

Shaping the future

The transformative potential of agentic AI and the strategic imperative for Google Cloud partners




Table of contents

Foreword		03
Executive summary		04
01	Understanding agentic AI: The next frontier	06
02	Solving real-world problems: The growing demand for agentic AI	09
03	The opportunity: Sizing the agentic AI market	13
04	The strategic imperative for partners: How services will evolve in an agentic world	19
05	Agentic AI in action: How partners can help transform key industries	26
06	Partnering with Google Cloud: Co-creating the agentic future	56
Conclusion		60
Appendices		61


Foreword

The emergence of agentic artificial intelligence (AI) marks a pivotal moment in technology, promising to reshape industries and redefine how businesses operate and create value. For Google Cloud partners, this represents not just a new technological wave but a profound opportunity to lead customers into a more intelligent, automated, and productive future. This report is intended for executive leaders within IT service organizations, providing actionable insights and a strategic framework to guide the development of an agentic AI strategy for 2025 and beyond.

The findings and perspectives shared within these pages are the culmination of rigorous, bottom-up analysis, where we examined the industries, roles, and tasks that stand to most benefit from agentic AI. This quantitative research was supplemented by findings from in-depth interviews with Google Cloud's own leaders and IT service partners at the forefront of AI development. These conversations covered topics such as sources of customer demand, emergent use cases, adoption barriers, and the opportunities for service partners to evolve their commercial approach. This mix of data-driven analysis and real-world market commentary yielded a focused and actionable view of agentic AI opportunities that will emerge going forward.


This report shows that agentic AI opens up two significant layers of opportunity—it elevates the value of traditional implementation engagements, and for partners ready to broaden their remit, it creates a clear runway to evolve into broader transformation partners. This report will show how partners can deepen customer relationships by helping them solve their most fundamental business challenges, alleviate long-standing industry pain points, and unlock new, sustainable revenue streams. We believe that by understanding the landscape, identifying high-potential use cases, and developing targeted capabilities, partners can achieve remarkable growth and differentiation.

Consider this report a comprehensive guide to navigating the exciting and dynamic world of agentic AI. Our goal is to spotlight the emerging demand signals for agentic AI and guide partners towards the pain points that will drive the most value for customers across industries.



Kevin Ichhpurani

President, Global Partner Ecosystem &
Channels, Google Cloud



Executive summary



Agentic AI represents a critical leap forward in AI advancement, offering novel capabilities to tackle deeply entrenched industry pain points that enterprises have faced for years. This wave of innovation is not just about improving productivity—agents will enable a fundamental reimagination of enterprise workflows, making businesses faster, more adaptive, and better attuned to customer needs.

In an ‘agentic world,’ AI agents will therefore unlock value in nearly every industry and workflow—whether it is streamlining error-prone claims adjudication in Insurance, proactively managing inventory misalignments in Industrials, delivering hyper-personalized experiences at scale in Retail, or accelerating analytical processes during pharmaceutical drug discovery. As a result, agentic AI can structurally redefine how enterprises generate value and compete.

Enterprises are already seeing the transformative potential of agentic AI, with several proof points quickly emerging. Today, more than 90% of enterprises report interest in deploying agentic AI solutions within the next three years, according to BCG’s most recent IT Buyer Pulse Check, and software companies’ public filings are mentioning agentic AI 12 times more today compared to a year ago. This trend is further underscored by the fact that the largest enterprise software vendors are treating agentic AI as a strategic priority, as evidenced by recent product launches and strategic acquisitions to support their agentic AI ambitions.

Given this context, the Partner Ecosystem has an immense opportunity to help enterprises realize the value that agents can bring. In total, agentic AI represents a potential ~\$1T market opportunity for agentic AI services globally, with \$350B to \$450B in the U.S. alone—a figure that surpasses the estimated \$320B U.S. spend on traditional services today—and if agentic AI follows a similar trajectory to cloud computing, we expect \$300B to \$600B of the global services opportunity to be realized by 2035-2040, though likely a conservative timeline given the current momentum and pace of adoption. Within this landscape, distinct value pools are fast emerging where partners can deliver outsized impact through agentic AI deployment. Therefore, partners must act quickly to capture the significant opportunity ahead.


Partners must also embrace their increasingly important role as long-term partners to enterprises, driving initiatives that closely align with ‘top-of-the-house’ business objectives. To that end, more onus will be placed on them to fundamentally redesign traditional business processes around agents, and as agentic deployments scale, partners will be expected to deliver new forms of ongoing support, including agentic fleet orchestration or agent evaluation/measurement. This evolution in responsibility will open up new opportunities for partners to explore alternative monetization models, including transaction-based, outcome-based, and SaaS-like recurring fee structures.

As a continued signal of a partner-first approach, Google Cloud will equip partners with resources to capitalize on this opportunity, providing differentiated market intelligence, access to our enterprise customer base, and a robust ecosystem of cutting-edge agentic tools. To that end, Google has already introduced a number of capabilities enabling the partner ecosystem to jump-start the agentic AI transformation journey for customers. At the beginning of the journey, customers lack awareness and want to see value early. Partners can run discovery workshops with Google in the room to help their customers evaluate agent use cases, as well as quickly create agent prototypes for pitches using Vertex AI and its family of cutting-edge multi-modal Gemini models.

During the build stage, many enterprises worry if their data pipeline is sufficiently ready. Partners can use BigQuery’s AI tools and modern connectors across systems to quickly activate agents with less structured data than needed previously. Finally, as agentic solutions scale, driving user adoption in a productive and sustainable manner is key. Partners can help establish the right change management and agent lifecycle habits for their clients, leveraging training and domain expertise available within the Google Cloud ecosystem.

These resources serve as a powerful starting point, further enhanced by industry-leading AI research from DeepMind, world-class TPU infrastructure, and a strong commitment to open-source collaboration through initiatives like the Agent Development Kit (ADK) and the Agent2Agent (A2A) protocol. This holistic approach, combined with co-innovation programs and partner support, is designed to fuel your success.

The future of agentic AI will be shaped by those who lead; partners who act decisively now will not just participate in this evolution—they will define the category.



01

Understanding agentic AI: The next frontier



Agentic AI unlocks the full power of AI

Recent breakthroughs in artificial intelligence (AI) have unlocked powerful new capabilities. Agentic AI represents the next step in applying these advances, creating systems that can autonomously reason, make decisions, and take action to achieve specific goals. By leveraging the latest AI models, agentic systems can go beyond simple task execution, like text generation, to independently manage complex processes and interact directly with business software.

This makes agentic AI particularly suited for real-world business challenges. An AI agent can monitor its environment, understand changing contexts, create plans, and execute them across different digital systems—all without constant human supervision. This allows them to handle complex, multi-step workflows and deliver tangible business outcomes in a way that previous automation and AI could not.

True paradigm shift, not just Automation 2.0

Various debates persist about whether agents are simply a new form of automation or something much greater. The facts offer clear support for the latter.

Enterprises experienced benefits from the first ‘evolution’ of business process automation, commonly known today as Robotic Process Automation (RPA). RPA bots gave enterprises a way to automate highly structured, repetitive tasks within stable environments. They followed predefined

scripts and often struggled when they were faced with changes or unexpected conditions.

Separately, the rise of generative AI (gen AI) introduced a powerful tool for unstructured, creative tasks. Its ability to perform with human-like proficiency drove rapid, widespread adoption; according to Google Cloud’s 2025 State of AI Infrastructure report, 98% of organizations are using or experimenting with it. However, using gen AI alone to transform core business processes and deliver clear ROI has proven challenging.

Agentic AI overcomes these limitations by combining the reasoning power of gen AI with the ability to take action in business systems. Agents can handle complex, multi-step workflows, learn from outcomes, and adapt to new information—all autonomously. This allows enterprises to move beyond simple task automation and begin redesigning core processes.

This step-change in capability has meaningful implications for how enterprises run. Taking a real-world example in the context of insurance claims adjudication, RPA has long struggled with exceptions, such as interpreting unstructured claims descriptions or dealing with ambiguity generated by incomplete documentation, often escalating cases when exact conditions aren’t met. Agentic AI transforms this. Agents can utilize gen AI to extract relevant details from multimodal inputs, then validate them against policy rules and member history, and retrieve supporting documents without relying on rigid logic. It moves beyond simple auto-adjudication by proactively coordinating with other systems to gather missing documents and assess risk from multimodal data. An agent can then

orchestrate the full, end-to-end resolution, from autonomously processing a wider scope of claims to initiating payment, significantly improving straight-through processing rates and freeing up specialists for high-judgment work.

Alternatively, consider personalization in marketing. Gen AI excels at creating personalized content elements (e.g., unique emails and ad call-to-actions or even distinct AI-rendered images to suit individual preferences). However, gen AI struggles to orchestrate entire individual customer journeys or dynamically adapt in real-time without marketing teams manually defining intricate journey rules and campaign parameters. Agentic AI can harness gen AI capabilities to autonomously manage hyper-personalization. It builds and refines individual profiles from online patterns, purchases data, and customer-provided data, uses gen AI for unique “N-of-1” content creation, and crucially, determines the optimal timing, sequence, and channels for interactions across the full journey. This enables truly adaptive, optimized marketing experiences at scale.


Examples like these abound, and they will continue to proliferate as agentic capabilities improve, allowing enterprises to fundamentally reimagine their business processes.

Examples	The ‘old’ world	The agentic world
Marketing journeys	Manual A/B testing campaigns with broad audiences and tweaked with imprecise results	Analyze real-time data to proactively create ‘N-of-1’ experiences, optimize ad campaigns, and forecast trends to maximize customer lifetime value
Go-to-market	Create outreach email template based on prompts, and hand off to sales representative to send to prospective customer	Drive full sales motion autonomously with creative discovery, tailored outreach, winning pitches, & tactful follow-ups to accelerate deal velocity
Customer service	Handle basic customer inquiries (the typical ~80%), using basic, pre-defined script responses	Agents anticipate customer needs, resolve complex issues in real-time, and build lasting customer loyalty at every service interaction
Supply chain	Planners use historical data and static rules to forecast demand, often leading to misaligned inventory and costly manual exception handling	A system of agents ingests real-time signals to forecast demand, autonomously triggers replenishment, and reroutes shipments to prevent disruptions

Exhibit 1.1: Agentic AI builds on RPA and gen AI to take actions autonomously to achieve business outcomes

02

Solving real-world problems: The growing demand for agentic AI



Unlocking the industry's deeply entrenched challenges

The power of the new agentic AI paradigm lies in its potential to help enterprises solve macro industry challenges—the so-called ‘trillion-dollar problems’ that exist for each industry.

In Retail, these pain points include rising input costs that threaten margin or large-scale customer service inefficiencies. In Healthcare, it's staffing shortages in clinics or persistent issues in patient billing. In Financial Services, it's error-prone and delayed insurance claims and Know Your Customer (KYC) processes, driving cost leakage and elevated compliance risks. The exhibit below crystallizes that several of these macro pain points exist, and that each industry faces its own unique concoction of these challenges.

Role of agents: Industries have grappled with long-standing pain points that have hindered enterprises from maximizing growth and productivity potential

Examples of key industry pain points

Non-exhaustive

Tech/TMEG	Financial Services/ Insurance	Auto/Industrial	Public Sector	Healthcare/ Life Sciences	Retail/CPG
Slow content creation Manual content creation reduces speed to market, personalization	Delayed, error-prone claims Manual processes and siloed data slow processing in insurance claims	Disruptive equipment failures Maintenance is reactive, issues identified post-failure	Delays in benefit delivery Eligibility checks require fragmented data & manual review	Chronic staffing shortages Shortages of medical, nursing, and admin. staff create delays and care coordination gaps	Inefficient customer service Support often inconsistent/delayed and causes churn
Media content misuse Limited rights oversight drives compliance & revenue gaps	Risky, manual KYC checks Inconsistent identity reviews cause compliance risk	Inadequate safety oversight Reactive safety monitoring increases incident risk	Permit delays stall activity Requires coordination across teams and manual validation	Error-prone billing/claims Manual billing processes lead to denials, delays, admin burden	Disjointed cust. experience Evolving customer expectations and fragmented experiences reduces loyalty and revenue
Prolonged service outages Manual triage of telco network events delays resolution	Slow, fragmented reporting Data fragmentation drives slow, error-prone compliance	Suboptimal load planning Static truckload & route plans miss consolidation opportunities	Delayed policy response Agencies lack capacity to analyze, act on new legislation	Slow, error-prone diagnoses Manual diagnostics prone to delays and inconsistency	Rising cost of inputs Prices not optimized to protect revenue under cost pressure
Slow creation of narratives Player immersion limited with few branching narratives	Generic wealth advice Static advice misses client needs without personalization	Inventory misalignment Manual allocation decisions lag real-time needs	Limited fraud oversight Fragmented data and manual review delay detection	Manual R&D slows discovery Documentation-heavy R&D slows drug discovery	Misaligned/unsold inventory Product inventory fluctuates frequently and unpredictably
Unoptimized pricing yield Fixed prices inflexible to real-time demand/availability	Slow loan underwriting Fragmented data slows approvals and limits access	Unpredictable lead sourcing Sourcing leads in B2B requires high effort, with unclear ROI	Slow dissemination of info Urgent public updates slowed by manual communication workflows	Inefficiencies in clinical trials Slow site startup delays clinical trials	Slow personalized marketing High cost and slow turnaround for highly personalized marketing

Note: Based on interviews with partners and Google industry solution leaders; Source: Partner and Google interviews, BCG analysis

Exhibit 2.1: Many long-standing pain points exist across industries

To address these pain points, enterprises can now use a constellation of agents that are tailored to industry-specific and business-specific needs. In Consumer Packaged Goods (CPG), an inventory planning AI agent could continuously forecast demand using live structured and unstructured data, and then autonomously orchestrate responses like triggering replenishment, rerouting shipments, or adjusting warehouse throughput requirements in real-time to prevent disruptions. In Healthcare, a care coordination agent can dynamically orchestrate personalized care plans, ensuring patient follow-ups are scheduled,

Role of agents: But in our ‘new world,’ AI agents can solve these pain points to unlock transformational value pools across industries

Agentic solutions to key industry pain points

Cross-industry use cases Non-exhaustive

Tech/TMEG	Financial Services/ Insurance	Auto/Industrial	Public Sector	Healthcare/ Life Sciences	Retail/CPG
Marketing/personalization agent Generates personalized content from analyzing audience behavior	Claims adjudication agent Automates claims validation, fraud checks, and payouts	Predictive maint. agent Predicts equipment failures, triggers maintenance tasks	Benefits processing agent Verifies eligibility and approves benefits across systems	Care coordination agent Coordinates full patient journey incl. care plans, readmission, etc.	Customer serv./ returns agent Resolves service issues and processes refunds/returns
Rights compliance agent Monitors content usage, flags violations, flags to legal team	KYC verification agent Verifies identity and flags risks for onboarding	Safety compliance agent Monitors safety data and initiates risk response	Permitting & licensing agent Validates applications, auto-issues licenses and permits	Rev. cycle mgmt. agent Automates key revenue cycle tasks – validates claims, manages submission, denials, and follow-ups	Customer experience agent Recommends personalized offers, support across customer journey
Network operations agent Detects and resolves outages by analyzing telemetry data, etc.	Reporting/compliance agent Prepares and submits regulatory reports	Logistics optimization agent Plans truck loads and routes alongside supply chain agent	Legislative analysis agent Identifies impact of new legislation, routes to departments	Medical diagnostics agent Analyzes data/imaging to generate diagnostic suggestions	Pricing analyst agent Optimizes prices based on market trends and demands
Narrative design agent Generates branching storylines, dialogue trees, character arcs	Wealth advisor agent Generates tailored financial plans + monitors portfolios	Inventory allocation agent Rebalances stock across sites using demand/supply signals	Fraud detection agent Monitors transactions, etc., to detect and act on potential fraud	Drug discovery R&D agent Finds targets and molecules from scientific data	Inventory planning agent Forecasts demand and triggers replenishment
Dynamic pricing agent Adjusts pricing based on demand, time, capacity, etc.	Loans underwriter agent Assesses borrow risk and issues approvals	B2B sales agent Sources leads based on sales goals and propensity to buy	Communications agent Drafts messages, routes to approval teams, publishes	Clinical trial ops agent Onboards sites and enrolls eligible patients	Marketing campaign agent Generates and publishes creatives, refines according to performance

Note: Based on interviews with partners and Google industry solution leaders; Source: Partner and Google interviews, BCG analysis

Exhibit 2.2: Agentic AI solutions can solve long-standing pain points across all industries

executed, and adapted based on incoming patient data, alleviating pressure from chronic staffing shortages. In Financial Services, a KYC verification agent can intelligently interpret a wide array of unstructured information to auto-complete the entire document verification process. Meanwhile, a loan issuance agent can move beyond simple checklist approvals, managing complex risk identification workflows to execute loan approval decisions autonomously with holistic judgment. Such examples exist in every industry, where agentic AI can offer a powerful new toolkit to address previously intractable pain points.

As a result, the agentic era will reshape how enterprises solve their core problems, augmenting human capabilities to drive substantial productivity and customer value. In

doing so, agents will enable human workers to shift their focus from routine execution to higher-value strategic tasks, and allow enterprises to deliver a new wave of improvement in both top-line performance (through faster speed-to-market and better customer engagement), and bottom-line efficiency (through routine task automation, optimized resource allocation, and proactive risk detection).

Agentic AI demand signals

Enterprises are already acknowledging the transformative potential of agentic AI to solve long-standing industry pain points and have started to take action. More than 90% of enterprises plan to incorporate agentic AI in the next three years, with the highest

interest currently in customer service and a broad set of analytics use cases (e.g., business intelligence, advanced predictive analytics), according to the most recent BCG IT Buyer Pulse Check.

Recognizing the emerging demand, leading software companies and investors are driving both public mindshare and capital investment in their agentic AI strategies. Leading SaaS vendors are investing significantly via different channels: launching pre-built agent solutions and platforms (e.g., Salesforce’s introduction of Agentforce), making strategic acquisitions (e.g., Informatica, Moveworks), and collaborating closely with others to build out the agentic ecosystem (e.g., Workday Agent Partner Network). These signals provide a strong indication of secular growth for agentic AI deployments going forward, creating a prime environment for Google Cloud’s partners to bring enterprise agentic AI implementations to market.

Software companies and investors are accelerating public mindshare and capital deployment around agentic AI




Exhibit 2.4: Software companies and investors are doubling down on agentic AI

Enterprise interest in agentic AI is already substantial...

Q: Which use cases of agentic AI is your company most interested in implementing? (N=565, BCG IT Buyers Pulse Check, Dec 2024)





Exhibit 2.3: More than 90% enterprises plan to incorporate agentic AI in the next 3 years

03

The opportunity: Sizing the agentic AI market



Agentic AI represents a ~\$1T market globally for partner services

The advent of agentic AI is not just a technological milestone; it's poised to create substantial new value across the global economy. For Google Cloud's partners, this translates into a significant market opportunity.

The total value pool for services in agentic AI (advisory, implementation, and maintenance) is estimated to reach \$350B to \$450B in the U.S. and ~\$1T globally in a 'full potential' scenario, assuming complete adoption.

Moreover, the full application layer value pool (the value of software accrued to enterprise agentic AI software products) is estimated at ~\$3T. This highlights the significant upcoming demand for specialized advisory, implementation, and maintenance services as agentic AI adoption grows across various industries and job functions.

Our methodology for sizing agentic AI value pools

To better understand the scale of this opportunity, we developed a detailed bottom-up model, with our primary aim to provide a clear perspective on the relative size of agentic opportunities across roles and industries (rather than to offer absolute forecasts of adoption timing).

Guiding principles – what our AI value pool sizing is:

- A 'full potential' scenario for agentic AI assuming complete adoption and value delivery.
- Assumes that AI agents successfully augment all tasks deemed addressable.
 - Embeds conservatism by assuming tasks are unlikely to be fully agentic, driven by the need for higher-level strategic thinking and/or exception handling.
 - Filters out companies with fewer than 500 FTEs, which System Integrators (SIs) are unlikely to serve.




Exhibit 3.1: Estimated global SI services value pool for agentic AI, by region and verticals

Guiding principles – what our AI value pool sizing is not:

- An estimate of any specific timeline on agentic AI adoption.
- A projection based on any specific SI monetization model or pricing level.

As such, this model quantifies the potential for agentic AI across approximately 330 task categories. Building from this foundation of addressable tasks, we then estimate the associated application layer market based on the value delivered to enterprises and, crucially, the associated SI services opportunity across industries. Please see the appendix for a detailed view of our methodology.

The potential timing for adoption

While the value pools discussed earlier outline a long-run ‘full potential’ value (expressed in 2025 dollars), we recognize that the actual trajectory and timeframes for realizing this potential are subject to considerable uncertainty.

The evolution of the cloud market provides a relevant—albeit imperfect—analogue for how the agentic AI market may play out, having also represented a structural technology shift for enterprises. Forrester, IDC, and Gartner suggest cloud spend has grown from ~\$20B in 2010 to >\$400B in real terms, growing at ~25% p.a. over a timeframe of ~15 years, and even faster at 30% p.a. between 2010 and 2015.




Exhibit 3.3: Historical growth in cloud revenue

While estimates on current agentic spend are unclear, Gartner¹ estimates generative AI software spend to be ~\$19B in 2024. In all likelihood, true agentic software spend today is lower than this, given its recency, but will see a swift acceleration to this baseline given its current infancy against the substantial enterprise traction.

As a base case, if you believe that the agentic market today can unfold at the same rate that cloud did from between 2010 and 2024, \$300B to \$600B (in 2025 dollars) of our total ~\$1T agentic services value pool estimate is achievable by 2035-2040. While this bracket is not a true point forecast, it does signal the magnitude of opportunity agentic AI brings, matching, if not exceeding, the current total U.S. SI services spend within 10 to 15 years of adoption.



This was the quarter of Agentforce... we have never seen products grow at these levels... in just a few months, we have seen this addressable market go to hundreds of billions of dollars.”

Salesforce CEO, Q4FY25 Earnings Conference Call



When cloud started, adoption took longer due to the size of transformation and upskilling needed. This infrastructure is already in place for agentic and has been fast-tracked by COVID.”

Google Cloud Partner

Additionally, the speed of AI momentum seen in the market suggests agentic adoption could follow a faster timeline than cloud, further pulling forward the timeline for partners and giving good reason to be optimistic that substantial value is at play for partners within the next 3 to 5 years. For instance, Stripe found that AI startups are building businesses faster. The top 100 AI companies on Stripe in 2024 took 24 months to \$5 ARR vs. 37 months for previous waves of SaaS companies.²

Regional market dynamics

The characteristics and make-up of the agentic AI market vary by region due to differences in industry exposure and the mix of service-based vs. manufacturing-based economies.

North America: North America represents the largest potential SI value pool globally, driven by a highly professionalized workforce. This market is driven by Technology/TMEG and the Financial Services & Insurance sectors, which together account for nearly 40% of the regional value pool. The technical maturity of enterprises in this region, along with organizational readiness, has laid a strong foundation for embracing AI-driven transformation.

Interest in agentic AI use cases spans both front- and back-office domains, given the concentration of large-scale enterprises with the resources, technical maturity, and innovation ambition to invest in AI. Non-exhaustively, we have observed interest in the personalization of customer experiences and the augmentation of backend processes.

EMEA: EMEA also features a relatively professionalized workforce. Compared to North America, the region exhibits a more balanced distribution across industries, though there is a hotspot in the Auto/Industrial sector, which accounts for around 20% of the total SI services value pool. This industrial diversity shapes the kinds of operational challenges SIs are brought in to solve.

Data maturity driven by heightened awareness of GDPR and customer data governance has created favorable conditions for AI exploration, especially in content generation and personalized engagement, though it also generates caution.

JAPAC: The SI value pool in JAPAC is more heavily weighted toward tasks performed by clerical, sales, and trade worker roles. This composition contrasts with more service-based economies and signals a stronger need for AI solutions that can address high-volume, task-based functions. Vertically, the region's value is concentrated in Retail/CPG, Auto/Industrial, and Financial Services, which



Many of our current opportunities are in North America, where new tech adoption seems to start first. This is driven by the speed and size of investments in the US and lower regulatory hurdles. Europe has been 3-6 months out, and specific pockets of demand have also picked up in APAC.”

Google Cloud Partner



Some sub-verticals across EMEA, APAC, and LATAM are cutting edge and quick to market. For example, banking and payments in Europe and LATAM is a progressive market with a lot of digital native customers... Or fashion players in Europe, who are exploring broader, customer experience-focused agentic AI use cases.”

Google Cloud Industry Solution Leader

together account for around 70% of the opportunity. This creates a natural fit for SIs to support operationally intensive industries, especially those with distributed workforces and manual frontline processes.

There are relatively high levels of data readiness and executive ambition across key JAPAC markets, particularly in Australia. Many retailers in Australia are actively exploring how AI can drive unified customer interactions and improve responsiveness across channels at scale.

LATAM: The LATAM SI value pool is smaller and is driven by more of a mixed labor base, with the types of roles being more diverse than other regions (~40% of the potential SI value pool is driven by augmentation of tasks performed by managers or professionals; with another ~40% by technicians or clerical and sales workers). LATAM has a higher concentration of value in the Retail/CPG, Public Sector, and Auto/Industrial verticals (~60%).

1: Gartner: "Gartner Forecasts Worldwide GenAI Spending to Reach \$644 Billion in 2025"

2: Stripe Annual Letter, 2024 (compared to Top 100 SaaS companies in 2018)



04

The strategic imperative for partners: How services will evolve in an agentic world



The sheer scale of the agentic AI opportunity makes it a strategic imperative for SIs to reshape how they deliver, monetize, and differentiate. A potential ~\$1T opportunity for agentic AI services is more than just a new service line—it reinforces the urgency for SIs to undertake a strategic rethink, act fast, and capture the opportunity ahead. Observing from the sidelines leaves untapped value on the table.

To navigate this transition successfully, there are **five key themes that SIs must consider when defining their strategic roadmap** in the agentic era.

01

Evolving the SI delivery model beyond just technical execution

In the traditional world, SIs helped customers solve large-scale, execution questions like “how can enterprises migrate on-prem SAP to the cloud?”. However, in the agentic world, customer expectations of SIs will evolve towards helping them leverage agents to solve fundamental business issues. This requires SIs to more deeply understand and potentially redesign business processes to realize the value of agents. This evolution has implications for how SIs deliver work over time:

- **Greater emphasis on upfront design work in the near term:** SIs need to invest time upfront to understand business processes, and potentially design new organization processes that embed agents into workflows in order to realize full value potential.
- **More prototypes—not slides:** Customers will increasingly expect ‘prototypes that work’ early on, which will cause SIs to shift from delivering slides to delivering working code—in turn, this will enable SIs to show value early, which strengthens their long-term positioning.
- **Finding the change leader to help drive implementation:** While technical implementation remains a key element, SIs over time will need to become more savvy in helping customers drive adoption as agents move into production environments. Crucially, this includes proactive change management: identifying change leaders, outlining adoption roadmaps, and managing production deployment risks. SIs may take this opportunity to build the capability to do this work themselves or work hand-in-hand with other consulting partners.
- **Evolution in maintenance support needs:** While core maintenance capabilities will remain relevant, we expect incorporation of new agent-specific elements that have never been seen before, such as orchestrating multi-agent workflows, evaluating and measuring agent performance, and refreshing agent knowledge to prevent model drift—all performed over time on a recurring basis.



Based on traction in the last 6 months, we fear clients will only use agents to enhance current workflows, not reimagine them. This is where SIs need to help them to unlock true, full value.”

Google Cloud Partner



We have moved away from slides and now bring agent prototypes, built with speed using Google tools and synthetic data. This proves ROI instantly for clients and helps get a POC in under 2 weeks vs. 4-6 weeks.”

Google Cloud Partner



As we create our library of agents, we expect customers to remain sticky and continue to invest in implementation and maintenance for continued access to 1000s of agents and best cross-industry practices to learn from. Both of these can save millions for enterprises.”

Google Cloud Partner

02

Rethinking commercial structures to match the nature of agentic value creation

Agentic AI will open up the opportunity for SIs to test different monetization structures, catering to varying levels of SI and enterprise risk appetite. This opportunity is a unique juncture to refresh pricing strategies, enabled by the step-change in both the magnitude of opportunity that agentic AI brings and the shift in work mix and knowledge base it requires. While we believe that traditional ‘time and material’ and ‘fixed fee’ models will continue as key structures for predictable, transparent pricing, SIs should use this opportunity to consider other monetization options that align with value delivery to customers:

- **Recurring software/SaaS-like pricing:** SIs may be able to embody their unique domain expertise into productized agentic solutions. If SIs are able to resell their agentic solution as reusable IP—potentially discoverable and accessible via marketplace—a recurring, SaaS-like monetization can drive customer stickiness and waves of additional value.
- **Non-recurring, transaction-based (e.g., one-time fee per agent deployed):** This model allows clients to purchase specific AI capabilities or ‘digital workers’ on demand for discrete tasks, offering cost control for well-defined, non-continuous needs, distinct from ongoing service subscriptions.
- **Outcome-based pricing:** Agents can inherently drive measurable business outcomes when deployed end-to-end. For such scenarios, SIs should consider outcome-based pricing when a feasible ROI exists with the risk upside vs. risk balance.




Exhibit 4.2: Various monetization models exist based on the SI's risk appetite

03

Using new tools to quickly reach sufficient data readiness for agentic deployment

Addressing how data is supplied to agentic AI is a foundational conversation with every customer, though the traditional concept of ‘data readiness’ is evolving. The emphasis is shifting towards architecting effective data pipelines that enable AI agents to reliably access and process the diverse information that they need, when they need it.

Thankfully, new gen AI-enabled tools have been built to address these needs, allowing SIs to quickly build knowledge maps and assemble data repositories that can be used to guide agents. Increasingly, these repositories can include unstructured data (e.g., documents, email, transcripts), which was often a challenge to integrate with traditional automation approaches.

As a result, SIs must be equipped to use these tools so that they can help customers map their existing data landscape, allowing them to efficiently and securely connect various datasets to AI agents. This includes mapping operational data flows, understanding AI/data regulations (which may differ by industry and geography), and pinpointing key data sources such as databases, APIs, and document repositories that can be used to guide the behavior of agents that rely on them.

Once these data pathways are better understood, SIs can guide clients across a spectrum of solutions. For immediate pilots/POCs, leveraging readily accessible client data, external data (e.g., from Google’s Kaggle), or even synthetic data, can offer tactical value and help quickly validate ROI. Modern approaches, including the use of advanced connectors (like the Model Context Protocol, MCP) and inter-agent communication protocols (like the Agent2Agent Protocol, A2A), can enable faster and more seamless data access from various enterprise systems.

Furthermore, agentic systems themselves are designed to improve over time in their decision-making capabilities as they process more information. While robust data governance, regulatory adherence, and data quality assurance remain important for building sustainable and scalable agentic AI solutions, the initial barrier to entry can be lowered by focusing on smart data pipeline design and gen AI’s ability to work with the data landscape as-is. Educating customers on these evolving data integration patterns and prerequisites will be vital.

04

Building the capabilities to scale agentic solutions from pilot to production

SIs will increasingly need to demonstrate core capabilities to map out existing business processes in their pre-agent state, redesign these workflows to incorporate agents (with or without humans in the loop), utilize best-in-class enterprise AI platforms such as Vertex AI to develop agents, as well as set up integrations with existing data systems that are required to execute that workflow. Balancing efficiency optimization and broader regulatory considerations is crucial in this redesign process and can be moderated by defining clear human-agent interaction points. SIs investing in proprietary knowledge, reusable IP (such as pre-built connectors to widely used existing enterprise software, custom model approaches, and design patterns), and strong change management capabilities will be able to accelerate the POC to deployment conversion cycle and create a long-term competitive advantage.

Furthermore, the ability to deploy these integrated solutions into production at scale is critical. This requires establishing processes that enable continuous integrations suitable for the pace of agentic AI evolution, rigorous testing methodologies to ensure agent reliability, and sound infrastructure planning. In the near-term, this means arming go-to-market (GTM) teams with tools to educate and provide assurance to customers, for instance, a library of vertical use cases, a set of ethical and AI regulation-compliant deployment guidances (which may differ by industry and geography), and ROI calculators with flexible pricing and prototypes/demos to bring to pitches. Over time, customers will also need support in implementing observability tools to monitor agent accuracy, efficiency, and ROI attribution, as well as emerging agentic maintenance work such as refreshing agent knowledge to prevent model drift for continued improvement in tangible value delivered.

05

Focusing on high-value, industry-specific and customer-specific use cases to drive differentiation

Agentic AI's applicability is vast, requiring SIs to prioritize use cases based on opportunity size and ease of deployment. While we extensively discuss the 'art of the possible' for each industry in the next chapter, SIs should consider two distinct horizons: pursuing quick wins with a clear near-term fit for agents (e.g., customer service), and building towards larger, transformative agentic ecosystems (e.g., N-of-1 personalization). This will enable SIs to build brand and capability for agentic services while building reusable IP and delivering increasing value for customers.





Exhibit 4.3: Critical to relentlessly prioritize agents that deliver value, focusing on quick wins to start

Google Cloud is dedicated to helping SIs answer these strategic imperatives and is committed to ensuring SIs have the right tools and ecosystem to address these key questions and foster innovation and success in the agentic era. How SIs can best leverage Google Cloud to expedite agentic AI deployment is further explored in Chapter 6.

Partners proactively addressing these themes can empower their customers to use AI agents in driving new avenues of growth and efficiency, as well as build a true competitive advantage for themselves to tap into the unprecedented agentic AI market opportunity.

05

Agentic AI in action: How partners can help transform key industries



Agentic AI is emerging as a far-reaching solution to augment and optimize tasks across every industry. Logistics agents can help mitigate inventory imbalances and routing inefficiencies, predictive maintenance agents can reduce unplanned equipment failures, and marketing agents can accelerate campaign execution and scale ‘N-of-1’ personalization. These, and other horizontally-applicable agents, such as customer service, compliance, and pricing agents, can drive impact across a wide range of sectors.

While agentic AI offers broad applicability, its true transformative power is unlocked when tailored to each industry’s specific nuances. Every sector faces unique structural challenges, regulatory demands, and operational workflows that agentic solutions must address. For instance, Financial Services and Insurance must navigate intricate compliance needs for claims processing; Retail

and CPG solutions need to adapt to margin pressures stemming from rising input costs; and in Healthcare, agentic tools must align with stringent patient data privacy and specific clinical protocols. Tailoring agentic AI solutions for these deeply embedded, industry-specific requirements is key to realizing significant transformation opportunities.

In the pages that follow, we explore the ‘art of the possible’ across six key industries, using selected examples to illustrate the transformative potential of agentic AI and to highlight how Google Cloud’s partners can lead the charge.

Tech/TMEG

Technology, Telecommunications, Media, Entertainment, and Gaming (Tech/TMEG) organizations are often early movers in enterprise AI adoption, with deep investments in data infrastructure and machine learning.




Exhibit 5.1: Regional split of global agentic AI value pool in Tech/TMEG and role mix in high-impact SI opportunities

But even in these advanced environments, execution bottlenecks persist. Personalized content delivery can still involve manual and time-intensive campaign workflows; network operations rely on fragmented telemetry systems; narrative design in gaming remains a labor-intensive process, requiring teams to manually script story branches and character interactions; and software development teams still face friction across code generation, testing, and deployment workflows.

In total, agentic AI presents a **\$125B to \$175B global opportunity** for SI services within the Tech/TMEG sector, and **\$60B to \$90B in the U.S.** Within this, high-impact value pools are likely to emerge in areas such as:

- **Content and campaign optimization in Media & Entertainment:** Empowering marketing teams to scale personalized content creation, streamline campaign execution, and optimize real-time performance across media and entertainment platforms
- **Intelligent network and operations management in Telco:** Supporting engineering and operations teams with autonomous monitoring, diagnostics, and resolution workflows to reduce service downtime and strengthen telco infrastructure
- **Scalable narrative design and personalization in Gaming:** Enabling game development teams to generate branching storylines, dynamic dialogue trees, and personalized character arcs based on player behavior, genre, and emotional tone
- **Software development acceleration in Tech:** Enabling engineering teams to

use code generation, test automation, and DevOps agents to streamline development cycles, reduce friction, and release higher-quality software faster

Sales is also a major source of value for agentic AI in Tech/TMEG and is widely applicable across other industries such as Retail/CPG, Auto/Industrial, and more, where customer-facing teams play a critical role in driving growth and retention. Agentic solutions like customer service agents can resolve support issues across channels, sales enablement agents can craft tailored outreach and proposals, and post-sale success agents can manage onboarding and renewals, helping organizations scale engagement and reduce friction across the customer lifecycle.

By augmenting high-impact workflows such as these, agentic AI enables Tech/TMEG organizations to unlock productivity, reduce manual overhead, and deliver more consistent and adaptive customer experiences.

Tech/TMEG: Media & Entertainment

In Media and Entertainment, agentic AI can transform workflows such as content creation and compliance. Marketing teams can deploy content marketing agents to produce and publish personalized, multi-format campaigns across channels, while localization agents can adapt those assets for global audiences. On the compliance side, rights management agents can track content usage, flag licensing violations, and trigger legal review processes. This way, agents can help reduce production bottlenecks, accelerate time-to-market, and mitigate regulatory risk.

In the near term, we have observed particular excitement around content marketing and personalization, where enterprises are testing new approaches. Today, marketing teams rely on a patchwork of tools and manual coordination to create, personalize, and deploy content across channels. Crafting copy, selecting visuals, segmenting audiences, and measuring campaign performance all require frequent handoffs and back-and-forth revisions, slowing down time to launch and limiting personalization. An agentic approach can change that, helping unlock significant value within marketing, a segment driving roughly ~12% of the total U.S. SI value pool within Tech/TMEG (Exhibit 5.1).

A **content marketing agent** can ingest campaign briefs, generate multi-format content personalized to different audience segments, and deploy assets across Content Management Systems (CMS) and ad platforms. It can continuously monitor performance, run A/B tests, and refine messaging in real time, closing the loop between strategy and execution. For media and entertainment brands managing complex, multi-channel campaigns, this unlocks a new operating model, one where teams can deliver highly personalized content with speed, precision, and consistency. Instead of working in cycles, marketing becomes a continuous, adaptive system—driving faster launches, increased audience engagement, and greater return on creative investment.




Exhibit 5.2: Agentic AI generates and deploys personalized content at scale in Media and Entertainment, continuously optimizing based on performance

Tech/TMEG: Telco

In Telco, agentic AI can improve operational continuity and pricing strategy. Engineering teams can use **network operations agents** to monitor telemetry data, detect anomalies, pinpoint root causes, and trigger resolution workflows, dramatically reducing service outages and minimizing downtime. Finance professionals can also explore use cases like **dynamic pricing agents**, which evaluate demand and usage to adapt pricing models in

real time, or **revenue leakage prevention agents** that detect billing anomalies and initiate recovery actions. These are just a few examples of how agents can address long-standing challenges across networks.

Network operations, in particular, offer a compelling and exciting near-term opportunity for transformation given their complexity, urgency, and direct impact on service reliability. Today, service outages, latency issues, and equipment failures can require manual monitoring and triage across complex telemetry and incident systems, often delaying resolution and hurting customer experience.

A **network operations agent** can ingest live telemetry data, detect anomalies, identify root causes, and coordinate resolution steps across systems and teams. Instead of waiting for failures to be reported and escalated, an agent can proactively manage outages and performance degradation, reducing time to resolution and improving overall reliability. This level of autonomy transforms network management from a reactive, manual process into a proactive, intelligence-driven system, reducing downtime, improving service reliability, and allowing engineering teams to focus on long-term optimization rather than day-to-day triage.




Exhibit 5.3: Agentic AI detects, diagnoses, and resolves network issues autonomously in Telco, reducing manual intervention across the incident lifecycle

Tech/TMEG: Gaming

In Gaming, agentic AI can enhance both creative development and monetization. Narrative teams can use **storyline generation agents** to develop branching plots and character arcs tailored to individual player decisions. Meanwhile, **game testing agents** can simulate gameplay to uncover bugs earlier in the development cycle. On the business side, finance teams might deploy **offer optimization agents** to test and adapt in-game bundles based on player behavior. These examples highlight the breadth of ways agentic AI can support gaming studios across storytelling, quality assurance, and engagement.

An area that we see as exciting and compelling in the near term is narrative design and personalization, where agentic AI can unlock creative scale and storytelling flexibility. Today, creating branching storylines and dialogue options can be a manual, time-consuming process that relies heavily on narrative designers and writers anticipating every player’s choice. As games grow more open-ended and player-driven, this model struggles to scale, limiting both personalization and replayability. Writers must manually script divergent arcs, build dialogue trees, and ensure consistency across story branches, often requiring repeated review and iteration. An agentic approach enables faster generation and dynamic adaptation of game narratives.

A **narrative design agent** can define story scaffolds based on genre and gameplay objectives, generate branching plotlines and character arcs, and write dynamic dialogue tailored to player choices and emotional tone. These agents can also adjust narrative flow in real time based on player decisions and simulate story paths to identify pacing gaps or logic breaks. By automating the development and refinement of narrative components, agentic AI allows creative teams to focus on deeper storytelling, enabling more immersive, personalized player experiences at scale, while also reducing the manual lift required to bring those experiences to life.




Exhibit 5.4: Agentic AI generates branching game narratives dynamically, personalizing player experience at scale

Tech/TMEG: Tech

In the broader Tech sector, agentic AI offers familiar yet powerful avenues to reduce software development cycle times and improve development velocity. Established applications include engineering teams relying on code generation agents to turn functional requirements into production-ready code, supported by test automation agents that create and execute QA test suites. DevOps agents can also be used to manage build pipelines, detect integration issues, and accelerate deployment. These examples illustrate how agentic AI can remove friction across the software development lifecycle, enabling

teams to release higher-quality software faster, with fewer bottlenecks. For SIs, these solutions are a core unlock within engineering, a segment directionally driving more than 50% of the total U.S. SI services value pool in Tech/TMEG (Exhibit 5.1).

Many of these use cases, while especially relevant in Tech, also extend across other industries. For example, DevOps automation agents are increasingly relevant in Retail or Financial Services for managing complex infrastructure that powers customer-facing applications. Similarly, code generation and test automation agents are equally valuable for healthcare companies building digital experiences or public sector teams modernizing legacy systems. These cross-cutting agents offer SIs a powerful entry point to drive repeatable, scalable impact, not only in traditional software companies, but in any enterprise looking to improve how technology is developed and deployed.

Taken together, these use cases show how agentic AI can enable a step-change in how Tech/TMEG organizations operate. By embedding reasoning and autonomy directly into execution workflows, agents are helping teams move beyond traditional tooling to deliver accelerated software development cycles, more resilient infrastructure, and hyper-personalized user experiences. For SIs, this represents a rare inflection point—an opportunity to re-architect core digital processes and help clients realize new levels of speed, scale, and adaptability across their businesses.

Retail/CPG

Retail and CPG companies face growing pressure to personalize the customer experience without compromising profitability. At the same time, they must navigate rising input costs, supply chain disruptions, and volatile demand. Many groups are still managing inventory, service, and promotions through static rules and disconnected systems.




Exhibit 5.5: Regional split of global agentic AI value pool in Retail/CPG and role mix in high-impact SI opportunities

Agentic AI presents a **\$150B to \$200B global opportunity** for SI services in the Retail and CPG sectors, including **\$20B to \$30B** in the U.S. Within this, the largest value pools are likely to emerge in areas such as:

- **Customer service and post-purchase support in Retail:** Deploying agents that handle returns, resolve inquiries, and manage escalations across channels to reduce churn and improve satisfaction
- **Inventory and demand planning in CPG:** Using agents that monitor real-time demand signals, coordinate restocking workflows, and optimize stock levels across locations to prevent overstock and outages
- **Personalized shopping and product discovery in Retail:** Enabling sales representatives to guide customers to relevant products, answer questions, and adapt recommendations based on behavior and preferences

Agentic AI also delivers value for engineering roles such as software developers and IT specialists by **automating system monitoring, detecting anomalies, and optimizing backend workflows** across e-commerce and supply chain platforms. These use cases are equally relevant in industries like Financial Services, Healthcare/Life Sciences, Tech/TMEG, and more, where similar technical infrastructure challenges persist.

By connecting real-time signals to coordinated action, agentic AI helps Retail and CPG organizations deliver more responsive, efficient, and personalized customer experiences at scale. Agents can adapt instantly to changing consumer behaviors, market conditions, or inventory constraints, enabling brands to engage customers with greater precision while protecting margins and streamlining execution across digital and physical channels.

Retail/CPG: Retail

In Retail, agentic AI can improve both customer-facing and operational workflows. Customer service agents can resolve inquiries and process returns end-to-end,

reducing friction and freeing up human sales representatives for higher-value interactions. Shopping advisor agents can guide customers through product discovery and purchase decisions using real-time preferences, browsing behavior, and context to personalize the experience. Customer experience agents can orchestrate personalized offers, messages, and support across the shopper journey based on real-time behavior and preferences. These use cases illustrate how agents can drive efficiency, responsiveness, and personalization at scale across the retail value chain.

Specifically, efforts to rethink customer service and returns are generating significant near-term interest, as the cost of poor experiences shows up in both loyalty and margin. Today, resolving a single issue can require multiple handoffs: a service representative reads and classifies the request, checks order systems, determines eligibility, and processes the return or refund manually. These steps often vary by channel, and complex cases require escalation, driving up service costs and customer frustration. An agentic solution can streamline this end-to-end.




Exhibit 5.6: Agentic AI in retail resolves customer service requests end-to-end by interpreting issues, retrieving answers, and executing resolution

A customer service agent can classify requests in real time, retrieve order details, determine eligibility based on policy rules, and process resolutions automatically, updating systems and notifying the customer without manual intervention. Complex or ambiguous cases are escalated as needed, while the bulk of requests are resolved quickly and consistently. For SIs, solutions like this, alongside sales rep agents, are critical to tap into the large sales value pool within Retail/CPG, roughly driving ~47% of the total U.S. SI services value pool in Retail/CPG (Exhibit 5.5).

Retail/CPG: Consumer Packaged Goods (CPG)

In CPG, agentic AI can help streamline supply chain execution and reduce operational inefficiencies. Inventory planning agents can forecast demand, monitor stock levels, and trigger replenishment to maintain optimal inventory across distribution networks. Procurement agents can evaluate pricing, assess vendor options, and initiate reorders to manage cost volatility and ensure continuity of supply. Meanwhile, vendor onboarding agents can automate partner setup by validating documentation, syncing systems, and coordinating tasks across internal teams. These applications show how agents can enhance agility, reduce manual workload, and improve throughput across the CPG value chain.

In the near term, inventory planning stands out as a particularly promising area for agentic transformation. Despite its central role in operational execution, it remains a manual and error-prone function in many CPG organizations. Planners often contend with incomplete demand signals, batch-based data updates, and fragmented coordination with procurement teams. This results in familiar tradeoffs: excess inventory that ties up working capital, or stockouts that lead to lost sales and diminished service levels.

An inventory planning agent can ingest real-time sales, promotion, and supply data to forecast demand and trigger replenishment actions automatically. It continuously monitors product availability and adjusts orders based on current conditions, improving sell-through and reducing capital tied up in misaligned inventory. Over time, this enables CPG brands to align supply with demand more precisely, even as buying patterns shift. For SIs, an inventory planning agent is a key value driver for their customers' operations functions, a segment broadly accounting for ~8% of the total US SI services value pool within Retail/CPG.




Exhibit 5.7: Agentic AI optimizes inventory in Retail/CPG by forecasting demand and triggering replenishment through integrated systems

As consumer expectations rise and value chain complexity intensifies, agentic AI offers Retail and CPG organizations a powerful lever to drive differentiation. In Retail, agents unlock new levels of responsiveness and personalization, transforming how brands engage shoppers across touchpoints. In CPG, agents bring precision and agility to supply chain decisions, enabling organizations to react in real time to demand signals, cost pressures, and disruptions. Across both segments, these solutions not only reduce operational burden but also enable faster, smarter execution, helping companies grow customer loyalty, improve margins, and compete more effectively in a market defined by speed and constant change.

Financial Services/Insurance

Financial Services and Insurance firms have invested heavily in digital infrastructure, yet they remain burdened by fragmented processes, legacy systems, and complex regulatory obligations. Some core workflows, such as claims processing, onboarding, and wealth advisory, can still depend on static rules, duplicated data, and high-touch manual review.

While newer entrants such as neo-banks often have a head start—built on modern, cloud-native architectures and more unified data systems—traditional institutions are also making steady progress in modernizing their tech stacks and consolidating data environments. This evolving foundation positions both types of firms to benefit from agentic AI, though their starting points may differ.




Exhibit 5.8: Regional split of global agentic AI value pool in Financial Services/ Insurance and role mix in high-impact SI opportunities

Agentic AI presents a **\$125B to \$175B global opportunity** for SI services in the Financial Services and Insurance sector, including **\$60B to \$90B in the U.S.** Within this, the largest value pools are likely to emerge in the 'Finance' segment, driving roughly ~61% of the U.S. SI services value pool in Financial Services/Insurance. Key areas of focus within this segment include:

- **Underwriting support in Insurance:** Empowering underwriters with tools that synthesize applicant, policy, and risk data to generate individualized risk profiles, flag edge cases, and streamline decision-making
- **Personalized financial planning in Wealth Management:** Supporting wealth managers with tools to generate tailored investment strategies, monitor portfolios, and adapt plans based on client data
- **Know Your Customer (KYC) verification in Banking:** Enabling compliance and onboarding teams to gather and verify identity documents, screen against watchlists, flag risk indicators, and auto-complete KYC workflows to reduce onboarding friction and compliance burden

In addition to finance-specific workflows, many high-value agentic use cases span engineering roles such as software developers, business analysts, and data analysts. Agents that support tasks like **code generation, report creation, and workflow optimization** are especially relevant for Financial Services/Insurance institutions looking to modernize internal systems, enhance analytical throughput, and scale regulatory reporting. These capabilities are also broadly applicable across industries such as Tech/TMEG and Healthcare/Life Sciences, enabling reusable agent patterns that drive cross-sector impact.

By transforming these high-friction workflows, agentic AI helps Financial Services and Insurance organizations increase accuracy, reduce compliance risk, and deliver faster, more responsive customer experiences. It enables institutions to move beyond rigid, manual processes toward adaptive systems that learn and improve over time, enhancing decision quality, accelerating time-to-value, and allowing teams to focus on higher-impact work.

Financial Services/Insurance: Insurance

In Insurance, agentic AI can streamline action-heavy workflows across claims, underwriting, and intake operations. Underwriting teams can deploy **underwriting support agents** to synthesize applicant, policy, and risk data, generate individualized risk profiles, and flag edge cases for human review, accelerating data gathering and improving consistency. Upstream in the process, **submission triage agents** can sort and prioritize incoming forms based on urgency and completeness, helping staff focus on high-value cases. Meanwhile, **broker insights agents** can analyze producer performance and engagement trends, helping insurers guide outreach, improve channel effectiveness, and refine distribution strategy. These examples show how agents can help reduce administrative burden, shorten cycle times, and improve both consistency and responsiveness across the insurance value chain.

In particular, underwriting is an exciting near-term area where agentic AI can ease manual workload and improve decision quality. Today, underwriters must gather and assess applicant and asset data from a range of fragmented sources, such as property records, inspection reports, and credit agencies, then apply static rules and checklists to assess risk. When data is missing or inconsistent, applications are escalated for manual review, creating delays and inconsistent evaluations.

An **underwriting support agent** can automate this process end-to-end. It can ingest data from internal and external sources, including satellite imagery or weather databases, apply underwriting logic to generate individualized risk scores, and flag outliers (such as roof damage or unexplained prior claims) for human review. This allows underwriters to focus on the highest-value decisions, significantly accelerating cycle times and enabling faster, more consistent, and data-driven quote generation at scale across property and casualty lines.




Exhibit 5.9: Agentic AI supports underwriting by synthesizing risk data and flagging edge cases in property & casualty insurance

Financial Services/Insurance: Wealth Management

In Wealth Management, agentic AI can enhance the delivery of personalized financial advice and improve how firms manage client engagement and portfolio oversight. Advisor teams can deploy **wealth advisor agents** to synthesize financial goals, account data, and transaction history to generate tailored investment plans, monitor portfolio drift, and dynamically adjust recommendations over time. **Data harmonization agents** can align inputs across CRM, policy, and risk systems, ensuring that advisors work from a complete and accurate client view. Meanwhile, **high-value client routing agents** can detect key service moments, such as large deposits or market events, and proactively reassign high-priority accounts to senior advisors. These examples show how agents can help scale personalization, improve advisor productivity, and drive stronger client retention in an increasingly competitive wealth management landscape.

Specifically, personalized financial planning stands out as a high-potential near-term area for agentic transformation in Wealth Management, where the ability to deliver timely, relevant advice remains constrained by legacy practices. Many advisors still operate with static planning templates, infrequent client check-ins, and fragmented data sources, making it difficult to tailor recommendations to evolving client needs. Time is often lost aggregating account data, tracking portfolio drift, and identifying rebalancing triggers, limiting both responsiveness and advisor productivity.

A **wealth advisor agent** can automate this process, aggregating account data, transaction history, and financial goals to generate tailored strategies. It monitors for portfolio drift, flags rebalancing opportunities, and proactively recommends updates. This empowers advisors to deliver precision-tailored guidance that evolves with each client’s goals, enabling more timely decisions, stronger engagement, and higher long-term value.




Exhibit 5.10: Agentic AI delivers personalized wealth planning and execution across the full client lifecycle in Wealth Management

Financial Services/Insurance: Banking

In Banking, agentic AI can improve both customer onboarding and core operational workflows. **KYC verification agents** help streamline identity checks by gathering documents, screening against watchlists, and completing onboarding steps automatically. Banks can also deploy **payments exception agents** to quickly resolve failed transactions, **model governance agents** to monitor and flag compliance risks, and **client segmentation agents** to tailor engagement strategies based on customer behavior. Together, these use cases show how agents can reduce regulatory friction, enhance service reliability, and improve customer experience.

In particular, we have observed that KYC verification is emerging as a high-impact opportunity for agentic transformation in banking, given its regulatory importance and operational complexity. Today, teams must manually collect identity documents, verify

them against internal records or registries, and screen applicants against watchlists and sanctions databases. This fragmented process slows down onboarding, increases risk exposure, and often results in rework when data is missing or flagged late.

A KYC verification agent can streamline this process while improving both customer experience and compliance rigor. It can ingest identity documents and form data from intake portals, validate them using internal and external sources, and screen applicants in real time against politically exposed persons (PEP), anti-money laundering (AML), and sanctions lists. By auto-completing onboarding for low-risk customers and routing only complex cases to compliance, the agent eliminates a major bottleneck in account activation, accelerating customer access, strengthening regulatory assurance, and freeing operations teams to focus on higher-value oversight.




Exhibit 5.11: Agentic AI automates KYC by verifying identities and flagging risk

Financial Services/Insurance: Capital Markets

In Capital Markets, agentic AI can enhance speed and accuracy across pricing, compliance, and operations. Securities underwriting agents can aggregate filings, market data, and risk inputs to streamline issuance and reduce manual documentation. Exception pattern analysis agents help identify and resolve issues like trade mismatches or delayed settlements, while security and anomaly detection agents monitor for unusual behavior or access risks. These agents support faster decision-making and help teams navigate complex regulatory requirements.

Many of these use cases extend beyond capital markets. Exception pattern analysis agents can be applied in banking, payments, or supply chain workflows to proactively detect and resolve process failures. Security and anomaly detection agents are equally relevant in data-sensitive sectors like healthcare and public sector. For SIs, these cross-industry applications offer a path to deploy proven agentic solutions more broadly, accelerating time-to-value and expanding impact across clients.

Taken together, these use cases demonstrate how agentic AI can reshape core operations across the Financial Services and Insurance landscape. By embedding intelligence into critical workflows—underwriting, wealth planning, onboarding, and beyond—agents enable faster execution, improved consistency, and more personalized client engagement. For SIs, this represents a meaningful opportunity to drive measurable impact by helping institutions modernize legacy processes, enhance decision quality, and scale the capabilities needed to compete in a more data-driven, real-time financial environment.

Healthcare/Life Sciences

Healthcare and Life Sciences organizations face persistent pressures to deliver more with limited resources. Providers are strained by staff shortages and administrative overhead, payers continue to struggle with slow, error-prone claims adjudication processes that increase administrative burden and delay reimbursements, and life sciences firms contend with fragmented trial coordination and long development cycles. These challenges slow down care, increase costs, and make it harder to bring new treatments to market.




Exhibit 5.12: Regional split of global agentic AI value pool in Healthcare/Life Sciences and role mix in high-impact SI opportunities

Agentic AI presents a **\$100B to \$150B global opportunity** for SI services in the Healthcare and Life Sciences sector, including **\$30B to \$50B in the U.S.** Within this, the largest value pools are likely to emerge in the ‘Scientist’ segment, accounting for roughly ~27% of the total U.S. SI services value pool in Healthcare/Life Sciences. Core areas within this segment are:

- **Revenue cycle automation for Healthcare providers and payers:** Enabling organizations to verify coverage, code encounters, validate and submit claims, detect fraud, manage denials, and trigger resolution workflows to improve payout speed, accuracy, and auditability

- **Care coordination and patient journey support for Healthcare providers:** Helping clinical teams streamline intake, triage, and care transitions by orchestrating data exchange and task management across systems
- **Drug discovery acceleration in Life Sciences:** Assisting scientists with literature review, target identification, and experimental design through autonomous research agents that surface insights and recommend next steps
- **Clinical trial operations in Life Sciences:** Enabling research teams to accelerate site activation, manage protocol adherence, and automate patient monitoring to reduce trial delays and dropout rates

Many agentic AI use cases (such as **digital experience optimization, customer service automation**, and credentialing) are broadly applicable across industries, including Retail, Public Sector, Healthcare, and more. These workflows often share a common foundation: structured, rule-based tasks that make them well-suited for agentic execution. However, in Healthcare and Life Sciences, these same workflows carry different stakes due to regulatory, privacy, and clinical safety requirements. As a result, even shared use cases may take different forms depending on whether they're applied in provider, payer, or life sciences settings.

Healthcare/Life Sciences: Payer

In the Healthcare Payer setting, agentic AI can streamline both operational efficiency and compliance-heavy workflows. Claims teams can deploy **revenue cycle**

management (RCM) agents to orchestrate tasks such as verifying coverage, coding encounters, submitting and tracking claims, detecting anomalies, and managing payouts. These agents integrate traditionally siloed functions, such as medical coding, billing, and claims follow-up, into a unified, intelligent workflow. On the risk side, **underwriting agents** can assess applicant data, flag anomalies, and generate risk scores to support faster, more consistent policy decisions. These examples show how agents can help reduce administrative burden, improve accuracy, and strengthen decision-making in payer operations.

Of these agents, we are observing growing excitement around RCM automation, as it addresses one of the most complex and high-impact workflows, where accuracy, speed, and compliance are critical, but often hindered by fragmented systems and manual handoffs. Today, much of this process is carried out manually: claims specialists gather documentation, validate coverage and policy terms, assign billing codes, run fraud checks, and liaise with billing and compliance teams. These steps span multiple disconnected systems, creating inefficiencies and increasing the risk of error.

An **RCM agent** can orchestrate these tasks end-to-end, retrieving insurance eligibility, validating policy rules, applying accurate coding, flagging fraud signals, handling denials, and initiating payment disbursement. It integrates previously disjointed activities into a seamless agentic workflow. Exceptions are escalated only for edge cases, enabling human specialists to focus on judgment-heavy or disputed claims. The result is a

step-change in operational efficiency and compliance assurance, delivering faster claim resolution, higher payout accuracy, and reduced overhead across the revenue cycle. From an SI perspective, this agentic solution is critical to unlock value for both the ‘Scientist’ and ‘Finance’ segments, together accounting for roughly ~47% of the US SI value pool for Healthcare/Life Sciences (Exhibit 5.12).




Exhibit 5.13: Agentic AI transforms revenue cycle management by automating coverage verification, coding, claim submission, and billing

Healthcare/Life Sciences: Provider

In the Healthcare Provider setting, agentic AI can enhance care delivery, reduce clinician workload, and improve operational efficiency. **Care coordination agents** can support full patient journey management, including tracking plans, surfacing follow-ups, and flagging readmission risks to guide proactive engagement. **Medical diagnostics agents** can assist clinicians by analyzing notes, labs, and imaging to suggest likely diagnoses and highlight risk patterns. At the administrative level, **credentialing automation agents** can verify licenses, track expirations, and streamline onboarding processes. These use cases help providers deliver more coordinated, timely, and compliant care.

Amid the opportunities above, care coordination is drawing the strongest near-term excitement, as providers look to tackle the root causes of care delays and clinician overload. Today, patient intake and care coordination require staff to manually validate eligibility, assess symptoms, schedule follow-ups, and ensure handoffs between providers. These workflows are not only time-consuming but also vulnerable to dropped tasks, contributing to preventable readmissions and poor patient experience.

In contrast, a **care coordination agent** can independently ingest intake forms, verify patient identity and insurance across systems, triage based on urgency and readmission risk, and coordinate appointments or follow-up actions. By continuously monitoring patient status and system availability, the agent ensures timely, well-coordinated care from intake through discharge. This elevates care continuity, reduces the risk of readmission, and allows clinical staff to spend more time where it matters most, delivering high-touch, patient-centered care.




Exhibit 5.14: Agentic AI orchestrates end-to-end patient intake by automating triage, verification, and scheduling

Healthcare/Life Sciences: Life Sciences

In Life Sciences, agentic AI can accelerate R&D, improve product quality, and reduce operational friction across the value chain. **Clinical trial ops agents** help research teams manage study trial startup by reviewing site documentation, confirming readiness criteria, and coordinating enrollment workflows, shortening trial timelines. **Manufacturing quality agents** monitor production data in real time to detect anomalies, flag deviations, and initiate corrective actions, reducing compliance risks and ensuring batch reliability. At the front end of the pipeline, **drug discovery agents** analyze literature, trial results, and biological datasets to identify promising targets and suggest candidate molecules. These agents support faster, more informed decisions, ultimately helping bring safer therapies to market faster.

In particular, clinical trial operations agents offer a substantial opportunity to demonstrate the impact of agentic AI augmentation, where agents can streamline the time-intensive process of getting trials up and running. Today, clinical research associates (CRAs) manually review site readiness documents, coordinate access to trial platforms, and track enrollment status across systems. Much of this effort is spent chasing down information, validating paperwork, and communicating updates to investigators and sponsors.

A **clinical trial operations agent** can take on these tasks: reviewing qualification documents, verifying readiness criteria, assigning study system access, and monitoring trial milestones. It can also scan Electronic Health Records (EHRs) and registries to identify eligible patients for recruitment and sync Clinical Trial Management Systems (CTMS), freeing up CRAs to focus on trial design and relationship management. By eliminating administrative bottlenecks and streamlining enrollment workflows, agentic trial operations agents accelerate trial startup, reduce delays in patient recruitment, and help sponsors bring therapies to market with greater speed, precision, and confidence.




Exhibit 5.15: Agentic AI streamlines clinical trial startup in Life Sciences by automating site readiness, onboarding, and patient enrollment

In Healthcare and Life Sciences, agentic AI represents a transformative opportunity to solve some of the sector’s most entrenched challenges that have long eluded prior automation efforts. In provider environments, agents can help unlock clinical capacity and improve care continuity by taking on coordination and administrative load. For payers, agents can bring new levels of precision, consistency, and speed to claims adjudication and risk assessment. And in life sciences, they can accelerate time-to-market by streamlining R&D and trial operations. As demands rise and systems grow more complex, agentic AI offers the potential to fundamentally rewire how healthcare and life sciences organizations operate, advancing speed, quality, and impact across the value chain.

Automotive/Industrial

In Auto and Industrial settings, coordination isn't just a challenge; it's also a competitive constraint. Supplier networks are complex, customer expectations are rising, and lean inventory strategies leave little margin for error. Despite digital tools, many critical workflows, like maintenance and B2B sales, still rely on reactive decisions and disconnected systems.




Exhibit 5.16: Regional split of global agentic AI value pool in Auto/Industrial and role mix in U.S. SI opportunity

Agentic AI presents a **\$200B to \$250B global opportunity** for SI services in the Auto and Industrial sector, including **\$50B to \$70B in the U.S.** Within this, the largest value pools are likely to emerge in areas such as:

- **Predictive maintenance in Auto:** Enabling software and mechanical engineers to analyze equipment telemetry, anticipate failures, and trigger proactive interventions to reduce unplanned downtime
- **B2B sales acceleration in Industrial:** Supporting sales representatives by qualifying leads, tailoring outreach, and automating follow-ups to shorten deal cycles and improve conversion
- **Inventory and logistics optimization in Auto and Industrial:** Helping supply chain and operations teams use real-time demand and supply signals to adjust stock levels, optimize routing, and minimize disruptions

By enabling faster, more coordinated action across machines, systems, and frontline teams, agentic AI helps Auto and Industrial organizations operate with greater efficiency, agility, and reliability. Some of these examples, such as inventory and logistics optimization, have broad applicability to other sectors, such as Healthcare, highlighted here to show the ‘art of the possible’ when applied with an industry-specific lens.

Automotive/Industrial: Automotive

In Automotive, agentic AI can drive impact across engineering, operations, and sales by enhancing reliability, efficiency, and decision-making. Engineering teams can deploy predictive maintenance agents to analyze machine or vehicle telemetry, identify early signs of equipment failure, and trigger proactive service interventions, reducing costly downtime. Logistics optimization agents can support operations by dynamically adjusting load plans based on capacity constraints and real-time delivery schedules. On the commercial side, intelligent sales enablement agents can pinpoint at-risk regions, monitor real-time performance data, and recommend targeted interventions to improve conversion and customer retention. These agentic AI solutions tackle longstanding challenges across the automotive value chain, from production floors to dealership networks.

Among the opportunities outlined above, predictive maintenance is surfacing as a leading use case, as even short periods of unplanned downtime can derail production schedules or delay order fulfillment. Today, technicians must review dashboards, logs, and sensor data manually to identify anomalies and decide whether intervention is needed, often after a breakdown has already occurred.

An agentic AI system can continuously monitor equipment telemetry, detect early signs of failure, determine likely root causes based on historical data, and initiate service requests or adjust production schedules accordingly. The agent can also learn from past events to improve its predictive accuracy over time. This proactive, closed-loop approach not only minimizes downtime and improves asset utilization but also unlocks capacity for engineering teams to focus on high-leverage innovation, turning maintenance from a cost center into a competitive advantage. This productivity unlock for engineers can also translate to significant value for SIs. Roughly ~39% of the total U.S. SI services value pool in Auto/Industrials is linked to ‘Engineer’ tasks (Exhibit 5.16), indicating the directional opportunity this agent can offer.




Exhibit 5.17: Agentic AI shifts equipment maintenance in Automotive from reactive troubleshooting to proactive failure prevention

Automotive/Industrial: Industrial

In Industrial sectors, agentic AI is poised to enhance productivity, compliance, and resilience across engineering, operations, and finance functions. Engineering teams can leverage safety compliance agents to monitor worksite telemetry, detect hazardous conditions, and trigger risk mitigation workflows, helping reduce incidents and improve regulatory adherence. Operations leaders can benefit from logistics and inventory agents that optimize route planning, balance material flows, and minimize costly surplus or shortages across distributed sites. Finance teams can deploy supplier risk mitigation agents to proactively flag performance issues, reroute orders, and maintain supply chain continuity during disruptions. These AI-driven solutions help industrial players overcome persistent inefficiencies while building smarter, safer, and more agile operations.

Within this broad landscape, momentum is building around B2B sales, the revenue engine for most industrial firms. Agentic solutions for sales tasks alone roughly drive ~23% of the total U.S. SI services value pool in Auto/Industrials (Exhibit 5.16). Many teams today continue to face challenges identifying and converting high-quality leads with speed and precision. Sales representatives often rely on static lists, basic firmographics, and manual follow-up processes that are slow to scale and difficult to personalize.

A B2B sales agent can monitor inbound inquiries, analyze firm-level data, qualify high-potential opportunities, generate tailored proposals, and coordinate order placement through connected systems. It can also track engagement and adapt messaging based on buyer behavior. By automating routine tasks, sales teams can concentrate on strategic engagement and closing complex deals, unlocking new levels of speed, precision, and scale that traditional sales workflows struggle to achieve.




Exhibit 5.18: Agentic AI automates B2B sales in high-tech industrial by prioritizing leads, personalizing pitches, and driving follow-through across systems

In Automotive and Industrial, where coordination gaps quickly turn into costly setbacks, the ‘art of the possible’ with agentic AI can deliver a step-change in execution. By turning fragmented data into real-time, cross-system action, agents help organizations move from reactive problem-solving to proactive, resilient operations. The result: faster decisions, smarter workflows, and value creation beyond the reach of traditional automation.

Public Sector

Public institutions are expected to deliver services with speed, consistency, and transparency, often while operating on legacy infrastructure and limited capacity. From benefits processing to public communications, many agencies still depend on manual workflows and fragmented platforms that delay responses and strain overstretched teams.




Exhibit 5.19: Regional split of global agentic AI value pool in public sector and role mix in high-impact SI opportunities

Agentic AI presents a **\$125B to \$175B global opportunity** for SI services in the Public Sector, including **\$25B to \$50B** in the U.S. Within this, the largest value pools are likely to emerge in areas such as:

- **Citizen engagement and service delivery:** Helping public affairs and communications teams draft, adapt, and publish timely messages across channels using communications agents, improving transparency, consistency, and responsiveness to constituent needs
- **Regulatory compliance and reporting:** Supporting fraud and audit investigators by automating the detection of anomalies, surfacing suspicious patterns in real time, and streamlining risk-based case reviews, strengthening institutional oversight and control
- **Administrative operations:** Enabling eligibility specialists and frontline caseworkers to process benefits faster with agents that verify documentation, auto-approve standard claims, and flag edge cases for escalation, helping reduce backlogs and lighten manual workload

In addition to frontline workflows, agentic AI offers significant value for IT specialists, data analysts, and technical support staff across public institutions. Agents can automate **system monitoring, flag anomalies, and route issues for resolution**, enhancing uptime and accelerating issue triage in overstretched IT environments. Data analyst agents can **synthesize inputs across legacy systems, generate reports, and surface insights for program evaluation or policy planning**. These use cases not only strengthen digital resilience in government but also carry relevance across other sectors with large-scale, distributed systems, such as Healthcare/Life Sciences, Auto/Industrial, and beyond.

A wide range of agentic AI use cases is emerging: **benefits processing agents** can help reduce delays by verifying eligibility, approving claims, and routing exceptions across fragmented systems. **Fraud detection agents** offer a more scalable approach to oversight by analyzing behavioral data and surfacing anomalies for review. **Communications agents** can accelerate how agencies craft and distribute time-sensitive public messages, routing content for approval and adapting it by channel.

Among these opportunities, fraud detection is likely to be broadly applicable across agencies, where today’s processes are often reactive and resource-intensive. Analysts must manually monitor transaction logs, flag anomalies based on static rules, and initiate investigations across siloed systems. The result is delayed detection, at times inconsistent risk scoring, and significant effort spent on false positives. By addressing these pain points, SIs can unlock productivity for Public Sector finance departments; we estimate tasks in the total ‘Finance’ segment account for roughly ~18% of the total U.S. SI services Public Sector value pool (Exhibit 5.19).

A **fraud detection agent** can continuously ingest data across systems, surface anomalous behaviors in real time, and trigger case workflows with supporting evidence already attached. It can escalate high-risk issues with evidence already attached, generate audit trails automatically, and continuously improve detection logic, enabling faster investigations, stronger oversight, and a more scalable model for enforcing accountability.




Exhibit 5.20: Agentic AI monitors behavior in real time, flags anomalies, and initiates investigations across teams

Another high-impact opportunity is in public communications, where the timely delivery of clear, consistent messages is critical, but often undermined by manual processes and approval bottlenecks. Today, drafting public updates, coordinating reviews across departments, and publishing to multiple channels is slow and error-prone.

A **communications agent** can monitor for relevant events or triggers, draft initial message variants tailored to the audience and channel, and route them through legal, policy, and leadership approval flows. Once approved, it can be published across platforms (e.g., web, SMS, email, social) and track engagement or sentiment in real time. For agencies managing emergencies or policy changes, this enables faster, clearer outreach at scale, reducing delays, minimizing confusion, and helping overextended teams maintain trust and transparency with the public.

From an SI perspective, this solution offers significant value to tap into. This agent could become a staple tool for Public Sector marketing departments, which is a segment that is in total driving approximately ~15% of the total U.S. SI services Public Sector value pool (Exhibit 5.19).





Exhibit 5.21: Agentic AI automates public communications by monitoring events, drafting content, and orchestrating multi-channel publishing

Agentic AI offers public institutions a path to dramatically improve service delivery and operational resilience, even in the face of resource constraints and legacy infrastructure. By automating high-volume, coordination-heavy workflows and embedding intelligence into decision processes, agents can help agencies deliver faster, more consistent services, communicate with clarity at scale, and uphold institutional transparency. This unlocks a new standard of responsiveness and accountability, empowering governments to meet rising public expectations while doing more with less.

Taken together, the use cases across these six industries reflect a growing shift in what enterprise customers expect: not just tools or platforms, but solutions that directly address their most pressing pain points, such as claims delays, inventory mismatches, fragmented service delivery, and more. Agentic AI presents a way to meet those expectations, but it won't implement itself. SIs are uniquely positioned to close the gap between ambition and execution by helping clients identify the right entry points, design agentic workflows that map to real business processes, and scale those solutions across organizations. In the next chapter, we explore how Google Cloud is supporting SIs to accelerate this journey, from cutting-edge platforms to co-innovation programs.



06

Partnering with Google Cloud: Co-creating the agentic future

The journey into the agentic AI era is one of immense opportunity, and Google Cloud is deeply committed to empowering our partners to lead this transformation. We believe that by combining your industry expertise and client relationships with our cutting-edge technology and AI innovation, we can collectively help businesses solve their most complex problems. Our partnership is designed to be a collaborative venture, providing you with the tools, resources, and support needed to build, deploy, and scale impactful agentic AI solutions.

Track record of open innovation and thought leadership

With Google Cloud, you start on the cutting edge of AI research. We remain committed to driving open innovations and AI thought leadership, actively contributing to and supporting open-source initiatives critical for the advantage of agentic AI:

- **Agent Development Kit (ADK)** provides an open-source framework to simplify and accelerate the creation of AI agents, which is model-agnostic and deployment-agnostic.
- **Agent2Agent (A2A)** protocol aims to establish open-source interoperability between agents, a collaborative effort already benefiting from contributions from over 50 partners.

At the heart of our innovation is Google's deep heritage and ongoing investment in AI research and technology infrastructure. Over a decade of pioneering work from Google AI and DeepMind has led to breakthroughs in large language models (LLMs) like Gemini and specialized hardware such as our Tensor Processing Units (TPUs) optimized for AI workloads. This foundation ensures that the agentic solutions you build on Google Cloud are powered by the most advanced and reliable underlying technologies available.

Anchored in this open, innovation-first approach, Google has built a cutting-edge, vertically integrated AI stack with choices at every layer to uniquely cater to an agentic world. The stack emphasizes interoperability. At its foundation, partners have freedom on AI accelerators and models. Access to GPUs, TPUs, and 200+ Google and third-party models allow partners to build for complex environments without worry of vendor lock-in. We believe this choice enables partners to carve out 'the Pareto frontier' for agentic AI, balancing performance and cost for enterprise needs.

We recognize the need to build a stack with security front-of-mind. Google's AI security operations platform (SecOps) can empower security teams to better detect, investigate, and respond to cyber threats.

Our commitment to partners' success extends beyond technology. Google Cloud offers a thriving ecosystem and comprehensive support to help our thousands of partners build and scale their agentic AI practices. We understand that our partners' success is our success, and we are committed to providing the resources and investment you need to move fast and capitalize on the significant downstream revenue opportunities that agentic AI presents with speed.

Cutting-edge platforms and tools for agentic AI development

Recognizing the unique demands for agentic AI, Google has the right tools across the AI stack to get you started. Our tools target elements that are most important to partners succeeding in the agentic era: solving data challenges and accessing ready-to-use agents to kickstart development and building complex AI agents, custom to enterprise needs.

Enterprise data readiness remains a key friction point at the agent pilot and production stage. **BigQuery** provides a scalable data warehousing platform for AI, automating the data lifecycle from ingestion to insights. Connectors can access external datasets and tools (e.g., Kaggle) or unique-to-Google offerings such as GraphCast, Earth Engine, and many more. This enables rapid agentic builds and a chance for partners to take agentic prototypes into action for pitches.

Vertex AI is Google’s unified platform to support every stage of the agentic AI lifecycle. Using **ADK** and **Vertex AI Agent Engine**, developers can create and deploy bespoke multi-agent systems. Interoperability and fluidity remain central. Partners can access 200+ LLMs via the **model garden** and build using other frameworks. Purpose-built **MLOps tools** enable partners to manage and optimize the entire agent lifecycle with ease.

Agentspace is a dedicated environment to manage and orchestrate agents. It provides customers a single interface to partner-created agents in **agent gallery**, a no-code UI via **agent designer** to build, and **connectors** with major apps like Google Drive and Confluence.

Across BigQuery, Vertex AI, Agentspace, and our wider AI tech-stack, we aim to arm partners with a powerful toolkit to confidently build, deploy and scale truly transformational AI agents.




Exhibit 6.1: Google Cloud is your unrivalled AI partner, offering business users and developers a full set of tools to build agents in a way that works for them

Building a responsible agentic future

As we step into this new frontier, our commitment to develop and deploy AI responsibly is unwavering. As articulated in our [AI Principles](#), we must be both **bold and responsible**.

Being **bold** means rapidly innovating and deploying AI in groundbreaking products for the benefit and use of people everywhere. We will continue to contribute to scientific advancements, helping address humanity's most pressing challenges. Being **responsible** means deploying AI that addresses user needs while safeguarding their safety, security, and privacy. This commitment to responsible AI development is integral to the trust we build with our customers and partners.

Conclusion

The path forward: A call to action for Google Cloud's partners

The era of agentic AI is no longer a distant vision; it is rapidly becoming a present-day reality, bringing with it an unparalleled opportunity to reshape how businesses operate and how value is delivered. As this report has illuminated, the market for agentic AI services is substantial, estimated at ~\$1T globally. As such, this is not merely an incremental shift but a fundamental transformation. For Google Cloud's partners, this moment is a clear call to action: to move boldly, invest strategically, and collaborate closely to help customers navigate this exciting new landscape.

The journey ahead requires more than just adopting new technologies; it demands a new way of thinking and operating. To help partners define their strategic roadmaps in the agentic era, we highlight six key guiding principles across the transformation journey:

1. **Identify pain points and build prototypes:** Focus on solving high-value, customer-specific pain points to drive differentiation, delivering working prototypes at pitches to prove value from day one.
2. **Reimagine core business processes:** Evolve the delivery to go beyond pure technical and systems execution, with initial focus on upfront consultative AI design to help customers restructure business processes by co-designing agentic workflows.
3. **Use new tools to solve data gaps:** Create opportunities to quickly overcome clients' perceived data readiness challenges by using generative AI tools and modern inter-agent communication protocols to activate agents with less structured datasets.
4. **Deploy agents at scale:** Manage change and embed agents into business workflow, and establish processes to track agent performance. Bring robust change management capabilities to accelerate deployment and deliver clear ROI.
5. **Manage the full agentic lifecycle:** Deliver new forms of ongoing support, such as orchestrating agentic fleets and refreshing agent knowledge. Create long-term competitive advantage by investing in reusable IP and reference integrations.
6. **Innovate commercially:** Consider different pricing models to align with the nature of agentic AI. Explore recurring, transaction-based, or outcome-based structures that reflect the measurable impact agents deliver.

Google Cloud stands ready to help partners lead the charge across these themes. Our cutting-edge technologies across the full AI tech stack, open-source contributions, thought leadership, and comprehensive partner support can accelerate your agentic plans and drive success in bringing high-value, impactful agentic AI solutions to the market.



Appendices

Acknowledgements

This report was made possible through the generous contributions of time, expertise, and insight from a wide range of individuals and organizations. We extend our sincere gratitude to all who played a role in shaping this research and its findings.

Partner and industry insights

We are grateful to the numerous partners and industry leaders who participated in our interview process. Your firsthand experiences, deep market knowledge, and candid perspectives on the challenges and opportunities presented by agentic AI were invaluable in grounding our analysis in real-world applications and ensuring the relevance of our conclusions. Your willingness to share your insights has significantly enriched this report.

Google contributors

This work reflects a significant collaborative effort. We thank our colleagues from various teams across Google Cloud for their expert contributions. Their deep understanding of Google Cloud's agentic AI capabilities and vision was invaluable.

Boston Consulting Group (BCG) collaboration

We acknowledge the dedicated BCG team for their insights, rigor, and collaboration on the TAM estimates. Their expertise in market analysis and business transformation was key to this whitepaper.

Data sources

The quantitative analysis presented in this report is built upon a foundation of reputable data. We acknowledge the various public and proprietary data sources, including market research firms, governmental labor statistics organizations, and specialized data providers, whose information was essential for our market sizing analysis. Specific sources are cited where appropriate.

Detailed methodology for sizing agentic AI value pools

U.S. value pools

Our buildup of the addressable market size was calculated as follows:

01

Estimate the total economic value of work activities

We started by mapping the U.S. workforce across 150 job titles and 7 job families (admin, engineer, finance, marketing, operations, sales, and scientist) using data from Revelio. Revelio aggregates public employee and compensation information (e.g., LinkedIn), triangulated with U.S. Bureau of Labor Statistics data. For each role, we multiplied the total FTE counts with the average compensation (salary and benefits) to estimate the total economic value of work activities.

This economic value was broken down into six industry categories and an ‘other’ bucket: Technology/Telecommunications, Media, Entertainment, Gaming (Tech/TMEG); Financial Services/Insurance; Auto/Industrial; Healthcare/Life Sciences; Public Sector; Retail/Consumer Packaged Goods (CPG); and Other (e.g., business consulting, advertising, hospitality).

02

Estimate the potential value of augmentation with agentic AI

To estimate the potential for agentic AI augmentation, we estimated the distribution of time spent, for each job title, on a variety of task groups (~330 in total), and subsequently assessed the potential for augmentation with agentic AI for these tasks.

Role	Task	Task frequency (Surveyed frequency converted to hours)	Task frequency total (Sum of total hours)	% of time allocation (Task hours / total hours)
Role 1	Task 1	Hourly (~4 hours/day)	5.25 hours	76%
	Task 2	2x daily (~1 hour per day)	5.25 hours	19%
	Task 3	Weekly (~.25 hours per day)	5.25 hours	5%

For time distribution estimation, we utilized O*NET, a U.S. Department of Labor-funded database, which provides task-level frequency data for ~800 job titles. This provided us with an analytical pathway to estimate a unique distribution of time spent on various tasks for each job category. For presentation purposes, we grouped these underlying tasks into eight categories, shown in Exhibit B.1 below

Task group	Description	Illustrative tasks
Data collection, synthesis & documentation	Gathering, recording, processing, and updating data to understand events, actions, and information	• Maintain operational records • Maintain current knowledge or expertise • Evaluate quality/accuracy of data • Collect consumer needs/opinions • Prepare documentation • Gather operational/development data • Read documents/materials
Analytics & performance evaluation	Examining data, processes, and quality standards to make informed assessments, ensure compliance, and quantify outcomes	• Analyze business/financial data • Analyze data to improve operations • Analyze market/industry conditions • Monitor financial/health data • Evaluate product/technology performance • Monitor regulatory compliance
Creative thinking & problem solving	Interpreting information and thinking creatively to resolve problems and make decisions	• Create visual displays • Design IT systems/applications • Implement procedures/processes • Develop business/marketing/education plans • Authorize transactions • Explain technical details
Resource & operations management	Monitoring, controlling, and managing people, systems, tools, and processes to coordinate resources effectively	• Recruit/hire • Execute financial transactions • Process data • Purchase goods/services • Perform administrative/clerical tasks • Replenish inventory • Operate office equipment
Work preparation & planning	Developing objectives by organizing, planning, prioritizing, and scheduling work	• Develop policies/procedures • Schedule appointments/activities • Plan work/events/programs • Draft legislation/regulations
Supervision, training & development	Supervising, training, and coaching others to guide and improve skillset and performance	• Direct operations • Supervise personnel • Train on procedures • Provide support/encouragement • Teach academic/vocational subjects • Assign/coordinate tasks

Task group	Description	Illustrative tasks
Communication & relationships	Communicating with people inside and outside the organization to consult, sell, maintain interpersonal relationships, and provide care	• Inform guests/clients • Communicate on operations/specifications • Advise on business/products • Respond to customer inquiries • Develop professional networks
Physical/manual tasks	Performing physical activities via operating vehicles/equipment, handing objects, inspecting resources, and maintaining equipment	• Clean/inspect tools/facilities • Dispose of waste • Prepare mixtures/food/beverages • Operate vehicles/equipment • Move/package objects • Arrange displays

After we determined this distribution of task time, we conducted a bottom-up assessment to understand which tasks are likely augmentable by agentic AI, applying a framework based on gating criteria and technical fit factors. For each filter, a task could be rated “Yes,” “No,” or “Partial.”

Gating criteria:

1. Is the task carried out primarily through digital interactions rather than physical or manual actions?
2. Does the task largely rely on objective, transactional exchanges rather than nuanced emotional interpersonal judgement?

Technical fit factors:

1. Does the task follow a structured workflow without excessive ambiguity or open-ended reasoning?
2. Are the data inputs available and observable by a potential agentic AI system?
3. Is the task’s outcome covered by rule-based logic rooted in existing documentation or precedents?

Assessement of agentic AI task addressability

Questionnaire applied to each of ~330 underlying tasks

Steps	Questions	Task marked 'yes' when:	Task marked 'partial' when:	Task marked 'no' when:
01 Gating factors preventing agentic AI use	Is the task carried out primarily through digital interactions rather than physical or manual actions? Does the task largely rely on objective, transactional exchanges rather than nuanced emotional interpersonal judgement?	Is completed largely digitally <i>e.g., gather info from electronic sources</i> Is largely transactional, objective <i>e.g., read documents or materials</i>	Addressability may vary by how a task is structured or the specific context it is performed in This variability occurs across different instances in the same job role	Task required in-person, manual dexterity <i>e.g., prepare food and beverage</i> Task depends on emotional/social cues <i>e.g., investigate criminal/legal matters</i>
02 Task is a technical fit for agentic AI	Does the task follow a structured workflow without excessive ambiguity or open-ended reasoning? Are the data inputs available and observable by a potential agentic AI system? Is the task's outcome covered by rule-based logic rooted in existing documentation or precedents?	Clear, repeatable steps/logic <i>e.g., gather infor from electronic sources</i> Uses structured data that AI retrieves reliably <i>e.g., estimate project development costs</i> Has logic/decision trees that can be derived using existing datasets <i>e.g., test performance of computers</i>	<i>e.g., doctors "providing health advice" may range from a routine phone check-in to a complex diagnosis requiring emotional intelligence and nuanced human communication</i>	Task has number of undefined processes <i>e.g., develop organizational goals or objectives</i> Task uses unstructured/inaccessible data <i>e.g., monitor individual behavior</i> Task requires expert intuition or judgement <i>e.g., maintain safety/security</i>
Included in value pool sizing?		All 'yes' answers indicate high agentic AI applicability (85-90% of task time)	All 'yes' and some 'partial' indicates mod. applicability (40-50% of task time)	'No' on any question indicates low agentic AI applicability

Exhibit B.2: Task agentic AI addressability assessment criteria

We assigned tasks to different AI augmentation levels, and to embed conservatism, account for some time being retained by humans for higher-level tasks, and/or exception handling:

- Fully augmentable (85-90% of task time) if the task passes all five criteria with "Yes".
- Partially augmentable (40-50% of task time) if the task passes all five criteria with "Yes" or "Partial".
- Not augmentable (0% of task time) if the task fails any one factor ("No").

This structure allowed us to compute the percentage of each job's time for which agentic AI is applicable for each of the 150 job roles, and the related portion of labor value. An illustrative sample of these roles is shown in Exhibit B.3.

Agentic AI applicability — examples | High agentic AI applicability for data and problem-solving roles such as customer service & software engineer




Exhibit B.3: Agentic AI applicability for illustrative U.S. job titles

03 Estimating the application layer value pool

Next, we estimated the agentic AI application layer value pool, which is the portion of value that flows to agentic AI software products and could come in many forms, such as license, subscription, or usage fees that a customer pays, distinct from the consulting, implementation, and maintenance fees paid to SIs. Based on partner interviews and secondary research, we estimate that the application layer could capture approximately 10-30% of the total compensation dollars that could be augmented by agentic AI. Job categories with high agentic addressability (greater than 50% of total job time) bias toward the upper end of this range. For other roles, we assumed the lower end of this range, based on benchmarks from BCG and Salesforce case studies, as well as market participant insights on past software and gen AI deployments, which suggest application layer capture in the 10-15% range.

This resulted in a total U.S. application layer value pool of approximately ~\$1T.




Exhibit B.4: Estimated U.S. application value pool for agentic AI

04 Estimating the SI services value pool

Finally, we estimated the amount of value that can be translated from the application layer specifically to SI services, assuming a ‘full potential’ steady state scenario. As a starting point, we exclude Small and Medium-Sized Businesses (SMBs) and lower middle-market firms (<500 FTEs).

From what remains, we translate the spend on agentic applications to spend on SI services to advise on, implement, and maintain those applications.

At the project level, the ratio of SI services to application spend can vary widely depending on the complexity, customization, and scope of the deployment, whether agentic or otherwise. Case studies based on information from BCG and Forrester show that SI services spend can be as low as 0.2x the application spend for relatively simple, off-the-shelf customer service agents. In contrast, highly tailored implementations that involve full-stack modernization or deep workflow integration, such as traditional CRM/ERP deployments and bespoke agents built from the ground up, can drive SI spend up to 4x+ the application layer. This variance reflects how different each technology deployment can be, depending on use case, industry, and enterprise readiness.

However, when we zoom out to the market level, this variance will be averaged away. Market estimates from sources such as Gartner and IDC suggest that SI services spend typically lands in the range of \$0.70-\$0.90 per \$1 of application software spend, tending towards the higher end of this range when isolating for larger companies.³ This reflects the

average mix of advisory, implementation, and ongoing support work that accompanies enterprise-grade software adoption.

With these data points in mind, we use a midpoint assumption that \$0.70-\$0.80 will be spent on SI services per \$1 of agentic AI application spend. This range reflects a market-level steady state, balancing high- and low-intensity deployments, while assuming a mix of recurring and non-recurring services. In practice, this means that some share of advisory and implementation work will ‘recycle’ across customers and projects over time, much like in the broader enterprise SaaS market today.

Importantly, we also exclude spend associated with job roles that have low agentification potential (<20%), as these roles are less likely to require material SI involvement. Together, these assumptions yield an estimated total SI services value pool of ~\$350B to \$450B in the U.S., based on an application layer value pool of ~\$1T.




Exhibit B.5: Bridge from agentic AI application layer to the SI services value pool

3: Gartner: “Forecast: Enterprise IT Spending by Vertical Industry Market, Worldwide, 2023-2029, 1Q25 Update” ; IDC: “IDC Worldwide ICT Spending Guide Enterprise and SMB by Industry - Forecast 2025 | Feb (V1 2025)”




Exhibit B.6: Estimated U.S. SI services value pool for agentic AI

Global value pools

We expanded our U.S. value pool calculations to a global scale by applying our established U.S. value pool ratios to International Labor Organization (ILO) data. This scale-up methodology involved a series of steps to estimate the labor value addressable by agentic AI, the application layer value pool, and ultimately the SI services value pool across different regions.

01 Calculating the global economic value of work activities

We determined each country's total labor value using ILO data, starting with the number of employees in each International Standard Classification of Occupations (ISCO) job role and industry. This covered approximately 3 billion employees, or 95% of global wage workers. We then multiplied the number of employees by the average salary. Salary figures were grossed up using US-calculated salary-to-compensation ratios to arrive at total compensation value.

02 Determining the potential global value of augmentation with agentic AI

We calculated the percentage of time augmentable by agentic AI for each global job role. This involved converting our U.S. 150-role taxonomy to nine ISCO roles. We then applied the U.S.-derived ISCO role agentification rates to the global dataset.

03 Estimating the global application layer and SI services value pool

We applied the same assumptions for application layer value capture and translation into SI services as in our detailed U.S. sizing methodology.

Through these steps, we calculated a global \$0.9-1.1 trillion SI services value pool. North America and EMEA represent the largest opportunities, driven by their significant application layer value pools and knowledge worker-oriented economies.

Global SI services value pool for agentic AI, by region




Exhibit B.7: Estimated global SI services value pool for agentic AI

Estimated agentic AI value pools




Exhibit B.8: Estimated global SI services value pool for agentic AI, by region

Disclaimers



This document has been prepared in good faith on the basis of information available at the date of publication without any independent verification. Authors do not guarantee or make any representation or warranty as to the accuracy, reliability, completeness, or currency of the information in this document, nor its usefulness in achieving any purpose. Readers are responsible for assessing the relevance and accuracy of the content of this document. It is unreasonable for any party to rely on this document for any purpose, and authors will not be liable for any loss, damage, cost, or expense incurred or arising by reason of any person or entity using or relying on information in this document.

To the fullest extent permitted by law, authors shall have no liability whatsoever to any party, and any person using this document hereby waives any rights and claims it may have at any time against authors with regard to the document. Receipt and review of this document shall be deemed agreement with and consideration for the foregoing. This document does not purport to represent the views of the companies mentioned in the document. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by authors.

This document includes information from the O*NET 29.3 Database by the U.S. Department of Labor, Employment and Training Administration (USDOL/ETA). Used under the CC BY 4.0 license. O*NET® is a trademark of USDOL/ETA. Authors have modified all or some of this information. USDOL/ETA has not approved, endorsed, or tested these modifications.

