

Introduction

Agentic AI has quickly become one of the most talked-about topics in enterprise technology. Vendors promise autonomous systems that can reason, plan, and act on behalf of the business, while executives are eager to understand how close these capabilities really are to production readiness. However, behind the excitement lies a reality that is often less clear: most enterprise AI solutions today are not truly agentic, and the path from a simple chatbot to a fully agentic system is neither short nor straightforward.

This paper aims to **clarify the journey from conversational AI to enterprise-grade agentic AI**, by breaking it down into a sequence of **distinct, incremental systems**, each adding new capabilities—and new complexity—on top of the previous one. Understanding these steps is critical, because **more advanced does not always mean more appropriate**. In many real-world use cases, simpler systems can deliver significant value without the operational risk, cost, and unpredictability associated with autonomous agents.

By clearly defining each stage—from **chatbots, to AI experts, to assistants, agents, and finally agentic AI ecosystems**—this guide provides a shared vocabulary and a practical framework for decision-makers. The goal is not to promote agentic AI as an inevitable destination, but to help organizations **choose the right level of intelligence for the right problem**, and to approach this emerging space with realism rather than hype.

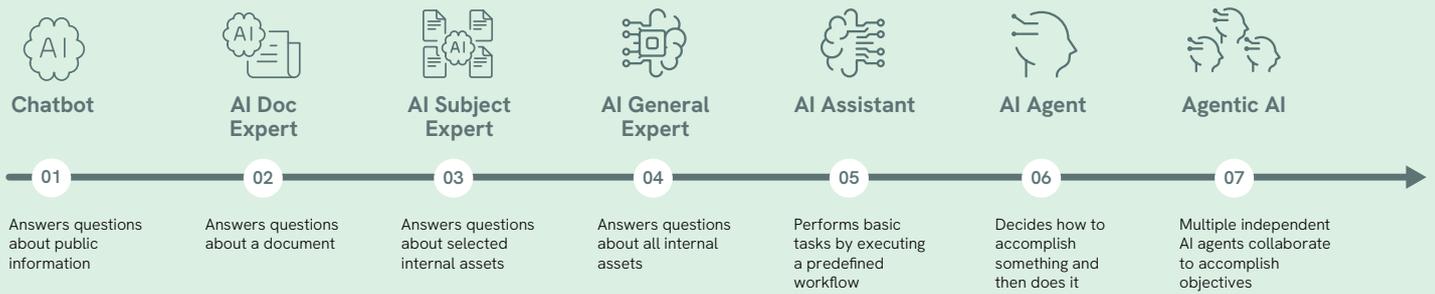
Every business is excited about agentic AI.

But very few companies have deployed true agentic AI for multiple reasons:

- 01.** The journey from a chatbot (like ChatGPT) to agentic AI is a long one with many steps along the way
- 02.** It's difficult to distinguish between the systems at each step (because they often appear very similar in their interface and operation)
- 03.** These systems are complex technologies and require some effort to understand
- 04.** Complexity increases dramatically with scale (PoCs might be easy, but production is hard)
- 05.** The marketplace is confusing, with conflicting claims and explanations
- 06.** Definitions are not yet standardized - most claims of "agents" are actually not agentic at all
- 07.** To take part in the hype, companies tend to exaggerate their capabilities and progress

As a result, it's easy to feel overwhelmed. This guide will help you understand and navigate this emerging space. Let's start with defining the agentic AI journey and explaining the steps along the way.

The following diagram summarizes the journey. Each step is explained in detail in the sections that follow.



While all the excitement seems to be focused on agentic AI, that doesn't mean it's the right destination for every application. There are needs that can be fulfilled by each of the systems on this diagram, and if (for instance) an AI General Expert accomplishes the goal, why add the complexity that comes with an AI assistant or agent? It's better to keep things simple when possible. So, while the market may make it sound like AI agents are going to solve everything, not all solutions need to be, or should be, agentic.

Understanding the journey requires understanding each step along the way.

The Steps on the Journey to Agentic AI

01 Chatbot



Chatbot: Large Language Model, or LLM

A software program that uses an LLM to answer questions and engage in conversation about general topics. Although chatbot carried a slightly different meaning prior to ChatGPT, this document uses chatbot to refer to any of conversational application built with an LLM.

The first step in the journey is the basic Chatbot, which serves as the foundation for all subsequent systems.

- Primary commercial examples: OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, x.AI's Grok
- Primary open-weight examples: Meta's LLaMa, Google's Gemma, Microsoft's Phi, Mistral's Mistral, DeepSeek by DeepSeek (and other Chinese models not widely used in the US).



02



AI Doc Expert

Building on the chatbot, organizations can layer additional capabilities to create increasingly specialized and intelligent systems.

AI Document Expert: LLM + document

Using a chatbot to work with the contents of a document in real time (in other words, the employee selects the document and gives it to the LLM). The chatbot serves as an expert intermediary between the employee and the document to answer questions or otherwise work with its contents. Typical tasks performed by an AIDE are:

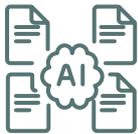
- **Summarize the document**
- **Create a list of the document's key points**
- **Answer a question that can be answered by what's in the document**
- **Analyze the document, for instance its tone, style, bias, etc.**
- **Translate**

Example: Questions about test specifications

An employee uploads a lengthy test specification for a particular part or assembly, and is able to ask questions about it and receive answers. For example, a question might be "what load testing is required for part x, and what maximum displacements are allowed for those loads?"

The AIDE will answer the employee's question and give citations to the exact section of the document that specifies this information.

03



AI Subject Expert

AI Subject Expert: LLM + Documents + Prompt

Software that combines a chatbot with one or more documents and a custom prompt to meet a specific need. Most AISEs engage in conversation about a particular topic based on a preselected set of documents that are uploaded for the model's use. Some AISEs also use stored prompts to elicit certain behaviors (e.g., help with brainstorming, write an email, evaluate if a document meets style guidelines). Note: OpenAI calls these "GPTs".

Example: Travel Policy AISE:

- LLM that uses your company's travel policy document(s) and is instructed to answer questions about the travel policy.
- employees can ask the AISE any travel policy question and get an immediate answer without having to scroll through a lengthy document



04



AI General Expert

AI General Expert (RAG): LLM + Search + Prompt

In its simplest form, an AIGE uses retrieval-augmented generation (RAG) to search for relevant knowledge and then provides that knowledge to a chatbot. This way, the chatbot can answer using the knowledge it was provided, making it able to speak intelligently about information that was not included in its training data. An AIGE can cite its sources for human verification, and greatly reduces the possibility of hallucinations by the LLM. For these reasons, RAG has become the standard for informing an LLM with enterprise knowledge.

Example: AI Expert

Can answer any question about your business; like the travel policy AISE above but on any subject, always up to date, and can securely handle knowledge with different permissions; the knowledge does not have to be contained in a curated set of documents but can reside anywhere in the systems that are searchable by the AIGE

05



AI Assistant

AI Assistant: RAG + workflow

An AI assistant takes RAG to the next level by making it multi-step. It follows a multi-step workflow which involves multiple actions, such as searches, branches, loops, etc. Workflows for AI assistants are predetermined (they are static and defined ahead of time, before the assistant starts to execute). AI Assistants can also use tools; tools are external applications separate from the LLM that give it additional capabilities; in fact, the search capability of an AI Assistant is often considered a tool.

Example: a status report assistant with the following predetermined workflow:

1. Request the desired timeframe for the report (e.g., week, month, quarter, etc.)
2. Open the status report template
3. Collect project X meeting recordings within the requested timeframe
4. Summarize the major issues, important milestones and action items into a bulleted list.
5. Insert the list into the status report template.
6. Identify all project X tickets with created or changed status in the requested timeframe
7. Insert the counts into the status report template
8. Summarize the customer's satisfaction with the project and highlight key opportunities to improve it.
9. Assess which of these areas may also offer opportunities for additional revenue.
10. Write a 3-5 sentence summary of the state and progress of the project
11. Add that summary to the template.
12. Save the status report to the repository.
13. Email the status report to the Project X distribution list, including the summary in the body of the email.

06



AI Agent

AI Agent: RAG + Tools + Planning

Similar to an AI assistant but instead of a predefined workflow, the LLM plans the workflow, and chooses from a suite of tools to execute the workflow. In this way, the LLM also decides when the work is completed (it may go back and repeat a part of the workflow, or modify it if it finds that its original plan isn't working). The workflow is created in real-time based on the request. With an AI agent, the steps to be performed is unknown and unpredictable. The use of tools expands what AI agents can do; the agent determines which tools to use in order to respond to the request.

Tools may help provide information to the LLM or give it the ability to interact with the real world or other systems. Examples of tools (in addition to the search example given earlier) include:

- **Calculator for math computation**
- **Email tool to send email**
- **Browser to interact with websites**
- **Credit card tool to make online payments**

A common example:

Deep Research agents like those provided by the main LLM vendors, or a Deep Internal Research agent that does the same kind of research but on a company's internal knowledge. The agents:

- conduct searches to collect information on a topic
- evaluate that information to determine what additional topics should be explored
- gather information on those related topics by running additional searches
- repeat the above process until the information gathered fully answers the request
- synthesize of the results into a report
- formalize the report, often as a PDF, with a table of contents, images, etc.
- They are agentic because they determine their own workflow and they choose when to use which tools (primarily, an internet search tool).



07



Agentic AI

Agentic AI: Two or more AI Agents that collaborate

Agentic AI is an ecosystem of multiple AI agents that work together to accomplish an objective. The AI agents share information with one another and coordinate control of the process.

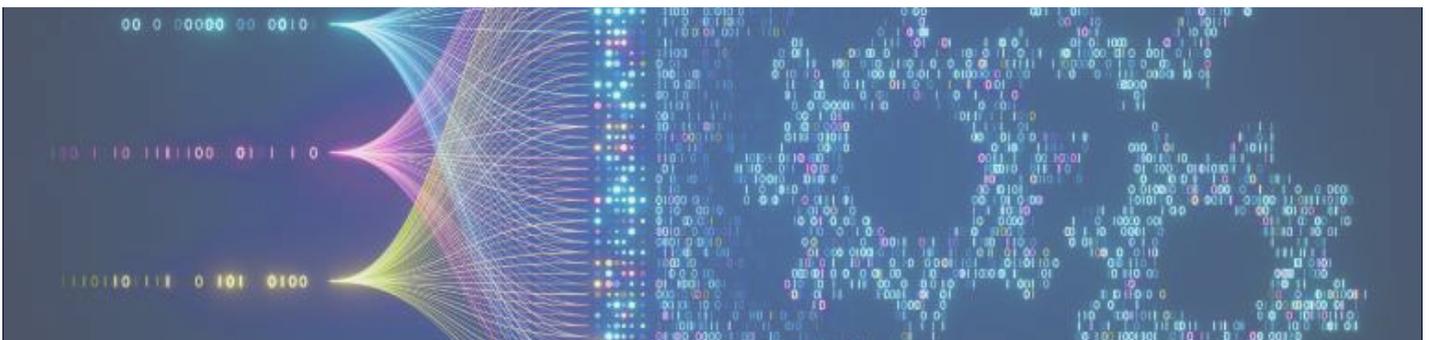
Agentic AI often includes a mechanism for AI agents to discover one another, so that they can determine which agents are most likely to be helpful in reaching the objective. Note: at the current state of technology, using an ecosystem of multiple specialized AI agents is often much more robust than trying to accomplish it with a single agent.

Example (that will soon be commonplace): A collection of agents dedicated to customer service (e.g., deployed by an airline):

- Customer satisfaction agent. Handles the primary customer conversation with access to tools such as a customer information lookup tool, a policy lookup tool, a flight information tool, a bookings lookup tool, etc.)
- Customer verification agent. Specializes in confirming a customer's identity through appropriate means (account number, phone, email, zip code, even 2-factor authentication)

And others like:

- Domestic reservation agent
- International reservation agent
- Group reservation agent
- Reservation change agent
- Refund agent



The Steps on the Journey to Agentic AI

The table below summarizes the systems described in this paper for quick reference.

System	Composition	Business Example
Chatbot	LLM	Chatbot to answer common-knowledge questions, summarize and translate documents, and write an email from a prompt
AIDE	LLM + document	Upload a document and chat with it
AISE	LLM + documents + a prompt	AI expert on your company's travel policy (public & static knowledge)
AIGE	LLM + search + a prompt	AI expert on anything in your business (secure & dynamic knowledge)
AI Assistant	RAG + predefined workflow	AI worker that creates a status report based on current information from all relevant sources (even email exchanges)
AI Agent	RAG + tools + planning	AI customer sales agent that can interact with a customer to place a new order
Agentic AI	Two or more AI agents that collaborate	AI customer service "department" that can place a new order, change an existing order, process a refund, check stock, track a shipment, etc.

Conclusion: Choosing the Right Level of AI

The journey from chatbot to agentic AI shows how systems grow in complexity, capability, and autonomy. While advanced AI can deliver significant value, simpler solutions - like AI Document Experts or AI General Experts - often meet business needs with lower risk and cost.

Decision-makers should focus on the system that aligns with the problem, desired autonomy, and organizational readiness. By choosing the right level of intelligence for the right task, organizations can benefit from AI effectively, without chasing complexity for its own sake.

For more information, please visit

[Enterprise Agentic AI platform](#)





About Sinequa by ChapsVision

Sinequa transforms how work gets done by providing employees with knowledgeable, accurate, and secure AI Agents that streamline workflows, navigate complex enterprise data, and deliver reliable, traceable insights. By combining enterprise search with generative AI in a configurable, easily managed framework, Sinequa enables organizations to deploy out-of-the-box Agents or tailor specialized workflows, ensuring every interaction is trusted, governed, and compliant while empowering employees to focus on high-value work.

For more information, visit www.chapsvision.com/platform/sinequa/ or www.sinequa.com