Just in Time

AI Agents

An Introduction to Agentic Systems, Market Impact, and Future Risks

November 2025

iason ESSENTIAL SERVICES FOR FINANCIAL INSTITUTIONS

Executive Summary

The rise of **Agentic AI** marks a significant step forward moment in the current AI revolution, marked by accelerating technological breakthroughs and massive capital deployment. Unlike previous AI developments, Agentic AI **holds the promise of significant practical utility and adaptability**, **already demonstrating early capabilities as a powerful automation tool**, which presents unprecedented opportunities alongside new complexities.

We provide here an overview on **Agents, starting from the basis**. A deep dive is then dedicated to possible **implementation strategies**, considering Agents' potential for widespread adoption in the short to medium term. Finally, the concluding sections outline first **global market perspectives** across all major sectors where the use of AI Agents can make a significant impact and, lastly, introduces the discussion on the **key risks** to monitor during agentic AI adoption. Further on this topic will follow on upcoming iason studies – *stay tuned*.

At a Glance

Keywords: Artificial Intelligence, GenAI, Agentic AI, Domain-Specific Intelligence

01

AI Agents Overview

Context of AI Agents

Beyond Standard LLMs: Domain – Optimized Agents

Types of Agents

AI Agents Overview 1/4

Context 1/2

Artificial Intelligence (**AI**) is continuously evolving, **from basic statistics and ML** to **powerful agents** capable of learning, interacting through natural language, and acting with more autonomy

This **evolution is reshaping** how we **work**, **communicate**, and **solve problems**: **faster** and **smarter** ***than ever before***

Artificial Intelligence (***AI***)

A field of computer science focused on building systems that mimic human thinking

Machine Learning (***ML***)

Branch of AI that enables systems to **learn from data** and **improve** their **performance** on specific tasks without explicit programming

Deep Learning (***DL***)

A **type of ML** using **multi-layer neural networks** to model **complex patterns**

Generative AI (***GenAI***)

Class of **DL** methods that generate new, realistic content (text, images, audio, or code) by modeling underlying data patterns

Large Language Models (***LLMs***)

GenAI **specialized in language**, capable of generating human-like text and performing linguistic tasks

Agents

Autonomous systems powered by LLMs that can make decisions and complete tasks, and learn from experience

AI Agents Overview 2/4

Context 2/2

The current **most vibrant space of development** is clearly **Generative AI**, an **advanced form of Deep Learning** that not only analyzes data but also **creates new content understanding underlying patterns**

Large Language Models (**LLMs**) are a cornerstone for GenAI applications, enabling machines to generate human-like text, answer questions, and assist with various linguistic tasks

While LLMs are a significant breakthrough, **GenAI extends beyond them**, covering broader applications such as:

- financial modeling
- synthetic data generation
- and automated decision-making

GenAI-based Agents **are now the frontier to unlock a *human-machine* interaction** like never experienced before, allowing at the same time to **craft *domain-specific* applications**

This evolution in AI unlocks new opportunities, particularly in Finance, where intelligent automation and predictive capabilities can be achieved through agentic systems, driving efficiency and innovation

AI Agents Overview 3/4

Beyond Standard LLMs: Domain – Optimized Agents

GenAI holds promise, but its ability to address complex, domain-specific challenges often **requires customization**, including fine-tuning on specialized data, integration with existing systems, and alignment with industry-specific regulations and workflows.

To be **effective**, GenAI **must be shaped into domain-specific** AI Agents

These Agents:

- **analyze data**
- **take decisions**
- **execute tasks** aligned with **business goals**

Offering greater:

- **precision**
- **reliability**

than generic models

	Standard LLMs		AI Agents
Purpose	General purpose		Task-oriented
Customization	Limited, often costly to customize		Tailored
Decision-Making	Lack of execution and reasoning capabilities		(Semi) Autonomous
Context Awareness	General purpose		Domain specific
Control & Audit	Limited		Reinforced by design *(e.g., human in the loop)*

AI Agents Overview 4/4

Types of Agents

AI Agents can be broadly classified into **several key types**, according to the level of capability, sophistication of decision-making processes, and degree of autonomy they possess. In the end, we acknowledge that an Agent can share many of the following properties at once.

Simple reflex Agents

Type of agents that select actions **based only on the current percept** (*what it sees or senses "now"*), **ignoring the rest of the percept history**. They function using **condition-action rules** [If condition, then action]. Moreover, they do **not store memory** of past events or consider future consequences

Autonomous Agents

Intelligent systems that operate **independently**, perceive the environment, and take decisions **without direct human intervention** to achieve **specific goals**.
They may learn from experience, adapt to changes, and pursue given objectives **proactively** and **reactively**

Learning Agents

Intelligent agents that can **improve their performance over time** by **learning from experience**. They learn how the environment behaves, how to make better decisions, or how to achieve goals more efficiently

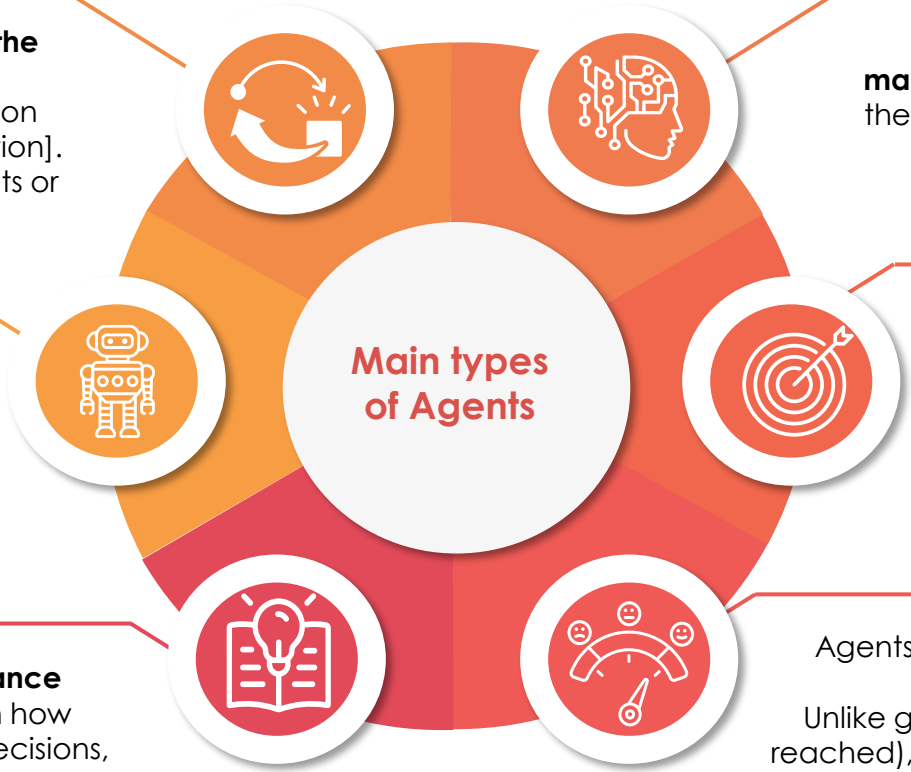

Model-based reflex Agents

More advanced form of reflex agents that **maintain internal state** to keep track of aspects of the **world they cannot currently observe**. They use a **model of the environment** to make more informed decisions

Goal-based Agents

Intelligent agents that **select actions** by considering **future consequences** and how those actions helps in **achieving a desired goal**. Unlike reflex agents (which act on immediate inputs), goal-based agents **reason** and **plan** before acting.

Utility-based Agents

Agents that choose actions to **maximize their level of satisfaction** measured by a **utility function**.
Unlike goal-based agents (which just check if a goal is reached), utility-based agents evaluate **how good each possible outcome is.**

02

Anatomy of an AI Agent

Overview

Deep Dive on the “Core”

Deep Dive on the “Behavior”

Knowledge Graphs

Anatomy of an AI Agent 1/4

Overview

An AI Agent can be conceptually understood through its **Core** architecture and resulting **Behavior**: the Core-Behavior framework enables researchers and developers to systematically examine how an **agent's internal architecture** translates into **observable actions.**

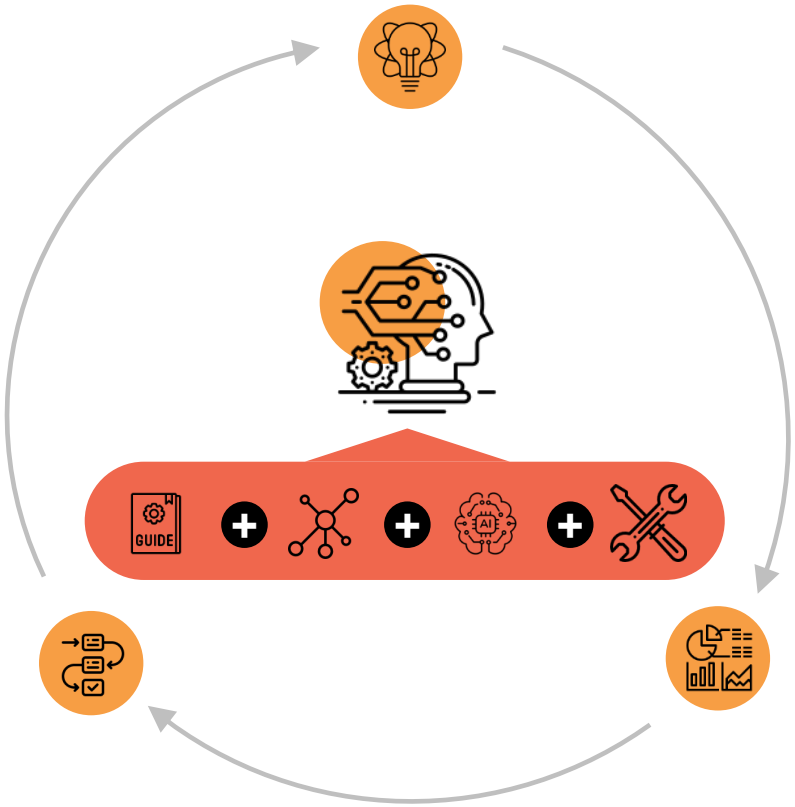

Core

Instructions
Proper Prompting and Action Plan for Agent's instruction (e.g., CoT [1], ToT [1], …)

Knowledge Base & Memory
Data, Context and Memory (Working and Persistent Memory)

AI Models
Usually LLM(s) acting as the agent's core intelligence, governing the agent's reasoning and decision-making

Tools
Mechanisms that allow the Agent to interact with external systems, perform complex operations, or access specialized resources

Behavior

Process
Analytical phase where the agent gathers and interprets environmental information

Plan
Represents the agent's reasoning/ elaboration process where it formulates strategies for action

Execute
Execution phase based on the planned strategy

(1) Chain of Thought (CoT), Tree of Thought (ToT)

Anatomy of an AI Agent 2/4

Deep Dive on the "Core"

AI Agents can mark a **major step in business automation**, using semi-autonomous decision-making to boost efficiency. At their **core**, a **layered architecture** combines machine learning, predictive modeling, and strategic frameworks to **analyze**, **forecast**, **plan**, and **act.**

Instructions

Each LLM underneath the Agent, depending on its tasks, must be instructed with (***1***) a proper **prompting**, and/or **fine-tuning,** to **guide** its behavior and (**2**) an "**action plan**" to define how it approaches problem-solving (e.g., Chain or Tree of Thoughts).
This allows to **control the specific LLM behavior**, even when fine-tuning (or training) the model is not an option or deemed necessary

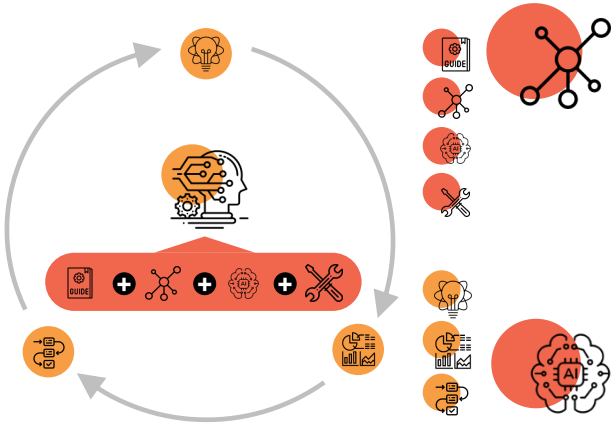

Knowledge Base & Memory

An Agent needs **more than a "database"** to properly work, especially in domain specific applications.
The so-called Knowledge Base must also include **context** of the data and tasks, and must allow the Agent to **account for past interactions** with the user (e.g., memory).
Complex Knowledge Bases can be tailored through **knowledge graphs**

AI Models

This is the centralized operational engine for agentic tasks, most commonly **LLMs.** The models functions as the **agent's core intelligence**, determining how it interprets inputs, perceives events, and reacts to various conditions. AI models govern the agent's **reasoning**, **decision-making**, and **response generation**

Tools

The Agent must not "*reinvent the wheel*", but rather have **access to properly developed/ proprietary models**, scripts, functions,…
This enforces also control and auditability; it ensures that the Agent acts more as expected, also in repeatable actions

Anatomy of an AI Agent 3/4

Deep Dive on the "Behavior"

At the heart of every effective AI Agent lies a **systematic approach to intelligent behavior**, structured around three core phases that mirror human cognitive processes: **Process**, **Plan**, and **Execute.**

Process

First step to be executed in order to transform **raw information** into **structured understanding**. In this phase the Agent analyze its environment and inputs, gather relevant data and proceeds to identify key constraints, objectives, and available resources

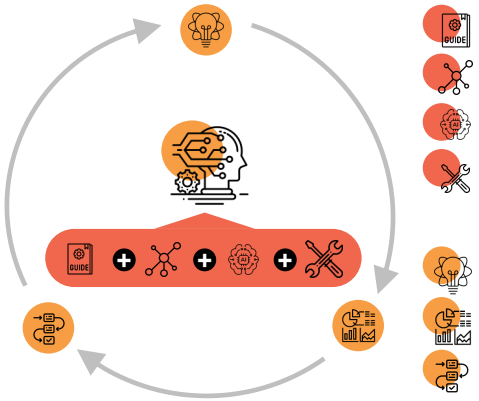

Plan

This is the Agent's **reasoning core**, translating a complex goal into an **actionable, sequential strategy**.

For routine tasks, the Agent retrieves and utilizes **predefined instructions** from its knowledge base; for novel or complex problems, it can **reason by itself** to logically decompose the main objective into manageable, sequenced sub-steps and selects the appropriate tools.

The depth of an **Agent's reasoning** capabilities is not uniform; it varies significantly based on its foundational component — namely, the use of **specialized LLMs (Reasoning LLMs vs. standard LLMs)** and the quality of instruction fine-tuning applied during its development or adaptation.

Execute

The agent **acts and adapts** based on its plan. It performs the planned actions (leveraging on LLMs only or exploiting the available tools), monitors results in real-time, and adjusts its approach when needed.
This could mean answering a query, perform a *what-if analysis*, or drafting a report. **Effective execution** depends on access to **external tools** and **resources**, which agent connects [1] to, for example, via APIs and integrations with SaaS products

(1) The accessibility is one of the key topics for agent risk management: e.g., APIs can be breached by malicious attacks, resulting in weak security in specific agents' architectures

Anatomy of an AI Agent 4/4

Knowledge Graphs

As explored in the deep dive on the Core, **Knowledge Graphs** can constitute a fundamental component of evolute AI agent's knowledge bases. Beyond organizing structured knowledge, they serve as an explicit representation of the **logical relations among data** and **reasoning paths** the agent must follow to take informed decisions. By encoding how concepts connect and interact, Knowledge Graphs can ground the **agent's behaviour** and **decision-making** processes. As AI agents increasingly need to integrate and reason over information from multiple heterogeneous data sources, **Knowledge Graphs provide a unified and interpretable framework** for both **knowledge representation** and **logical reasoning.**

A **Knowledge Graph** is established through:

1. **Ontology** definition
2. **Nodes** creation
3. **Edges** among nodes
4. **Properties** for more context

4. Attach Properties

These are **attribute-value pairs** that can be attached to nodes and edges to provide **rich, contextual information** that AI agents can leverage for sophisticated reasoning and decision-making

3. Establish Edges

Define the **relationships of interest** between Nodes (for example, an edge can establish the relationship between a client and a company as a "Client Relationship")

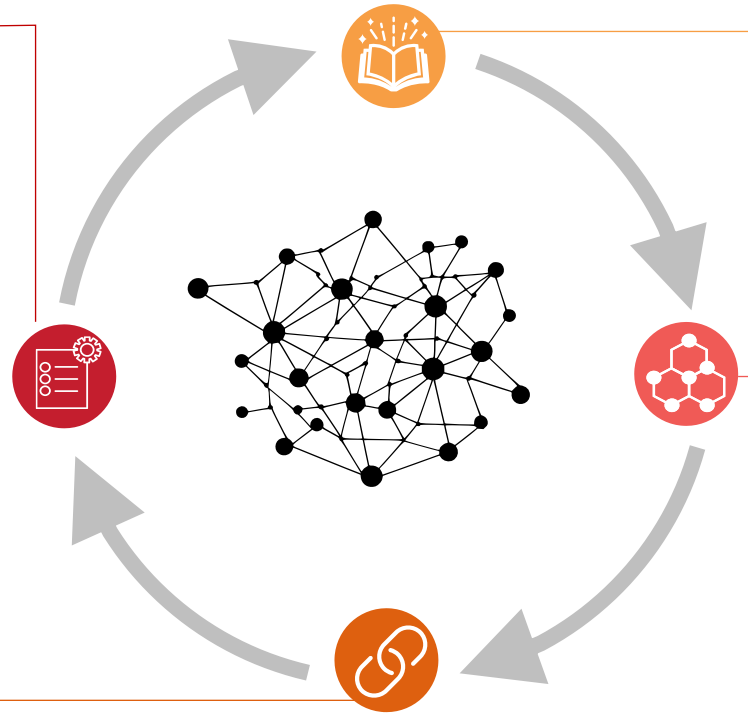

1. Define Ontologies

Ontologies are used to create a formal representation of the graph's entities. They are typically based on a **taxonomy** and on Resource Description Framework

2. Create Nodes

Nodes represents **elements** or **entities** of the Knowledge Base.
Anything can act as a node (e.g., Customers, Trades, Counterparties, …)

03

Implementation Strategy

A Building Block Approach

Balancing Control and Flexibility

Implementation Strategy 1/2

A Building Block Approach

Creating an **agentic framework** is key to properly take advantage of the current state of GenAI. The **implementation of an Agent** follows a common path, where different small Agents can be combined modularly in **domain-specific Agents** (which can interact/collaborate with each other).

Without a methodical and well-structured approach, organizations face **different risks**:

- **Wasting resources** on fragmented and non-scalable solutions
- **Creating critical dependencies** on specific LLM vendors
- **Losing control** over automated decision-making processes
- **Limiting interoperability** between different systems

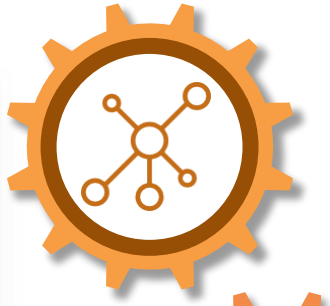

Knowledge Base

The first foundational element is a **highly structured Knowledge Base**, **specifically built** to address the unique requirements of use-case at hand. The Knowledge Base can include technical documentation, internal policies, historical cases, outputs from other agents or tools, and can be filtered or categorized by metadata or domains

Sub-Agents & Tools

The second building block centers around **specialized agents** and the **tools** they use to perform specific tasks. **Sub-agents** interact with each other and can be **built on generalist** or **fine-tuned LLMs**, **provided with commands** (e.g., *prompts and chain of thoughts*) and **tools dedicated to the specific purposes** these Agents are designed for

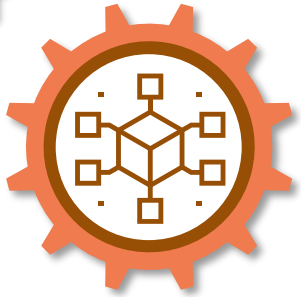

Graph

The final building block is represented by a **logical structure encoded through a graph** that guides the system's behavior and coordinates the interactions between agents, tasks, and data.
This **boost interpretability**, **interoperability** and **control** (e.g., *guardrails* included by-design)

Implementation Strategy 2/2

Balancing Control and Flexibility

An optimal **Agent-based architecture** must strike a **balance between control** (such as security and compliance) **and flexibility** to enable reasoning and innovation. The ideal trade-off depends on the business value intended to be unlocked by the organization and its risk-tolerance.

Guided Agent
Delegates the most critical or hallucination-prone analyses to specific tools

Autonomous Agent
A more independent agent that does not rely on predefined specialized tools

Greater traceability, auditability, and replicability	**Security and Control**	Less transparent decisions, higher risk of hallucinations
Lower costs, thanks to the efficiency of dedicated tools	**Operating Cost**	Higher, due to increased autonomous reasoning
More defined pathways, consistent responses	**Reproducibility**	Greater variability
More constrained (e.g., bounded to specific tools)	**Flexibility and Adaptability**	Handles unexpected questions and scenarios
Longer timelines due to tools development	**Implementation Time**	Faster deployment
Requires constant tools updates to meet new requirements	**Management of Dynamic Environments**	Lower sensitivity to changes in data and sources

04

Market Perspectives

State of the Art and Potential Evolution

Financial Services Application

Market by Numbers

Market Perspectives 1/4

State of the Art and Potential Evolution

The world of **AI** is undergoing a **radical transformation**, mostly driven by the rise of **semi-autonomous AI agents.** It is reasonable to think that the next decade will mark the beginning of the era of **more concrete adoption** of these systems in key sectors of the **global economy.**

	Current Adoption	Future Perspective
Finance & Banking	The **financial sector** shows the **highest AI implementation** among regulated industries, focusing on anti-money laundering, fraud detection systems, risk modelling, and trading algorithms[1]	The **financial sector** can see **rapid growth** in the next decade, with **AI Agents, for example, helping in managing portfolios**[2] through advanced predictive analysis of global markets and economic data together with a massive adoption in **Risk Management**
Customer Services	**Customer service** is one of the most developed area for AI agents, featuring **extensive chatbot deployment** across industries, with AI agents offering superior natural language capabilities compared to traditional scripted chatbots [3]	**AI Agents** will **transform customer service** by handling complex conversations and creative problem-solving, coordinating with human teams and integrating with CRM systems for maximum impact [4]
Sales	**AI Agents** already excel at **follow-up** automation, **scheduling**, and **sentiment analysis**, with companies using advanced chatbots for initial prospect screening before human handoff [5]	**AI Agents** may evolve into **comprehensive sales assistants**, capable of conducting autonomous sales conversations, negotiating basic contracts, and providing real-time strategic recommendations
Human Resources	**AI Agents** are transforming **HR** by streamlining time-intensive processes like **recruitment** and **assessments**, with chatbots managing employee inquiries about company policies [6]	Future **HR assistants** will autonomously conduct interviews, manage performance evaluations, and personalize development paths, with advanced analytics enabling better predictions

(1) See *"Artificial Intelligence in Financial Services – White Paper"; World Economic Forum, January 2025*
(2) See *"Alpha Agents: Large Language Model based Multi-Agents for Equity Portfolio Constructions"; T.Zhao et al, BlackRock Inc, August 2025*
(3) See *"Complete Guide to chatbot of Customer Service on 2025"; Botpress, March 2025*
(4) *Gartner predicts Agents AI will autonomously resolve 80% of Customer Service Issues without human intervention by 2029; Stamford Conn. March 2025*
(5) See *"AI Sales Agents, Use Cases, Industries and Key Tools"; Botpress, June 2025*
(6) See *"Reimagining recruitment: traditional methods meet AI intervention – A 20-year assessment"; A. Rukadikar et al. Cogent Business & Management, January 2025*

Market Perspectives 2/4

Financial Services Application

One of the most compelling **use cases** for **AI** and **AI Agents** is undoubtedly within the **financial services industry**, both in terms of innovation level and availability of resources to invest in implementing new technologies. Below are just some of the possible applications.

Automated Risk Management[1] through:

- **Real-time monitoring**: agents with capabilities and tools for financial data analysis, behavioral patterns, market conditions
- **Proactive alerts**: agents handling early warning for defaults and credit deterioration
- **Dynamic risk management** through real-time simulations, shock scenario evaluation and what-if analysis eased

Advanced Customer Service[4] through:

- **Internal Task Automation**: Operational agents automate standardized and repetitive tasks such as invoice recording, bank reconciliation and document processing
- **Customer Service and Hyper-Personalization**: providing 24/7 assistance AI advisory agents offer automated financial advice, creating personalized investment plans based on risk profiles.

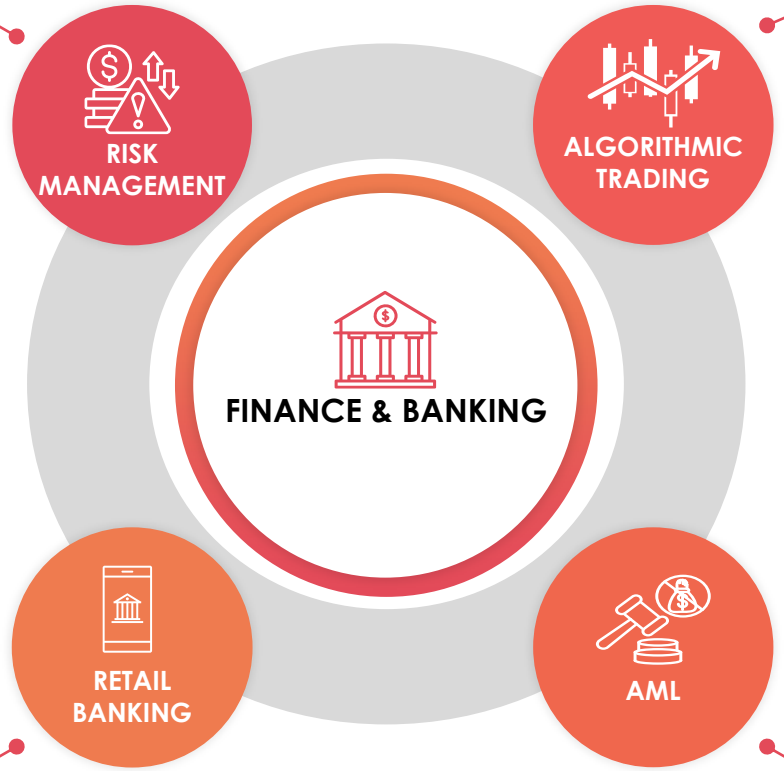

Optimized Algorithmic Trading[2] through:

- **Markets and Credit studies**, drafted automatically by agents, providing traders with more and faster insights
- **Agents** powered with **Reinforcement Learning** algorithms to dynamically optimize trading strategies in real-time, offering traders a faster trading ideas generation process
- **High-Frequency Trading (HFT)** boosting thanks to Agents which incorporate in trading signal processing news consumption in real-time

Proactive AML Compliance[3] through:

- **AI Agents** that analyze billions of transactions daily to detect anomalous patterns indicative of financial crimes
- **Automated reporting**: SAR generation with supporting documentation
- **Real-Time Verification and Onboarding**: improve efficiency in processing critical documents such as KYC forms

(1) Many players in the European Financial Industry have adopted a semi-autonomous digital worker to support risk management functions
(2) Goldman Sachs deployed a multi-agents trading systems for real-time automated equity trading execution that absorb massive data volumes, identify patterns and execute trades autonomously
(3) Wells Fargo AI Agents analyze transactions in real-time, detecting anomalies and suspicious patterns with greater accuracy and reducing false positives
(4) Bank of America's *Erica* AI-powered virtual assistant provides 24/7 customer support and personalized guidance leveraging advanced analytics and machine learning to offer proactive financial insights

Market Perspectives 3/4

Market by Numbers – Best Case Scenario

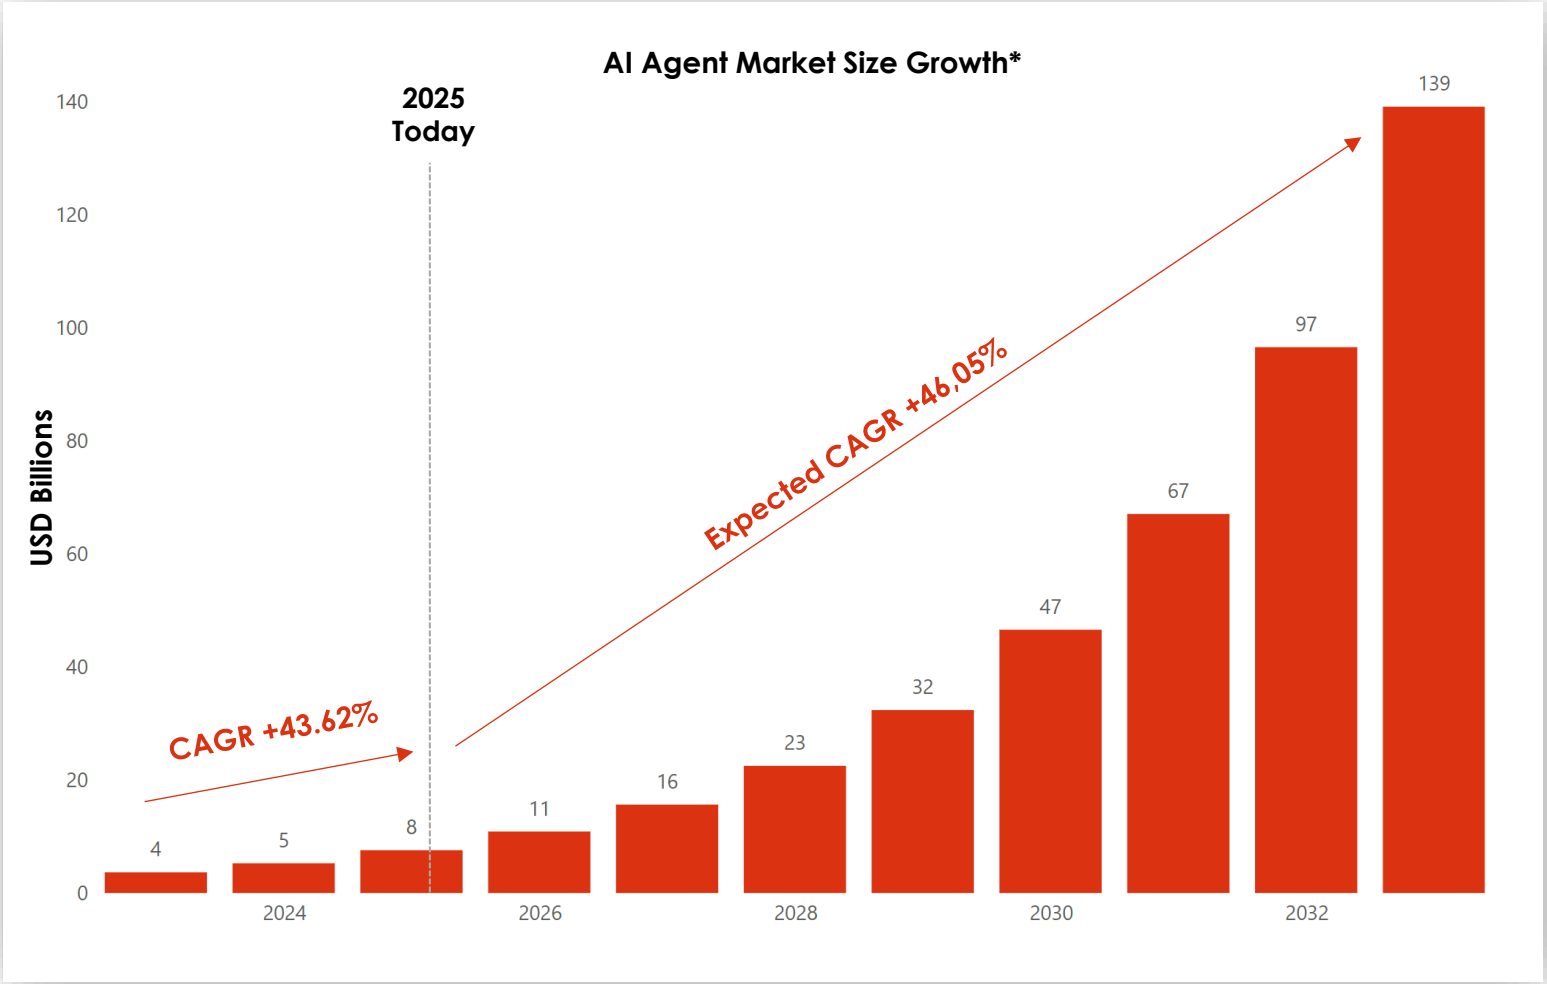

* Data: Market.US, 2025, Global Market Analysis, Size, Share, Growth, Trends and Forecast, 2019-2033

Since 2023, the market for AI agents has garnered increasing attention, with a **market value growth** of nearly **$4 billion** between 2023 and 2025.

The potential applications across various industries, driving business **efficiency and productivity**, are considered key factors in the evolution of business processes in the coming years. Companies will increasingly adopt AI agent solutions with an **expected CAGR of 46.05%** over the next five years, reaching **$50.31 billion.**

While in 2023, the largest share of the market was represented by Single Agent Systems (over 73,49% of the market share), it is reasonable to expect that **multi-agent systems** will experience significant growth in the coming years, with CAGR estimated between **35% and 46%** through **2030–2034**.

Market scenarios projected here are based on **optimistic assumptions** regarding macroeconomic measures, linear development of regulation, and increasing demand for automation in business operations.

Market perspectives 4/4

Market by Numbers – Conservative Scenario

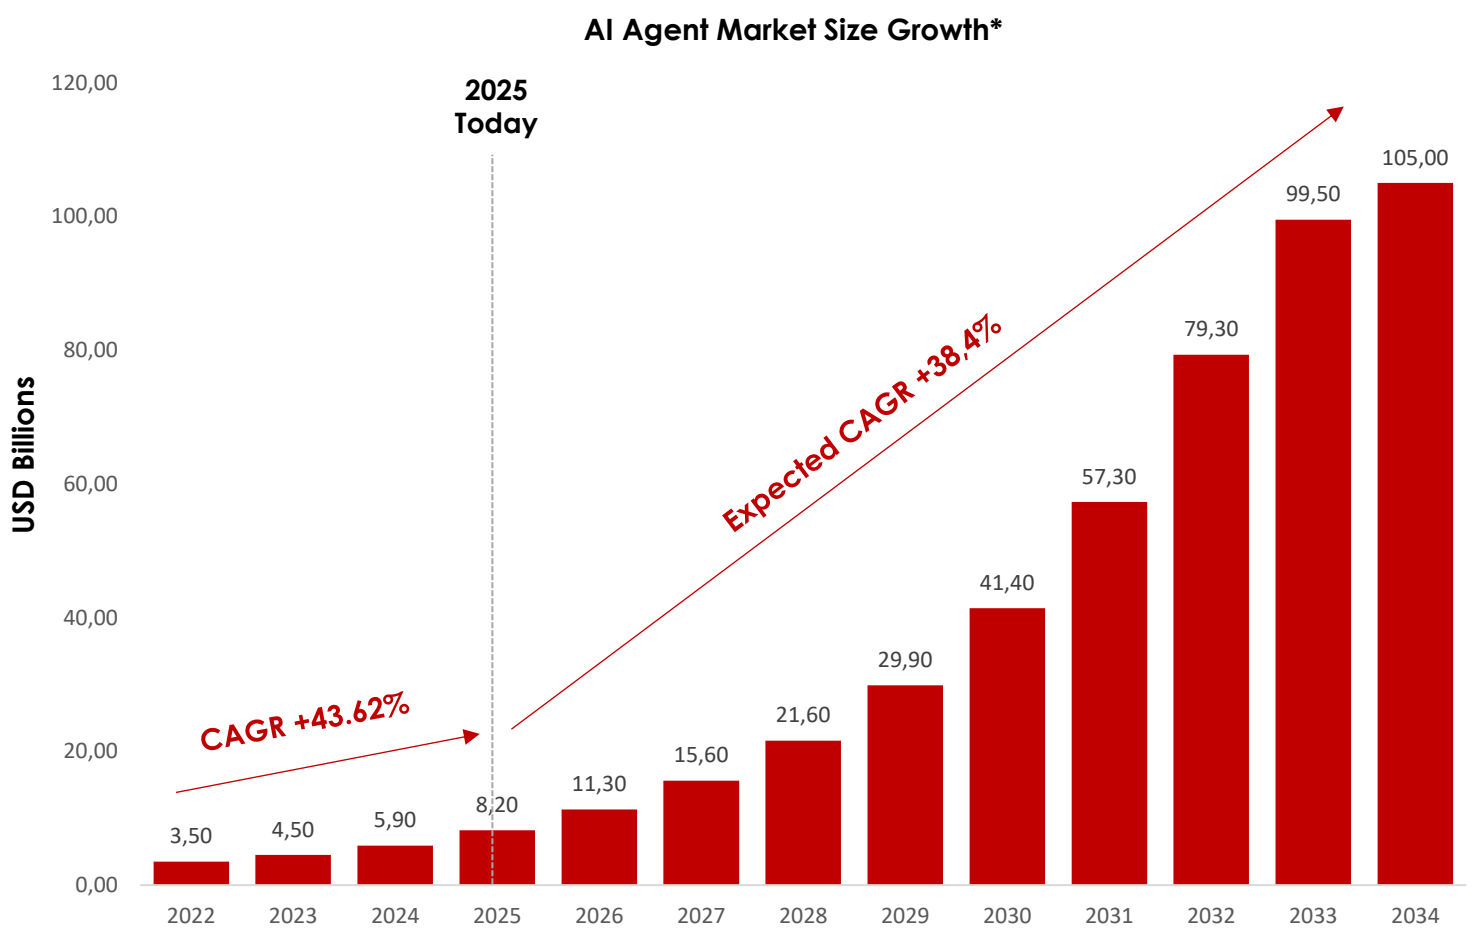

According to a recent report of Global Market Insights, the rapid growth of the AI Agents market, driven by substantial investments in the sector in recent years, will continue at a sustained pace in the years to come, but is expected to settle at a **more moderate CAGR** (Compound Annual Growth Rate) in the final years of the next decade. Overall, the study estimates CAGR growth by 2034 of 38.4% leading to **$105 billion.**

The more moderate growth rate is primarily attributed to two key restraining factors identified in the report:

- **Technical limitations**: such as the lack of contextual understanding and accuracy.
- **Market adoption barriers**: mainly due to high initial implementation costs and limited awareness among business users, which may slow training adoption and affects scale and ROI.

* Data: Global Market Insights: AI Agents Market Size – By Agents, by Technology By Deployment Mode, By Application, By End Use, Growth Forecast, 2025 - 2034

05

Key Risks

Risks Associated with AI Agents Adoption

Key Risks

Risks Associated with AI Agents Adoption

The adoption of **AI agent–based** technologies, although offering all the benefits described in the previous sections, if not carefully designed and guided by **professionals** capable of **leading the change**, carries with it a **series of risks** that should not be underestimated.[1]

Hallucination

AI Agents can generate incorrect, unverifiable, or misleading outputs (***hallucinations***). This may lead to wrong decisions, loss of customer trust, or reputational damage if proper **human-in-the-loop** oversight mechanisms are not in place.

Third-Party risk

Many AI Agents rely on models or cloud infrastructures provided by **third parties**. This exposes organizations to risks related to:

- **service reliability** and continuity,
- **regulatory compliance** (data protection, localization),
- **technological lock-in**, making it difficult to switch to alternative solutions.

Obsolescence

The **fast pace of AI evolution** means that models, frameworks, or platforms can quickly become outdated. Companies risk investing in technologies that may soon be unsupported or less competitive, leading to **additional costs** for updates or migrations.

Security

AI Agents can be manipulated with **malicious inputs** (e.g., prompt injection, data poisoning) or become vectors for **cyberattacks** (e.g., APIs can be breached), exposing the organization to security threats.

For continued insights and Generative AI updates, follow iason. Future publications are planned to address the critical risks organizations face when implementing AI Agent solutions and detail the necessary mitigation strategies for successful, full-scale adoption.

(1) For a full picture of risks associated with Gen AI adoption see: "Artificial Intelligence: Financial Industry Market Overview"; *G. Campaniolo, J. Figuriello, N.Mazzoni, Just in Time, July 2025*

ESSENTIAL SERVICES FOR FINANCIAL INSTITUTIONS

iason is an international consulting firm that has been supporting both financial institutions and regulators in topics related to Risk Management, Finance and ICT since 2008

Strategy

Strategic advisory on the **design** of **advanced frameworks** and **solutions** to fulfil both **business** and **regulatory needs** in Risk Management and IT departments

Methodology & Governance

Implementation of the designed **solutions** in bank departments **Methodological support** to both **systemically important financial institutions** and **supervisory entities**

Solution

Advanced **software solutions** for **modelling, forecasting, calculating** metrics and **integrating** risks, all on cloud and distributed in Software-as-a-Service (**SaaS**)

KEEP IN TOUCH

iason

Company Profile

iason is an international firm that consults Financial Institutions on Risk Management. **iason** integrates deep industry knowledge with specialised expertise in Market, Liquidity, Funding, Credit and Counterparty Risk, in Organisational Set-Up and in Strategic Planning.

Valerio Ciminelli

www.iasonltd.com