



A PRACTICAL GUIDE TO AI AGENTS

Key agentic AI concepts, use cases and considerations to drive ROI

TABLE OF CONTENTS

| | |
|---|-----------|
| Introduction | 3 |
| CHAPTER 1 | 5 |
| The Business Case: What Can AI Agents Offer | |
| CHAPTER 2 | 7 |
| The Platform: How To Implement AI Agents | |
| CHAPTER 3 | 12 |
| Ethics And Bias: Overcoming Challenges | |
| CHAPTER 4 | 14 |
| Bringing It All Together: Rolling Out AI Agents | |
| CHAPTER 5 | 16 |
| Snowflake For Data Agents | |
| INDUSTRY USE CASES | 18 |



INTRODUCTION

AI agents are like computer programs that can perform tasks independently using artificial intelligence. They can significantly boost business productivity by handling complex, multi-step operations quickly. To work effectively, these agents need access to an organization's data, which can be both structured data — patient records at hospitals or transaction details at a department store for example — and unstructured data, like video files or social media posts.

The potential for AI agents is massive but the quality of work depends heavily on the data it uses and how accurately it can find the right information. Organizations often struggle with this because how AI understands data doesn't always match how the data is stored.

Still, AI agents are expected to become a valuable resource in the workforce, helping with roles like customer support, field technicians, data analysis and more. It's also important to understand that there are two different types of AI agents: personal and company.

Personal AI agents are customized digital assistants that learn from and adapt to an individual user's preferences, habits, and needs, acting primarily in that person's interest. They typically have access to individual data only.

Company AI agents, on the other hand, usually work with shared company data and are designed to serve organizational goals and operate within specific business parameters, often handling multiple users and following corporate policies.

Both types of agents will free time for employees to focus on more important and complex tasks. Specialized AI agents are particularly useful because they can combine data and tools to provide more accurate and reliable insights by finding and extracting answers from the best data sources for each task.

For AI agents to be widely used in businesses, they need secure access to company data and a way to manage that access, similar to how employee access is managed. They must follow company rules about data, efficiently use different data sources, and find accurate information to deliver valuable results.

However, deploying these agents at a large scale presents challenges, such as ensuring the accuracy of their outputs, maintaining security and trust, and managing access to diverse data sources.

WHAT'S COVERED IN THIS BOOK

This ebook provides an introductory overview of agentic AI, a form of artificial intelligence designed to function autonomously. It covers how agentic AI differs from traditional AI and generative AI, highlighting its ability to make decisions, perform tasks and adapt to its environment without constant human intervention. We'll cover the potential benefits of agentic AI, including enhanced efficiency, improved decision-making, speedier time to market, higher customer satisfaction and more opportunities for innovation.

We'll also lay out a strategic roadmap for implementing agentic AI, focusing on the need to assess your data foundation and align with business goals, while prioritizing security, compliance, and ethical concerns. This ebook will also explore various real-world and industry use cases of agentic AI, showcasing its diverse applications and potential impact across different sectors.

ANTHROPIC

Summaries for each chapter were written with the help of Anthropic Claude 3.5 Sonnet.

GENERATIVE AI

Can create content and communications personalized to a customer's inquiries by analyzing large data collections and adapting the tone and style of its outputs to match brand guidelines or customer needs.

AI AGENTS

Act as active problem solvers that can handle interactive and multi-step engagements with customers using a natural language interface.

CHAPTER 1:

THE BUSINESS CASE: WHAT CAN AI AGENTS OFFER

ANTHROPIC

Chapter summary written by Anthropic Claude 3.5 Sonnet:

AI agents — particularly data agents in enterprise settings — are specialized tools that focus on efficiently combining data and tools to deliver accurate, data-grounded insights for specific business functions.

These agents operate through key actions including data retrieval, tool/agent calling and output delivery, with critical requirements being accuracy, efficiency and proper governance of data.

The future of AI agents involves multiple agents working together autonomously to manage complex tasks, analyze information from various sources, make decisions, and continuously learn from their actions, though this advancement also brings forth various challenges and ethical considerations.

AI agents capable of completing complex tasks without human intervention are quickly emerging as a major paradigm shift for businesses to drive a new wave of productivity, automation and creativity. These autonomous, goal-oriented systems are capable of more complex, multi-step tasks such as performing a root cause analysis and generating a report by accessing data and leveraging different tools. For these advanced analytical workflows with many edge use cases, access to private data — whether it be database or PDF files stored in object storage — is key.

AI agents can bring ROI to every part of the organization to accelerate existing workflows or enable new ones where previously technical talent gaps prevented more users from being part of data-driven decision-making. For most organizations, data is their core differentiator and being able to use that data effectively in decision-making is not a new goal — rather, it's become a more clear necessity.

Many companies have already made investments: in a survey of 1,100 large-enterprise executives, 10% of organizations currently use AI agents, more than half plan to use them in the next year, and 82% **plan to integrate them** within the next three years. Here are a few examples:

- **Finance:** Automate financial forecasting by integrating real-time revenue (structured) and market reports (unstructured), enabling faster decision-making and risk assessment.
- **Operations:** Optimize supply chain logistics by analyzing internal inventory levels (structured) and shipping contracts (unstructured) to reduce delays and costs.
- **Engineering:** Identify patterns in bug reports and customer requests by analyzing customer support chat transcripts (unstructured) and customer product usage patterns (structured) to proactively improve product quality and prioritize fixes.

- **Marketing:** Personalize campaign targeting by analyzing customer purchase history, and behavioral data (structured) and analyzing images in ads (unstructured) to improve engagement and conversion rates.
- **Product:** Enhance feature adoption analysis by integrating user telemetry, A/B test results (structured) and summaries from customer calls (unstructured) to prioritize product improvements.
- **Human resources:** Enable employees to use AI agents to compare healthcare plans (unstructured) while taking into account monthly salary and tax implications (structured).

For the remainder of this book we will use the term **structured data** to make reference to data organized in tables whether it's a database or an Excel sheet. And we will use the term **unstructured data** to encompass complex file formats like PDFs, images or video.

Hopefully, you've seen the pattern that most decision-making requires multiple data points and types to make decisions. AI agents are replacing manual data stitching with automated, accurate insights — boosting productivity and enabling faster strategic decisions that ultimately will help unlock new revenue opportunities. Here is one thing to keep in mind as you start the journey to give employees a decision intelligence agent: delivering accurate answers while ensuring data privacy becomes increasingly complex as more data sources are connected and more AI agents are built on top of them. To address this from the start, it's important to evaluate solutions with a platform mindset where shared data governance, processing and retrieval resources drive cost-efficiencies.

In 2024, **37% of venture capital (VC) funding** was to AI startups. Autonomous agents and digital co-workers saw **the biggest growth in VC deal activity** in 2024.

The companies that effectively leverage the opportunities, mitigate the risks and become early adopters of this groundbreaking technology are poised to reap substantial benefits. Here are just a few:

- **Enhanced efficiency** and cost savings by automating routine and time-consuming tasks and streamlining workflows
- **Improved decision-making** with more accurate insights and real-time adaptability
- **Speedier time to market** with the ability to process information and carry out tasks in a fraction of the time
- **Higher customer satisfaction** and retention by providing personalized, more responsive support
- **More opportunities for innovation** by freeing up workers to focus on more complex and creative tasks

In a recent survey, executives say **AI agents will increase automation in their workflows (71%) and improve customer service and satisfaction (64%), and that the potential productivity improvements outweigh the risks (57%).**



CHAPTER 2:

THE PLATFORM: HOW TO IMPLEMENT AI AGENTS

ANTHROPIC

Chapter summary written by Anthropic Claude 3.5 Sonnet:

AI agents — particularly data agents in enterprise settings — are specialized tools that focus on efficiently combining data and tools to deliver accurate, data-grounded insights for specific business functions. These agents operate through key actions including data retrieval, tool/agent calling and output delivery, with critical requirements being accuracy, efficiency and proper governance of data.

The future of AI agents involves multiple agents working together autonomously to manage complex tasks, analyze information from various sources, make decisions, and continuously learn from their actions, though this advancement also brings forth various challenges and ethical considerations.

THE EVOLUTION FROM LLMS TO AI AGENTS

The first program using natural language processing was developed by an MIT scientist in the 1960s. In 2022, the introduction of OpenAI's ChatGPT transformed the world of AI. Today, LLMs drive a diverse range of applications, from generating creative content to playing a pivotal role in national defense. LLMs laid the groundwork for natural language understanding and created the ability for AI to process and generate text-based on patterns in existing data. However, they are fundamentally passive and cannot take direction or respond adaptively. They have no decision-making or autonomous capabilities. AI agents fill these gaps by expanding the role of AI, from passive responder to active participant.



THE EVOLUTION OF AI: FROM DATA TO ACTION

Traditional AI began in the 1950s with a focus on rule-based systems and explicit programming to solve specific problems through logical reasoning and decision trees. Gen AI gained prominence starting around 2010. It represented a shift toward machine learning models that could create new content by learning patterns from vast amounts of data, with transformers and LLMs leading to breakthroughs in the creation of text, images and audio.

AI agents are the latest evolution in this technological journey. They combine both approaches by using gen AI capabilities within autonomous systems that can plan, reason and take actions to achieve specific goals. This shift potentially marks the beginning of more versatile and independent artificial intelligence systems. Let's take a look at how AI agents work.

By 2028, agentic AI **will make 15%** of daily work decisions.



HOW AN AI AGENT WORKS

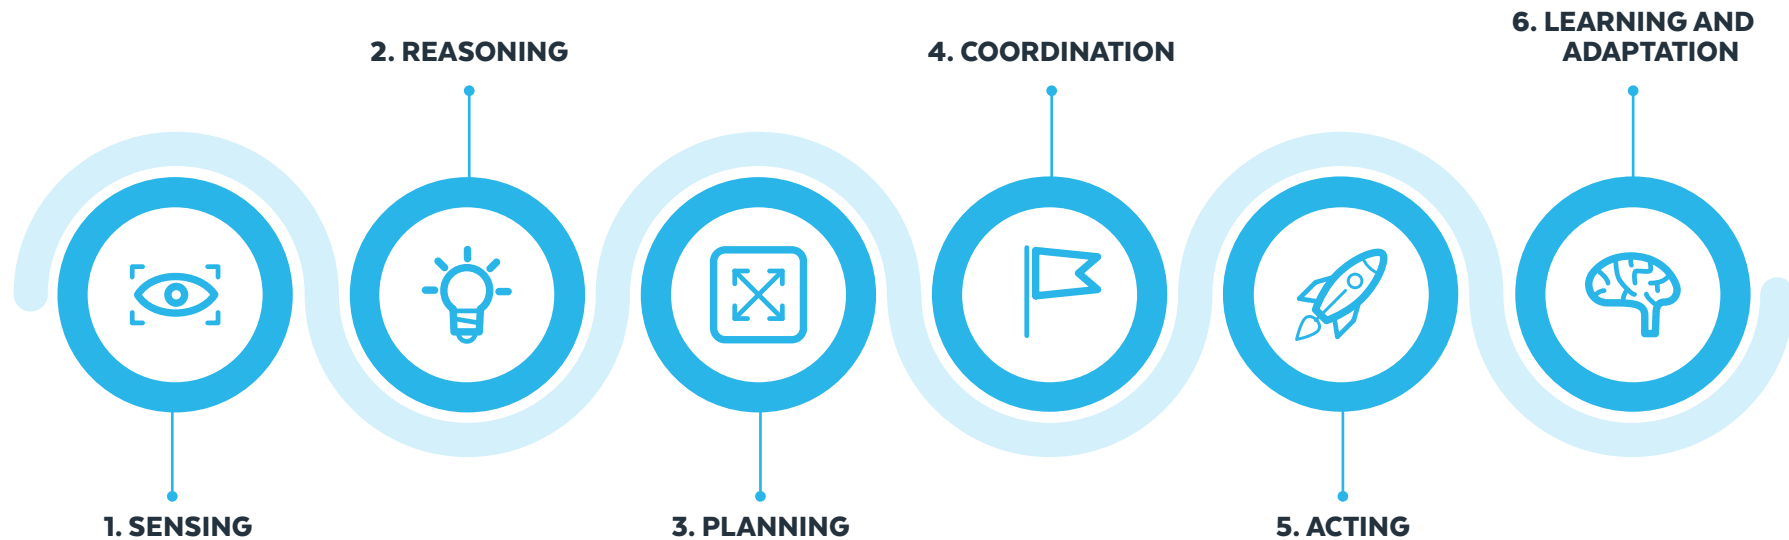


Figure 1: AI agent

- 1. Sensing:** Agents define the task or problem to be solved, and ascertain and gather data from the most relevant, reliable, accurate and up-to-date sources available.
- 2. Reasoning:** Agents process and interpret the data, often relying on an LLM, to understand the context and requirements and make informed decisions.
- 3. Planning:** Based on the insights from reasoning agents, develop plans of actions to achieve the objective.
- 4. Coordination:** Agents share their plans with users or systems to ensure alignment and collaborative decision-making.
- 5. Acting:** Agents implement the plan and execute the necessary actions.
- 6. Learning and adaption:** Agents assess the outcomes and incorporate feedback to reflect, determine success and areas for improvement, and refine processes for future tasks.

THE DATA AGENT

While general-purpose AI assistants play a key role in our personal lives — say, for example, who doesn't want to automate flight search for the next vacation — in the enterprise, accessing internal data is what drives value for people.

This particular type of AI agent can be referred to as data agents. Data agents are a category of AI agents focused on efficiently combining data and tools to deliver data-grounded insights with an extreme

focus on the accuracy of the results — because the decisions business teams make on the data can be the difference between making or losing millions. Just like with any type of agent, specialization is key, and these will need to be scoped to specific business functions and units. For example, sales teams may have a different data agent than those working in research and development.

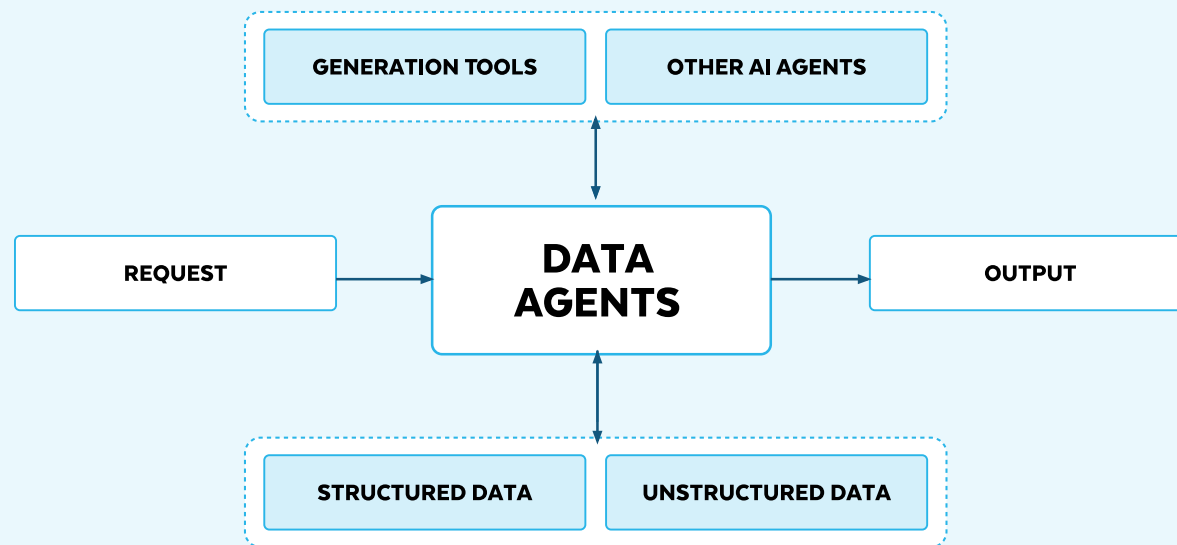


Figure 2: Data agent workflow

The key actions of data agents:

- **Data retrieval:** Data agents must be able to plan and route requests to the correct structured or unstructured data source to retrieve data from to complete a request. And when doing so, it must enforce data access and governance policies.
- **Calling tool or calling other agents:** To perform tasks such as generating a histogram or creating a report, data agents will need to combine retrieval outputs with other tools or even collaborate with other agents that have access to multiple tools.
- **Output delivery:** Delivery of output might involve sending a message through a chat interface, writing the result to a table that feeds results to a dashboard, or it could mean emailing a PDF report that includes asset performance, portfolio updates and a summary of market sentiment for each fund manager.

Key requirements for data agents

Without data, there are no data agents. A team of developers could have great tools to generate charts or reliable integrations between agents and downstream systems, but without the correct data, what are those good for? With data retrieval being foundational, here are three critical elements that organizations need to address:

- 1. **Accuracy:** If the data retrieved is not accurate, the entire workflow becomes useless. Why would a sales leader trust a data agent that makes fancy charts if the underlying data is incorrect?
- 2. **Efficiency:** Fast responses that require pulling data are only possible if the agent quickly identifies the right set of sources and retrieves the necessary results. Having too much information accessible can explode compute cost and not having enough access can impact response accuracy.
- 3. **Governance:** To scale to multiple use cases and reach production, organizations need a scalable approach to define and enforce access and privacy controls. This ensures a unified framework for hundreds, if not thousands, of agents to retrieve data securely.

From an architectural perspective, the primary tools needed are retrieval tools for structured and unstructured data. And these tools must have access to data that is properly prepared and has the necessary governance and access controls in place.




Figure 3: Retrieval stack for structured and unstructured data

In the diagram above, we’ve oversimplified to show how a single agent works. But the reality is that — similar to business structures — there need to be “managers” and there need to be “cross-functional teams” for AI agents as well.

This is the future, where multiple AI agents can work together to autonomously manage complex tasks such as processing account changes, resolving billing issues and providing personalized incentives for higher customer satisfaction. These agents can do this by analyzing information scattered in different formats and areas of a business — such as databases, web pages and emails — to discern the most reliable, up-to-date information.

Then, they use reasoning by employing analysis, pattern recognition and decision-making algorithms to generate options for a solution. After choosing the best option, they coordinate the plan and execute it. Finally, they learn from the results of their actions and improve for the next time. This ability to autonomously manage, execute and learn from complex tasks marks a significant advancement from previous iterations of AI.

At the same time AI agents are transforming modern business operations, they are also bringing up a host of challenges and ethical implications that need consideration.

CHAPTER 3:

ETHICS AND BIAS: OVERCOMING CHALLENGES

ANTHROPIC

Chapter summary written by Anthropic Claude 3.5 Sonnet:

AI agents face significant technical challenges including integration with existing systems, infrastructure requirements and security concerns, while also requiring robust governance frameworks. Ethical considerations are paramount, necessitating proper guardrails, continuous monitoring, and evaluation frameworks to prevent biases and unethical decisions. Effective human-AI collaboration and transparency in decision-making processes are crucial for successful implementation, though achieving the right balance between automation and human oversight remains an ongoing challenge.

Organizations thinking about deploying AI agents face significant ethical and technological challenges that have to be accounted for. Ethical concerns start with data privacy and algorithmic bias because these systems can process sensitive information without the proper context and then perpetuate existing societal inequalities.

Technical challenges must also be overcome, including ensuring AI systems remain reliable and consistent in their decision-making, particularly when faced with unique situations not built into their training data.

Integration with legacy systems and existing workflows can also be difficult. Organizations must carefully manage the transition to AI-enhanced processes while maintaining operational continuity.


These organizations also have to be prepared to face the challenge of maintaining transparency in their AI systems' decision-making processes. This element is critical to building trust with users and meeting regulatory requirements. Let's dive deeper into some of these challenges and how to overcome them.

Challenges in implementation: Integration with data and other systems, compute infrastructure, security and governance

AI agents need to integrate seamlessly with existing data ecosystems and enterprise systems to work effectively. But that can be challenging because of the differences in data formats, silos and interoperability issues. These agents need vast compute infrastructure, often demanding significant investment in GPUs or cloud resources, to access, process and analyze real-time data efficiently.

There are also security and governance issues. Like employees, AI agents must comply with data privacy regulations, while preventing unauthorized access or data leakage. Organizations also have to enforce strong governance policies to monitor model behavior, manage access controls and mitigate biases, ensuring AI aligns with business objectives and regulatory standards.

For AI agents to work at scale, they need a secure connection with enterprise data and unified governance to manage their access, similar to existing controls for your employees. They must follow data policies, access multiple sources efficiently and retrieve accurate information to deliver reliable, high-value outcomes.



Ethical considerations: The importance of guardrails, evaluation and observability

AI agents operate in dynamic environments where their decisions have real-life consequences on people, businesses and society at large. That is a great responsibility. Without proper guardrails, these agents risk amplifying biases, making unethical decisions or generating misleading content.

Organizations that deploy AI agents must implement robust evaluation frameworks to continuously test them against fairness, bias and safety criteria. Organizations must also ensure the real-time monitoring of models so they do not stray from established ethical norms or produce unintended consequences. Organizations must adopt transparent and accountable AI governance strategies, including red-teaming exercises, diverse data audits and clear escalation paths when AI actions violate their ethical guidelines.

Human-AI collaboration to ensure appropriate handoffs and feedback

It's important to remember AI agents are not standalone solutions that can be turned on and trusted to work on their own. They work best when designed to work effectively with humans. But one major challenge to that scenario is defining when and how AI agents should hand off tasks to humans, especially in high-stakes scenarios such as customer service, healthcare or financial advisory.

Ensuring smooth transitions requires AI agents to detect when they are uncertain, recognize complex queries beyond their capabilities, and escalate appropriately. Continuous human feedback loops are important because they improve AI performance over time.

However, organizations often struggle to collect, analyze and implement this feedback efficiently. Striking the right balance between automation and human intervention remains an ongoing challenge but one that needs to be built in up front.

Transparency and explainability: Using agents to ensure appropriate implementation

AI agents often function as black boxes, making it difficult for users to understand how decisions are made. This lack of transparency erodes trust, especially in regulated industries like finance, healthcare and law.

AI agents need to be designed so they can provide clear rationales for their decisions, highlighting key data points and reasoning pathways. But balancing explainability with performance can be challenging because more-transparent models may sacrifice accuracy or efficiency.

AI agents themselves can help with this challenge by generating human-readable explanations, identifying biases and ensuring compliance with defined regulatory requirements. Achieving transparency and explainability in AI requires tools like model interpretability frameworks alongside automated audit log reviews as well as manual reviews where models are asked to explicitly explain the steps they took to get to the answer. As organizations get better at defining the review process, AI can be used to automate and scale explainability practices.

| | GUARDRAILS | EVALUATION | OBSERVABILITY |
|--------------------|---|--|--|
| What is it | Business policies are given to LLMs to filter harmful content, block denied topics, and remove specific words and phrases | Test scores provide a quantitative measure of trust in responses, evaluating accuracy, consistency, safety and reliability using benchmarks and real-world tests | Continuous evaluation of AI app responses to track performance, behavior and interactions with data sources to ensure reliability, accuracy and compliance |
| Why it's important | Ensures AI responses align with organizational policies and values | Ensures AI delivers relevant, safe and high-quality responses | Helps detect issues and changes in models that can impact AI responses |

CHAPTER 4:

BRINGING IT ALL TOGETHER: ROLLING OUT AI AGENTS

ANTHROPIC

Chapter summary written by Anthropic Claude 3.5 Sonnet:

Implementing agentic AI successfully requires a strong technological foundation capable of handling data processing and model deployment while aligning with business objectives. Clear communication from leadership and proper employee training are essential for effective integration and collaboration with AI systems. Organizations must establish robust governance frameworks to address security, compliance and ethical concerns while maintaining appropriate human oversight.

PLANNING FOR AGENTIC AI: A STRATEGIC ROADMAP

- 1. Assess your technology foundation.** Agentic AI requires a robust technology infrastructure with physical and virtual resources that can handle the demands of data processing, data retrieval in natural language, model training and deployment. Investing in a solid foundation can lead to better ROI, results and adaptability to integrated future AI innovation.
- 2. Align with business goals.** Agentic AI isn't just a technological innovation — it's a powerful tool for companies to accelerate business results. Develop and manage the system in tandem with business strategies and specific use cases to maximize value.
- 3. Communicate.** Company leaders should clearly explain the reasoning and goals behind using AI agents and foster a culture of innovation throughout the organization.
- 4. Invest in change management.** Employees need to be trained to collaborate effectively with agentic systems and encouraged to provide productive feedback that can make the systems better meet organizational needs.
- 5. Prioritize security, compliance and ethical concerns.** Companies will need governance frameworks that establish parameters for the use of agentic AI and maintain human oversight. They also need to stay up to date with evolving regulations and standards to safeguard the use of confidential data.



5 PRINCIPLES FOR AGENTIC AI ARCHITECTURE

AI agents can be a game changer for the enterprise. But for successful deployment and ROI, company leaders need to start with a solid foundation that can handle its significant resource demands. Here are five principles for developing and implementing a comprehensive platform:

1. **Scalability.** The data and AI platform must be able to handle growing computational demands and workloads as the number of LLMs, AI agents, models and human users grow.
2. **Flexibility.** The architecture should support an agentic AI system as part of a complete technology stack where it can integrate and communicate seamlessly with other systems as well as adapt to the evolving AI landscape. These include data sources, LLMs and collaboration tools.
3. **Data accessibility.** AI agents need easy access to reliable, accurate sources of data. These include databases, IoT devices and external APIs. Does your data infrastructure include easily integratable first-, second- and third-party data — whether structured or unstructured? Is the data in those sources updated in real or near real time?
4. **Trust.** Using guardrails, evaluations and observability provides reliability, accountability and opportunity for continuous improvements in both automated responses and AI agent to human handoff.
5. **Security and compliance.** The infrastructure must include strong security measures to protect data and models with granular access controls and proactive log monitoring.



CHAPTER 5:

SNOWFLAKE FOR DATA AGENTS

ANTHROPIC

Chapter summary written by Anthropic Claude 3.5 Sonnet:

AI agents are being deployed across enterprise departments to automate tasks and improve efficiency, from customer service to sales, finance and IT operations. Snowflake offers a unified platform that enables organizations to process both structured and unstructured data while maintaining governance and security controls. The technology has shown impressive results, including 14% more resolved customer service issues per hour and 90% reduction in sales prospecting time, with the AI agent market expected to reach \$47.1 billion by 2030.

Agentic outputs are only as good as the quality of the underlying data and the accuracy of the retrieval systems that help ground them. Yet organizations struggle to pave a path to production due to an AI and data mismatch. LLMs excel at unstructured data, but many organizations lack mature preparation practices for this type of data; meanwhile, structured data is better managed, but challenges remain in enabling LLMs to understand rows and columns (the text-to-SQL task).

“Traditional data management methods are increasingly insufficient given the exponential data growth,” said Jim Rowan, Head of AI Deloitte Consulting LLP. “Many enterprises face overwhelming data sources, from structured databases to unstructured social media feeds. Manual processes can be time-consuming and error prone. Agentic AI automates these processes, helping ensure data integrity and offering real-time insights. Leveraging advanced machine learning and natural language processing, these intelligent agents can efficiently manage and analyze vast data amounts. The integration of Snowflake’s AI Data Cloud and launch of Cortex Agents, along with Deloitte’s experience, can optimize these processes for efficiency and innovation.”

Snowflake customers now have a unified platform for processing and retrieval of both structured and unstructured data with high accuracy out of the box. End-to-end unified governance, from ingestion to application, enables teams to deliver a new wave of data agents. Customers can build scalable solutions while enforcing access and privacy controls.

“At Luminate, we’re revolutionizing how we deliver precise, data-driven insights to our clients through generative AI applications. **Snowflake’s unified data and AI platform provides our developer team with scalable processing and retrieval for both structured and unstructured data — the critical building blocks for developing, deploying and orchestrating data agents powering our applications.** Using Snowflake Cortex AI brings advanced AI within the same security and governance perimeter as our data and saves us countless development hours, allowing us to unlock the full potential of entertainment industry data with agentic AI.”

— **Glenn Walker**
Chief Data Officer, Luminate Data




Figure 4: AI applications framework

INDUSTRY USE CASES

Large enterprises from various industries are dispatching AI agents across departments to drive business results. Agents can continuously monitor and analyze vast amounts of streaming data and decide, plan and execute in real time.

- **Customer service** AI agents can handle customer inquiries and provide faster, more personalized resolutions. In one use case, AI agents helped a large call center resolve 14% more issues each hour, while time spent handling issues went down 9%.
- In **sales and marketing**, AI agents can shorten sales cycles by providing real-time guidance to sales teams, boosting conversion rates and deal times. Agent AI can minimize prospecting times by 90%, according to one report.
- In the **finance department**, agent AI is helping organizations cut costs by automating repetitive tasks like invoice processing and providing real-time insights for decisions that drive growth.
- In **IT**, agent AI is cutting costs by streamlining processes and reducing manual workload. Assistants can write and debug code, monitor systems, resolve tickets and generate reports.

USE CASE: ACCENTURE AND SNOWFLAKE'S ROLE IN FACILITATING AGENTIC AI

Accenture has leveraged Snowflake's platform to create a groundbreaking AI insurance claims agent, combining Snowflake's tools to automate key parts of the claims process with Accenture's expertise in agentic AI. This AI agent can review documents, summarize information, make claims decisions and generate personalized claims letters to clients, explaining the reasoning behind an approval or denial.

According to Mike Lao, a leader of Accenture's data and AI team, "Underlying the Gen AI Claims Agent is Snowflake's AI Data Cloud, which includes capabilities like Document AI that can process various documents, such as driver's licenses, receipts and forms. We also use Snowflake's Cortex Analyst, which helps the AI Claims Agent analyze data and policy documents to make informed decisions."

Accenture's focus on claims stems from insights into the financial services industry, where claims approval does not always require 100% accuracy. Interestingly, clients are often willing to accept partial payouts in exchange for faster processing. As Accenture found, there's room for flexibility in claims payments — decisions can be appealed internally, avoiding regulatory intervention. This contrasts with the more stringent accuracy requirements of areas such as financial advice, which demands fiduciary responsibility. Claims, in comparison, are less regulated, offering insurance companies more leeway to adjust payouts within their margins. The ability to speed up payments can significantly enhance both the insurer's bottom line and customer satisfaction.

To ensure reliability and client satisfaction, Accenture has included human ingenuity in the AI Claims Agent. Mike Lao emphasized, "Keeping a human in the loop is important. We are using AI to reinvent work processes to scale AI adoption in organizations and meet client expectations."

The Accenture team in Manila, Philippines, brings deep expertise in claims payments supporting global insurance companies. This project marks the first step toward broader adoption of AI agents across various financial services sectors. As Kaushik GD, Snowflake's Head of Financial Services in APJ, noted, "In the future, we expect to see AI agents assisting with financial planning and personal investment advice. However, this will take time, as both technical solutions and regulatory acceptance evolve."

As businesses continue to integrate agentic AI, they are unlocking new efficiencies and innovations, positioning themselves at the forefront of their industries.

| INDUSTRY | USE CASES |
|---|---|
|  Media | <ul style="list-style-type: none">• Identify ad performance, shift spending accordingly, and test new creatives to improve ROI• Recommend content to users based on their consumption habits and provide future programming ideas based on audience analytics |
|  Financial services | <ul style="list-style-type: none">• Review market conditions and adjust a client's portfolio allocations to maximize trends and mitigate risks• Identify early warning signs of delinquency in a borrower and implement a tailored repayment plan to avoid default |
|  Healthcare and life sciences | <ul style="list-style-type: none">• Remotely monitor and analyze patient data in real time through wearable devices, alerting providers to concerning changes or potential future issues• Design and test new molecules for diseases to identify potential drug candidates |
|  Manufacturing | <ul style="list-style-type: none">• Monitor equipment performance, predict failures and dispatch maintenance teams in real time• Optimize inventories on the fly in response to real-time demand fluctuations or weather disruptions |
|  Retail | <ul style="list-style-type: none">• Predict demand trends and adjust stock levels and prices in real time• Provide timely personalized product recommendations and faster issue resolution for shoppers |
|  Technology | <ul style="list-style-type: none">• Actively monitor networks for unusual activity, detecting and mitigating potential threats in real time• Prototype designs or codes, iterating to improve results |
|  Telecom | <ul style="list-style-type: none">• Predict traffic loads and manage bandwidth allocation accordingly• Identify customers at risk of leaving and carry out retention strategies |

Whether you're just starting to experiment with generative AI or already have your first AI applications in production, Snowflake enables you to extend the benefits of AI to your entire organization by streamlining development, reducing complexity and helping secure your data every step of the way.

Discover how Snowflake can help your organization be more innovative and competitive through generative AI. Visit Snowflake's gen AI [webpage](#) for more information and hands-on resources to help you get started.





ABOUT SNOWFLAKE

Snowflake makes enterprise AI easy, efficient and trusted. Thousands of companies around the globe, including hundreds of the world's largest, use Snowflake's AI Data Cloud to share data, build applications, and power their business with AI. The era of enterprise AI is here.

Learn more at snowflake.com (NYSE: SNOW)



© 2025 Snowflake Inc. All rights reserved. Snowflake, the Snowflake logo, and all other Snowflake product, feature and service names mentioned herein are registered trademarks or trademarks of Snowflake Inc. in the United States and other countries. All other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Snowflake may not be associated with, or be sponsored or endorsed by, any such holder(s).