SalesChoice

AGENTIC AI

The Next Frontier of Enterprise Transformation

An Executive White Paper on Growth, Value Creation, Frameworks, Governance, Data Foundations, and Strategic Opportunity

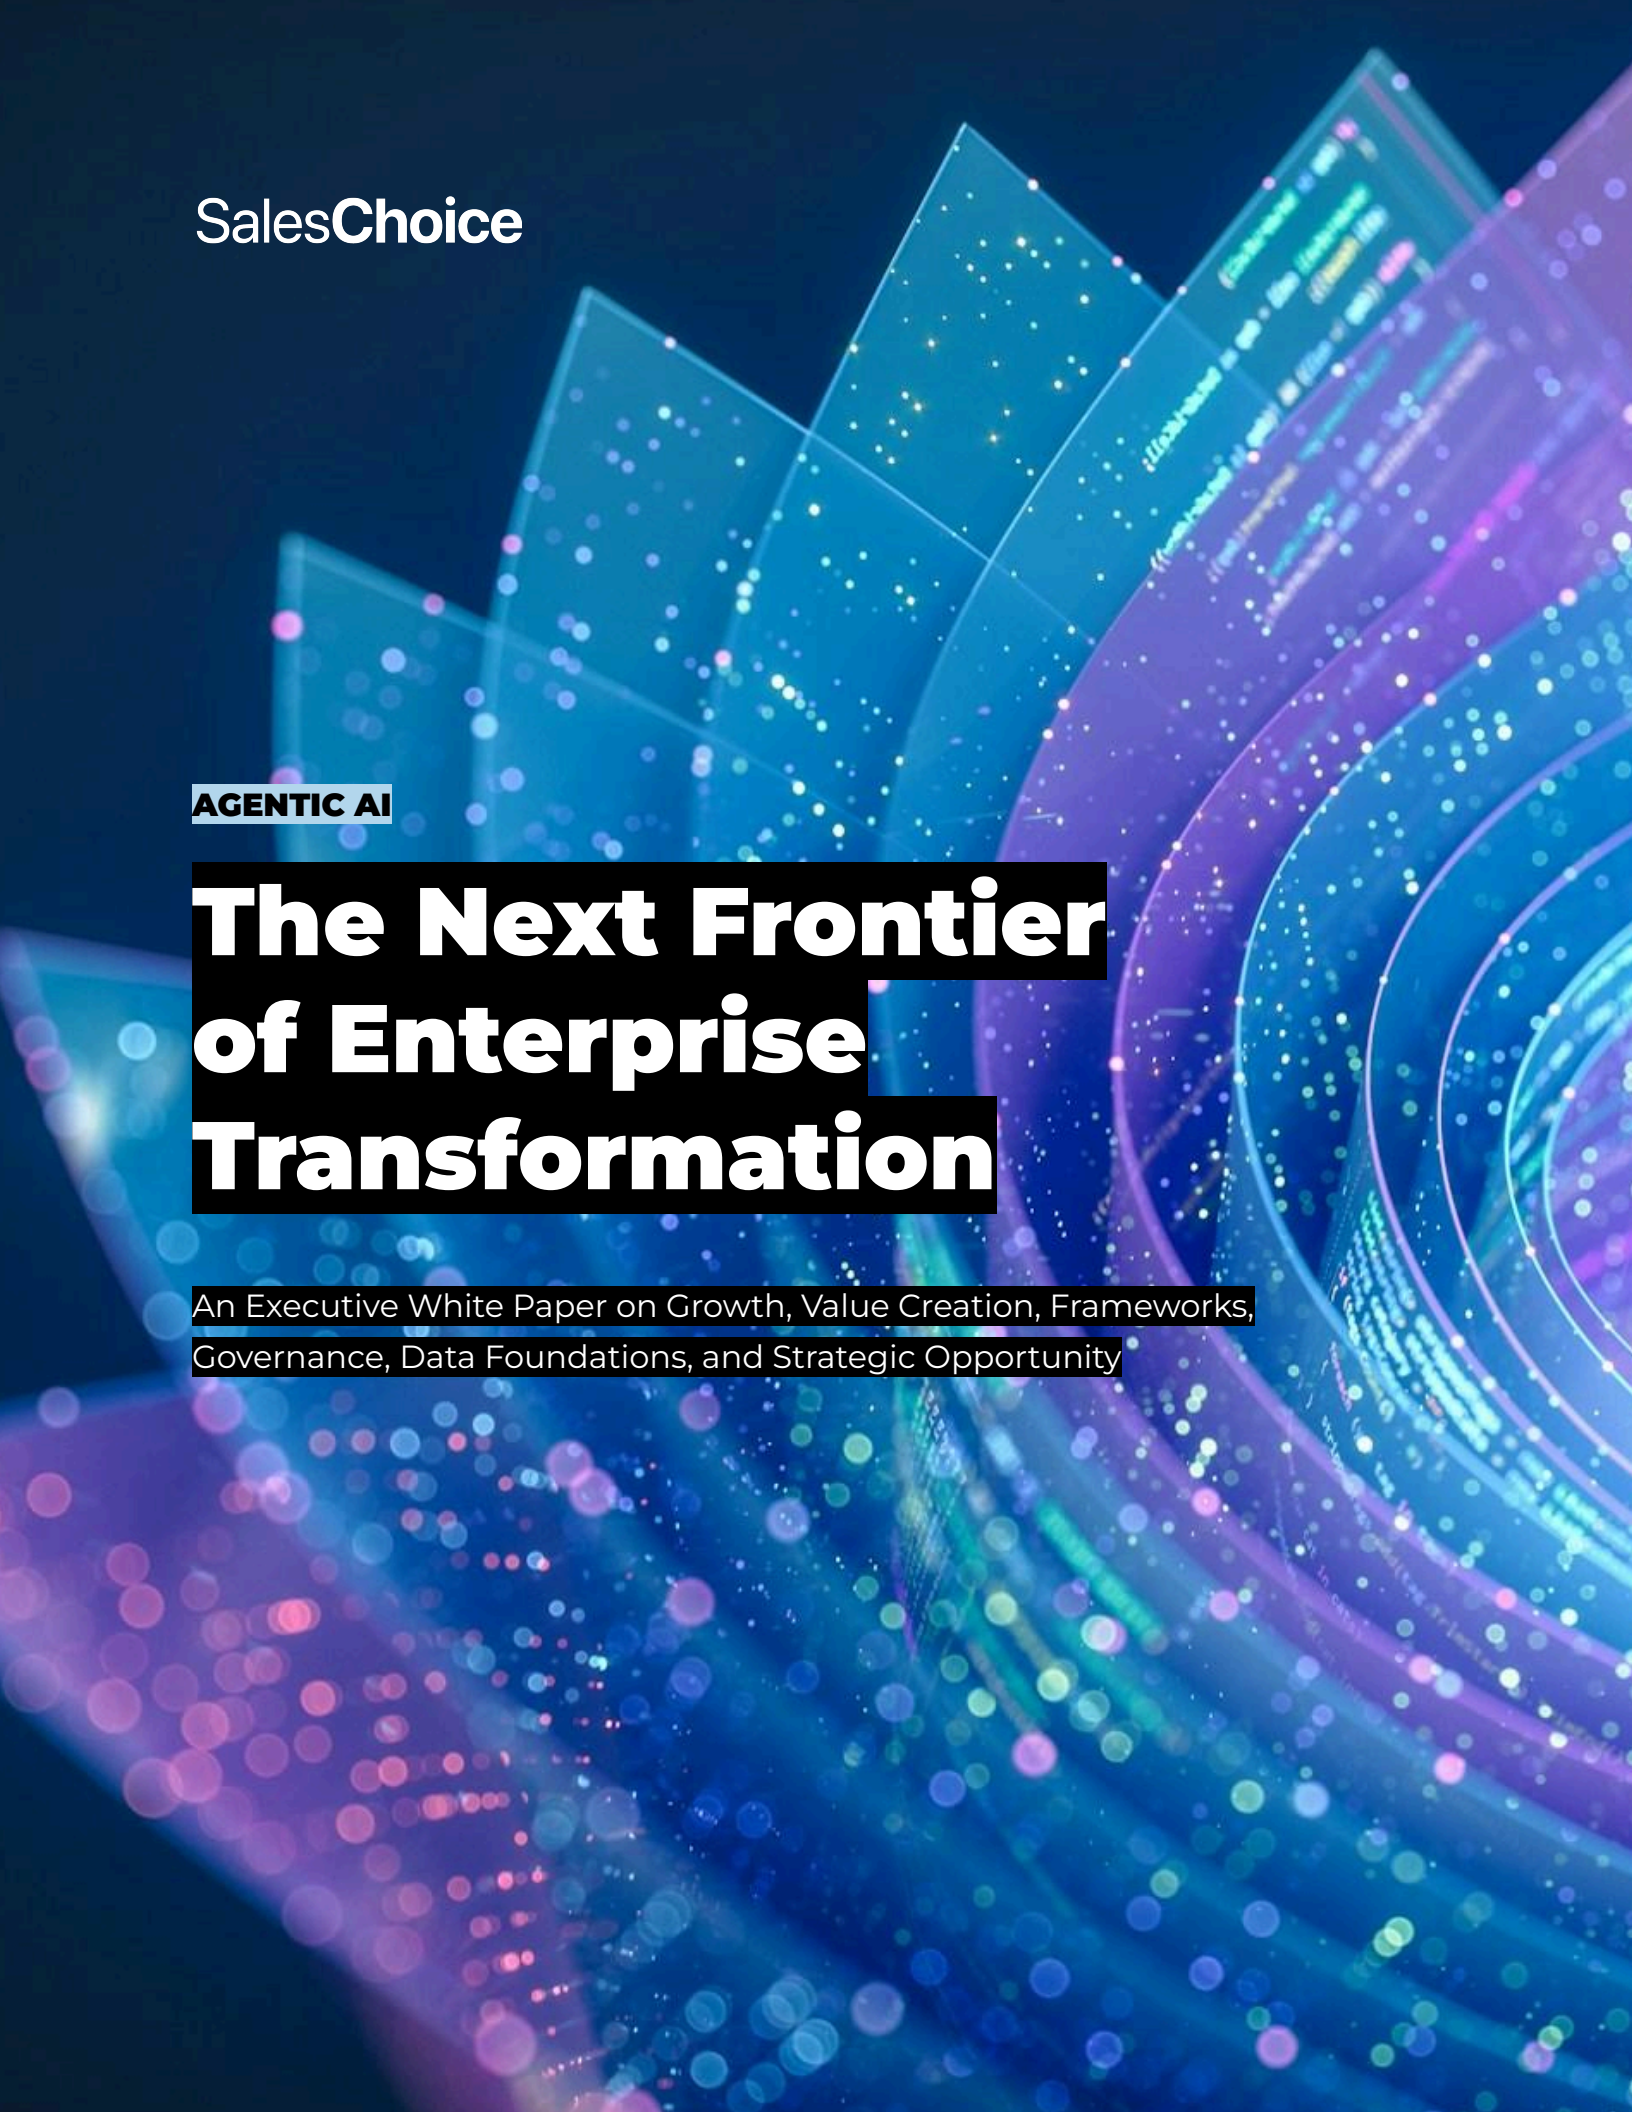

Contents

Executive Summary ____ 3

What Is Agentic AI? ____ 4

Why is Agentic AI growing rapidly? ____ 5

Why does Agentic AI matter strategically? ____ 5

Why is data readiness so crucial? ____ 8

Agentic AI Frameworks: leading options ____ 11

How can you get started with an Agentic roadmap? ____ 12

What enterprise companies are advancing with Agentic AI? ____ 14

What is the strategic role of SalesChoice in Agentic AI? ____ 15

What are Agentic governance risks and critical success factors? ____ 16

What governance questions should boards and CEOs plan for? ____ 18

Conclusion ____ 19

Author ____ 20

References ____ 21

Executive Summary

Enterprise AI is shifting from analytics and isolated assistance to Agentic AI: autonomous systems that pursue goals, reason across steps, interact with tools, and execute workflows with limited human intervention. This evolution promises significant gains in productivity, speed, accuracy, customer response, and cost efficiency by managing processes and coordinating decisions, not just producing outputs.

However, this autonomy introduces substantial governance and execution risks. Agentic AI's strength is limited by its underlying data environment.

Traditional governance models designed for human-directed AI are insufficient, as agentic systems can initiate actions and operate across multiple systems in real-time.

Success requires combining innovation with strong data discipline, transparent controls, human oversight, measurable ROI, and ethical standards to responsibly manage autonomous AI.

What Is Agentic AI?

Agentic AI refers to artificial intelligence systems that function as autonomous or semi-autonomous agents capable of interpreting goals, decomposing tasks, selecting tools, reasoning through multiple steps, learning from feedback, and completing workflows with minimal direct supervision.

Traditional AI systems generally require explicit prompts and deliver bounded outputs. **Agentic systems, by contrast, are oriented toward objectives and can determine how best to achieve them.**

An enterprise example would be instructing a system to improve quarterly sales forecast accuracy with specific actions in a relevant context. A conventional AI tool might generate a report. An agentic system would analyze pipeline data, detect risk patterns, engage with CRM tools, recommend interventions, monitor results, and continuously refine its approach over time. It could even surface real-time coaching insights, relevant to the sales motion in play.

Table 1.0: Traditional AI vs. Agentic AI

Dimension	Traditional AI	Agentic AI
Primary Role	Responds to prompts	Pursues goals
Workflow Scope	Single-task outputs	Multi-step execution
Tool Usage	Limited	Extensive
Memory	Session-based	Persistent/contextual
Human Oversight	High	Moderate to high
Business Impact	Productivity	Operating model transformation

Why is Agentic AI growing rapidly?

The agentic AI market is experiencing explosive growth, projected to grow roughly 10x between 2024 and 2030, with some forecasts suggesting it will grow at a compound annual growth rate (CAGR) of over 40% to 175%, depending on the analyst. Fortune Business Insights projects that the global Agentic AI market will grow from approximately USD $9.1 billion in 2026 to USD $139 billion by 2034 and nearly $200 billion by 2034 (Fortune Business Insights, 2026).

The growth of Agentic AI is being driven by several structural market forces:

- Enterprises increasingly require measurable returns from AI investments and are moving beyond experimentation toward the automation of real workflows.
- At the same time, organizations are managing complex digital ecosystems composed of disconnected software platforms.
- Agentic AI can bridge these environments by operating across process value chains like CRM (Customer Relationship Management), ERP (Enterprise Resource Planning), HRM (Human Resource Management), and PRM (Partner Relationship Management), etc.
- Recent advances in large language models (LLMs), tool use, memory layers, orchestration frameworks, and enterprise APIs have also made autonomous systems more feasible.

What are the drivers of Agentic Growth?

- **Shift from Chatbots to Action:** Unlike generative AI that produces text, agentic AI acts as autonomous software that can plan, reason, and execute complex, multi-step tasks across systems.
- **Multi-Agent Systems:** Rise in platforms allowing developers to create teams of agents that work together (e.g., AutoGen, CrewAI), enhancing efficiency.
- **High ROI:** Early adopters report up to 50% reduction in support costs and significant improvements in employee productivity
- **Major Tech Player Investment:** including Anthropic, Apple, Cohere, Google, Meta, Microsoft, NVIDIA, OpenAI, and Salesforce, to name a few, are investing heavily in agentic frameworks.

Why does Agentic AI matter strategically?

Agentic AI transforms the economic equation of enterprise AI. Most organizations do not suffer from a shortage of dashboards or reports. They suffer from delays in execution, fragmented ownership, and inefficient coordination across functions.

The dominant AI use cases underway in the market are: code development, virtual assistants, and customer service, expected to reach a market of over $48.2 B by 2030. The stakes are high. According to CB Insights, customer service is the #1 area of AI agent adoption in their Q4'25 enterprise survey. 115+ companies operate in the market (putting it in the top 10 of all CBI-tracked markets), and six private companies are generating $100M+ in revenue.

While the majority of private market revenue leaders were founded before the genAI wave, AI-native entrants are gaining ground fast. Companies like Sierra (founded in 2023) and Crescendo (founded in 2024) have already broken into the top 10 by revenue generation.

The fastest growing region in adoption is the USA, but the Asia Pacific is emerging as the fastest growing region according to Intel Market Research. While growth is high, Gartner Research indicates that the majority of organizations are in the early emergence and testing phase, with only 17% having deployed agents as early as 2026.

Sales organizations can detect deteriorating pipeline quality, reprioritize opportunities, and initiate follow-up actions. In finance, they can reconcile anomalies and accelerate close processes. In customer service, they can resolve issues, escalate exceptions, and continuously optimize service quality.

In human resources, they can support onboarding, workforce planning, and engagement processes. Table 2.0 summarizes some of the functional enterprise use cases.

Table 2.0: Functional Enterprise Use Cases

Function	Agentic AI Use Case	Value Created
Sales	Forecasting and opportunity prioritization	Revenue growth
Finance	Reconciliation and anomaly detection	Lower costs
Customer Service	Automated case resolution	Faster response
HR	Onboarding and workforce analytics	Better engagement
Operations	Workflow orchestration	Productivity gains

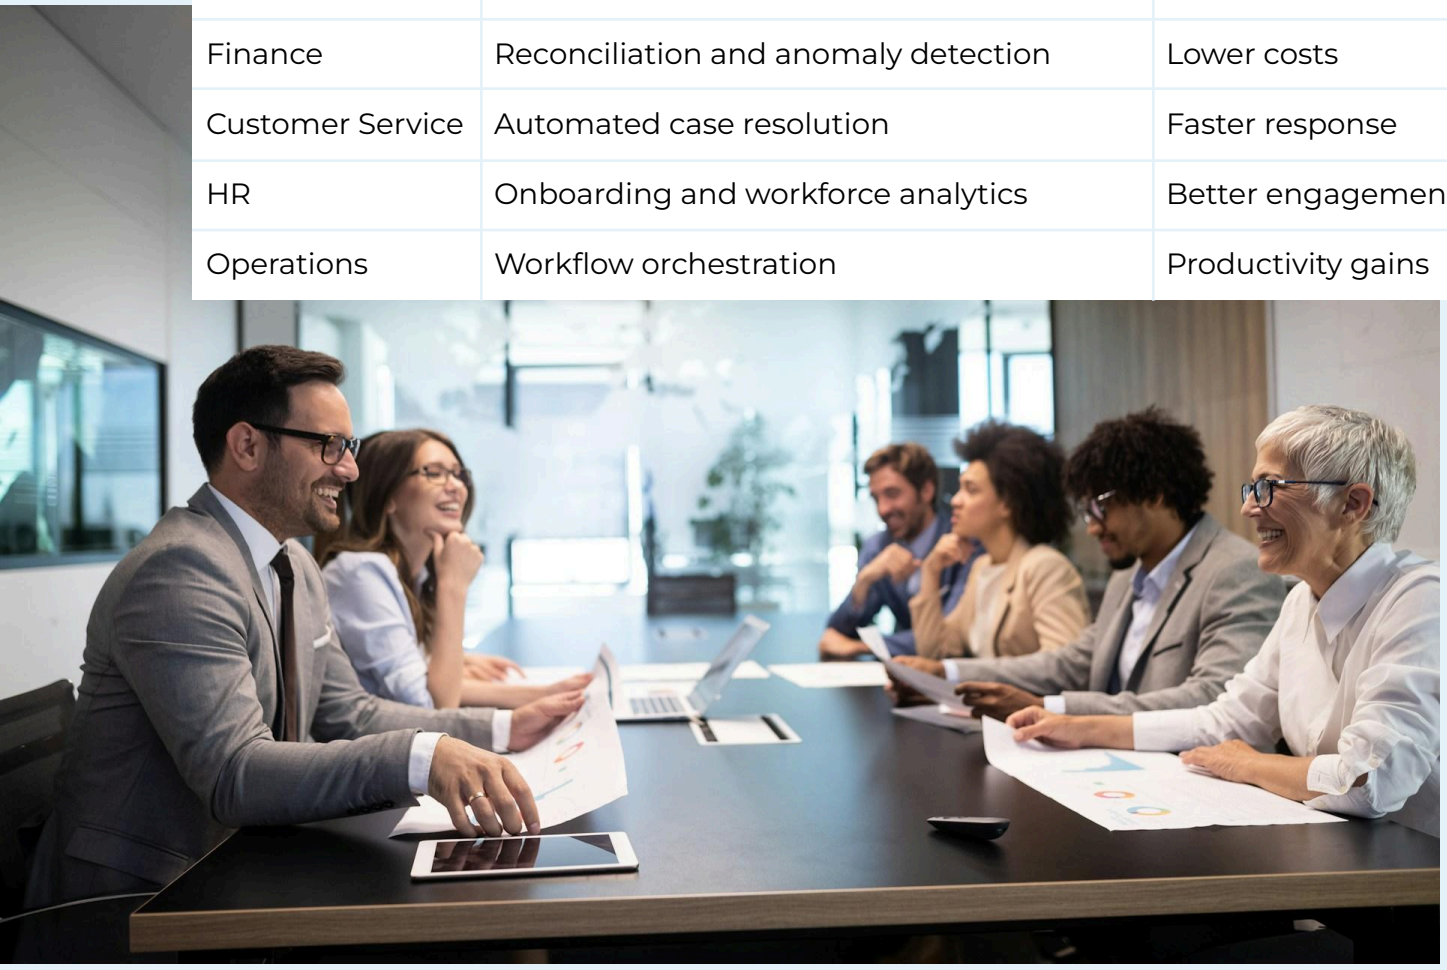

Why is data readiness so crucial?

McKinsey research (2026) on Agentic found that "*nearly two-thirds of enterprises worldwide have experimented with agents, but fewer than 10 percent have scaled them to deliver tangible value. Shaky data is often to blame; eight in ten companies cite data limitations as a roadblock to scaling agentic AI. Addressing this issue is a core element of building a solid capability foundation—and that's what distinguishes companies that create value from AI from those that don't - and data readiness is core to execution success.*"

Our SalesChoice research has also found that the majority of mid-market companies and many large enterprise organizations have weak master data records, fragmented systems, duplicate records, poor metadata, and inconsistent definitions, which collectively are undermining autonomous AI performance. Hence, the emergence of Agentic AI is marginalized.

The most successful adopters will combine innovation with strong governance, clear ROI measurement, a modern data foundation, and a human-centered operating model. Humans in the loop is also crucial for Agentic AI foundations to set up a living organization to learn from, as routine operating processes are redesigned and simplified, so functions can, in time, be performed independently

Many organizations approach Agentic AI as a model or software decision. In reality, the first strategic question should be whether the enterprise data estate is ready for autonomous decision systems.

Agents depend on data to perceive context, make decisions, and trigger actions. If customer records are duplicated, product hierarchies are inconsistent, pipeline stages are unreliable, or financial records are delayed, the agent's decisions will also be flawed. The traditional phrase "**garbage in, garbage out**" becomes more serious in Agentic AI. Poor data quality does not merely produce a weak dashboard. It can produce **incorrect autonomous actions at scale**. Enterprises, therefore, need to strengthen foundational data disciplines before expecting agentic systems to deliver value. **Context is increasingly becoming one of the most valuable assets in the AI Age.**

Table 3.0: Data Issues That Undermine Agentic AI

Data Weakness	Operational Impact
Duplicate customer records	Incorrect actions / poor personalization
Inconsistent definitions	Conflicting decisions
Missing metadata	Weak context awareness
Delayed data feeds	Slow or outdated actions
Poor CRM hygiene	Unreliable sales agents
Fragmented systems	Incomplete reasoning

Success with agentic AI depends on a data architecture that can support increasing levels of autonomy, coordination, and real-time decision-making. This often looks like modular, interoperable frameworks that give agents reliable access to the data they need to operate safely.

While GenAI has already shown the need for data access control, lineage, and traceability, agentic platforms place greater operational pressure on these foundations. Because agentic AI coordinates multiple models and data sources continuously, often without human intervention, it requires tighter, more automated governance to ensure reliability and control at scale.

To enable a scaled transformation into an agentic organization, companies must start by building foundational data capabilities. This requires not just a technology reboot, but also an organizational one. That's because a company's data strategy and operating model are just as important as its underlying data quality and architecture. McKinsey (2026) outlined four coordinated steps that link strategy, technology, and people:

- **Step 1: Identify high-impact workflows to "agentify."** Organizations can identify a small number of high-value, end-to-end workflows where increased autonomy could unlock impact. Building on established approaches to building data products, leaders can prioritize agentic use cases based on value potential, feasibility, and strategic fit before scaling more broadly.
- **Step 2: Modernize each layer of the data architecture for agents.** Rather than rebuilding everything from scratch, leaders can modernize existing platforms to support interoperability and governance across systems. While some may be tempted to lean on advancements in AI to shortcut data architecture best practices, the strongest organizations build modular, evolutionary architectures with components that can be replaced as new technologies emerge.
- **Step 3: Ensure that data quality is in place.** Organizations must move from periodic data cleanup to continuous, real-time quality management. They can do this while ensuring that both structured and unstructured data, as well as agent-generated outputs, meet consistent standards for accuracy, lineage, and governance.
- **Step 4: Build an operating and governance model for agentic AI.** Scaling agentic AI requires rethinking how work gets done. Human roles are shifting from execution to supervision and orchestration of agent-driven workflows. In a hybrid human–agent work environment, clear governance is essential to allow agents to operate transparently and safely at scale.

Notation: See McKinsey Research for more agentic data insights here.

Agentic AI Frameworks: leading options

Several frameworks have emerged to help enterprises design and deploy agentic systems. LangChain and LangGraph offer flexibility and orchestration depth.

CrewAI is useful for fast experimentation. Microsoft AutoGen and Semantic Kernel are strong in Microsoft-centric environments. OpenAI Agents support rapid deployment. Custom frameworks may suit highly strategic use cases. Table 4 provides a comparison of leading Agentic frameworks with strengths, weaknesses, and best fit perspectives.

Table 4.0: Framework Comparison

Framework	Strengths	Weaknesses	Best Fit
LangChain / LangGraph	Flexible ecosystem	Complexity at scale	Custom enterprise builds
CrewAI	Fast experimentation	Less enterprise maturity	Innovation teams
AutoGen	Strong orchestration	Higher engineering needs	Advanced teams
Semantic Kernel	Microsoft integration	Narrower ecosystem	Azure enterprises

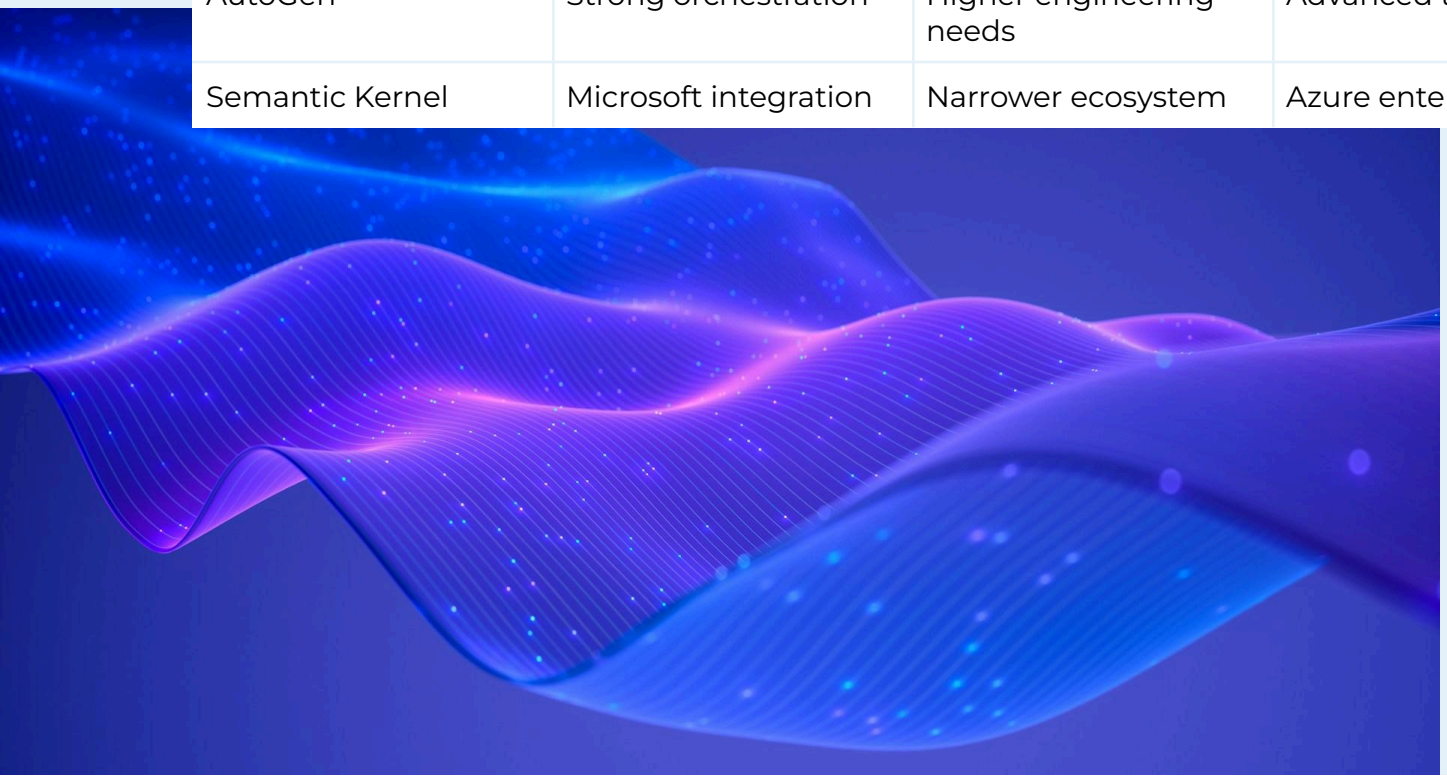

How can you get started with an Agentic roadmap?

Organizations should begin with a disciplined transformation roadmap that places data readiness before agent scale.

The first phase should focus on strategic use-case selection and enterprise data assessment. Companies should identify workflows where measurable value exists, while simultaneously evaluating data quality, master data maturity, system integration, taxonomy consistency, and governance controls.

The second phase should involve data remediation and architecture modernization. This includes cleansing records, harmonizing definitions, improving metadata, connecting systems through APIs, and establishing real-time data pipelines.

The third phase is crucial as it involves rethinking the current workflow and ensuring that agentic is not just being put into old, outdated operating processes and procedures. Taking the time to evolve and take strategic leaps forward in how work is being done is often a missing step in the majority of AI deployments, which is reflected in the low value realization outcomes on AI investments.

For most organizations, the path to generating value from agentic AI starts not with redesigning everything at once, but with deliberately rewiring a few critical workflows in high-impact domains. Leading companies identify domains such as knowledge management or marketing, analyzing existing data to pinpoint where the greatest value lies and where increased autonomy could materially change business outcomes. This requires first mapping end-to-end workflows and then identifying steps where agentic capabilities could add value and the corresponding data needed for those agents to do their tasks. This helps companies prioritize use cases based on their value, potential, and feasibility

Only after the data foundation is stable and the new workflow is designed should organizations move into pilot deployments. Initial pilots should focus on two or three use cases where value can be demonstrated within three to six months. Successful pilots should then be integrated across enterprise systems and monitored through clear ROI metrics.

The final phase involves scaling through an AI Center of Excellence that governs standards, controls, talent, and investment priorities. Depending on the size of the organization, an AI COE governance structure can be established at the onset of investing in Agentic AI enablements or aligning with existing COE operating practices.

Table 5.0 summarizes an illustrative 12-month Agentic Roadmap of key activities to advance an Agentic strategy.

Table 5.0: Recommended 12-Month Agentic Roadmap

Quarter	Priority
Q1	Use-case selection and enterprise data readiness assessment to identify foundational priorities and investments for agentic agent success.
Q2	Data cleansing, integration, governance, and architecture upgrades to have access to Agentic Frameworks and toolkits.
Q3	Redesign Workflow Logic and launch pilots with strong data controls and measure ROI
Q4	Align with the AI Center of Excellence or establish a governing structure to continue to monitor all Agent activity and associated use cases. Continue to monitor the ROI of Agentic solutions to ensure executive alignment and stakeholder buy-in.

What enterprise companies are advancing with Agentic AI?

Microsoft has embedded autonomous capabilities within Copilot and the Azure ecosystems. Salesforce has introduced Agentforce within CRM workflows.

Amazon Web Services has advanced Bedrock agent capabilities. Major financial institutions, including JPMorgan Chase, Goldman Sachs, Citi, and Morgan Stanley, have discussed using AI agents for productivity and operations (Business Insider, 2026). It is early days in using agentic frameworks across large enterprises, but the experimentation is underway, and we will learn from our mistakes - often the reality of being human.

Table 6.0: Illustrative Corporate Leaders

Company	Agentic Focus Area	Reported Value
Microsoft	Productivity automation	Faster workflows
Salesforce	CRM agents	Improved customer operations
AWS	Cloud orchestration	Scalable automation
JPMorgan	Internal productivity	Efficiency gains

What is the strategic role of SalesChoice in Agentic AI?

SalesChoice is well-positioned to support our customers' Agentic AI strategy because we bring over 15 years of proven predictive intelligence, forecasting accuracy, opportunity scoring, and workforce sentiment insight. In revenue operations, SalesChoice's SalesInsights™ product innovation, an app on the Salesforce App Exchange, can strengthen agents through predictive opportunity scoring, pipeline risk detection, forecast confidence modeling, and next-best action recommendations. **We are evolving our core capabilities to be more "agentic alert" - our phrase for contextual intelligence.**

In workforce performance, MoodInsights™ provides a differentiated capability by helping leaders monitor employee sentiment in real time, identify burnout risks, and improve organizational resilience.

Generative AI analyzes in real-time all the verbatim submissions, and anonymizes with high-impact clustering relevance for more effective executive decision-making

Table 7.0: SalesChoice Strategic Agentic Contribution

Capability	Strategic Value
Opportunity Scoring	Better pipeline prioritization
Forecast Accuracy	Stronger planning decisions
Predictive Analytics	Higher ROI from agents
MoodInsights™	Improved engagement and resilience
Executive Dashboards	Better enterprise decisions

SalesChoice also has an advanced AI and Data Sciences advisory practice where we help global companies design their Agentic AI Strategies, select the optimal architectural solutions to modernize their technology,, and help advance successful MVPs and pilot orchestration. Our depth of expertise in digital and change transformation continues to derisk our customers' AI Investments.

What are Agentic governance risks and critical success factors?

The governance challenge in Agentic AI is structurally different from earlier AI waves. If a chatbot generates an imperfect summary, the consequence may be limited. If an autonomous agent reprices products incorrectly, escalates the wrong customer cases, makes flawed procurement recommendations, or acts on inaccurate data across thousands of transactions, the impact can be immediate and material.

Agentic AI, therefore, shifts governance from a model-risk issue to an **enterprise operating risk issue.**

Leaders must ensure that autonomous systems remain aligned to business objectives, regulatory obligations, financial controls, and human values.

Many Agentic AI initiatives will fail because organizations underestimate governance complexity. Poor data quality, unclear ownership, weak security controls, opaque decision logic, and absence of ROI metrics remain common failure points. Successful organizations define business outcomes in advance, strengthen data governance, retain humans in the loop for consequential decisions, continuously monitor agent performance, and invest in change management. Table 8.0 summarizes common Agentic governance risks and potential business impacts.

Table 8.0: Common Agentic Governance Risks and Business Impacts.

Governance Risk	Description	Potential Business Impact
Poor Data Quality	Inaccurate or incomplete data drives flawed decisions	Revenue leakage, operational errors
Lack of Accountability	No clear owner of AI actions or outcomes	Slow remediation, legal exposure
Opaque Logic	Inability to explain actions taken	Regulatory scrutiny, trust erosion
Security Exposure	Excessive system access or weak controls	Cyber incidents, fraud
Ethical Bias	Unfair or discriminatory outcomes	Reputation damage, fines
Runaway Automation	Autonomous actions scale errors rapidly	Service disruption
Weak ROI Tracking	No measurement of outcomes	Wasted investment

What governance questions should boards and CEOs plan for?

Agentic AI should be positioned and viewed as both a growth agenda and a control agenda. Governance discussions should move beyond technical curiosity toward strategic oversight. Board Directors and CEOs should ensure there is optimal risk management with agentic solutions and answer these five questions on a quarterly basis.

1	Where are autonomous agents currently making decisions?
2	What financial or operational value has been realized?
3	What incidents, exceptions, or control breaches occurred?
4	Is data quality sufficient for scaled autonomy?
5	Are employees and customers maintaining trust in AI-enabled operations?

Organizations that cannot answer these questions based on SalesChoice research and experiences are scaling faster with stronger risk management practices.

Conclusion

Agentic AI represents the next frontier of enterprise transformation. It shifts artificial intelligence from generating outputs to achieving outcomes. Yet no autonomous system can outperform the quality of the enterprise data on which it depends. In the agentic age, technology leaders are finding that their data foundations increasingly define competitive positioning. Yet despite the promise of using data to generate real value from agentic AI, many organizations still struggle to make data accessible and governable for digital agents. The time has come to drive data transformations that pave the way for an agentic future.

Organizations that succeed will recognize that Agentic AI is not only an AI strategy. It is equally a data, governance, and operating model strategy. The imperative for leaders is clear. Build the data house first. Cleanse, govern, integrate, and modernize the enterprise information environment. It is crucial to rethink how work should be performed versus automating inefficient processes more rapidly. Then deploy agents into workflows where value is measurable and trust can be maintained.

Without trust in AI, investments in agentic AI won't scale. In this context, SalesChoice can play a meaningful strategic role for customers by enabling smarter revenue decisions, stronger forecast accuracy, and healthier human performance at scale. SalesChoice's robust AI Advisory services, with its proven AI Strategy, governance, and solution offerings, further help organizations build sustainable AI solutions and deliver a positive ROI.

Notation: See SalesChoice's Case Studies here.

Author

Dr. Cindy Gordon, CEO & Founder

Dr. Gordon is the CEO and Founder of SalesChoice. Author of 16 books, the most recent being *The AI Precipice: An Executive Guide to Innovation and Safety*, and hundreds of articles and research reports. She is the recipient of more than twenty AI awards. She is a former senior partner or executive with scaling experience from Accenture, Citicorp, and Xerox. She has also been a Venture Capitalist and is an active board advisor of diverse AI companies. Dr. Gordon has been advancing AI digital transformation strategies, governance, and building innovative AI solutions to ensure AI Value Realization is sustained for over 20+ years. She has a doctorate in Applied AI and teaches AI at the United States AI Institute (USAII) and other academic institutions.

Follow Cindy on **LinkedIn →**

References

Amazon Web Services. (2026). *Amazon Bedrock agents and enterprise AI orchestration.* Retrieved April 18, 2026, from https://aws.amazon.com/

Business Insider. (2026). *Wall Street banks accelerate AI agent strategies across operations and productivity functions.* Retrieved April 18, 2026, from https://www.businessinsider.com/

CB Insights.(2026). Agent Predictions-2026. Retrieved April 19, 2026, from: https://www.cbinsights.com/reports/CB-Insights_Agent-Predictions-2026.pdf?utm_campaign=5026645-Q1%202025%20Enterprise%20Tech&utm_medium=email&_hsenc=p2ANqtz-9-VSelvO9YWakXdel5Liu17ZY8tDiJIypRi6UkNKVlav_gxWpzZoS6VAmNURZHQAkZ1heU0hdC6jK6XwAoiZIJ58ic9Q&_hsmi=385295217&utm_content=385295217&utm_source=hs_automation

Deloitte. (2025). *Scaling trustworthy AI in the enterprise.* Retrieved April 18, 2026, from https://www.deloitte.com/

Fortune Business Insights. (2026). *Agentic AI market size, share, and industry analysis 2026–2034.* Retrieved April 18, 2026, from https://www.fortunebusinessinsights.com/

Gordon, C (2026). Strategy, Governance and Operating Models for an AI Center of Excellence. Retrieved April 20,2026, from https://www.saleschoice.com/resources/whitepapers/whitepaper-scaling-ai-responsibly/

McKinsey & Company. (2025). *The state of AI: Global survey on enterprise adoption and value realization.* Retrieved April 18, 2026, from https://www.mckinsey.com/

McKinsey & Company (2026). Building the Foundations for Agentic AI at Scale, Retried April 19 from https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/building-the-foundations-for-agentic-ai-at-scale

Microsoft. (2026). *Copilot and Azure AI enterprise automation announcements.* Retrieved April 18, 2026, from https://www.microsoft.com/

Salesforce. (2026). *Agentforce and enterprise AI agent strategy updates.* Retrieved April 18, 2026, from https://www.salesforce.com/

Let's Talk AI

THAT WORKS FOR YOU

Get in touch to learn how SalesChoice can help you:

- ✔ Plan your AI journey to future-proof your investments
- ✔ Forecast sales with increased confidence
- ✔ Clean and modernize your Salesforce
- ✔ Capture the opportunity of AI with real-time coaching
- ✔ Unleash the latent power of your workforce
- ✔ Make data-driven decisions you trust
- ✔ Build agentic AI routines for smarter workflows
- ✔ Ensure explainable and responsible AI

Book a call at **SalesChoice.com**

Prefer to email? Say hi at **hello@saleschoice.com**