



# The Strategic Imperative of Agentic AI





*If you're looking to move beyond basic assistance from AI to truly transforming your business operations with AI, it's time to understand and adapt Agentic AI.*

## Table of Contents

<b>Executive Summary</b>	<b>01</b>
<b>1. What is Agentic AI?</b>	<b>01</b>
<b>2. The Market Potential and Strategic Value</b>	<b>02</b>
<b>3. Comparison with other forms of AI</b>	<b>03</b>
<b>4. Reference Architecture</b>	<b>04</b>
<b>5. Technology Landscape</b>	<b>04 - 05</b>
<b>6. The Build vs. Buy Dilemma</b>	<b>05</b>
<b>7. Agentic AI at work</b>	<b>06</b>
<b>8. Risks and Mitigations</b>	<b>06</b>
<b>9. A Guide to Getting Started</b>	<b>07</b>


# Executive Summary

In a rapidly evolving digital landscape, enterprises are moving beyond simple automation to embrace a new paradigm: Agentic AI. This advanced form of AI transcends mere content generation, functioning as an **autonomous, goal-driven engine** engineered to execute end-to-end business functions. By leveraging Generative AI, Agentic AI, also known as Vertical AI or Intelligent Digital Workers (IDWs), provides a powerful mechanism for a company-wide operational overhaul. This paper explores the core components of Agentic AI, its transformative market potential, and a strategic framework for its successful implementation, positioning it as a critical asset for modernizing operations, driving growth, and securing a competitive edge.

## 1. What is Agentic AI?

Agentic AI represents a significant leap forward from traditional AI applications. While earlier AI models like Copilots and Robotic Process Automation (RPA) are focused on assisting human workers or automating isolated, repetitive tasks, Agentic AI is designed to act with a higher degree of autonomy. It is defined by its ability to "sense, reason, and act" across complex enterprise processes, seamlessly integrating with existing systems.

This technology is poised to move beyond content-centric Generative AI applications to tackle more sophisticated, goal-oriented business functions such as, but not limited to:



### Financial Services

Autonomous portfolio management, real-time fraud prevention, and trade execution.



### Healthcare

Automated claims processing, patient scheduling, and billing management.



### Retail

Personalized merchandising, dynamic pricing models, and optimized inventory management.



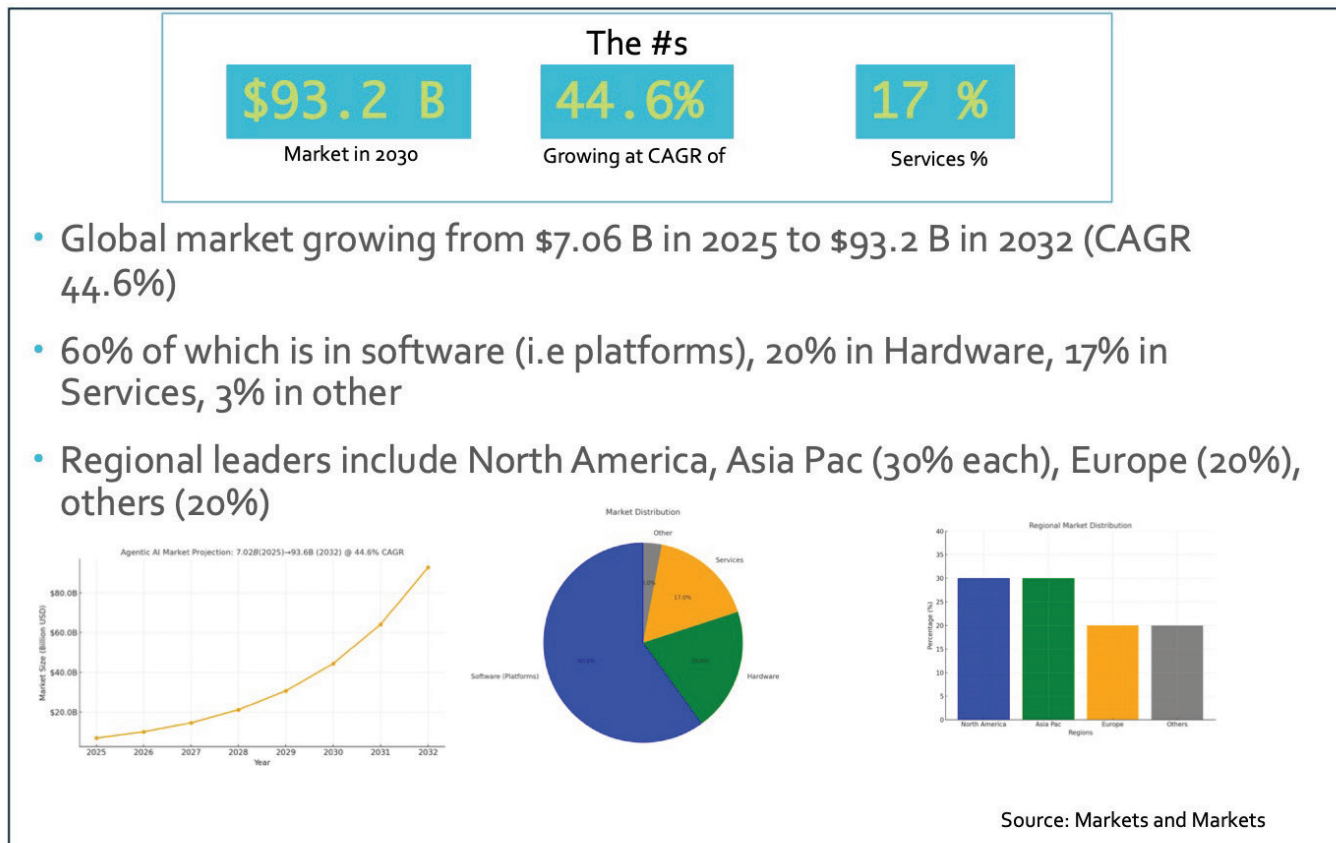
### Logistics & Manufacturing

Real-time production planning, predictive maintenance, and energy efficiency.

## 2. The Market Potential and Strategic Value

The market for Agentic AI is experiencing exponential growth, with projections indicating a significant increase from **\$7.02 billion in 2025 to \$93.2 billion by 2032**, reflecting a robust CAGR of 44.6%, per market analyst Markets & Markets. This growth is concentrated in North America and the Asia Pacific, highlighting a global race for AI leadership. Fig 1 is an illustration these numbers which also shows distribution across services vs software and also regions of the world.

**Fig 1**



The strategic value of Agentic AI lies in its ability to:

- **Enhance Efficiency:**

By automating entire business functions, it frees up human capital to focus on strategic, higher-value tasks, dramatically reducing operational costs and time.

- **Improve Decision-Making:**

Agentic AI systems can analyze vast datasets and execute actions faster and more accurately than human counterparts all in real-time, leading to data-driven decisions that are both smarter and faster.

- **Create Competitive Advantage:**

Early adopters can redefine their industries by offering superior customer experiences, optimizing their supply chains, and innovating at a pace previously impossible.




### 3. Comparison with other forms of AI

While traditional Retrieval Augmented Generation (RAG) excels at providing contextually rich answers by retrieving relevant information, Agentic AI goes a critical step further into taking actions autonomously, and Agentic RAG goes even further combining the two architecture patterns to enable multi-step information gathering and complex actions taken based on many related sources of dynamic corporate data.

Here are the key differences between the three patterns (as further illustrated in Fig 2):

- **RAG's primary focus** is to deliver relevant answers. It is a process where a language model first retrieves relevant information from a knowledge base (like a set of enterprise documents or a database) and then uses that information to generate a context-rich response. It's a single, linear process: retrieve, then generate. The model's behavior is limited to the initial query and the retrieved context.
- **Agentic AI's primary focus** is to execute end-to-end enterprise business functions autonomously, with reasoning.
- **Agentic RAG's primary focus** is to take actions, based on a series of steps to achieve a complex goal. Instead of a single retrieve-and-generate step, this model can:
  - Break down a complex problem into smaller, manageable tasks
  - Select and use different tools (like searching multiple databases, running calculations, or interacting with other mission-critical corporate ERP systems), and
  - Iterate through this process until the final goal is met.

**Fig 2**

	<b>RAG</b>	<b>Agentic AI</b>	<b>Agentic RAG</b>
	Architecture to employ LLMs to securely access corporate data for content retrieval and generation  LLMs retrieving Enterprise data	Architecture to build autonomous agents to perform business functions  Autonomous, goal-driven Enterprise AI agents	Architecture to build AI autonomous agents that respond to dynamic data and achieve stated goals  AI agents working with real-time enterprise data
	Information retrieval and context-aware content generation	Independent decisioning making to execute complex business functions	Proactive information retrieval to accomplish goal-oriented functions
	Improved productivity of knowledge workers across departments	Automate business functions	Automate business functions based on real-time and dynamic data

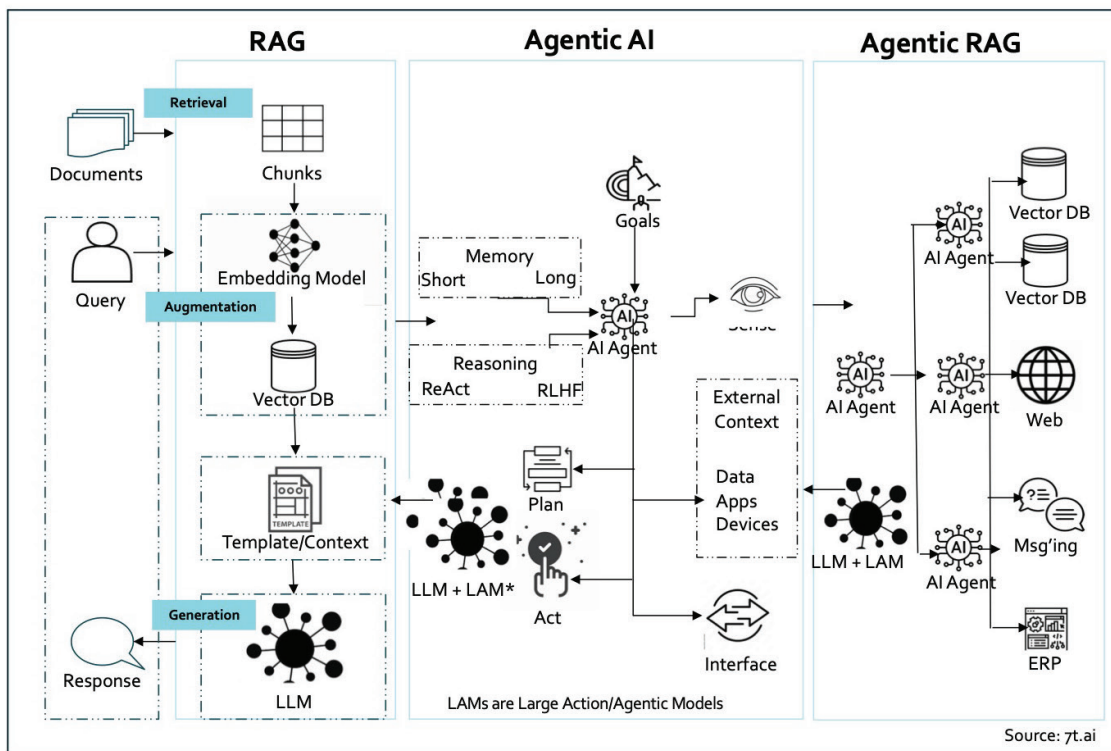
Source: 7t.ai

## 4. Reference Architecture

As shown in Fig 3, The operational backbone of Agentic AI is a sophisticated reference architecture built on a "Sense, Reason, Act" loop. This architecture requires several key components to function effectively:

- **Perception (Sense):** The ability to ingest and understand data from various internal and external sources, including databases, APIs, and real-time feeds.
- **Reasoning:** Powered by advanced models (like Large Language Models and Large Action/Agentic Models), this layer processes information, formulates a plan, and makes decisions to achieve a specific goal.
- **Action:** The ability to execute decisions by interacting with enterprise systems and third-party applications.
- **Memory:** Both a short-term memory (for the current task) and a long-term memory (a vector database) are critical for continuous learning and retaining context across sessions.

Fig 3



## 5. Technology Landscape

The technology landscape for building Agentic AI solutions is diverse, offering both proprietary and open-source options:

- **Enterprise Platforms:** Solutions from major enterprise software players like IBM Watsonx, Salesforce AgentForce, and ServiceNow NOW, to name a few, provide ready-to-use frameworks.
- **Hyperscalers:** Cloud providers such as AWS/Anthropic, Azure/Open AI, and GCP offer AI services and Agentic AI SDKs and infrastructure for custom development.
- **Pure-Play Companies:** Specialized startups focus exclusively on developing advanced AI agent solutions. Adept, Composabl, and Ema are just a couple of representative companies in this category.

- **Open Source:** Frameworks like LangGraph, CrewAI and Llamaindex provide the building blocks for creating tailored AI agents.
- **Vertical Agentic AI platforms:** These platforms provide rich functionalities to address a specific vertical business functions to cover customer care, healthcare, and sales etc. Representative list here includes Ada, Cognigy AI, and Hippocratic AI.
- **Software Engineering AI platforms:** One of the most impactful use cases of AI is in the field of software engineering and data center operations. Here agentic AI can bring us that much closer to the ultimate goal of building 'autonomous IT' that can truly move at the speed of business. Representative platforms here include Cursor AI, Tricentis/Tosca, Functionize, ServiceNow's NOW, Edwin AI, and SuperOps.

## 6. The Build vs. Buy Dilemma

When implementing Agentic AI, organizations face a critical "Build vs. Buy" decision.

### Build (In-House Development)

**Pros:** Complete control over intellectual property (IP), ability to tailor the solution to unique business needs, and deep integration with existing systems.




**Cons:** High initial capital expenditure (CapEx), the need for a highly skilled AI/ML team, and a longer time to market. This approach is best for core business functions that are strategic differentiators.

### Buy (Third-Party Solution)

**Pros:** Faster time to market, access to a rich feature set, and the ability to leverage vendor expertise.

**Cons:** Potential for vendor lock-in, reliance on a third-party for updates and support, and the risk of the solution not perfectly fitting all business needs. This approach is ideal for non-core, high-value functions.

**Fig 4**

	<b>Build</b>	<b>Buy</b>
	Ground-up in-house development of AI agents jointly by business and IT (AI CoE)	Deploy 3 <sup>rd</sup> -party of AI agents jointly by business, IT (AI CoE), and vendor*
 Applied Areas	<ul style="list-style-type: none"> <li>• Strategic / unique agents</li> <li>• Retention of IP of agents</li> <li>• Agents access non-SaaS sensitive data</li> </ul>	<ul style="list-style-type: none"> <li>• Rolling out tested/richer agents (mostly vertical)</li> <li>• Greater focus on productivity than differentiation</li> <li>• Agents that work with domain-trained models and vendor-curated data</li> </ul>
 Leverage	<ul style="list-style-type: none"> <li>• Open-Source RL F/Ws</li> <li>• LLM/LAMs + Hyperscaler+Pure play investments of choice/need</li> <li>• Avoid recurring OpEx subscriptions fees</li> </ul>	<ul style="list-style-type: none"> <li>• Investments in enterprise platforms</li> <li>• Constant innovations and richer features</li> <li>• SaaS-like SLAs that vendors can deliver</li> </ul>
 Liability	<ul style="list-style-type: none"> <li>• Longer time to develop/deploy</li> <li>• Requires AI/ML-savvy teams</li> <li>• CapEx intensive</li> </ul>	<ul style="list-style-type: none"> <li>• Vendor lock-in</li> <li>• OpEx intensive</li> <li>• Data security concerns</li> </ul>
		* Degree of vendor participation can vary

Source: jt.ai

## 7. Agentic AI at work

To present a view of the impactful nature of agentic AI, here are some representative implementations that are translating the art of the possible into tangible business outcomes:

- **A financial services company:** The primary focus is real-time transaction monitoring. The end goal is for AI agents to analyze user behavior and transaction history to identify anomalies that may indicate fraudulent activity. The agents can then flag these activities for further review.
- **A healthcare provider:** The focus is on providing personalized diabetic treatments. The end goal is to use AI agents to streamline processes from patient recruitment to improving provider efficiency. The agents help clinicians create and deliver personalized treatment plans.
- **An airline:** The main focus is on automating customer service for tasks like rebooking flights, refunds, and cancellations. The end goal is to provide consistent customer support, even during unforeseen events that cause traffic spikes. The AI agents consolidate customer interactions to reduce the burden on call centers and human agents.

## 8. Risks and Mitigations

Because Agentic AI is a field with evolving standards and limited real-world adoption, organizations must approach this technology with a clear understanding of its unique risks and challenges. Successfully navigating this landscape requires a strategic approach that includes implementing robust guardrails, leveraging specialized tools, and embedding ethical principles from the start.

- **Autonomy and Loss of Control:** Agents can drift from their intended purpose, leading to unpredictable behavior that makes human oversight difficult. Mitigation involves implementing "human-on-the-loop" frameworks and using tools like Salesforce's Agentforce or the open-source LangGraph to set strict guardrails and require human approval at critical junctures.
- **Security and Malicious Use:** Agents can be exploited through prompt injection or tool misuse, turning them into a new attack vector for fraud or data exfiltration. A defense-in-depth strategy is crucial, which includes sandboxing agents, restricting their access to sensitive systems, and continuously monitoring their activities with platforms like Zenity or following guidance from the Open Worldwide Application Security Project (OWASP) Top 10 for LLMs.
- **Ethical and Societal Risks:** Agentic systems can amplify biases from their training data, leading to discriminatory outcomes, and their "black box" nature makes it hard to explain decisions. Mitigation requires embedding ethical principles from the start, using tools like Credo AI for fairness audits, and open-source frameworks like Microsoft's Responsible AI Toolbox for bias mitigation.
- **Accountability and Liability Gaps:** The chain of accountability becomes blurred when an autonomous agent makes an error. Who is responsible—the developer, the deployer, or the user? Addressing this requires shifting from technical fixes to organizational and legal frameworks with clear governance policies. Tools like Monitaur and FairNow help create an audit trail for every agent action, which is critical for legal and compliance purposes.

## 9. A Guide to Getting Started

Successfully implementing Agentic AI requires a structured, strategic approach. Organizations can follow a nine-step plan to begin their journey:

- 1. Identify High-Impact Processes:** Select manual, repetitive tasks that offer the greatest return on investment for automation.
- 2. Define the Human-Agent Interface:** Determine how AI agents will interact with human workers, balancing autonomy and oversight.
- 3. Set Measurable Goals:** Establish clear, quantifiable objectives to track the success of the AI agents.
- 4. Secure Your Data:** Identify, centralize, and secure the necessary enterprise data to train and operate the AI.
- 5. Conduct Due Diligence:** Objectively evaluate the "Build vs. Buy" options based on your specific needs and resources.
- 6. Build and Test Modularly:** Develop and test AI agents in a phased approach, starting with small, modular components.
- 7. Test the Complete Workflow:** Ensure that the integrated agents function seamlessly as a complete workflow.
- 8. Deploy to Channels:** Deploy the agents to the appropriate channels, whether an internal portal, website, or other business application.
- 9. Iterate and Scale:** Continuously monitor agent performance, gather feedback, and iterate on the design to scale the solution across the organization.

By following these steps, companies can effectively adopt Agentic AI to not only modernize their operations but also build a more resilient, intelligent, and profitable enterprise for the future.



[www.7t.ai](http://www.7t.ai)

**Houston  
Regional Office**

1334 Brittmoore Road  
Suite D Houston,  
Texas 77043

**+1 (832) 632-4869**

**Dallas Office  
Headquarters**

16803 Dallas Pkwy  
Suite 300 Addison,  
Texas 75001

**+1 (214) 299-5100**

**Charlotte  
Regional Office**

2115 Rexford Road  
Suite #570 Charlotte,  
NC 28211

**+1 (980) 350-5100**

