row is not found. I need to search Front Row (software) .

Act 3: Search[Front Row (software)]

Obs 3: Front Row is a discontinued media center software ...

Thought 4: Front Row (software) is controlled by an Apple Remote or the keyboard function keys. So the answer is keyboard function keys.

Act 4: Finish[keyboard function keys]

Source: Prompt Engineering Guide (<https://www.promptingguide.ai/>)

Reflexion Agents

What are Reflexion Agents?	AI agents that resemble ReAct agents but incorporate self-reflection and long-term memory. Following each trial, the agent evaluates its successes and failures, retains feedback, and applies this knowledge to enhance future performance.
Logic Structure	(Thought → Act → Observe) → (Self-reflect → Update Long-term Memory) → Future Trial
Advantages	Can interact with the environment and access various tools, adapts based on feedback and good at use cases involving interactive tasks such as coding, Q&A augmented by retrieval
Limitations	Does not have persistent memory across interactions and does not learn from past experience

Question: Aside from the Apple Remote, what other device can control the program Apple Remote was originally designed to interact with?

Answer: iPod

Thought: Let's think step by step. Apple Remote was originally designed to interact with Apple TV. Apple TV can be controlled by iPhone, iPad, and iPod Touch. So the answer is iPhone, iPad, and iPod Touch.

Answer: iPhone, iPad, iPod Touch

Standard

Reason only

Act only

Act 1: Search[Apple Remote]

Obs 1: The Apple Remote is a remote control introduced...

Act 2: Search[Front Row]

Obs 2: Could not find [Front Row]. Similar: ...

Act 3: Search[Front Row (software)]

Obs 3: Front Row is a discontinued media center software ...

Act 4: Finish[yes]

ReAct

Thought 1: I need to search Apple Remote and find the program it was originally designed to interact with.

Act 1: Search[Apple Remote]

Obs 1: The Apple Remote is a remote control introduced in October 2005 by Apple ... originally designed to control the Front Row media center program ...

Thought 2: Apple Remote was originally designed to control the Front Row media center program. I need to search Front Row next and find what other device can control it.

Act 2: Search[Front Row]

Obs 2: Could not find [Front Row]. Similar: ['Front Row Seat to Earth', 'Front Row Motorsports', 'Front Row (software)', ...]

Thought 3: Front Row is not found. I need to search Front Row (software) .

Act 3: Search[Front Row (software)]

Obs 3: Front Row is a discontinued media center software ...

Thought 4: Front Row (software) is controlled by an Apple Remote or the keyboard function keys. So the answer is keyboard function keys.

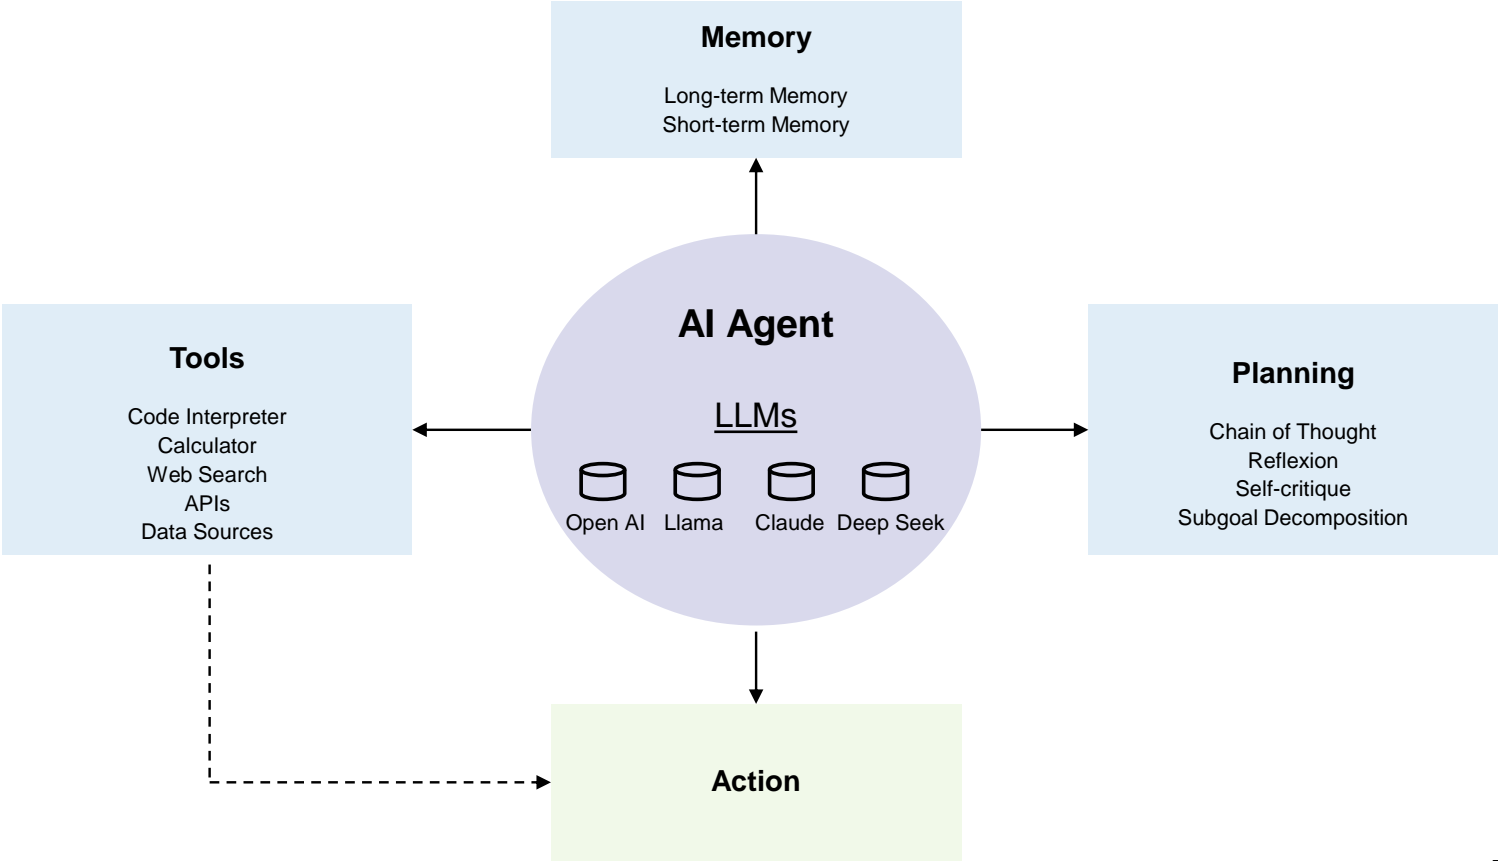
Act 4: Finish[keyboard function keys]

Source: Prompt Engineering Guide (<https://www.promptingguide.ai/>)

The Emerging AI Agent OS

Key Elements of the Agent OS

The emerging AI Agent OS consists of a few core components: 1) the LLM, 2) Tools, 3) Access to Memory and 4) the ability to plan



Source: Lilian Weng's Lil'Log (<https://lilianweng.github.io/>)

Breaking Down the Agent OS

The Four Core Components Of Agent OS	
Foundational Models	Every agent is powered by an AI or Large Language Model to take instructions, call tools, apply reasoning, execute tasks, structure response and deliver outputs
Tools	Define the set of actions an AI agent can take and can include code interpreters, calculators, calendars, web search, enterprise search, public/private APIs and databases/data sources
Memory	Refers to the process used to acquire and store and retrieve information. There are two major categories of memory: 1) short-term and 2) long-term. Short-term memory (aka working memory) is enabled by approaches such as retrieval augmented generation (RAG) to provide the agent necessary information and context to execute the task at hand. Long-term memory refers to the ability of the Agent to recall information over extended episodes usually by leveraging a vector/operational data store
Planning	Refers to the ability of the agent to decompose the problem into subgoals to enable efficient execution of tasks. Furthermore, planning includes the ability to reflect and refine behavior through self-criticism, self-reflection and learning from mistakes

Source: Lilian Weng's Lil'Log (<https://lilianweng.github.io/>)

Morgan Stanley

MORGAN STANLEY RESEARCH

Multi-agent Frameworks

What Is an Agent Framework?

Multi-agent AI frameworks is software that simplifies the process of building systems by providing high levels of abstraction including ready-made components to handle tasks such as orchestration, memory management, tool usage and reasoning

Agent Frameworks Showing Significant Signs of Market Traction Today



A python-based, open source multi-agent framework that is designed to orchestrate agents in a collaborative environment. Agents are defined by their roles and collaborate intelligently to achieve common objectives



Created by LangChain, LangGraph is an open source framework that allows developers to build complex agentic applications using a graph-based architecture to model and manage the relationships involved in an AI workflow



An open source programming framework developed by Microsoft, Autogen allows developers to build agents and enables agents to work together to accomplish tasks by replicating the structure of human teams and organizations

Source: LangGraph, CrewAI, Microsoft websites

Morgan Stanley

MORGAN STANLEY RESEARCH

Agent Protocols

**What are
Agent
Protocols?**

Protocols enable uniform standards for AI models as well as agents to connect to external services and for fostering agent to agent communication and collaboration

Leading Agent Protocols

MCP

Developed by Anthropic, Model Context Protocol is a protocol that provides a standard interface for LLMs to connect to external tools, services and data such as databases, APIs and files. With MCP, developers can forego custom and repetitive logic to connect each LLM to external services.

A2A

The Agent-to-Agent protocol is a cloud-based standard developed by Google that enables AI agents to communicate with each other across different systems. This allows agents from different vendors to collaborate and coordinate action to execute a given workflow

ACP

An open agent orchestration standard developed by IBM & BeeAI, Agent Communication Protocol enables structured communication, discovery and coordination between AI agents. ACP differs from A2A in that it focuses on local runtime environment as opposed to being cloud-based

Morgan Stanley

MORGAN STANLEY RESEARCH

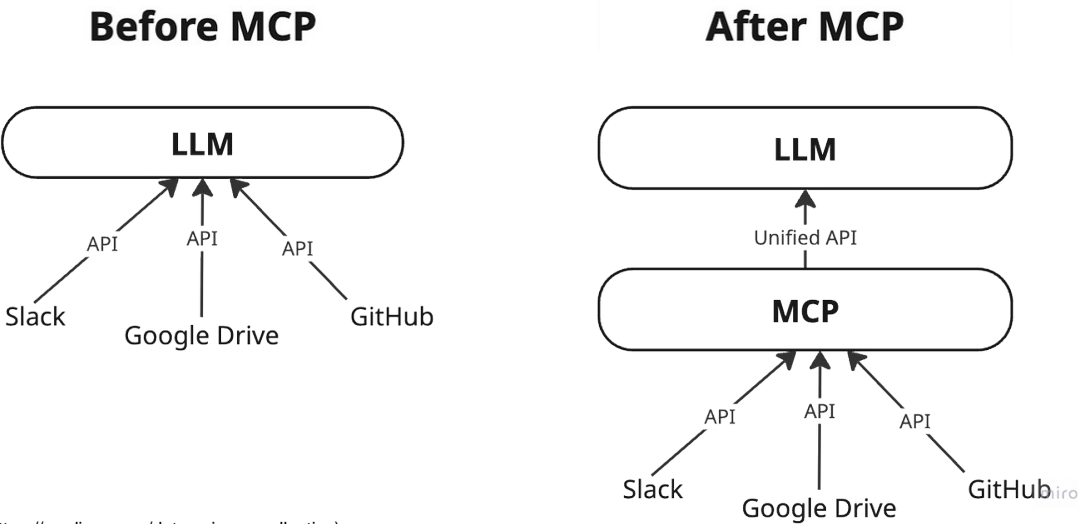
Comparing and Contrasting The Major Agent Protocols

	Model Context Protocol (MCP)	Agent Communication Protocol (ACP)	Agent-to-Agent Protocol (A2A)
Developed by	Anthropic	IBM & BeeAI	Google
Primary focus	Enabling LLMs to get real-time context by creating a standard interface to tools, data & APIs	Enabling local communication, coordination and orchestration between agents	Enabling cross-platform agent communication, collaboration and orchestration
Architecture	Client-server (host/server model)	Decentralized, local runtime	HTTP-based client/server with Agent Cards
Scope	Vertical integration: connection to tools/services to AI models	Agents running locally (say on a desktop or single server)	Horizontal integration: agent-to-agent communication & coordination
Ecosystem Support	Wide acceptance across the ecosystem including Microsoft, Open AI, Google, AWS, Cloudflare, Paypal, MongoDB, Atlassian and Glean	Linux foundation, BeeAI and IBM Research	Over 50 companies have announced support including Atlassian, Salesforce, SAP, MongoDB, LangCahin, Elastic, Datadog and Confluent
Security Model	App-layer auth, OAuth2, scoped APIs	Runtime sandboxing, private network security	OAuth2, scoped endpoint exposure
Best for	LLM apps with external data/tool needs	Edge AI, embedded systems, offline agents	Multi-agent workflows across platforms in the cloud
Use Case	Connecting an LLM to internal APIs	On-device coordination of multiple small agents	Distributed enterprise agents collaborating

Source: Edwin Lisowski Medium Blog (<https://medium.com/@elisowski>)

MCP: Unlocking a New Wave of Agentic Application Development

What’s the Problem?	For AI agents to be effective, they need to take action in real-time. This means getting relevant context, connecting to data and accessing a suite of tools. Prior to MCP, developers would have to build custom integrations to connect to the necessary tools and services for each LLM they were running
Solution	MCP provides a universal standard for LLMs to connect apps, data and services without the need for coding
An Analogy	Just as personal computers utilize USB as a standard way to connect to a host of services (storage, chargers, headphones, monitors), AI models use the MCP protocol as a standard way to extend their knowledge and capabilities
Implication	The emergence of MCP makes it easier to build AI agents that are useful and capable in real world environments helping to usher in new a wave of agentic applications



Source: Data Science Collective (<https://medium.com/data-science-collective>)

Morgan Stanley

MORGAN STANLEY RESEARCH

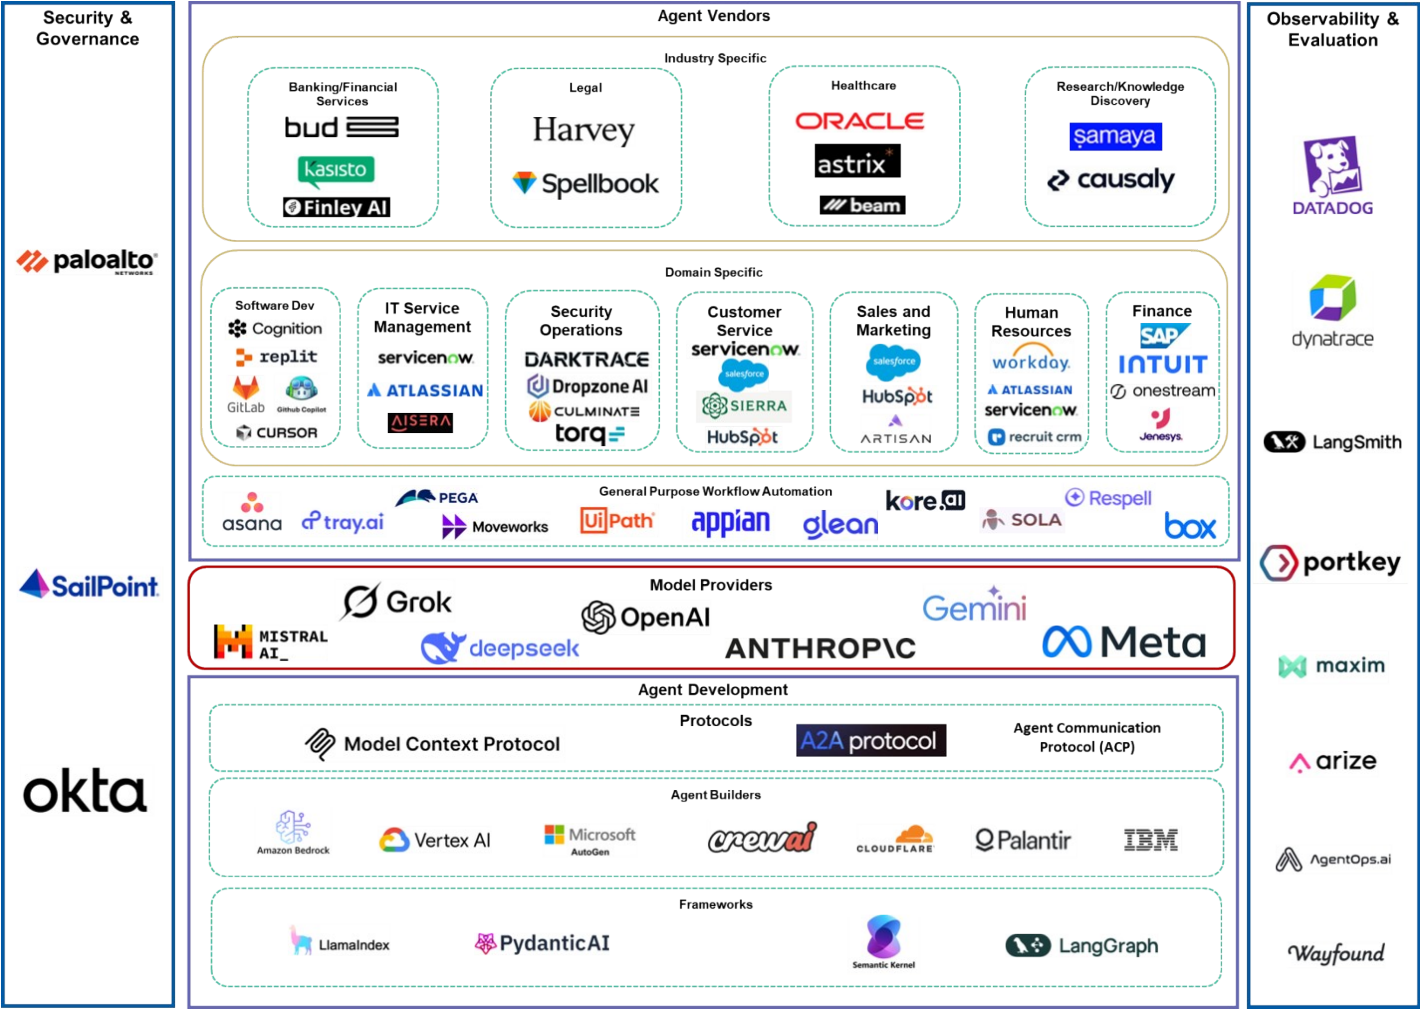
Part IV

Market Landscape

Morgan Stanley

MORGAN STANLEY RESEARCH

Current Market Landscape



Morgan Stanley

MORGAN STANLEY RESEARCH

Most Useful Domains



Source: Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

Part V

Company Spotlights

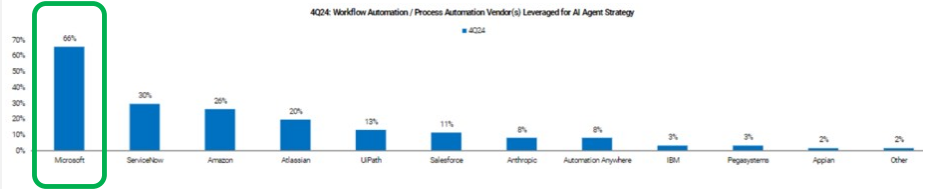
Hyperscalers – Microsoft

Microsoft Agents

- In March 2023, Microsoft announced Microsoft 365 Copilot building the foundation for its agents, which plug into existing business applications, security tools, and data management.
- Autonomous agents focused on key business functions (sales, customer support, security, etc.) entered a paid preview program on February 1, 2025, and pricing is based on the complexity of responses and actions generated
- **Microsoft's agents are priced based on consumption, which is tiered by the 'level of autonomy'**
 - E.g. A hypothetical agent might answer questions from customers on a website. Yesterday, it consumed 500 classic answers and 2,000 generative answers. Therefore, it would cost 4,500 messages, equivalent to \$45 for that day.
 - A hypothetical agent in Copilot Chat uses data stored in Microsoft Graph to answer employee questions about HR policies. Yesterday, the agent consumed 200 generative answers and 200 tenant Graph grounding for messages. Therefore, it would cost 6,400 messages or \$64 for that day.
 - A hypothetical autonomous agent responds to and routes inbound sales orders from customers. Yesterday, it consumed 100 generative answers, 100 tenant Graph grounding for messages, and 800 autonomous actions. Therefore, it would cost 23,200 messages or \$232 for that day.
- The consumption will be charged back to customers as MACC credits (Microsoft Azure Consumption Commitment).
- Microsoft benefits from its comprehensive AI portfolio and its breadth of Agent specific offerings, including: GitHub Copilot's asynchronous coding agent; Azure AI Foundry Agent Service for orchestrating multiple specialized agents, Microsoft 365 Copilot Studio / Agent Build, Agents for Microsoft Teams, M365 Agents toolkit for Visual Studio, M365 Agents SDK, Data Agents in Microsoft Fabric, SRE agents (among others)

		Microsoft 365 Copilot Chat	Microsoft 365 Copilot
		Free + Consumption	\$30 p/upm
Chat	Copilot Chat – Web grounded (powered by GPT-4o)	Included	Included
	Copilot Chat – Work grounded (work data in your tenant's Microsoft Graph and 3rd party data via Graph connectors)	Included	Included
	Copilot Pages	Included	Included
	File upload ¹	Included	Included
	Code Interpreter ²	Included	Included
	Image generation ³	Included	Included
Agents ²	Create agents using Copilot Studio ³ , including SharePoint agents	Included	Included
	Discover and pin agents	Included	Included
	Use agents grounded in Web data	Included	Included
	Use agents grounded in work data (work data in your tenant's Microsoft Graph and 3rd party data via Graph connectors)	Included – Metered	Included – Metered
	Use agents that act independently using autonomous actions	Included – Metered	Included – Metered
Personal assistant	Copilot reasons over personal work data (e.g. Outlook, OneDrive, meeting transcripts)	Included	Included
	Copilot in Teams	Included	Included
	Copilot in Outlook	Included	Included
	Copilot in Word	Included	Included
	Copilot in Excel	Included	Included
	Copilot in PowerPoint	Included	Included
	Copilot Actions	Included	In preview
Copilot Control System	Pre-built M365 agents (Interpreter, Facilitator, Project Manager, Employee Self-Service)	Included	In preview
	Enterprise Data Protection (EDP)	Included	Included
	IT management controls	Included	Included
	Agent management	Included	Included
	SharePoint Advanced Management	Included	Included
		Copilot Analytics to measure usage and adoption ⁴	Included
		Pre-built reports and advanced analytics to measure ROI	Included

1. Limits apply. 2. Applies to employee-facing agents only. 3. Learn more about the full capabilities of Copilot Studio aka.ms/CopilotStudioCapabilities. 4. Basic reporting in Microsoft Admin Center available for Copilot Chat.



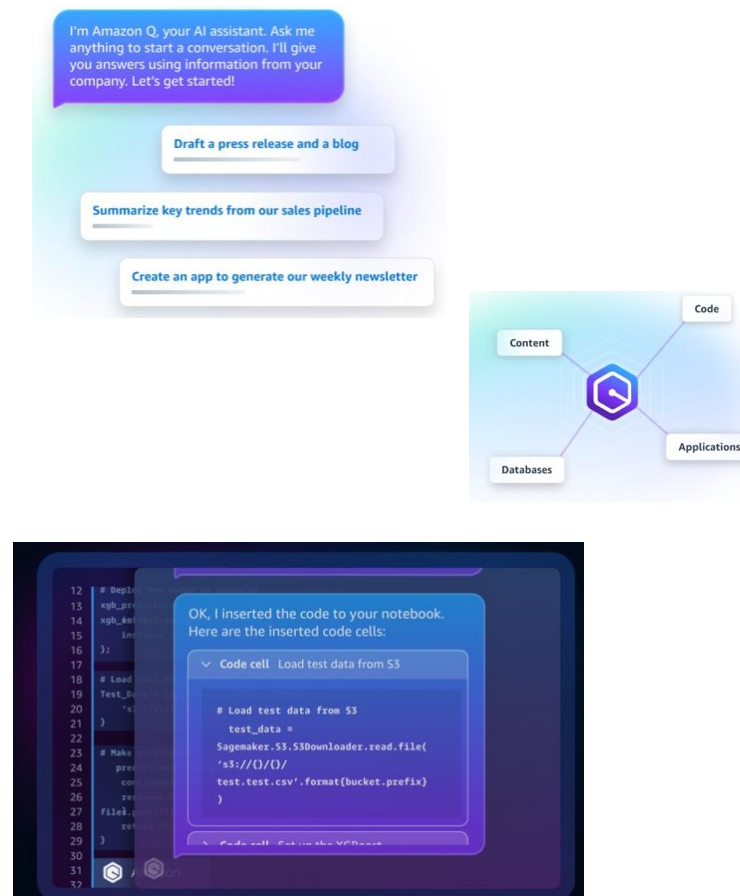
Morgan Stanley

MORGAN STANLEY RESEARCH

Hyperscalers – Amazon

Amazon Q

- An all-in one AI assistant, AMZN's Q spans across the business, integrating with a user's AWS account, business intelligence tools, contact center and business apps. **Q spans across 40 popular data sources including S3, Salesforce, Google Drive, Gmail, Microsoft 365 and others.** After connecting these data sources, Q indexes all the information, understanding core concepts and organizational idiosyncrasies, allowing Q to provide comprehensive and personalized answers specific to each user.
 - Agentic Capabilities:** Q's agentic capabilities also extend into AWS services like Connect and QuickSight, enabling services like real-time agent assist for contact centers or creation of executive summaries of key business data for analysts. In 1Q25 Q launched a new agentic coding experience that can execute complex workloads autonomously.
 - Pricing:** \$20 and \$25 a month, respectively, for the business and builder tiers.
- Amazon Q Business:** Users can leverage their company's content, data and systems by using Q to perform a series of tasks including; search across company systems and data, app creation to accelerate automation, actions in other applications (Jira, Salesforce, PagerDuty, etc.).
- Amazon Q Developer:** Users can leverage Q developer for building, operating and transforming software, helping developers with coding, testing, deploying, trouble shooting and optimizing AWS resources. AMZN has seen up to 40% improvement in developer productivity for those who adopt Q Developer.
- AMZN developed Nova Act, a new AI model trained to perform actions within a web browser that allows developers to break down complex workflows into reliable atomic commands.



Source: Company Data, Morgan Stanley Research

66

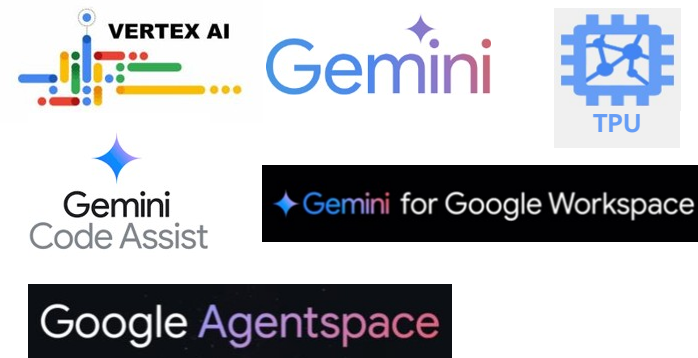
Morgan Stanley

MORGAN STANLEY RESEARCH

Hyperscalers – Google

Gemini

- GOOGL's Gemini LLM is deeply integrated into a variety of enterprise offerings, including Vertex AI, Google Workspace, Agentspace, and Gemini Code Assist
- Key Capabilities:**
 - Vertex AI Agent Builder**
 - Easier Agent Creation:** Enterprises can use an agent development kit (ADK) to build conversational or task-specific agents with no-code/low-code prompts that can connect to their data, API, and workflows securely
 - Agent2Agent (A2A) Protocol:** Universal communication standard that allows agents across different ecosystems to communicate with each other, allowing business agents to work together across different tasks
 - Gemini for Google Workspace**
 - Integration Across Drive:** Gemini can link to an enterprise's Google Docs, Sheets, Slides, etc. and can generate text, charts, emails, meeting notes etc.
 - Agentspace**
 - Centralized Hub for Agents:** Integrates with an enterprise's database and acts as a one-stop shop for employees to search across enterprise data, use/create agents, and generate video/image/text/code
 - Gemini Code Assist**
 - AI Coding Assistant:** Code generation, completion, debugging that can integrate with enterprise databases
- Vertex AI Agent Builder and Gemini Code Assist are available currently through Google Cloud; Gemini for Google Workspace is available in Google Workspace Business and Enterprise plans, and Agentspace is currently only available to select customers



Source: Company Data, Morgan Stanley Research

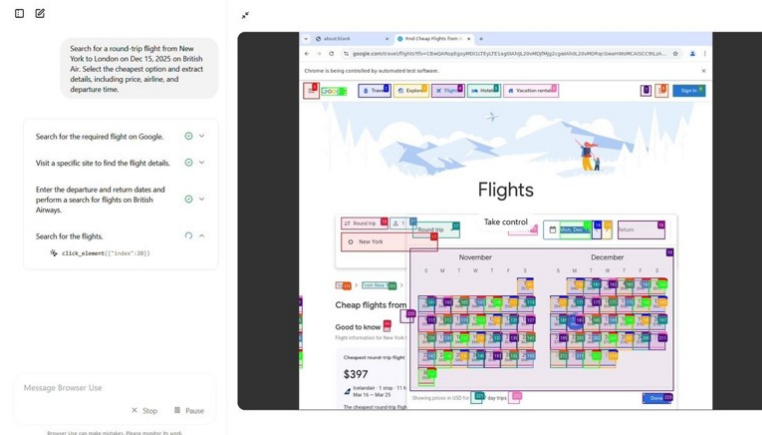
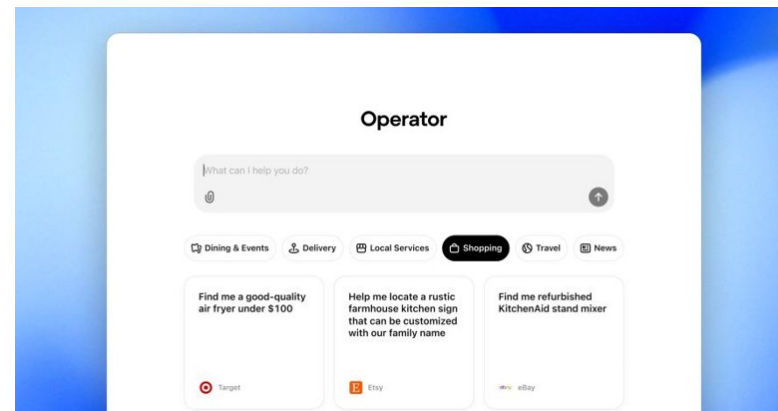
Morgan Stanley

MORGAN STANLEY RESEARCH

3rd Party Model Providers – OpenAI

Operator

- Launched in January 2025 as a research preview, Operator is an AI Agent that combines advanced vision capabilities with browser automation to perform web-based tasks independently.
- Operator is powered by a new model called Computer-Using Agent (CUA), which combines 1) GPT-4o's multimodal vision for analyzing webpage screenshots, 2) reinforcement learning for decision-making, and 3) virtual input simulation (mouse clicks, keyboard typing, scrolling).
- Operator interacts directly with graphical user interfaces (GUIs), mimicking human browsing behavior, instead of using APIs. This allows it to work across websites without requiring custom integrations.
- Key Capabilities:**
 - Task Automation:** Handles repetitive web activities with 38% success rate on OSWorld benchmarks (vs 72% human baseline). Examples include grocery ordering, flight booking, etc.
 - Workflow Personalization:** Users can 1) save recurring task sequences (i.e., weekly grocery orders), 2) set site-specific preferences (i.e., airline seat selection), 3) run multiple concurrent tasks through separate browser sessions.
 - Security Protocols:** Auto-switches to manual mode for logins/payments, blocks screenshots during sensitive interactions, and requires user confirmation before finalizing orders/transactions.
- Currently available only to U.S.-based ChatGPT Pro subscribers, which are paying \$200/month.
- Token Pricing:** GPT-4o charges \$2.5/1M input tokens and \$10/1M output tokens.



Source: Company Data, Morgan Stanley Research

3rd Party Model Providers – Anthropic

Computer Use

- Launched in October 2024 as a public beta, Computer Use is an agentic system that enables Claude 3.5 Sonnet models to interact with graphical user interfaces (GUIs) using mouse movements, keyboard inputs, and screen analysis – mimicking human-computer interactions.
- The agent combines vision-language models for screen interpretation with reinforcement learning for decision-making, allowing it to navigate software interfaces without API integration.

Key Capabilities

- Task Automation and Workflow Execution:** The agent autonomously executes multi-step digital tasks, including 1) web browsing, 2) application management, content creation, and administrative workflows.
 - Unlike traditional RPA, it adapts to unstructured interfaces and dynamic workflows without predefined scripts.
- Technical Implementation:** The system works through a four-stage loop, including 1) prompt analysis to interpret user instructions, 2) chooses between integrated capabilities/tools like terminals, text editors, or browser controls, 3) performs tasks within a sandboxed Docker container or VM, taking screenshots to verify outcomes, and 4) repeats steps until task completion or a 10-iteration limit.

- Token Pricing:** Claude 3.7 and 3.5 Sonnet charge \$3/1M input tokens and \$15/1M output tokens.

Pricing

As a subset of tool use requests, computer use requests are priced the same as any other Claude API request.

We also automatically include a special system prompt for the model, which enables computer use.

Model	Tool choice	System prompt token count
Claude 3.5 Sonnet (new)	auto	466 tokens
	any , tool	499 tokens
Claude 3.7 Sonnet	auto	466 tokens
	any , tool	499 tokens

In addition to the base tokens, the following additional input tokens are needed for the Anthropic-defined tools:

Tool	Additional input tokens
computer_20241022 (Claude 3.5 Sonnet)	683 tokens
computer_20250124 (Claude 3.7 Sonnet)	735 tokens
text_editor_20241022 (Claude 3.5 Sonnet)	700 tokens
text_editor_20250124 (Claude 3.7 Sonnet)	700 tokens
bash_20241022 (Claude 3.5 Sonnet)	245 tokens
bash_20250124 (Claude 3.7 Sonnet)	245 tokens

If you enable thinking with Claude 3.7 Sonnet, the tokens used for thinking will be counted against your `max_tokens` budget based on the `budget_tokens` you specify in the thinking parameter.

Source: Company Data, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH









SaaS Vendors – Atlassian Corp (TEAM)

Rovo

- First introduced in April 2024 and launched to GA in October 2024, Rovo is Atlassian's first AI-powered SKU designed as an enterprise search platform and virtual teammates (Agents) that can be called on or created by any team member to collaborate and automate tasks.
- Rovo operates across three fundamental dimensions: Find (Rovo Search), Learn (Rovo Chat), and Act (Rovo Agents).
 - Rovo Search:** Get relevant information from your Atlassian tools (Cloud and Data Center), 3rd party SaaS applications, and custom data sources. Use pre-built connectors that seamlessly integrate with Jira, Confluence, Loom, and Bitbucket, along with essential platforms like Google Drive, Microsoft SharePoint, Figma, Slack, Microsoft Teams, and GitHub.
 - Rovo Chat:** With deep understanding of your organization, Rovo offers personalized answers and provides expert advice. Answers questions about your team, projects, and company.
 - Rovo Agents:** *Out-of-the-Box Agents* can easily step into tasks like writing release notes, translating content into multiple languages, grooming your backlog and more. *Build Custom Agents* with no-code/low-code or the Forge development platform. Also explore a diverse range of Agents created by 3rd party partners on the Atlassian Marketplace.
 - Rovo for GitHub Copilot** provides instant code suggestions tailored to a team's design rules, Jira tasks, and more all within the IDE.
- Rovo's truly "Agentic" nature is evident** by its ability to execute complex tasks autonomously by leveraging AI models while integrating seamlessly into existing workflows. Its purpose-built agents can perform tasks with minimal human intervention, while learning and adapting in real-time.

Source: Company Data, Morgan Stanley Research

Out-of-the-box Agents

 <p>BY ATlassian</p> <p>Jira Theme Analyzer</p> <p>Build innovative products, faster. Analyzes and groups Jira issues into themes, reducing what used to take hours to seconds.</p> <p>Learn more →</p>	 <p>BY ATlassian</p> <p>Issue Organizer</p> <p>Speed up Jira backlog management. Auto-groups work into epics, creates new issues from Confluence, and prioritizes high-value tasks.</p> <p>Learn more →</p>	 <p>BY ATlassian</p> <p>Comms Crafter</p> <p>Write clear and engaging content. Creates and revises content that perfectly aligns with your company's brand voice and tone.</p> <p>Learn more →</p>	 <p>BY ATlassian</p> <p>Ops Guide</p> <p>Amp up your incident management. Helps your team efficiently handle on-call tasks and swiftly detect, respond to, and resolve incidents.</p> <p>Learn more →</p>
 <p>BY ATlassian</p> <p>Product Requirements Guide</p> <p>Craft the perfect PRD. Offers handy templates, expert tips, and best practices for writing concise documents.</p> <p>Learn more →</p>	 <p>BY ATlassian</p> <p>Global Translator</p> <p>Reach new audiences. Translates all your content into multiple languages to reach new users and geographies in less time.</p> <p>Learn more →</p>	 <p>BY ATlassian</p> <p>Release Notes Drafter</p> <p>Create release notes in a snap. Quickly generates release notes from multiple Jira issues, saving time and effort.</p> <p>Learn more →</p>	 <p>BY ATlassian</p> <p>Service Request Helper</p> <p>Power through support tickets. Speeds up issue resolution by summarizing problems, identifying support teammates, and assigning tickets.</p> <p>Learn more →</p>

Pricing

- Initially introduced with a \$20/user/month price tag
- Now access is free for Premium and Enterprise customers,
 - Coming to Standard soon.
 - \$5/month for non-Atlassian users
 - Consumption limit in place to control costs, with overconsumption triggering a bill – though not punitive
- Also included as part of new Teamwork Collection bundle along with Jira, Confluence, and Loom

/U

SaaS Vendors – ServiceNow, Inc (NOW)

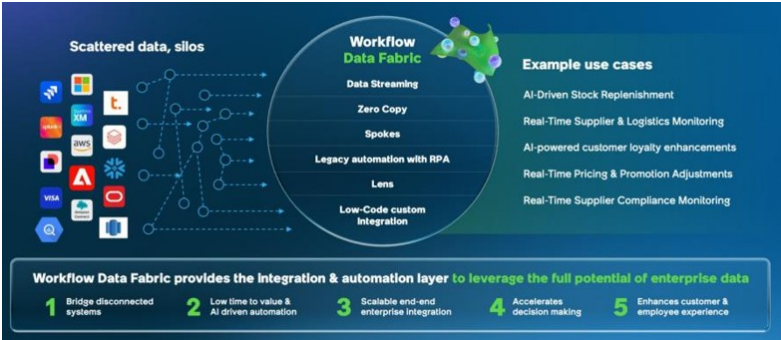
Pro Plus – AI Agents

- ServiceNow announced the general availability of its Yokohama release in early March, introducing AI Agent Studio, Workflow Data Fabric, and providing customers with teams of preconfigured AI agents, enabling users to automate workflows and improve productivity levels
- The company offers its solutions in three packages: 1) Standard, 2) Pro, and 3) Pro Plus, with access to AI Agents being offered within the Pro Plus package. The Pro Plus offering provides customers with an initial set of agentic capacity with potential to purchase / add further capacity should organizations exceed initial usage levels
- **AI Agent Studio:** Provides users the ability to efficiently design custom AI agents that integrate into their existing enterprise-wide workflows and data. The solution leverages a no-code language-based interface, enabling customers to build a team of AI agents
 - Examples of pre-configured AI agents include:
 - Security Operations Expert
 - Autonomous Change Management
 - Proactive Network Test and Repair
 - Operational Technology Knowledge Generation
 - ITSM Incident Categorization
 - Post-Incident Review Generation
- **Workflow Data Fabric:** A core element of ServiceNow's Agentic offering as it's powered by the company's automation engine and serves as an integrated data layer that unifies enterprise data to more efficiently process workflows and power AI agents with real-time, secure access to data from a variety of sources
- **Agentic AI is "The New Frontier" and ServiceNow is at the Core** as its preconfigured AI agents address a broad set of complex workflows while its data connectivity (via Workflow Data Fabric) provides unique differentiation vs peers as it provides AI agents incrementally more data to leverage thereby making them more impactful across enterprise employee workflows over time

Agentic Pricing Not Specified, Though Management Commentary Suggests a Value-Based Metering System

Agentic Commentary	
Investor Day	"We have well over a 1,000 customers using our AI platform in Agentic AI. Well over 250 million in ACV. It's the fastest growing product family in the history of ServiceNow, and it will continue to be that. And as Amit said, it's a hybrid pricing model. And what that means is, when we delivered Pro Plus on day one, there was a seat-based license like we've always had, but there was also a consumption based that we built into that licensing model. And that's super important as we move forward."
1Q25 Transcript	"Q1 was another substantial acceleration for ServiceNow AI. The number of Pro Plus deals more than quadrupled year-over-year, including 39 deals with three or more Now Assist products. Average ACV deal sizes grew by one-third quarter-over-quarter as Pro Plus products were included in 15 of our top 20 deals."
Morgan Stanley TMT Conference	"The more you use it and the more it accommodates interactions with multiple systems across your business or your network, the more we will make on the consumption-based pricing model."
4Q24 Transcript	"As the agents become increasingly productive, they will drive the consumption pricing meter. And that, of course, will be in addition to the seat-based licensing foundation. You'll get both." "If they (customers) go beyond the capacity and the Assist, which is the metered capacity, if they go beyond that, they're still buying Assist packs. So it's not like one assist per time in terms of consumption."

Workflow Data Fabric Serves as a Key Piece of ServiceNow's Agentic Architecture



Source: Company Data, Morgan Stanley Research. Note: Pre-configured agents are a small sample of others that are offered.

Morgan Stanley

MORGAN STANLEY RESEARCH

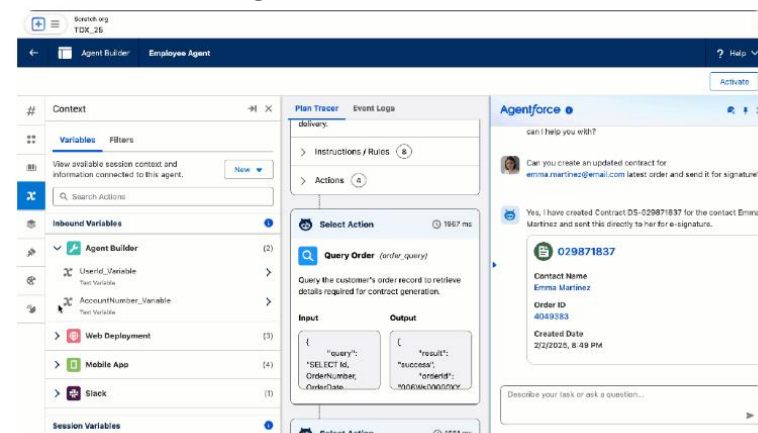
SaaS Vendors – Salesforce (CRM)

Agentforce

- Agentforce (General Availability in October 2024) is a suite of autonomous, out-of-the-box agents, allowing customers to automate administrative tasks while reorienting employee bandwidth to higher priority work.
 - Service agents:** automatically help a customer troubleshoot an issue with personalized, branded advice based on their purchase history and company knowledge documents
 - SDR agents:** send a personalized introductory email to a lead, respond to a question on pricing, and access an account executive's calendar to schedule a follow-up call
 - Marketing agents:** generate campaigns, briefs, audience segments, and emails based on your company's goals
 - Ecommerce agents:** manage websites and generate personalized promotions
- The latest version, **Agentforce 2dx** (announced March 2025) expands beyond reactive chat interfaces to enable proactive AI agents that work behind the scenes without constant human oversight, with capabilities including multimodal experiences and embedded AI in any workflow. **AgentExchange**, a trusted marketplace and community where partners can build and monetize agentic AI, will accelerate the deployment of AI agents.
- Strong Early Customer Traction:** As of F1Q26, Salesforce signed >4,000 paid Agentforce deals, up from ~3k the previous quarter, with Agentforce now crossing a \$100M ARR run rate since going GA.
- Pricing:** After initially being rolled out at a fixed price of **\$2 per call**, Salesforce recently added 'Flex Credits' to the Agentforce pricing model, a consumption-based model that ensures customers only pay for the actions Agentforce performs. Credits are available for purchase in packs of 100k at a \$500 price point— each action consumes 20 credits (or \$0.10). The pay per action model better connects price to value and allows customers to scale usage of Agentforce into additional workflows (beyond call centers) at a more favorable price.
- The combination of broad application functionality, deep data sets to power the AI models and an Agentic Computing architecture to take action frames CRM's strong positioning for GenAI.

Source: Company Data, Morgan Stanley Research

Agentforce Work Interface



Pricing

Flex Credits	Conversations	Standard Success Plan
Pay per action as you scale Agentforce.	Drive efficiency with digital labor 24/7.	Set your company up for success with self-guided, always-on resources.
\$500 USD/Per 100k Credits	\$2 USD/Per conversation	Included in all licenses
<ul style="list-style-type: none"> Out-of-the-box agent skills Agent Wizard Agent Builder Prompt Builder 	<ul style="list-style-type: none"> Out-of-the-box agent skills Agent Wizard Agent Builder Prompt Builder 	<ul style="list-style-type: none"> Access to Knowledge Articles, documentation, and Trailhead Global Trailblazer Community
Contact us	Contact us	Contact us

/2

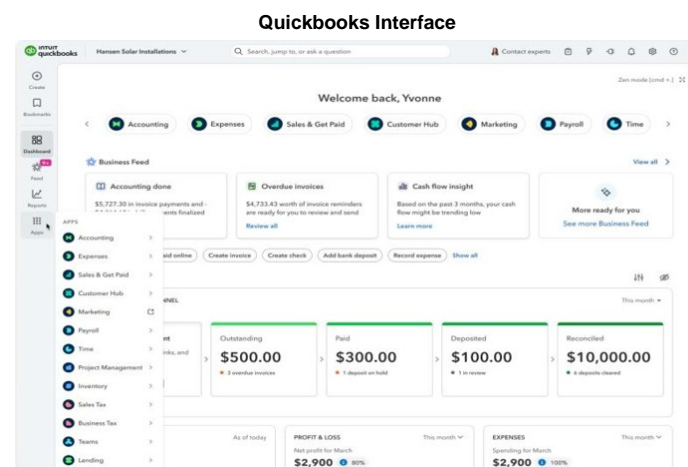
Morgan Stanley

MORGAN STANLEY RESEARCH

SaaS Vendors – Intuit (INTU)

Intuit AI Agents

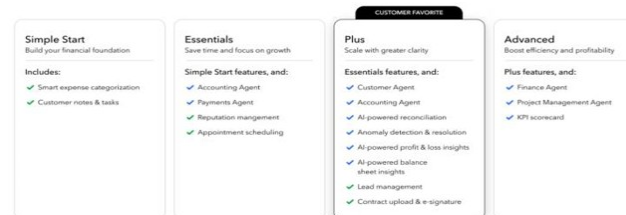
- During Intuit's 3Q FY25 Earnings call, CEO Sasan Goodarzi introduced their new suite of specialized AI Agents, focusing on automating tasks such as invoicing, payment tracking, transaction reconciliation, and project management. These agents are designed to act proactively, offering insights and executing tasks on behalf of users, streamlining operations for SMBs.
- Why is Intuit a Wildcard?** Given small businesses propensity to buy bundled solutions, we like Intuit's opportunity to upsell agentic capabilities into their broad install base (10M+ small and mid-market businesses as of Jan '25). That said, given the solution is still not generally available, and we have not seen it in the marketplace (pricing dynamic still unknown), we keep Intuit in the wildcard bucket of our Agentic Winners Framework.
 - Accounting Agent:** keeps books clean and accurate by automatically categorizing transactions, detecting and resolving anomalies, and reconciling accounts—flagging what needs approval or attention.
 - Payments Agent:** helps facilitate faster payments, and bills paid on time by monitoring cash flow and optimizing invoicing and payment collection.
 - Finance Agent:** analyzes financial data and automatically creates forecasts, highlights variances, and pinpoints the source to help you stay on track to your goals.
 - Customer Agent:** your clients can source leads from their inbox, draft personalized email responses, and track every customer opportunity in the sales cycle.
- Our Thoughts on Pricing:**
 - Pricing for Value:** Intuit estimates SMBs value their time at about \$75 per hour
 - Add-on module pricing:** plan to introduce AI capabilities as add-on modules in the coming months



Offerings

All in one. Get it done.

Amp up efficiency with AI agents and optional access to trusted experts.*



更多调研关注公众号: 调研纪要更新

添加微信: jiyaox

Morgan Stanley

MORGAN STANLEY RESEARCH

SaaS Vendors – Adobe (ADBE)

Adobe Experience Platform Agent Orchestrator & Purpose-built Agents

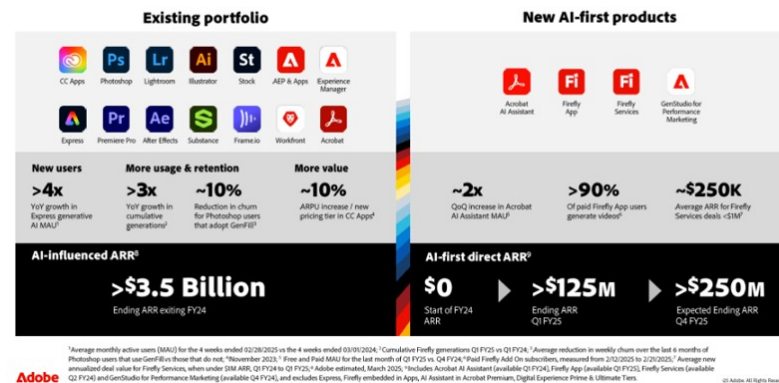
- At the 2025 Summit conference, Adobe unveiled the **Adobe Experience Platform Agent Orchestrator**, which enables businesses to build, manage, and orchestrate AI agents from Adobe and third-party ecosystems.
- Adobe released **10 purpose-built agents** with a deep understanding of customer data and ability to support multiagent collaboration and adaptive reasoning. These agents include: *Account Qualification Agent, Audience Agent, Content Production Agent*, Data Insights Agent, Data Engineering Agent, Experimentation Agent, Journey Agent, Product Advisor Agent, Site Optimization Agent*, Workflow Optimization Agent**
- Extended collaboration with Microsoft.**
 - Adobe Express Agent for Microsoft 365 Copilot** enables users to create high-quality images directly within Copilot, PowerPoint and Word.
 - Adobe Marketing Agent for Microsoft Copilot 365** empowers marketers to refine audience targeting, extract actionable insights and improve cross-team collaboration.
- Brand Concierge:** An agentic app that draws on a company's unique brand attributes and customer data to enhance brand loyalty by engaging customers through multimodal interactions across text, voice, and images.
- Pricing:** The company has not disclosed specific pricing for these agents. Their statements around general pricing strategy for AI SKUs at the Summit in March 2025 suggest they will likely follow their established pattern of tiered, value-based pricing approaches for these new AI capabilities.
- Adobe is capitalizing on its domain expertise across content, data and workflows to drive competitive advantage with GenAI. While monetization takes time, the product road-map and positive partner conversations provide comfort on Adobe's competitive moat.

*Included as part of the GenStudio.

Source: Company Data, Morgan Stanley Research

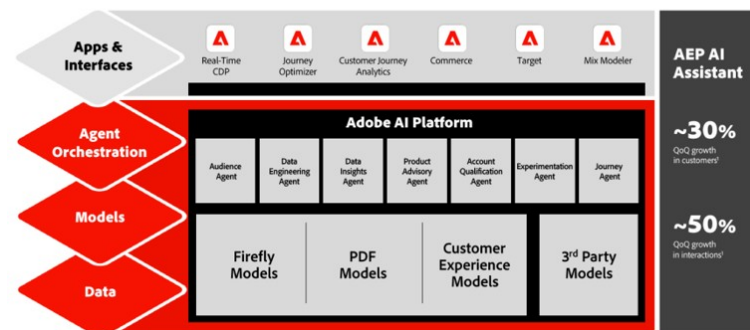
More than \$3.5B AI-Influenced Revenue

Adobe's AI monetization



AI Agents incorporated to the AEP

Adobe AI Platform powers AEP & Apps



14

Morgan Stanley

MORGAN STANLEY RESEARCH

SaaS Vendors – Workday (WDAY)

Illuminate

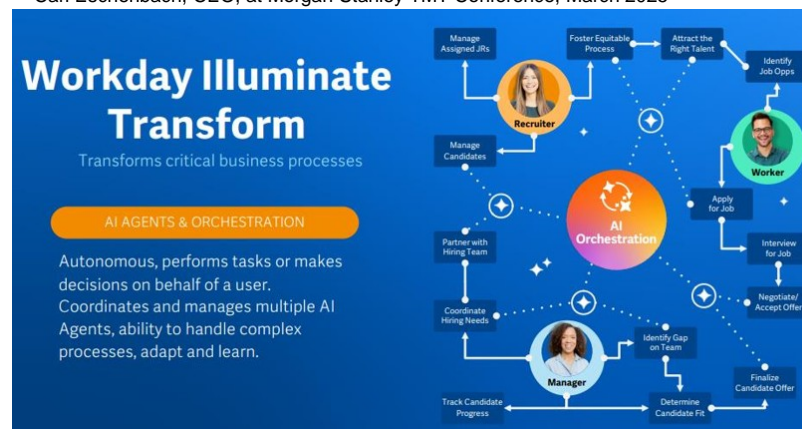
- At Workday's Rising conference in September 2024, Workday introduced **Illuminate**, which is their platform of AI products and models built on Workday's HR and finance datasets. Illuminate is focused on three key areas: 1) Accelerating tasks with GenAI, 2) Assisting users with Copilots like Workday Assistant, and 3) Transforming business processes through role-based agents.
- Workday has rolled out multiple role-based agents so far:
 - Recruiting Agent:** Provides AI-powered talent sourcing and screening capabilities, boosting recruiting solution selling price by 150%
 - Expenses Agent:** Provides users with flexible ways to submit and approve expenses, while giving leaders the ability to set controls and analyze spend
 - Succession Agent:** Assists HR teams with succession planning by identifying and developing potential candidates for key positions within the organization
 - Contracts Agent:** Continuously analyzes contracts across the enterprise, surfaces obligations and opportunities buried in unstructured data, and drives business actions to capture value and mitigate risk
 - Payroll Agent:** Identifies and updates invalid payroll data, automates audit workflows, surfaces insights, recommends fixes, monitors compliance, and delivers system updates
 - Financial Auditing Agent:** Increases efficiency during audits by connecting complex business documents to monitor transactions, reconcile balances, and review internal controls, allowing audit firms to develop apps that connect directly to their Workday customers
- Underlying these role-based agents is Workday's Agent System of Record, a centralized platform to manage an organization's entire fleet of AI agents (both from Workday and third parties). This system helps organizations maintain control, governance, and visibility while leveraging AI agents.

"We monetize it through one of three ways. Number one, our agents, whether it's our Recruiter Agent or – we have a new agent that we've announced that's really an internal mobility agent for employees to start to think about how they move around. We have the succession agent, we have a whole bunch of agents. We announced the policy agent. We announced a financial audit agent. So we're going to build our own agents. And that's a way for us to monetize the Workday platform.

At the same time as others out there start to build agents, we've developed something as part of the Agent System of Record an agent gateway. So when people want to get access and onboard into the enterprise, a new digital agent or digital employee, we're going to monetize this new gateway as people come through and get access to our platform. So it's another way for us to monetize it.

And then the third way we're monetizing it is into an agent or the agent gateway. But it's our products that include AI. An example is Extend Pro. Another example would be what we're doing with something like talent optimization, which is a SKU, it's not an agent, but it's 100% driven or through something like Evisort, which is a CLM solution that a 100% AI driven as well. So we have many different ways to monetize AI going forward and we think we're uniquely positioned."

- Carl Eschenbach, CEO, at Morgan Stanley TMT Conference, March 2025



Source: Company Data, Morgan Stanley Research

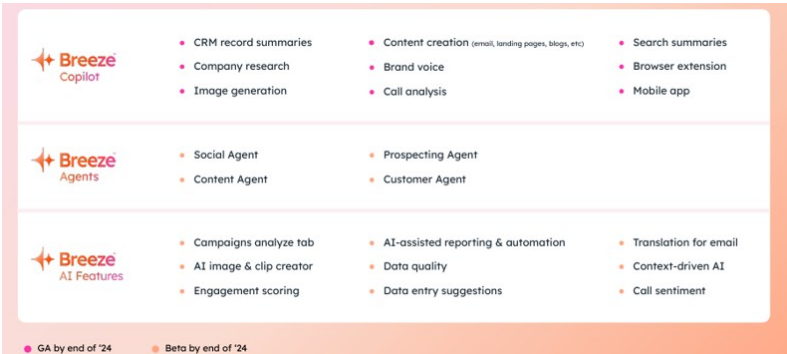
SaaS Vendors – HubSpot (HUBS)

Breeze

- HubSpot introduced their **Breeze platform**, which delivers AI functionality to customer-facing teams to boost productivity, scale growth, and unlock actionable insights, in September at 2024’s INBOUND Conference.
- HubSpot currently has **four Breeze Agents** rolled out in the market, all still in Beta aside from Customer Agent (went GA with credit-based pricing, for Pro and Enterprise tiers as of June 2nd), acting as AI-powered experts designed to automate workflows from planning through to execution:
 - **Content Agent:** Utilizes CRM data to streamline content creation for engaging and differentiated marketing campaigns. It helps marketers quickly produce different types of content that attract visitors and generate leads and drive conversion, including high-quality landing pages, podcasts, case studies, and blogs.
 - **Social Media Agent:** AI-powered expert that enhances your social media strategy, analyzing your social performance, company details, and marketing best practices to develop a tailored, multi-channel content strategy that amplifies your social presence.
 - **Prospecting Agent:** The ‘always on’ agent drives pipeline generation 24/7 by researching and executing personalized outreach strategies for ideal prospects within HubSpot’s Smart CRM.
 - **Customer Agent:** Engages with customers on websites in full conversation and at any time to provide instant answers– the agent utilizes customer-specific knowledge base articles, URLs, help sites, and blogs to deliver accurate and personalized responses to common customer issues.
- HubSpot is well positioned to address the agent opportunity given: **1) deep data sets** combining rich structured and unstructured data; **2) a unified platform approach** that provides context across data, engagement (hubs), and action (AI-powered execution); and **3) a vibrant agent ecosystem** developed through the company’s Agent.ai network

Source: Company Data, Morgan Stanley Research

AI Powers the HubSpot Customer Platform



Pricing

After driving consistently strong adoption/utilization with agents, HubSpot is moving to direct monetization as of June 2nd, with Customer Agent now available via credit-based pricing (in Pro/Enterprise tiers).

Highest level of HubSpot subscription	Included HubSpot Credits in subscription per month
<i>Starter</i>	500
<i>Professional</i>	3,000
<i>Enterprise</i>	5,000

Morgan Stanley

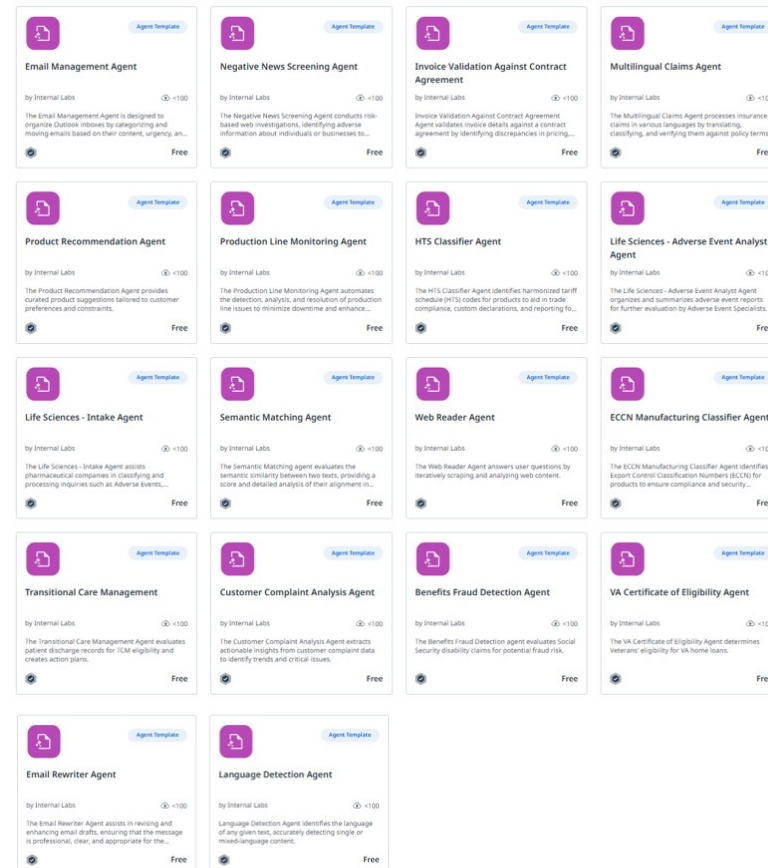
MORGAN STANLEY RESEARCH

SaaS Vendors – UiPath, Inc (PATH)

Agent Builder – UiPath Maestro (Orchestrator) – AutoPilot

- UiPath's product innovation remains healthy with AI-focused solutions including Agent Builder, Agentic Orchestration, and AutoPilot announced in 2024.
- The company has made further progress with recent announcements of UiPath Platform for Agentic Automation, and bi-directional integrations with Microsoft Copilot Studio and other 3rd party agents through UiPath Maestro.
- **AutoPilot:** Provides a set of AI-powered solutions that can automate and streamline a variety of general and department-specific tasks. Core capabilities include automation generation, self-service and assisted experiences, cross-application compatibility, data processing, contextual awareness, process optimization, and compliance. The functionality is available in the below UiPath products
 - **Developers:** Studio, StudioX, Apps, and Studio Web
 - **Testers:** Test Manager and Studio
 - **Everyone:** Assistant
- **Agent Builder:** Announced at the company's FORWARD user conference in 2024, Agent Builder enables users to create agents from scratch or from the company's pre-configured agents' catalog. The offering launched with "hundreds of customers" in private preview in December and management noted thousands of agents have been created generating >250,000 agent runs to date (as of 1Q26)
- **UiPath Maestro:** Combines RPA, AI, and human collaboration in a unified orchestration platform. Designed for complex, long-running enterprise workflows, Maestro enables organizations to model, execute, and optimize processes that span many systems (175 average in large enterprises) while reducing implementation timelines.

UiPath's Pre-Built Agentic Workflow Templates and Prompts



Source: Company Data, Morgan Stanley Research

SaaS Vendors – GitLab (GTLB)

GitLab Duo

- First introduced in May 2023 as part of the GitLab 16 launch and released to general availability in August 2024, GitLab Duo is a comprehensive suite of AI-powered tools integrated into GitLab's platform, designed to assist users throughout the software development lifecycle, from code creation and security assessment to collaboration and analytics.
- GitLab Duo distinguishes itself in the market in two core aspects
 - Goes beyond code assistance, **integrating across the platform** to enable greater automation across the SDLC.
 - **Operates with a privacy-first approach**, emphasizing that it does not use customers' proprietary code as training data. This is particularly important for enterprise customers concerned about IP protection.
- GitLab has established strategic partnerships to power its AI capabilities, leveraging Google's Vertex AI for tuning foundation models and Anthropic's Claude for specific features like code suggestions. **GitLab selects the most suitable LLM for each individual feature.**
- GitLab Duo is currently offered in two tiers, **Duo Pro** which is largely an AI code assistant and **Duo Enterprise** which provides summarization, security, and analytics capabilities.
- **GitLab Duo Workflows** – often described as an “agentic workflow system” represents the next evolution, moving beyond reactive, prompt-based AI assistants and functioning as an autonomous team member that actively contributes to software development processes.
 - Duo Workflows actively monitors projects, anticipates issues, automatically identifies and resolves vulnerabilities, optimizes applications for performance, and streamlines development environments.
 - *Currently in private beta stage.*



GitLab Duo Pricing		
Duo Pro	Duo Enterprise	
\$19 per user/month	Contact Sales	
Available for Premium and Ultimate customers	Available for Ultimate customers	
Features Include:	Everything from Duo Pro, plus:	
Organizational User Controls	Summarization and Templating tools	Advanced Troubleshooting
User permissions for AI capabilities	Discussion summary	Root Cause Analysis
Code Suggestions	Merge request summary	AI Analytics
Code generation	Code review summary	AI Impact and Productivity Reporting
Code completion	Security and Vulnerability tools	Personalized GitLab Duo
Available in many popular IDEs and supports >20 programming languages	Vulnerability explanation	Self-Hosted (optional)
Chat	Vulnerability resolution	Model Personalization (planned)
Code explanation	Advanced Chat with	
Test generation	Merge Request context	
Code refactoring	Issue and Epic context	

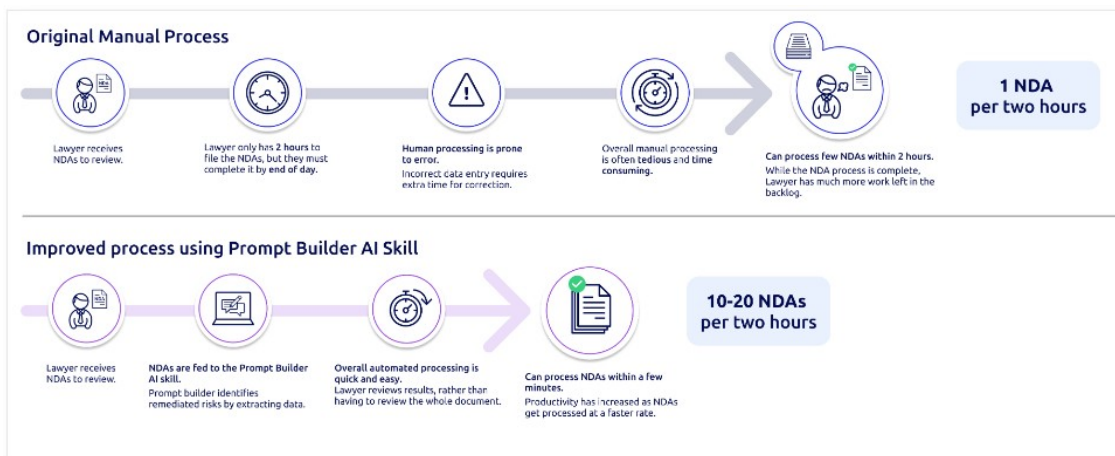
Morgan Stanley

MORGAN STANLEY RESEARCH

SaaS Vendors – Appian (APPN)

AI Skills

- Initially launched in May 2023 as part of the Appian 23.2 release, and significantly enhanced through its 24.2 release in 2024, Appian AI Skills represents the company's implementation of agentic AI within its low-code process automation platform. It exemplifies practical agentic AI, balancing purpose-built enterprise processes automation with human oversight.
- Through its collaboration with AWS Bedrock, Appian allows users to leverage the latest Gen to:
 - Incorporate AI agents into mission-critical processes.
 - Train custom AI models.
 - Build your own generative AI prompts in processes.
 - Leverage the latest models while maintaining privacy.
 - Automate common use cases like content generation and processing, PII extraction, and data extraction.



Pricing

Pricing is based on a token limit system within Appian's subscription tiers:

- Standard tier: 100M tokens/month
- Advanced tier: 200M tokens/month
- Premium tier: 500M tokens/month

The exact cost per token is not publicly disclosed.

Appian's plan pricing is customized and requires contacting sales for a quote.

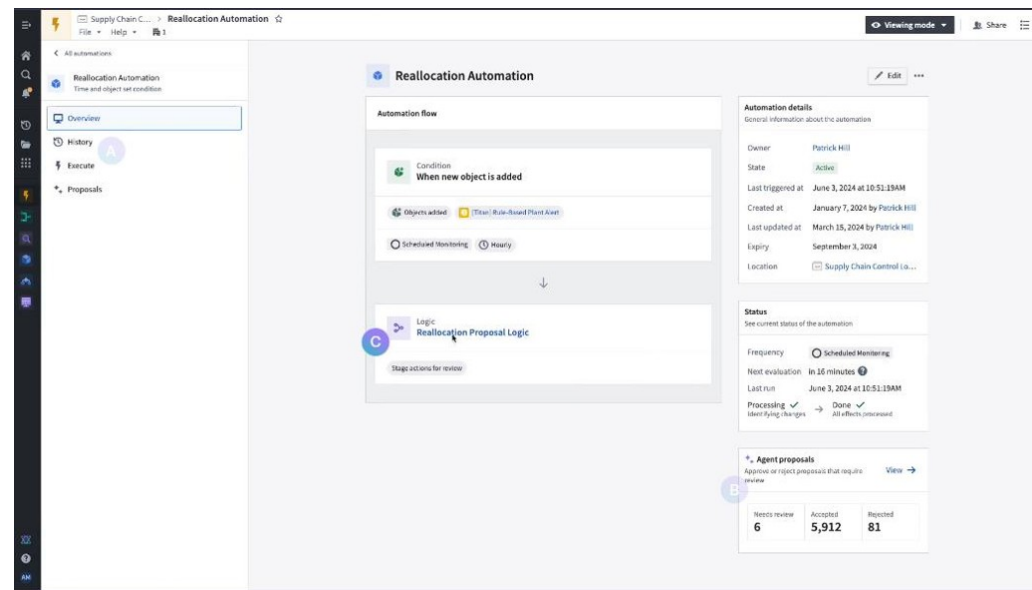
Morgan Stanley

MORGAN STANLEY RESEARCH

SaaS Vendors – Palantir (PLTR)

Palantir AIP

- Palantir's Artificial Intelligence Platform (AIP) which today includes agent functionality was **first announced in April 2023**.
- **AI Agent Functionality:** AIP enables users to build and deploy AI Agents using customized LLMs that **interact with enterprise data and integrate with the customers' "ontology"** (digital twin) to improve the decision-making capabilities of the agents. Palantir's agent use cases focus primarily on data heavy, analytics use cases, but can be broadly customized for different operational use cases.
- **Integration:** AIP integrates with Palantir's Foundry, leveraging existing data pipelines and operational workflows to automate processes. Based on recent demos and customer conversations, most agent workflows continue to have a human in the loop rather than being fully autonomous.



Pricing

- Pricing details are not publicized but our channel conversations suggest that customers can in many cases get initial developer access via pilots (free trials), while enterprise pricing is then provided on a custom basis based on existing contracts.

Source: Company Data, Morgan Stanley Research

80

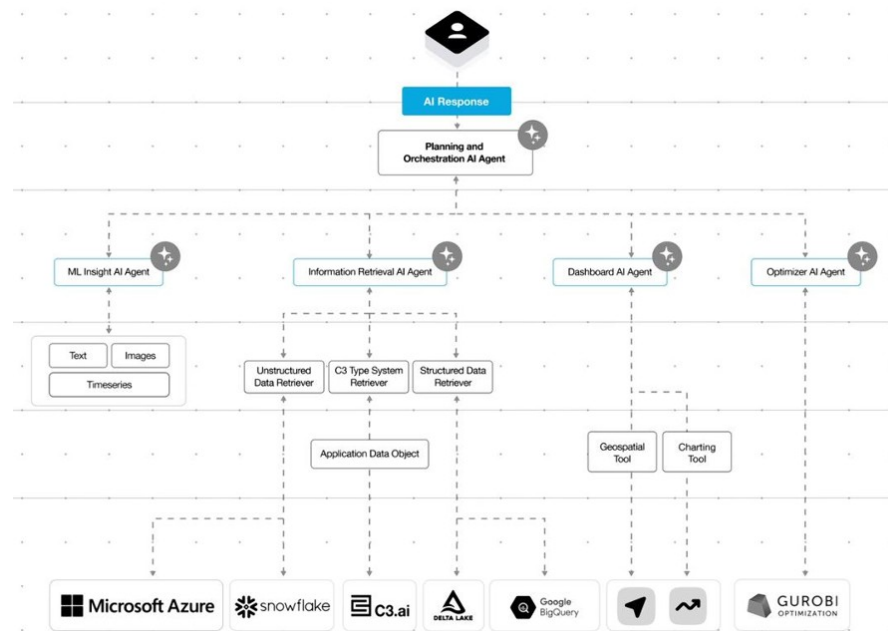
Morgan Stanley

MORGAN STANLEY RESEARCH

SaaS Vendors – C3.ai (AI)

C3 AI Agents

- C3 AI's AI agent functionality was **announced initially as part of the C3 Generative AI suite in March 2023**.
- **AI Agent Functionality:** C3 AI agents leverage orchestration software to create and deploy multiple specialized AI agents that autonomously retrieve, analyze, and summarize enterprise data via a natural language interface, driving interactive decision support and workflow automation.
- **Integration:** Agent capabilities are built into the C3 AI Platform and integrate with different pre-built enterprise AI applications, to extend existing models and connectors.



Pricing

- C3 AI's pricing is partly consumption-based with pilots starting at \$250,000 to launch an AI application, which is followed by charges per vCPU/vGPU hour with volume discounts.
- The service is available via cloud marketplaces for self-service deployment.

Source: Company Data, Morgan Stanley Research

SaaS Vendors – SAP (SAPG)

Joule

SAP's AI strategy centers around Joule, its generative AI copilot launched in September 2023. Joule is embedded throughout SAP's cloud enterprise portfolio, allowing users to interact with business data using natural language across applications from HR to finance, supply chain, and customer experience.

Joule serves as an intelligent assistant that understands business context and can perform tasks like analyzing financial data, optimizing supply chains, and enhancing customer experiences. It provides personalized insights and recommendations based on users' roles and permissions within the SAP ecosystem.

In October 2024, SAP enhanced Joule with collaborative AI agents - specialized AI systems that work together on complex business workflows. These agents handle specific tasks across finance, service, and sales with customers able to build their own agents alongside SAP's library. The agents can automate processes like invoice matching, customer service inquiries, and sales opportunity qualification.

The SAP Knowledge Graph, introduced in Q1 2025, provides deeper business understanding by mapping relationships across SAP's data landscape. This solution grounds AI in business semantics, connecting entities like purchase orders and invoices while reducing the risk of inaccurate AI results.

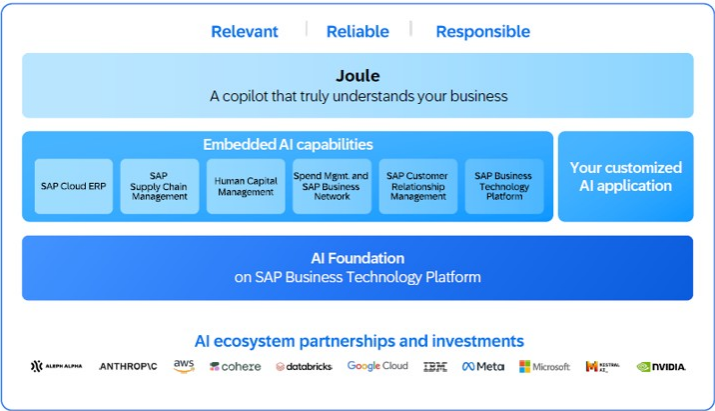
For developers, SAP offers AI-powered tools including:

- Generative AI capabilities in SAP Build for code explanation
- ABAP Large Language Model for customizations
- Vector database features in SAP HANA Cloud

SAP delivered over 130 AI use cases at the end of 2024 and has committed to delivering 400+ by end of 2025, focusing on measurable business outcomes while maintaining principles of relevance, reliability, and responsibility.

Source: Company Data, Morgan Stanley Research

Out-of-the-box Agents



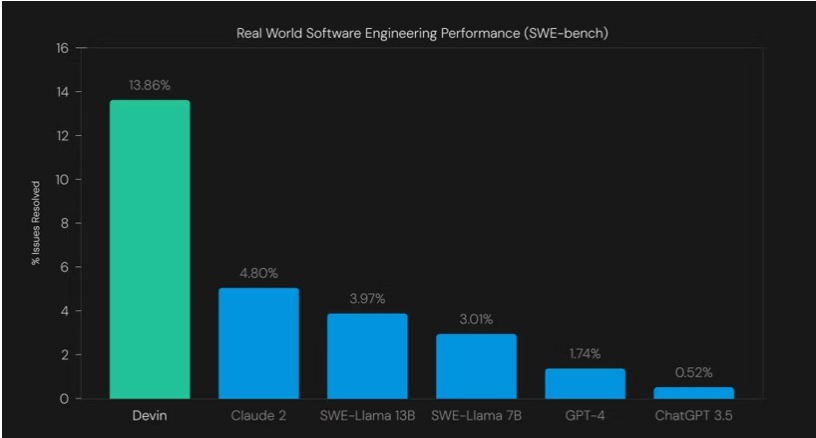
- 1 **Joule**
Collaborative AI agent experience across our portfolio
500+ skills added so far
- 2 **AI Capabilities**
Embedded or customized AI by partner, customer, or SAP
100+ Generative AI use cases
- 3 **AI Foundation**
Integrated enterprise data, security, and compliance
100+ Partners on AI foundation

Startups – Cognition Labs

Devin

- Launched in March 2024, Devin is an autonomous AI assistant created by Cognition Labs, **branded as the “first AI software engineer”** and designed to complete software development tasks with minimal human intervention.
- Cognition Labs (also known as Cognition AI) is a startup founded by a small team led by CEO Scott Wu, CTO Steven Hao, and CPO Walden Yan.
 - The company received \$21 million in Series A funding led by Peter Thiel’s Founders Fund – at \$350 million valuation.
 - Raised another \$175M shortly after in a Series B round led by Founders Fund – at \$2 billion valuation.
 - Recently the company secured additional funding (undisclosed amount) valuing the company close to \$4 billion as reported by [Bloomberg](#).
- **Core Functionalities:**
 - Autonomous planning and execution of complex engineering tasks
 - Coding, debugging, and problem-solving with ability to recall relevant context and fix mistakes
 - Learning unfamiliar technologies by researching online resources to complete tasks
 - Building and deploying applications end-to-end from simple text prompts
 - Finding and fixing bugs in existing codebases independently
 - Training and fine-tuning its own AI models when necessary

Claimed Performance



Pricing

- After an initial early access period, Devin became generally available in December 2024 with **pricing starting at \$500 per month**,
- Subscription includes no seat limits, access to its Slack integration, IDE extension, and API, along with onboarding and support.
 - There is a pay-as-you-go component for additional ACUs or “Devin Activity Units” at \$2/ACU, up to a limit. Each ACU is roughly equivalent to 15 minutes of active Devin work.

Source: Company Data, Morgan Stanley Research

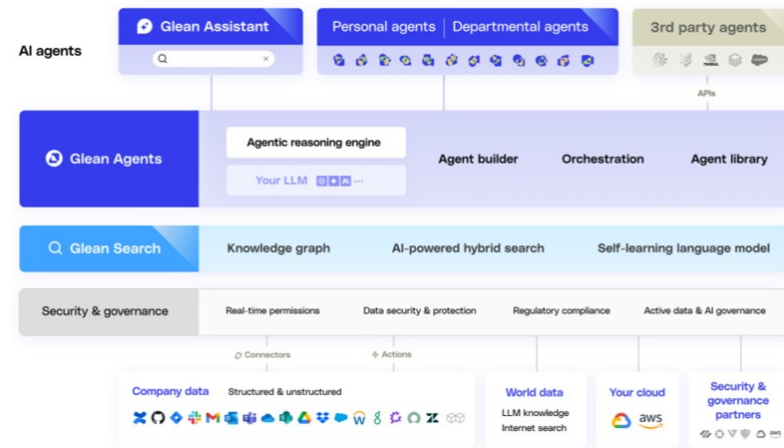
Morgan Stanley

MORGAN STANLEY RESEARCH

Startups – Glean

Work AI Platform

- Founded in 2019 by former Google search engineers led by Arvind Jain, Glean has evolved from an AI-powered enterprise search solution to **a comprehensive “Work AI” platform**.
- Beginning with a \$15 million Series A in 2019, Glean has secured over \$500 million in total investment, including a \$200 million Series D round in February 2024 at a \$2.2 billion valuation and a \$260 million Series E round in September 2024 that **valued the company at \$4.6 billion**.
- Glean’s core offering is its Work AI platform, which connects to and understands enterprise data, enabling AI-powered search, content generation, and task automation.
- **Key Platform Components:**
 - **AI-Powered Search:** Uses vector search and LLMs to enable natural language queries. Delivered personalized results based on the user’s role, current work, and collaborators.
 - **Glean Assist:** AI assistant that allows users to ask questions, find information, and perform advanced tasks. In February 2025, Glean enhanced this with “multi-step prompts”, enabling reasoning through complex tasks.
 - **Glean Agents:** Launched in February 2025, it provides a horizontal agent environment for building, deploying, orchestrating, and governing AI Agents.
 - These agents can automate personal tasks and business processes, with access to both enterprise data and internet data.



Pricing

For organizations, there is *no public pricing structure or free plan*, requiring potential customers to schedule a demo and discussion with their sales team for a customized quote.

For Individuals, there are 5 pricing editions:

1. \$12/month for monthly access
2. \$129 for 1 year - \$10.75/month
3. \$245 for 2 years – \$10.20/month
4. \$348 for 3 years - \$9.67/month
5. \$439 for 4 years – \$9.15/month

Source: Company Data, Morgan Stanley Research

84

Morgan Stanley

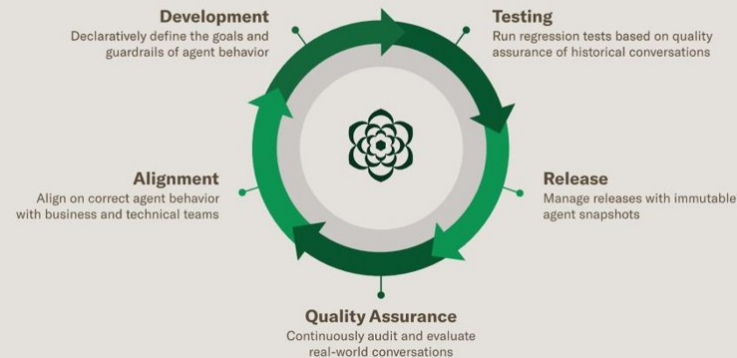
MORGAN STANLEY RESEARCH

Startups – Sierra

Agent OS

- Founded in February 2023 by Bret Taylor (current chairman of OpenAI's BoD and previously Co-CEO of Salesforce) and Clay Bavor (previously led Google Labs), **Sierra is a conversational AI platform** that enables businesses to deploy sophisticated, brand-aligned AI agents.
- In October 2024, Sierra raised \$175 million in a Series B round led by Greenoaks Capital at a **\$4.5 billion valuation**, more than quadrupling its valuation since January 2024 when it raised \$110 million from Sequoia Capital and Benchmark at a ~\$1 billion valuation.
 - Sierra has **reportedly crossed \$20 million in annualized revenue**, per Reuters.
- Sierra targets the customer experience software market by addressing systemic inefficiencies in traditional contact centers. The platform, branded as **Agent OS**, combines LLMs with proprietary orchestration layers to create AI agents capable of handling complex, multi-step customer interactions.
- Key Platform Components:**
 - Agent Development Environment:** Provides no-code tools for defining **Goals** (reducing call handling time or increasing upsell rates), **Guardrails** (prevent unauthorized actions or violations), and **Knowledge Grounding** (integration with 100+ enterprise systems via APIs).
 - Multi-LLM Architecture:** Combines foundation models from OpenAI, Anthropic, and Meta. Primary LLMs handle conversation flow, while secondary models verify accuracy and monitor for hallucinations.
 - Voice Interaction Engine:** Launched in October 2024, it processes tone, pacing, and emotional inflection, naturally handles conversational overlaps/interruptions, and analyses sentiment in real-time to trigger de-escalation protocols.
 - Action Execution Framework:** Agents can initiate business processes like subscription cancellations, order returns/exchanges, and appointment rescheduling.

Sierra Agent Development Life Cycle



Pricing

Sierra doesn't publicly disclose specific pricing tiers, instead the company emphasizes an "outcome-based pricing" approach, which is tied to tangible business impacts, such as a resolved support conversation, a saved cancellation, an upsell, a cross-sell. If the conversation isn't resolved, in most cases there's no charge.

Source: Company Data, Morgan Stanley Research

Startups – Aisera

Aisera Agentic AI Platform

- Founded in 2017 by CEO Matt Thomas, Aisera is an enterprise **AI-driven service automation platform**, delivering solutions that blend generative AI with sophisticated workflow orchestration.
- Since its funding, Aisera has raised \$180 million across multiple funding rounds, though its valuation remains undisclosed.

Key Platform Components:

- Intentless Orchestration Engine:** Enables natural language interactions without predefined intent mappings. It acts autonomously to understand, reason, and proactively fulfill requests through AI agents.
- Multi-LLM Gateway:** Employs a model routing system that selects between Azure AI, OpenAI, Google Vertex AI, Llama 3, AWS Bedrock, or proprietary AiseraLLM based on query complexity.
- AI Agent Development:** Build AI Agents leveraging AiseraLLM, task-specific agents, domain-specific LLMs, and AI Studios.
 - LLM Studio:** Integrate and execute any LLM of your choice in your Gen AI application with Aisera's Gateway (LLMOps)
 - Prompt Studio:** Create, test, and refine prompts to mitigate hallucinations and irrelevant outputs.
 - Event Studio:** Take actions from system events by dynamically triggering action flows based on the priority, sequence, and nature of the event.
 - Workflow Studio:** Create workflows using a low-code/no-code interface, leveraging a library of 3K+ pre-built automations.
 - Copilot Studio:** Create context-aware and time-saving skill flows that complete tasks easily.
 - Hyperflow Studio:** Complete tasks by identifying which APIs to call and dynamically orchestrating their execution in sequence on the fly.

Aisera's Agentic AI Platform

Securely Deploy AI Agents in Weeks, Not Months

Pricing per Azure Marketplace

AI Service Desk [Save to my list](#)

Aisera

[Free trial](#)

[Overview](#) [Plans + Pricing](#) [Ratings + reviews](#)

Plan	Description	Price + payment options	Billing term	Subtotal
Up to 10000 users	Free trial This plan is for AI Service Desk and supports up to 10000 users on 1 ticketing/knowledge system	First month free, then \$1,200,000.00/one-time payment	1-year	\$1,200,000.00
Up to 1000 Users	Free trial This plan is for AI Service Desk and supports up to 1000 users on 1 ticketing/knowledge system	First month free, then \$200,000.00/one-time payment	1-year	\$200,000.00
Aisera AI Service Desk	Generative AI and Automation solution.	\$5.00/one-time payment	1-year	\$5.00
AI Service Desk	Generative AI and Automation solution.	\$5.00/one-time payment	1-year	\$5.00

Source: Company Data, Morgan Stanley Research

Startups – CrewAI

CrewAI

- Founded in 2023 by João Moura in São Paulo, Brazil, CrewAI is a **multi-agent AI platform providing a framework for building, deploying, and managing AI agents**, with both high-level simplicity and precise low-level control, for creating autonomous AI agents tailored to any scenario:
 - **CrewAI Crews:** Optimize for autonomy and collaborative intelligence, enabling you to create AI teams where each agent has specific roles, tools, and goals
 - **CrewAI Flows:** Enable granular, event-driven control, single LLM calls for precise task orchestration and supports Crews natively.
- CrewAI began as an open-source framework (*still is*), but after a successful private beta, it launched an Enterprise edition in October 2024 after raising \$18 million across two funding rounds including a Series A led by Insight Partners. Valuation is unknown.
- **Key Features and Capabilities:**
 - **Role-Based Agents:** Create specialized agents with defined roles, expertise, and goals.
 - **Intelligent Collaboration:** Agents work together, sharing insights and coordinating tasks to achieve complex objectives.
 - **Flexible Tools:** Equip agents with custom tools and APIs to interact with external services and data sources.
 - **Task Management:** Define sequential or parallel workflows, with agents automatically handling task dependencies.
 - **Event-Driven Orchestration:** Define precise execution paths responding dynamically to events.
 - **Fine-Grained Control:** Manage workflow states and conditional execution securely and efficiently.
 - **Native Crew Integration:** Effortlessly combine with Crews for enhanced autonomy and intelligence.
 - **Deterministic Execution:** Ensure predictable outcomes with explicit control flow and error handling.

Streamline Your **Workflows** with AI-Powered Development

Our AI multi-agentic platform provides a wide range of tools for every step in the agent life cycle.

Plan

Build

Deploy

Monitor

Iterate

Evaluator

Crew Studio

Auto API Creation

Management UI

Crew Testing

Chat with Docs

Templates

Deploy CLIA

Quality Reporting

Training UI

CrewAI Courses

Integrations

Auto Generate UI

Crew Dashboards

Auto Train

CrewAI Pricing Tiers, Features, and Intended Users						
	Free	Basic	Standard	Pro	Enterprise	Ultra
Pricing	Free	\$99/month	\$6,000/year	\$12,000/year	\$60,000/year	\$120,000/year
Features	1 Crew, 50 monthly task executions	2 Crews, 100 monthly executions, total of 5 developers	5 Crews, 1K monthly executions, unlimited seats	10 Crews, 2K monthly executions, unlimited seats	50 Crews, 10K monthly executions, unlimited seats	100 Crews, 500K monthly executions, unlimited seats
Intended Users	Techies searching for a new AI tool	Small businesses with limited workflows	Small businesses that are scaling to median	Median to large businesses, growing	Large organizations and enterprises	Fortune 500/Enterprises

Source: Company Data, Morgan Stanley Research

Morgan Stanley

MORGAN STANLEY RESEARCH

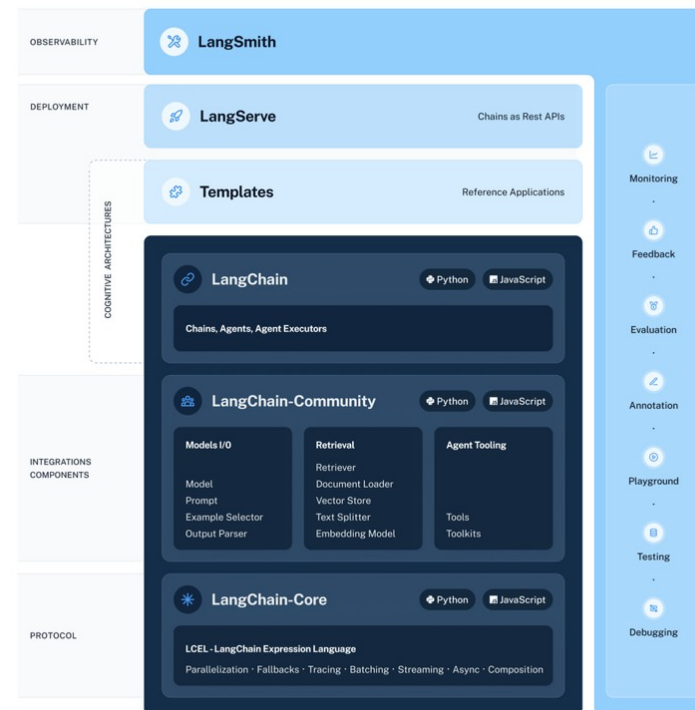
Startups – LangChain

LangChain – LangSmith – LangGraph

- Founded in October 2022 by CEO Harrison Chase and co-founder Ankush Gola, LangChain originated as an open-source project to facilitate development of applications connected to LLMs. After gaining significant traction, it incorporated in April 2023 and has since raised ~\$35 million in private funding.
- **Key Platform Components:**
 - **Model I/O:** Interfaces with language models through standardized abstractions for LLMs, chat models, and embedding models
 - **Prompt Management:** Creates templated and dynamically selected prompts with example selectors to improve model outputs
 - **Retrieval Framework:** Enables information retrieval through document loaders, text splitters, embeddings, and vector stores for effective RAG
 - **Memory Systems:** Implements chat message history and conversation persistence between chain runs
 - **Chains:** Provides common building block compositions for creating workflows that combine multiple components
 - **Agent Development:** Builds AI agents that can reason about issues, break them down into sub-tasks, and choose appropriate tools to complete goals
- **Ecosystem:**
 - **LangChain Libraries:** Core open-source packages available in Python and JavaScript with various integrations
 - **LangChain Expression Language (LCEL):** Declarative syntax for orchestrating components, introduced in Q3 2023
 - **LangServe:** Deployment tool launched in Oct 2023 to host LCEL code as production-ready APIs
 - **LangSmith:** Developer platform for debugging, testing, evaluating, and monitoring LLM applications with pricing starting at \$39/seat/month
 - **LangGraph Platform:** Cloud service for deploying agentic applications at scale with usage-based pricing (\$0.001/node executed)

Source: Company Data, Morgan Stanley Research

LangChain Ecosystem



Startups – Anysphere

Cursor

- Founded in 2022 by four MIT classmates (Sualeh Asif, Arvid Lunnemark, Aman Sanger, and Michael Truell), the company launched its flagship product Cursor in 2023. Later that year it raised \$8 million in seed funding led by OpenAI’s Startup Fund
- Anysphere has raised over \$1 billion in funding to date, including a \$900 million funding round in May 2025 led by Thrive Capital at a \$9 billion valuation
- **Cursor** is an AI-powered integrated development environment (IDE) that enhances developer productivity through 1) intelligent code generation, 2) codebase understanding, 3) multi-model integration, 4) smart reviews, 5) VS code compatibility, and 6) Agent Mode
- **Agent Mode:** Autonomous AI coding agent that independently explores, plans, and executes complex codebase changes with full tools. Capabilities include:
 - **Autonomous Operation:** Independently explores your codebase, identifies relevant files, and makes necessary changes
 - **Full Tool Access:** Uses all available tools to search, edit, create files, and run terminal commands
 - **Contextual Understanding:** Builds a comprehensive understanding of your project structure and dependencies
 - **Multi-Step Planning:** Breaks complex tasks into manageable steps and executes them in sequence
- **Ecosystem & Adoption:**
 - **Enterprise Adoption:** Used by engineering teams at major tech companies including Stripe, OpenAI, Spotify, and Uber
 - **Developer Impact:** Reportedly crossed 1 billion lines of generated code per day
 - **Acquisition:** Acquired AI coding assistant Supermaven in November 2024 to expand capabilities

Pricing

Choose the plan that works for you

MONTHLY

YEARLY (SAVE 20%)

Hobby

Free

Includes

✓ Pro two-week trial

✓ 2000 completions

✓ 50 slow requests

DOWNLOAD

OTHERS

Pro

\$16 /month

Everything in Hobby, plus

✓ Unlimited completions

✓ 500 requests per month

✓ Unlimited slow requests

✓ Max mode

GET STARTED

Business

\$32 /user/month

Everything in Pro, plus

✓ Enforce privacy mode org-wide

✓ Centralized team billing

✓ Admin dashboard with usage stats

✓ SAML/OIDC SSO

GET STARTED

Source: Company Data, Morgan Stanley Research



Disclosure Section

The information and opinions in Morgan Stanley Research were prepared by Morgan Stanley & Co. LLC, and/or Morgan Stanley C.T.V.M. S.A., and/or Morgan Stanley Mexico, Casa de Bolsa, S.A. de C.V., and/or Morgan Stanley Canada Limited. As used in this disclosure section, "Morgan Stanley" includes Morgan Stanley & Co. LLC, Morgan Stanley C.T.V.M. S.A., Morgan Stanley Mexico, Casa de Bolsa, S.A. de C.V., Morgan Stanley Canada Limited and their affiliates as necessary.

For important disclosures, stock price charts and equity rating histories regarding companies that are the subject of this report, please see the Morgan Stanley Research Disclosure Website at www.morganstanley.com/researchdisclosures, or contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY, 10036 USA.

For valuation methodology and risks associated with any recommendation, rating or price target referenced in this research report, please contact the Client Support Team as follows: US/Canada +1 800 303-2495; Hong Kong +852 2848-5999; Latin America +1 718 754-5444 (U.S.); London +44 (0)20-7425-8169; Singapore +65 6834-6860; Sydney +61 (0)2-9770-1505; Tokyo +81 (0)3-6836-9000. Alternatively you may contact your investment representative or Morgan Stanley Research at 1585 Broadway, (Attention: Research Management), New York, NY 10036 USA.

Analyst Certification

The following analysts hereby certify that their views about the companies and their securities discussed in this report are accurately expressed and that they have not received and will not receive direct or indirect compensation in exchange for expressing specific recommendations or views in this report: Josh Baer, CFA; Brian Nowak, CFA; Elizabeth Porter, CFA; Chris Quintero; Sanjit K Singh; Keith Weiss, CFA.

Global Research Conflict Management Policy

Morgan Stanley Research has been published in accordance with our conflict management policy, which is available at www.morganstanley.com/institutional/research/conflict/policies. A Portuguese version of the policy can be found at www.morganstanley.com.br

Important Regulatory Disclosures on Subject Companies

The analyst or strategist (or a household member) identified below owns the following securities (or related derivatives): Josh Baer, CFA - Amazon.com Inc(common or preferred stock), Peloton Interactive, Inc(common or preferred stock), Roblox Corporation(common or preferred stock), Uber Technologies Inc(common or preferred stock); Brian Nowak, CFA - Cloudflare Inc(common or preferred stock); Elizabeth Porter, CFA - Microsoft(common or preferred stock); Chris Quintero - Expedia Inc(common or preferred stock); Oscar R Saavedra - Alphabet Inc(common or preferred stock), Amazon.com Inc(common or preferred stock); Sanjit K Singh - Alphabet Inc(common or preferred stock), Pinterest Inc(common or preferred stock); Keith Weiss, CFA - Uber Technologies Inc(common or preferred stock).

As of April 30, 2025, Morgan Stanley beneficially owned 1% or more of a class of common equity securities of the following companies covered in Morgan Stanley Research: Adobe Inc., Airbnb Inc, Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, Appian Corp, AppLovin Corp, Asana Inc, Atlassian Corporation PLC, Autodesk, BILL Holdings Inc, Blackline Inc, Booking Holdings Inc, Box Inc, Bumble Inc., C3.ai, CCC Intelligent Solutions Holdings Inc, Chewy Inc, Cloudflare Inc, Confluent, Inc., Couchbase, Inc., Coursera, Inc., Criteo SA, CrowdStrike Holdings Inc, CyberArk Software Ltd, Datadog, Inc., DigitalOcean Holdings Inc, DocuSign Inc, DoorDash Inc, DoubleVerify Holdings Inc, E2open Parent Holdings Inc, eBay Inc, Elastic NV, Electronic Arts Inc, Etsy Inc, Expedia Inc., Five9 Inc, Fortinet Inc, Gen Digital Inc, GoDaddy Inc, Intuit, Jamf Holding Corp, Klaviyo, Inc, LegalZoom.com Inc, Lightspeed POS Inc, Liveramp Holdings Inc, Lyft Inc, Match Group Inc, Meta Platforms Inc, Microsoft, MongoDB Inc, NICE Ltd., Okta, Inc., Opendoor Technologies Inc, Oracle Corporation, PagerDuty, Inc., Palantir Technologies Inc., Palo Alto Networks Inc, Peloton Interactive, Inc., Pinterest Inc, Qualys Inc, Rapid7 Inc, Reddit Inc, Revolve Group Inc, RingCentral Inc, Roblox Corporation, Salesforce, Inc., Samsara Inc, Semrush Holdings Inc -A, ServiceNow Inc, Shopify Inc, Shutterstock Inc, Snowflake Inc., Sprinklr Inc, Sprout Social Inc, Take-Two Interactive Software, Tenable Holdings Inc, Toast, Inc., Trade Desk Inc, Twilio Inc, Uber Technologies Inc, UiPath Inc, Varonis Systems, Inc., Vertex Inc., Webtoon Entertainment Inc, Workday Inc, WW International Inc, Yelp Inc.

Within the last 12 months, Morgan Stanley managed or co-managed a public offering (or 144A offering) of securities of Akamai Technologies, Inc., Alphabet Inc., AppLovin Corp, BILL Holdings Inc, Booking Holdings Inc, Box Inc, Chewy Inc, CoreWeave, Datadog, Inc., DoorDash Inc, Klaviyo, Inc, Meta Platforms Inc, OneStream Inc, Sabre Corp, SailPoint Inc, ServiceTitan Inc, Uber Technologies Inc, Unity Software Inc, Varonis Systems, Inc., Webtoon Entertainment Inc, Zeta Global Holdings Corp.

Within the last 12 months, Morgan Stanley has received compensation for investment banking services from Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, AppLovin Corp, Autodesk, BILL Holdings Inc, Booking Holdings Inc, Box Inc, CCC Intelligent Solutions Holdings Inc, Chewy Inc, CoreWeave, CyberArk Software Ltd, Datadog, Inc., DigitalOcean Holdings Inc, DocuSign Inc, DoorDash Inc, Klaviyo, Inc, Meta Platforms Inc, OneStream Inc, Qualys Inc, Sabre Corp, SailPoint Inc, Salesforce, Inc., ServiceTitan Inc, Toast, Inc., Uber Technologies Inc, Udemy Inc, Varonis Systems, Inc., Webtoon Entertainment Inc, Zeta Global Holdings Corp.

In the next 3 months, Morgan Stanley expects to receive or intends to seek compensation for investment banking services from Adobe Inc., Airbnb Inc, Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, Amplitude Inc., Appian Corp, AppLovin Corp, Asana Inc, Atlassian Corporation PLC, Autodesk, BigCommerce Holdings, Inc., BILL Holdings Inc, Blackline Inc, Booking Holdings Inc, Box Inc, Bumble Inc., C3.ai, CCC Intelligent Solutions Holdings Inc, Check Point Software Technologies Ltd., Chewy Inc, Cloudflare Inc, Confluent, Inc., CoreWeave, Couchbase, Inc., Coursera, Inc., Criteo SA, CrowdStrike Holdings Inc, CyberArk Software Ltd, Datadog, Inc., DigitalOcean Holdings Inc, Docebo Inc., DocuSign Inc, DoorDash Inc, DoubleVerify Holdings Inc, Duolingo Inc, eBay Inc, Elastic NV, Electronic Arts Inc, Etsy Inc, Expedia Inc., FIGS, Inc., Five9 Inc, Fortinet Inc., Freshworks Inc, Gen Digital Inc., GitLab Inc, GoDaddy Inc, HubSpot, Inc., Instacart, Integral Ad Science Holding Corp., Intuit, Jamf Holding Corp, JFrog Ltd., Klaviyo, Inc, LegalZoom.com Inc, Lightspeed POS Inc, Liveramp Holdings Inc, Lyft Inc, Match Group Inc, Meta Platforms Inc, Microsoft, MongoDB Inc, Nextdoor Holdings Inc, NICE Ltd., Okta, Inc., OneStream Inc, Opendoor Technologies Inc, Oracle Corporation, PagerDuty, Inc., Palantir Technologies Inc., Palo Alto Networks Inc, Peloton Interactive, Inc., Pinterest Inc, Playtika Holding Corp, Qualys Inc, Rapid7 Inc, Reddit Inc, Revolve Group Inc, RingCentral Inc, Roblox Corporation, Sabre Corp, SailPoint Inc, Salesforce, Inc., Samsara Inc, Semrush Holdings Inc -A, SentinelOne, Inc., ServiceNow Inc, ServiceTitan Inc, Shopify Inc, Shutterstock Inc, Snap Inc., Snowflake Inc., Sprinklr Inc, Sprout Social Inc, Take-Two Interactive Software, Tenable Holdings Inc, Toast, Inc., Trade Desk Inc, Twilio Inc, Uber Technologies Inc, Udemy Inc, UiPath Inc, Unity Software Inc, Varonis Systems, Inc., Vertex Inc., Webtoon Entertainment Inc, Wix.Com Ltd, Workday Inc, WW International Inc, Yelp Inc, Zeta Global Holdings Corp, Zillow Group Inc, Zoom Video Communications Inc, Zscaler Inc.

Within the last 12 months, Morgan Stanley has received compensation for products and services other than investment banking services from 8x8 Inc, Adobe Inc., Airbnb Inc, Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, AppLovin Corp, Asana Inc, Atlassian Corporation PLC, Autodesk, Blackline Inc, Booking Holdings Inc, Box Inc, CCC Intelligent Solutions Holdings Inc, Check Point Software Technologies Ltd., Chewy Inc, Cloudflare Inc, Confluent, Inc., CoreWeave, DigitalOcean Holdings Inc, DocuSign Inc, DoorDash Inc, Dynatrace Inc, eBay Inc, Electronic Arts Inc, Expedia Inc., Five9 Inc, Fortinet Inc., Freshworks Inc, Gen Digital Inc., Integral Ad Science Holding Corp., Intuit, Jamf Holding Corp, LegalZoom.com Inc, Match Group Inc, Meta Platforms Inc, Microsoft, NICE Ltd., OneStream Inc, Oracle Corporation, Palantir Technologies Inc., Palo Alto Networks Inc, Peloton Interactive, Inc., Pinterest Inc, Playtika Holding Corp, Reddit Inc, Revolve Group Inc, RingCentral Inc, Sabre Corp, SailPoint Inc, Salesforce, Inc., ServiceTitan Inc, Snap Inc., Snowflake Inc., Take-Two Interactive Software, Tenable Holdings Inc, Toast, Inc., Uber Technologies Inc, Udemy Inc, UiPath Inc, Unity Software Inc, Workday Inc, WW International Inc, Zeta Global Holdings Corp, Zoom Video Communications Inc, Zscaler Inc.

Within the last 12 months, Morgan Stanley has provided or is providing investment banking services to, or has an investment banking client relationship with, the following company: Adobe

Inc., Airbnb Inc, Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, Amplitude Inc., Appian Corp, AppLovin Corp, Asana Inc, Atlassian Corporation PLC, Autodesk, BigCommerce Holdings, Inc., BILL Holdings Inc, Blackline Inc, Booking Holdings Inc, Box Inc, Bumble Inc., C3.ai, CCC Intelligent Solutions Holdings Inc, Check Point Software Technologies Ltd., Chewy Inc, Cloudflare Inc, Confluent, Inc., CoreWeave, Couchbase, Inc., Coursera, Inc., Criteo SA, CrowdStrike Holdings Inc, CyberArk Software Ltd, Datadog, Inc., DigitalOcean Holdings Inc, Docebo Inc, DocuSign Inc, DoorDash Inc, DoubleVerify Holdings Inc, Duolingo Inc, eBay Inc, Elastic NV, Electronic Arts Inc, Etsy Inc, Expedia Inc., FIGS, Inc., Five9 Inc, Fortinet Inc., Freshworks Inc, Gen Digital Inc., GitLab Inc, GoDaddy Inc, HubSpot, Inc., Instacart, Integral Ad Science Holding Corp., Intuit, Jamf Holding Corp, JFrog Ltd., Klaviyo, Inc, LegalZoom.com Inc, Lightspeed POS Inc., Liveramp Holdings Inc, Lyft Inc, Match Group Inc, Meta Platforms Inc, Microsoft, MongoDB Inc, Nextdoor Holdings Inc, NICE Ltd., Okta, Inc., OneStream Inc, Opendoor Technologies Inc, Oracle Corporation, PagerDuty, Inc., Palantir Technologies Inc., Palo Alto Networks Inc, Peloton Interactive, Inc., Pinterest Inc, Playtika Holding Corp, Qualys Inc, Rapid7 Inc, Reddit Inc, Revolve Group Inc, RingCentral Inc, Roblox Corporation, Sabre Corp, SailPoint Inc, Salesforce, Inc., Samsara Inc, Semrush Holdings Inc -A, SentinelOne, Inc., ServiceNow Inc, ServiceTitan Inc, Shopify Inc, Shutterstock Inc, Snap Inc., Snowflake Inc., Sprinklr Inc, Sprout Social Inc, Take-Two Interactive Software, Tenable Holdings Inc, Toast, Inc., Trade Desk Inc, Twilio Inc, Uber Technologies Inc, Udem Inc, UiPath Inc, Unity Software Inc, Varonis Systems, Inc., Vertex Inc., Webtoon Entertainment Inc, Wix.Com Ltd, Workday Inc, WW International Inc, Yelp Inc, Zeta Global Holdings Corp, Zillow Group Inc, Zoom Video Communications Inc, Zscaler Inc.

Within the last 12 months, Morgan Stanley has either provided or is providing non-investment banking, securities-related services to and/or in the past has entered into an agreement to provide services or has a client relationship with the following company: 8x8 Inc, Adobe Inc., Airbnb Inc, Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, AppLovin Corp, Asana Inc, Atlassian Corporation PLC, Autodesk, BigCommerce Holdings, Inc., Blackline Inc, Booking Holdings Inc, Box Inc, CCC Intelligent Solutions Holdings Inc, Check Point Software Technologies Ltd., Chewy Inc, Cloudflare Inc, Confluent, Inc., CoreWeave, CyberArk Software Ltd, Datadog, Inc., DigitalOcean Holdings Inc, DocuSign Inc, DoorDash Inc, Dynatrace Inc, eBay Inc, Electronic Arts Inc, Etsy Inc, Expedia Inc., Five9 Inc, Fortinet Inc., Freshworks Inc, Gen Digital Inc., GoDaddy Inc, HubSpot, Inc., Integral Ad Science Holding Corp., Intuit, Jamf Holding Corp, LegalZoom.com Inc, Match Group Inc, Meta Platforms Inc, Microsoft, MongoDB Inc, NICE Ltd., OneStream Inc, Opendoor Technologies Inc, Oracle Corporation, PagerDuty, Inc., Palantir Technologies Inc., Palo Alto Networks Inc, Peloton Interactive, Inc., Pinterest Inc, Playtika Holding Corp, Qualys Inc, Reddit Inc, Revolve Group Inc, RingCentral Inc, Sabre Corp, SailPoint Inc, Salesforce, Inc., ServiceNow Inc, ServiceTitan Inc, Shopify Inc, Snap Inc., Snowflake Inc., Take-Two Interactive Software, Tenable Holdings Inc, Toast, Inc., Twilio Inc, Uber Technologies Inc, Udem Inc, UiPath Inc, Unity Software Inc, Varonis Systems, Inc., Wix.Com Ltd, Workday Inc, WW International Inc, Zeta Global Holdings Corp, Zillow Group Inc, Zoom Video Communications Inc, Zscaler Inc.

An employee, director or consultant of Morgan Stanley is a director of Alphabet Inc., eBay Inc, Elastic NV, Tenable Holdings Inc. This person is not a research analyst or a member of a research analyst's household.

Morgan Stanley & Co. LLC makes a market in the securities of 8x8 Inc, Adobe Inc., Airbnb Inc, Akamai Technologies, Inc., Alphabet Inc., Amazon.com Inc, Amplitude Inc., Appian Corp, AppLovin Corp, Asana Inc, Atlassian Corporation PLC, Autodesk, BigCommerce Holdings, Inc., Blackline Inc, Booking Holdings Inc, Box Inc, Bumble Inc., C3.ai, CCC Intelligent Solutions Holdings Inc, Check Point Software Technologies Ltd., Compass, Inc., Confluent, Inc., CoreWeave, Couchbase, Inc., Coursera, Inc., Criteo SA, CyberArk Software Ltd, Datadog, Inc., DigitalOcean Holdings Inc, Docebo Inc., DocuSign Inc, DoubleVerify Holdings Inc, Duolingo Inc, Dynatrace Inc, E2open Parent Holdings Inc, eBay Inc, Elastic NV, Electronic Arts Inc, Etsy Inc, Expedia Inc., FIGS, Inc., Five9 Inc, Fortinet Inc., Freshworks Inc, Gen Digital Inc., GoDaddy Inc, Integral Ad Science Holding Corp., Intuit, Jamf Holding Corp, JFrog Ltd., Karooooo Ltd, Klaviyo, Inc, LegalZoom.com Inc, Liveramp Holdings Inc, Match Group Inc, Meta Platforms Inc, Microsoft, Nextdoor Holdings Inc, OneStream Inc, Opendoor Technologies Inc, Oracle Corporation, PagerDuty, Inc., Palo Alto Networks Inc, Playtika Holding Corp, Qualys Inc, Rapid7 Inc, Revolve Group Inc, RingCentral Inc, Sabre Corp, SailPoint Inc, Salesforce, Inc., Semrush Holdings Inc -A, ServiceNow Inc, ServiceTitan Inc, Shopify Inc, Shutterstock Inc, Sprinklr Inc, Sprout Social Inc, Take-Two Interactive Software, Tenable Holdings Inc, Udem Inc, Varonis Systems, Inc., Vertex Inc., Wix.Com Ltd, Workday Inc, Yelp Inc, Zeta Global Holdings Corp, Zillow Group Inc, Zoom Video Communications Inc, ZoomInfo Technologies Inc, Zscaler Inc.

The equity research analysts or strategists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality of research, investor client feedback, stock picking, competitive factors, firm revenues and overall investment banking revenues. Equity Research analysts' or strategists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

Morgan Stanley and its affiliates do business that relates to companies/instruments covered in Morgan Stanley Research, including market making, providing liquidity, fund management, commercial banking, extension of credit, investment services and investment banking. Morgan Stanley sells to and buys from customers the securities/instruments of companies covered in Morgan Stanley Research on a principal basis. Morgan Stanley may have a position in the debt of the Company or instruments discussed in this report. Morgan Stanley trades or may trade as principal in the debt securities (or in related derivatives) that are the subject of the debt research report.

Certain disclosures listed above are also for compliance with applicable regulations in non-US jurisdictions.

STOCK RATINGS

Morgan Stanley uses a relative rating system using terms such as Overweight, Equal-weight, Not-Rated or Underweight (see definitions below). Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold and sell. Investors should carefully read the definitions of all ratings used in Morgan Stanley Research. In addition, since Morgan Stanley Research contains more complete information concerning the analyst's views, investors should carefully read Morgan Stanley Research, in its entirety, and not infer the contents from the rating alone. In any case, ratings (or research) should not be used or relied upon as investment advice. An investor's decision to buy or sell a stock should depend on individual circumstances (such as the investor's existing holdings) and other considerations.

Global Stock Ratings Distribution

(as of May 31, 2025)

The Stock Ratings described below apply to Morgan Stanley's Fundamental Equity Research and do not apply to Debt Research produced by the Firm.

For disclosure purposes only (in accordance with FINRA requirements), we include the category headings of Buy, Hold, and Sell alongside our ratings of Overweight, Equal-weight, Not-Rated and Underweight. Morgan Stanley does not assign ratings of Buy, Hold or Sell to the stocks we cover. Overweight, Equal-weight, Not-Rated and Underweight are not the equivalent of buy, hold, and sell but represent recommended relative weightings (see definitions below). To satisfy regulatory requirements, we correspond Overweight, our most positive stock rating, with a buy recommendation; we correspond Equal-weight and Not-Rated to hold and Underweight to sell recommendations, respectively.

Stock Rating Category	Coverage Universe		Investment Banking Clients (IBC)			Other Material Investment Services Clients (MISC)	
	Count	% of Total	Count	% of Total IBC	% of Rating Category	Count	% of Total Other MISC
Overweight/Buy	1493	40%	379	46%	25%	698	41%
Equal-weight/Hold	1650	44%	372	45%	23%	782	46%
Not-Rated/Hold	4	0%	0	0%	0%	2	0%
Underweight/Sell	602	16%	74	9%	12%	235	14%
Total	3,749		825			1717	

Data include common stock and ADRs currently assigned ratings. Investment Banking Clients are companies from whom Morgan Stanley received investment banking compensation in the last 12 months. Due to rounding off of decimals, the percentages provided in the "% of total" column may not add up to exactly 100 percent.

Analyst Stock Ratings

Overweight (O). The stock's total return is expected to exceed the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Equal-weight (E). The stock's total return is expected to be in line with the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Not-Rated (NR). Currently the analyst does not have adequate conviction about the stock's total return relative to the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Underweight (U). The stock's total return is expected to be below the average total return of the analyst's industry (or industry team's) coverage universe, on a risk-adjusted basis, over the next 12-18 months.

Unless otherwise specified, the time frame for price targets included in Morgan Stanley Research is 12 to 18 months.

Analyst Industry Views

Attractive (A): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be attractive vs. the relevant broad market benchmark, as indicated below.

In-Line (I): The analyst expects the performance of his or her industry coverage universe over the next 12-18 months to be in line with the relevant broad market benchmark, as indicated below.

Cautious (C): The analyst views the performance of his or her industry coverage universe over the next 12-18 months with caution vs. the relevant broad market benchmark, as indicated below.

Benchmarks for each region are as follows: North America - S&P 500; Latin America - relevant MSCI country index or MSCI Latin America Index; Europe - MSCI Europe; Japan - TOPIX; Asia - relevant MSCI country index or MSCI sub-regional index or MSCI AC Asia Pacific ex Japan Index.

Important Disclosures for Morgan Stanley Smith Barney LLC Customers

Important disclosures regarding the relationship between the companies that are the subject of Morgan Stanley Research and Morgan Stanley Smith Barney LLC or Morgan Stanley or any of their affiliates, are available on the Morgan Stanley Wealth Management disclosure website at www.morganstanley.com/online/researchdisclosures. For Morgan Stanley specific disclosures, you may refer to www.morganstanley.com/researchdisclosures.

Each Morgan Stanley research report is reviewed and approved on behalf of Morgan Stanley Smith Barney LLC. This review and approval is conducted by the same person who reviews the research report on behalf of Morgan Stanley. This could create a conflict of interest.

Other Important Disclosures

A member of Research who had or could have had access to the research prior to completion owns securities (or related derivatives) in the Atlassian Corporation PLC, Microsoft, Uber Technologies Inc, Workday Inc. This person is not a research analyst or a member of research analyst's household.

Morgan Stanley Research policy is to update research reports as and when the Research Analyst and Research Management deem appropriate, based on developments with the issuer, the sector, or the market that may have a material impact on the research views or opinions stated therein. In addition, certain Research publications are intended to be updated on a regular periodic basis (weekly/monthly/quarterly/annual) and will ordinarily be updated with that frequency, unless the Research Analyst and Research Management determine that a different publication schedule is appropriate based on current conditions.

Morgan Stanley is not acting as a municipal advisor and the opinions or views contained herein are not intended to be, and do not constitute, advice within the meaning of Section 975 of the Dodd-Frank Wall Street Reform and Consumer Protection Act.

Morgan Stanley produces an equity research product called a "Tactical Idea." Views contained in a "Tactical Idea" on a particular stock may be contrary to the recommendations or views expressed in research on the same stock. This may be the result of differing time horizons, methodologies, market events, or other factors. For all research available on a particular stock, please contact your sales representative or go to Matrix at <http://www.morganstanley.com/matrix>.

Morgan Stanley Research is provided to our clients through our proprietary research portal on Matrix and also distributed electronically by Morgan Stanley to clients. Certain, but not all, Morgan Stanley Research products are also made available to clients through third-party vendors or redistributed to clients through alternate electronic means as a convenience. For access to all available Morgan Stanley Research, please contact your sales representative or go to Matrix at <http://www.morganstanley.com/matrix>.

Any access and/or use of Morgan Stanley Research is subject to Morgan Stanley's Terms of Use (<http://www.morganstanley.com/terms.html>). By accessing and/or using Morgan Stanley Research, you are indicating that you have read and agree to be bound by our Terms of Use (<http://www.morganstanley.com/terms.html>). In addition you consent to Morgan Stanley processing your personal data and using cookies in accordance with our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html), including for the purposes of setting your preferences and to collect readership data so that we can deliver better and more personalized service and products to you. To find out more information about how Morgan Stanley processes personal data, how we use cookies and how to reject cookies see our Privacy Policy and our Global Cookies Policy (http://www.morganstanley.com/privacy_pledge.html).

If you do not agree to our Terms of Use and/or if you do not wish to provide your consent to Morgan Stanley processing your personal data or using cookies please do not access our research. Morgan Stanley Research does not provide individually tailored investment advice. Morgan Stanley Research has been prepared without regard to the circumstances and objectives of those

who receive it. Morgan Stanley recommends that investors independently evaluate particular investments and strategies, and encourages investors to seek the advice of a financial adviser. The appropriateness of an investment or strategy will depend on an investor's circumstances and objectives. The securities, instruments, or strategies discussed in Morgan Stanley Research may not be suitable for all investors, and certain investors may not be eligible to purchase or participate in some or all of them. Morgan Stanley Research is not an offer to buy or sell or the solicitation of an offer to buy or sell any security/instrument or to participate in any particular trading strategy. The value of and income from your investments may vary because of changes in interest rates, foreign exchange rates, default rates, prepayment rates, securities/instruments prices, market indexes, operational or financial conditions of companies or other factors. There may be time limitations on the exercise of options or other rights in securities/instruments transactions. Past performance is not necessarily a guide to future performance. Estimates of future performance are based on assumptions that may not be realized. If provided, and unless otherwise stated, the closing price on the cover page is that of the primary exchange for the subject company's securities/instruments.

The fixed income research analysts, strategists or economists principally responsible for the preparation of Morgan Stanley Research have received compensation based upon various factors, including quality, accuracy and value of research, firm profitability or revenues (which include fixed income trading and capital markets profitability or revenues), client feedback and competitive factors. Fixed Income Research analysts', strategists' or economists' compensation is not linked to investment banking or capital markets transactions performed by Morgan Stanley or the profitability or revenues of particular trading desks.

The "Important Regulatory Disclosures on Subject Companies" section in Morgan Stanley Research lists all companies mentioned where Morgan Stanley owns 1% or more of a class of common equity securities of the companies. For all other companies mentioned in Morgan Stanley Research, Morgan Stanley may have an investment of less than 1% in securities/instruments or derivatives of securities/instruments of companies and may trade them in ways different from those discussed in Morgan Stanley Research. Employees of Morgan Stanley not involved in the preparation of Morgan Stanley Research may have investments in securities/instruments or derivatives of securities/instruments of companies mentioned and may trade them in ways different from those discussed in Morgan Stanley Research. Derivatives may be issued by Morgan Stanley or associated persons.

With the exception of information regarding Morgan Stanley, Morgan Stanley Research is based on public information. Morgan Stanley makes every effort to use reliable, comprehensive information, but we make no representation that it is accurate or complete. We have no obligation to tell you when opinions or information in Morgan Stanley Research change apart from when we intend to discontinue equity research coverage of a subject company. Facts and views presented in Morgan Stanley Research have not been reviewed by, and may not reflect information known to, professionals in other Morgan Stanley business areas, including investment banking personnel.

Morgan Stanley Research personnel may participate in company events such as site visits and are generally prohibited from accepting payment by the company of associated expenses unless pre-approved by authorized members of Research management.

Morgan Stanley may make investment decisions that are inconsistent with the recommendations or views in this report.

To our readers based in Taiwan or trading in Taiwan securities/instruments: Information on securities/instruments that trade in Taiwan is distributed by Morgan Stanley Taiwan Limited ("MSTL"). Such information is for your reference only. The reader should independently evaluate the investment risks and is solely responsible for their investment decisions. Morgan Stanley Research may not be distributed to the public media or quoted or used by the public media without the express written consent of Morgan Stanley. Any non-customer reader within the scope of Article 7-1 of the Taiwan Stock Exchange Recommendation Regulations accessing and/or receiving Morgan Stanley Research is not permitted to provide Morgan Stanley Research to any third party (including but not limited to related parties, affiliated companies and any other third parties) or engage in any activities regarding Morgan Stanley Research which may create or give the appearance of creating a conflict of interest. Information on securities/instruments that do not trade in Taiwan is for informational purposes only and is not to be construed as a recommendation or a solicitation to trade in such securities/instruments. MSTL may not execute transactions for clients in these securities/instruments.

Morgan Stanley is not incorporated under PRC law and the research in relation to this report is conducted outside the PRC. Morgan Stanley Research does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC. PRC investors shall have the relevant qualifications to invest in such securities and shall be responsible for obtaining all relevant approvals, licenses, verifications and/or registrations from the relevant governmental authorities themselves. Neither this report nor any part of it is intended as, or shall constitute, provision of any consultancy or advisory service of securities investment as defined under PRC law. Such information is provided for your reference only.

Morgan Stanley Research is disseminated in Brazil by Morgan Stanley C.T.V.M. S.A. located at Av. Brigadeiro Faria Lima, 3600, 6th floor, São Paulo - SP, Brazil; and is regulated by the Comissão de Valores Mobiliários; in Mexico by Morgan Stanley México, Casa de Bolsa, S.A. de C.V. which is regulated by Comisión Nacional Bancaria y de Valores. Paseo de los Tamarindos 90, Torre 1, Col. Bosques de las Lomas Floor 29, 05120 Mexico City; in Japan by Morgan Stanley MUFG Securities Co., Ltd. and, for Commodities related research reports only, Morgan Stanley Capital Group Japan Co., Ltd; in Hong Kong by Morgan Stanley Asia Limited (which accepts responsibility for its contents) and by Morgan Stanley Bank Asia Limited; in Singapore by Morgan Stanley Asia (Singapore) Pte. (Registration number 199206298Z) and/or Morgan Stanley Asia (Singapore) Securities Pte Ltd (Registration number 200008434H), regulated by the Monetary Authority of Singapore (which accepts legal responsibility for its contents and should be contacted with respect to any matters arising from, or in connection with, Morgan Stanley Research) and by Morgan Stanley Bank Asia Limited, Singapore Branch (Registration number T14FC0118); in Australia to "wholesale clients" within the meaning of the Australian Corporations Act by Morgan Stanley Australia Limited A.B.N. 67 003 734 576, holder of Australian financial services license No. 233742, which accepts responsibility for its contents; in Australia to "wholesale clients" and "retail clients" within the meaning of the Australian Corporations Act by Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813, which accepts responsibility for its contents; in Korea by Morgan Stanley & Co International plc, Seoul Branch; in India by Morgan Stanley India Company Private Limited having Corporate Identification No (CIN) U22990MH1998PTC115305, regulated by the Securities and Exchange Board of India ("SEBI") and holder of licenses as a Research Analyst (SEBI Registration No. INH000001105); Stock Broker (SEBI Stock Broker Registration No. INZ000244438), Merchant Banker (SEBI Registration No. INM000011203), and depository participant with National Securities Depository Limited (SEBI Registration No. IN-DP-NSDL-567-2021) having registered office at Altimus, Level 39 & 40, Pandurang Budhkar Marg, Worli, Mumbai 400018, India; Telephone no. +91-22-61181000; Compliance Officer Details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: tejarshi.hardas@morganstanley.com; Grievance officer details: Mr. Tejarshi Hardas, Tel. No.: +91-22-61181000 or Email: msic-compliance@morganstanley.com. Morgan Stanley India Company Private Limited (MSICPL) may use AI tools in providing research services. All recommendations contained herein are made by the duly qualified research analysts; in Canada by Morgan Stanley Canada Limited; in Germany and the European Economic Area where required by Morgan Stanley Europe S.E., authorised and regulated by Bundesanstalt fuer Finanzdienstleistungsaufsicht (BaFin) under the reference number 149169; in the US by Morgan Stanley & Co. LLC, which accepts responsibility for its contents. Morgan Stanley & Co. International plc, authorized by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority, disseminates in the UK research that it has prepared, and research which has been prepared by any of its affiliates, only to persons who (i) are investment professionals falling within Article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2005 (as amended, the "Order"); (ii) are persons who are high net worth entities falling within Article 49(2)(a) to (d) of the Order; or (iii) are persons to whom an invitation or inducement to engage in investment activity (within the meaning of section 21 of the Financial Services and Markets Act 2000, as amended) may otherwise lawfully be communicated or caused to be communicated. RMB Morgan Stanley Proprietary Limited is a member of the JSE Limited and A2X (Pty) Ltd. RMB Morgan Stanley Proprietary Limited is a joint venture owned equally by Morgan Stanley International Holdings Inc. and RMB Investment Advisory (Proprietary) Limited, which is wholly owned by FirstRand Limited. The information in Morgan Stanley Research is being disseminated by Morgan Stanley Saudi Arabia, regulated by the Capital Market Authority in the Kingdom of Saudi Arabia, and is directed at Sophisticated investors only.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (DIFC Branch), regulated by the Dubai Financial Services Authority (the DFSA)

or by Morgan Stanley & Co. International plc (ADGM Branch), regulated by the Financial Services Regulatory Authority Abu Dhabi (the FSRA), and is directed at Professional Clients only, as defined by the DFSA or the FSRA, respectively. The financial products or financial services to which this research relates will only be made available to a customer who we are satisfied meets the regulatory criteria of a Professional Client. A distribution of the different MS Research ratings or recommendations, in percentage terms for Investments in each sector covered, is available upon request from your sales representative.

The information in Morgan Stanley Research is being communicated by Morgan Stanley & Co. International plc (QFC Branch), regulated by the Qatar Financial Centre Regulatory Authority (the QFCRA), and is directed at business customers and market counterparties only and is not intended for Retail Customers as defined by the QFCRA.

As required by the Capital Markets Board of Turkey, investment information, comments and recommendations stated here, are not within the scope of investment advisory activity. Investment advisory service is provided exclusively to persons based on their risk and income preferences by the authorized firms. Comments and recommendations stated here are general in nature. These opinions may not fit to your financial status, risk and return preferences. For this reason, to make an investment decision by relying solely to this information stated here may not bring about outcomes that fit your expectations.

The trademarks and service marks contained in Morgan Stanley Research are the property of their respective owners. Third-party data providers make no warranties or representations relating to the accuracy, completeness, or timeliness of the data they provide and shall not have liability for any damages relating to such data. The Global Industry Classification Standard (GICS) was developed by and is the exclusive property of MSCI and S&P.

Morgan Stanley Research, or any portion thereof may not be reprinted, sold or redistributed without the written consent of Morgan Stanley.

Indicators and trackers referenced in Morgan Stanley Research may not be used as, or treated as, a benchmark under Regulation EU 2016/1011, or any other similar framework.

The issuers and/or fixed income products recommended or discussed in certain fixed income research reports may not be continuously followed. Accordingly, investors should regard those fixed income research reports as providing stand-alone analysis and should not expect continuing analysis or additional reports relating to such issuers and/or individual fixed income products. Morgan Stanley may hold, from time to time, material financial and commercial interests regarding the company subject to the Research report.

Registration granted by SEBI and certification from the National Institute of Securities Markets (NISM) in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Investment in securities market are subject to market risks. Read all the related documents carefully before investing.

INDUSTRY COVERAGE: Software

COMPANY (TICKER)	RATING (AS OF)	PRICE* (06/05/2025)
Chris Quintero		
BILL Holdings Inc (BILL.N)	O (01/16/2025)	\$45.68
Blackline Inc (BL.O)	O (09/29/2024)	\$58.22
E2open Parent Holdings Inc (ETWO.N)	E (01/17/2024)	\$3.23
OneStream Inc (OS.O)	O (01/16/2025)	\$29.07
Vertex Inc. (VERX.O)	O (01/17/2024)	\$41.87
Elizabeth Porter, CFA		
Amplitude Inc. (AMPL.O)	E (01/13/2025)	\$12.97
Autodesk (ADSK.O)	O (08/23/2024)	\$298.22
Freshworks Inc (FRSH.O)	E (10/18/2021)	\$15.61
GoDaddy Inc (GDDY.N)	E (07/19/2021)	\$180.39
HubSpot, Inc. (HUBS.N)	O (03/21/2023)	\$602.61
Klaviyo, Inc (KVYO.N)	E (10/16/2023)	\$33.95
LegalZoom.com Inc (LZ.O)	U (07/28/2022)	\$9.20
Liveramp Holdings Inc (RAMP.N)	E (01/13/2025)	\$33.12
Semrush Holdings Inc -A (SEMR.N)	O (01/13/2025)	\$10.11
Sprinklr Inc (CXM.N)	E (07/19/2021)	\$8.85
Sprout Social Inc (SPT.O)	E (11/17/2020)	\$21.86
Wix.Com Ltd (WIX.O)	O (01/13/2025)	\$153.25
Zeta Global Holdings Corp (ZETA.N)	E (08/01/2024)	\$13.64
ZoomInfo Technologies Inc (GTM.O)	E (02/01/2024)	\$9.97
Josh Baer, CFA		
Asana Inc (ASAN.N)	U (05/20/2025)	\$14.72
BigCommerce Holdings, Inc. (BIGC.O)	E (05/11/2021)	\$5.12
Box Inc (BOX.N)	E (05/21/2024)	\$37.93
CCC Intelligent Solutions Holdings Inc (CCCS.O)	O (11/13/2024)	\$8.92
Coursera, Inc. (COUR.N)	E (06/05/2025)	\$8.86
DigitalOcean Holdings Inc (DOCN.N)	O (01/16/2025)	\$27.77
Docebo Inc. (DCBO.O)	E (05/12/2025)	\$26.85
DocuSign Inc (DOCU.O)	E (01/16/2024)	\$92.90
Lightspeed POS Inc. (LSPD.N)	E (02/18/2021)	\$11.34

Sabre Corp (SABR.O)	E (03/16/2021)	\$2.70
ServiceTitan Inc (TTAN.O)	E (01/06/2025)	\$114.55
Toast, Inc. (TOST.N)	O (12/16/2021)	\$42.82
Udemy Inc (UDMY.O)	E (06/05/2025)	\$7.55
Keith Weiss, CFA		
Adobe Inc. (ADBE.O)	O (07/31/2023)	\$415.20
Akamai Technologies, Inc. (AKAM.O)	E (04/29/2020)	\$76.14
Atlassian Corporation PLC (TEAM.O)	O (01/13/2020)	\$218.60
Check Point Software Technologies Ltd. (CHKP.O)	E (10/16/2023)	\$231.11
Cloudflare Inc (NET.N)	O (12/02/2024)	\$176.33
CoreWeave (CRWV.O)	E (04/22/2025)	\$135.05
CrowdStrike Holdings Inc (CRWD.O)	O (01/09/2024)	\$462.94
CyberArk Software Ltd (CYBR.O)	E (01/09/2024)	\$392.18
Fortinet Inc. (FTNT.O)	O (10/07/2022)	\$103.50
Gen Digital Inc. (GEN.O)	E (06/07/2024)	\$29.13
Intuit (INTU.O)	O (02/26/2025)	\$766.64
Jamf Holding Corp (JAMF.O)	E (10/13/2024)	\$10.54
Microsoft (MSFT.O)	O (01/13/2016)	\$467.68
Okta, Inc. (OKTA.O)	O (12/02/2024)	\$104.18
Oracle Corporation (ORCL.N)	E (01/15/2019)	\$171.14
Palo Alto Networks Inc (PANW.O)	O (10/10/2017)	\$197.11
Qualys Inc (QLYS.O)	U (02/09/2021)	\$139.70
Rapid7 Inc (RPD.O)	E (08/11/2015)	\$23.17
SailPoint Inc (SAIL.O)	E (03/10/2025)	\$18.02
Salesforce, Inc. (CRM.N)	O (12/21/2023)	\$267.14
Samsara Inc (IOT.N)	E (03/23/2023)	\$47.25
SentinelOne, Inc. (S.N)	E (12/02/2024)	\$18.12
ServiceNow Inc (NOW.N)	E (10/21/2024)	\$1,017.60
Shopify Inc (SHOP.O)	O (04/19/2024)	\$105.03
Snowflake Inc. (SNOW.N)	E (02/29/2024)	\$210.10
Tenable Holdings Inc (TENB.O)	E (12/02/2024)	\$32.56
Varonis Systems, Inc. (VRNS.O)	O (10/16/2023)	\$50.14
Workday Inc (WDAY.O)	E (02/19/2025)	\$250.10
Zscaler Inc (ZS.O)	E (01/12/2023)	\$300.88
Meta A Marshall		
8x8 Inc (EGHT.O)	U (06/14/2024)	\$1.77
Five9 Inc (FIVN.O)	E (10/10/2022)	\$28.55
NICE Ltd. (NICE.O)	O (10/16/2023)	\$175.96
RingCentral Inc (RNG.N)	E (08/08/2023)	\$26.80
Twilio Inc (TWLO.N)	O (02/24/2025)	\$120.80
Zoom Video Communications Inc (ZM.O)	E (10/11/2022)	\$81.33
Roy D Campbell		
Karooooo Ltd (KARO.O)	O (04/27/2021)	\$57.86
Sanjit K Singh		
Appian Corp (APPN.O)	E (12/03/2021)	\$31.80
C3.ai (AI.N)	U (01/04/2021)	\$25.34
Confluent, Inc. (CFLT.O)	E (01/16/2025)	\$24.18
Couchbase, Inc. (BASE.O)	E (08/16/2021)	\$19.46
Datadog, Inc. (DDOG.O)	E (01/16/2025)	\$121.75
Dynatrace Inc (DT.N)	E (02/13/2024)	\$54.94
Elastic NV (ESTC.N)	O (12/16/2024)	\$86.29
GitLab Inc (GTLB.O)	O (10/09/2024)	\$48.77
JFrog Ltd. (FROG.O)	O (12/21/2023)	\$43.18
MongoDB Inc (MDB.O)	O (04/12/2023)	\$225.38

PagerDuty, Inc. (PD.N)	E (01/24/2024)	\$15.38
Palantir Technologies Inc. (PLTR.O)	E (02/04/2025)	\$119.91
UiPath Inc (PATH.N)	E (09/07/2022)	\$13.26

Stock Ratings are subject to change. Please see latest research for each company.

* Historical prices are not split adjusted.

INDUSTRY COVERAGE: Internet

COMPANY (TICKER)	RATING (AS OF)	PRICE* (06/05/2025)
Brian Nowak, CFA		
Airbnb Inc (ABNB.O)	U (12/06/2022)	\$137.29
Alphabet Inc. (GOOGL.O)	O (08/11/2015)	\$168.21
Amazon.com Inc (AMZN.O)	O (04/24/2015)	\$207.91
Booking Holdings Inc (BKNG.O)	E (01/09/2019)	\$5,570.52
DoorDash Inc (DASH.O)	O (02/21/2024)	\$215.84
Expedia Inc. (EXPE.O)	E (01/09/2019)	\$173.38
Instacart (CART.O)	E (01/29/2024)	\$46.24
Lyft Inc (LYFT.O)	E (10/24/2019)	\$15.53
Meta Platforms Inc (META.O)	O (03/20/2023)	\$684.62
Nextdoor Holdings Inc (KIND.O)	U (05/12/2025)	\$1.61
Pinterest Inc (PINS.N)	E (03/28/2022)	\$33.76
Reddit Inc (RDDT.N)	O (12/08/2024)	\$112.24
Snap Inc. (SNAP.N)	E (07/22/2024)	\$8.20
Uber Technologies Inc (UBER.N)	O (06/04/2019)	\$84.67
Matthew Cost		
AppLovin Corp (APP.O)	O (04/10/2025)	\$414.14
Compass, Inc. (COMP.N)	E (07/20/2022)	\$6.16
Criteo SA (CRO.O)	E (01/26/2016)	\$26.28
DoubleVerify Holdings Inc (DV.N)	E (06/25/2024)	\$15.03
Electronic Arts Inc (EA.O)	E (08/04/2021)	\$147.88
Integral Ad Science Holding Corp. (IAS.O)	E (04/16/2024)	\$8.24
Opendoor Technologies Inc (OPEN.O)	E (07/24/2023)	\$0.66
Playtika Holding Corp (PLTK.O)	E (11/27/2022)	\$4.88
Roblox Corporation (RBLX.N)	O (11/04/2024)	\$94.20
Shutterstock Inc (SSTK.N)	E (07/28/2022)	\$17.86
Take-Two Interactive Software (TTWO.O)	O (02/01/2018)	\$231.03
Trade Desk Inc (TTD.O)	O (06/01/2023)	\$71.11
Unity Software Inc (U.N)	O (09/02/2024)	\$24.79
Webtoon Entertainment Inc (WBTN.O)	E (07/22/2024)	\$8.78
Yelp Inc (YELP.N)	U (01/10/2019)	\$36.80
Zillow Group Inc (Z.O)	E (04/18/2018)	\$69.80
Nathan Feather		
Bumble Inc. (BMBL.O)	E (03/08/2021)	\$5.35
Chewy Inc (CHWY.N)	O (10/31/2023)	\$47.49
Duolingo Inc (DUOL.O)	O (04/23/2025)	\$522.99
eBay Inc (EBAY.O)	O (04/18/2024)	\$77.74
Etsy Inc (ETSY.O)	U (04/18/2024)	\$62.72
FIGS, Inc. (FIGS.N)	E (02/29/2024)	\$5.09
Match Group Inc (MTCH.O)	E (04/18/2024)	\$31.50
Peloton Interactive, Inc. (PTON.O)	E (03/14/2022)	\$6.94
Revolve Group Inc (RVLV.N)	E (10/20/2024)	\$22.57
WW International Inc (WGHTQ.PK)	NR (05/20/2025)	\$0.27

Stock Ratings are subject to change. Please see latest research for each company.

* Historical prices are not split adjusted.

© 2025 Morgan Stanley