


NK RESEARCH

AI Agents

— September 2026

**Beyond the model:
Mapping the value chain
of the agentic shift**



AI AGENTS: WHAT ARE THEY?

The terms get used loosely, so it is worth separating them precisely. **Agentic AI** is the **underlying technology**, the architecture and techniques that let a model plan, take actions and evaluate its own results. An **AI agent** is the product **built on top of that technology**: the thing a business or an individual actually deploys **to get work done**. In short, agentic AI is the engine, and an AI agent is the car built around it.

It helps to contrast this with what came before. A **classic chatbot**, or “copilot,” behaves like a junior analyst who will **answer exactly the question that is asked**, and only that question, and then stop and wait. You prompt it, it replies once, and if you want anything further, you have to ask for it explicitly, prompt by prompt. Useful, but **entirely dependent on the user driving every step**.

An **AI agent** is closer to an analyst you can hand a mandate to. Instead of a single question-and-answer exchange, you give it a goal, and **it works out the plan itself**: which steps are needed, in what order, and which tools to use along the way (a web search, a piece of code, a spreadsheet, an internal database). It **executes** those **steps**, **checks its own output against the original goal**, and **loops back to correct itself** before it is satisfied. Only then does it come back to you with a **finished result**. The defining shift is the **move from a single prompt-response exchange to a closed loop** that runs largely on its own: goal, plan, act, check, deliver, repeating the middle steps as many times as needed **with far less need for a human to micromanage** each one.

The shift from chatbot to agent matters **commercially** as much as it matters **technically**, because it is exactly the line the market is now using to separate **durable businesses** from **fragile ones**. A large share of the **software built over the past three years** has been what people in the industry call an “**LLM wrapper**”, a product that takes an **existing model** such as GPT or Claude, **adds a user interface**, a handful of **prompts** and some **light automation**, and **sells the packaging**. A wrapper adds convenience but no differentiated capability of its own, which means it is exposed the moment the underlying model vendor decides to offer the same convenience natively. A **genuine agent**, by contrast, is **doing work** — orchestrating tools, taking actions, producing outcomes — that is **harder to replicate and harder to commoditise**.

Four traits recur in every serious definition of an agent, and they are worth holding onto because they double as the criteria the market is now using to judge which products are “real” agents and which are marketing dressed up as one:



Autonomy. The system acts toward a goal with minimal step-by-step supervision, rather than waiting for an instruction at every turn.



Tool use. It can reach outside the model itself, calling on search engines, running code, querying databases or operating other software, rather than only generating text from what it already knows.



Memory and context. It retains state across multiple steps of a task, so that step five still remembers what was decided in step one, rather than treating every exchange as a blank slate.



Goal-directedness. It is organised around producing an outcome (a finished document, a working piece of software, a completed transaction) rather than simply producing a plausible-sounding reply.

WHY IS EVERYONE TALKING ABOUT THIS, RIGHT NOW?

The centre of gravity of the AI conversation has moved. For the past two years, the **dominant question was which laboratory had built the best model**. That question has not disappeared, but it has been joined by a different one: **who can actually get agentic AI to run reliably, securely and profitably inside a real business**. Infrastructure, execution and cost discipline are increasingly differentiating business models within the AI ecosystem. That reframing is the biggest reason the subject has stayed at the top of the financial and technology press through the summer of 2026, and it shows up in the news cycle along several distinct threads.



Reliability is under genuine scrutiny

Coverage in outlets such as the Financial Times⁽¹⁾ has repeatedly asked **how reliable (or unreliable) autonomous agents actually are**. This is not idle scepticism: businesses are being asked to hand agents real budgets, real customer interactions and real infrastructure access, and the press is right to **ask whether the technology has earned that trust yet**. Several stories this summer have gone further, framing parts of the current AI safety testing regime as itself becoming **a source of risk rather than only a safeguard**, a sign that the reliability conversation is maturing past simple boosterism.



Security concerns are pushing budgets higher, fast

A second thread concerns **security**, and it cuts in an unusual direction: agents are not only a target for attackers, they are increasingly **a tool attackers themselves are experimenting with**. Public demonstrations of agents' hacking-like capabilities — automated reconnaissance, automated exploitation, even reports of frontier models being used to generate deceptive developer profiles designed to trick other engineers into integrating malicious code — have made the case for **higher cybersecurity spending** almost on their own. For enterprises, this converts an abstract technology risk into a concrete, budgeted line item, which is one of the clearest ways a trend shows up in earnings calls rather than just in headlines.



There is now a visible revenue signal

Perhaps the most commercially persuasive thread is the emergence of **hard revenue evidence**. Companies that have pivoted a product from simple search or static software toward an agent-based interaction have, in a number of disclosed cases, reported **meaningfully higher monthly revenue**. That is the kind of argument that moves a discussion from the innovation team's slide deck to the CFO's forecast, and it explains why so many software businesses are racing to describe themselves as “agentic” this year, whether or not the underlying product has genuinely changed.



Physical infrastructure has become the visible bottleneck

The fourth thread is the least “AI” of the four, and arguably the most important for markets: **electricity**. Training and running agentic workloads at scale requires **data-centre capacity** that, in several major markets, the existing electricity grid simply cannot support. The broader pattern is that US grid capacity has not kept pace with data-centre demand, forcing operators toward **on-site generation, long-term power purchase agreements** and, increasingly, **direct investment in energy infrastructure**. This is the mechanism through which AI has pulled classic industrial names (power equipment, generators, cabling) into what used to be a purely technology story.



The market reaction has been concentrated and large

All of this has shown up plainly in prices. The **Philadelphia Semiconductor Index (SOX) was up roughly 68% year-to-date⁽²⁾** as this note was prepared, more than four times the return of the NASDAQ 100 over the same period, itself up around 17%⁽¹⁾. That gap is a useful shorthand for the state of **the market's conviction**: it is not simply rewarding “technology” or “growth” as a broad category, **it is rewarding the specific layer of the stack** (compute and the infrastructure around it) that this note identifies as the current centre of gravity.



And the societal debate keeps the topic in the news on its own

Underneath the market-specific threads sits a broader public argument that keeps AI in the news regardless of quarterly results: **whether the technology is safe, and whether it will eliminate large numbers of jobs**. Neither question has a settled answer, and both feed directly into the reliability and regulatory debates above. For an asset manager, the practical takeaway is less about resolving those questions than about **recognising that they are a permanent feature of the backdrop**: a source of headline risk and periodic volatility that is unlikely to go away even as the commercial picture becomes clearer.

(1) Financial Times, Perplexity monthly revenue jumps 50% in pivot from search to AI agents, Apr 8 2026. (2) Data as of 8 September 2026.

THE COMPANIES AT THE HEART OF THE PHENOMENON

It is worth first simply mapping **who the main actors are and where they sit in the value chain**, from the raw compute that makes any of this possible, through the infrastructure that hosts it, to the software and applications that end users actually touch. Thinking in terms of this chain: **who supplies what to whom**.

01

The silicon layer: chipmakers and memory

At the base of the chain sit the **companies that make agentic AI computationally possible at all: Nvidia, AMD, Intel, Taiwan Semiconductor Manufacturing Company (TSMC) and Broadcom design** and fabricate the processors that both train and run these models. **Micron** sits slightly to the side of that group: it does not make AI accelerators itself, but **it makes the high-bandwidth memory those accelerators depend on**.

Micron is currently exposed to increased demand for high-bandwidth memory used in AI infrastructure: demand for memory has run consistently ahead of supply, and competition in that specific niche is thin, which has kept pricing power as a clear tailwind. Even so, Micron's product has not changed and is not itself an 'AI' product in any direct sense.

02

The orchestration layer: cloud, infrastructure and the model labs

One layer up sit the companies that turn raw compute into something usable at scale. **Microsoft and Amazon** (through AWS) provide the **cloud infrastructure** that the majority of agentic workloads actually run on, and both have pushed hard into governing agents, not just providing access to them, through concepts such as a "control plane" that lets an enterprise manage many agents at once. **Palantir** has positioned itself around connecting agentic AI directly to **enterprise data and operational workflows**. **Equinix, Lambda and Digital Realty** own and operate the **physical data-centre capacity** that all of the above ultimately depends on, and both have converted what used to be a capital-expenditure story into one of signed, revenue-generating contracts. **Anthropic and OpenAI** sit at the **model-development end** of this layer: their commercial argument increasingly rests not on the model in isolation but on the agentic products built around it.

03

The physical and industrial layer: electrification of the data centre

A less obvious but increasingly important group of actors sits outside the technology sector altogether. **Schneider Electric and Prysmian** supply the **electrical equipment and cabling** that every new data centre needs, and both have seen substantial share-price appreciation over the past year as a direct consequence of the data-centre build-out, even though **neither company has changed its underlying product**. **Schneider** is still selling essentially the same category of equipment it has sold for two decades; it is the **volume of data-centre construction** that has changed. Engine and generator manufacturers such as **Caterpillar and Cummins** belong to the same category: their equipment increasingly serves as **backup and supplementary power generation for data centres** in regions where grid capacity cannot keep up with demand. These companies are experiencing indirect exposure to higher levels of data-centre-related investment: companies pulled into the story by demand for what they already made, rather than companies that reinvented themselves around AI.



CONTINUED

THE COMPANIES AT THE HEART OF THE PHENOMENON

04

The enterprise software layer

Above the infrastructure sits a large and genuinely contested layer: enterprise software: **Salesforce, Adobe, ServiceNow, HubSpot, Capgemini, Oracle and C3.AI** all sit here, and this is the layer where the “potential beneficiaries and challenged business models” question is hardest to answer cleanly. These companies' fortunes **depend** less on whether they use AI and **more on whether they can convert their existing position** (the data they hold, the workflows they already run inside their customers' businesses) into a genuine execution layer for agents, rather than being replaced by agents that can reproduce their interface from scratch.

05

The application layer: turning models into finished products

At the top of the chain sit the companies that turn all of the above into something a non-technical end user can actually operate. Replit is the clearest current example: it takes **a business idea** and, through its agent, **turns it directly into a working, deployed piece of software**, orchestrating the model, the code generation, the testing and the deployment itself.

06

The layer under the most direct pressure: single-purpose content and service tools

Finally, a group of companies sits at the layer **most exposed to displacement**: businesses whose entire value proposition was a **single, narrow task that a general-purpose agent can now perform** as a side capability rather than a dedicated product. **Reverso** (translation), **Chegg** (homework answers), **Fiverr** and **Upwork** (marketplaces for small freelance tasks), **Otter.ai** (meeting transcription) and **Midjourney** (image generation in isolation from any broader workflow) all fall into this category. **Accenture**, whose traditional systems-integration and staffing-heavy model is being squeezed as agents take over more of the execution work that used to require large outsourced teams.




POTENTIAL BENEFICIARIES AND CHALLENGED BUSINESS MODELS

Not every company touched by AI is winning from it. As agentic AI moves from experimentation to real-world execution, **value is concentrating in a narrow part of the chain**, and draining away from another.

The map sorts the sector accordingly: how different **segments** may be **affected** by the current development of agentic AI, and who still has the **balance sheet and the time to adapt**. It is a working view, not a permanent verdict.

The pattern running through it is consistent: market participants currently appear to focus on **execution capabilities and operational deployment**, not intelligence alone. Owning the compute, the infrastructure or the workflow matters more than owning a model.

From silicon to execution: companies with different levels of exposure to AI, and those facing greater challenges



(1) Data as of 8 September 2026. Illustrative share price data, not indicative of future performance.

ILLUSTRATIVE EXAMPLES
CASE STUDIES



NVIDIA
▼

Nvidia is a significant participant in the AI infrastructure ecosystem due to its role in supplying processors used in the training and deployment of AI systems. Its position rests on being the primary supplier of the processors that both train and run agentic AI systems, at a moment when demand for that capacity is running structurally ahead of supply.

The evidence for this is not simply share-price performance but order backlogs and long-term supply agreements that confirm demand is already observed, not merely projected. Nvidia's growing footprint in AI infrastructure investment more broadly (including its recent acquisition of Hugging Face, the open-model hosting platform) suggests the company is positioning itself across more of the value chain than compute alone, reinforcing rather than diversifying away from its core exposure to the agentic AI build-out.



Prysmian

SCHNEIDER ELECTRIC
& PRYSMIAN
▼

Schneider Electric and Prysmian provide examples of companies whose products are used in data-centre construction and electrical infrastructure projects.

Neither company has really changed its product. Schneider still sells the same kind of electrical distribution equipment it has sold for two decades. Prysmian still sells cable. What has changed is the scale of data-centre construction, and with it, the volume of electrical infrastructure that construction now requires. That is almost entirely demand pulled in from a sector neither company operates in directly. It is a useful reminder: some of the most durable exposure to a technology cycle sits one or two steps removed from the technology itself, in the unglamorous infrastructure built to support it.



REVERSO
▼

Reverso illustrates the potential challenges faced by specialised software solutions when broader AI-enabled platforms incorporate similar functionalities.

It was first squeezed by DeepL, an early and genuinely specialised machine-translation product, and has since been made close to redundant by general-purpose large language models, which now translate competently as one capability among dozens rather than as a dedicated product a user has to seek out. The lesson generalises: a business built around a single narrow task is vulnerable not only to a better version of that same tool, but to a broader agent that happens to perform the task well as a side effect of doing something else entirely.

ILLUSTRATIVE EXAMPLES

CASE STUDIES (CONTINUED)



REPLIT



Replit is presented as an example of an application-layer platform that integrates multiple AI capabilities into an operational workflow.

Earlier coding tools, GitHub Copilot among them, offered line-by-line suggestions to a developer who stayed in control. Replit's agent goes further. It takes a business idea and turns it directly into a working, deployed product. It orchestrates every stage itself: selecting the model, generating the code, writing and running tests, configuring the environment, executing the program and deploying it. It hands back an operational result, not a code snippet.

The economics matter as much as the product. Nvidia and other chipmakers generate the raw intelligence. Replit supplies the application layer that turns that intelligence into something usable, even for firms with no in-house development team. Every drop in the cost of that intelligence, driven by lower-cost or open models such as DeepSeek and Qwen, expands Replit's addressable market. It lowers the bar for which businesses can afford custom software at all. The real advantage in this layer sits in orchestration and deployment, not in the underlying model. That is why we classify Replit as an orchestrator-adjacent potential beneficiary, not just another AI-branded coding tool.




ADOBE



Adobe illustrates why the “To Be Determined” category exists. The future positioning of established software vendors will depend on strategic, technological and competitive factors that remain uncertain at this stage. Through much of the period to mid-August 2026, Adobe traded as a clear challenged business. The market was pricing in the risk that agentic tools would erode its core creative-software franchise.

Yet Adobe generates on the order of \$10 billion a year in free cash flow. A business generating that kind of cash has real resources. Within two to three years, it can fund a genuine reinvention around agents, rather than simply watch its interface get replicated. Whether Adobe's management actually deploys that cash toward reinvention or instead manages the business for cash extraction as the core franchise erodes, is the real question. It is what separates the eventual beneficiaries from the eventual challenged businesses in this category. And it is a question best answered by ongoing diligence, not a single snapshot classification.



TO CONCLUDE



- 01 The agentic AI cycle is still young, but a pattern is already visible. Value is not spreading evenly across the technology sector. It is concentrating in whoever executes the work, not whoever merely provides the intelligence behind it.
- 02 That is true at every layer of the chain. Enablers win because compute, not demand, is the constraint. Orchestrators win because they own the workflow, not just the model. Even outside the technology sector, some companies show that durable exposure can sit one or two steps removed from AI itself, in the infrastructure quietly built to support it.
- 03 The challenged businesses side of the ledger tells the same story in reverse. Reverso, Chegg and similar tools were not out-competed by a better version of themselves. They were made redundant by agents that perform their task as one capability among many. A narrow, single-purpose product is the most exposed position in this cycle, whatever sector it sits in.
- 04 Between the two sits a genuine grey zone. Adobe is the clearest example. Cash-rich incumbents are not protected by their cash alone. They are protected by what they choose to do with it. Reinvention and managed decline can look identical for a year or two. Only time will tell them apart.
- 05 The broader lesson, for investors watching this cycle from the outside, is to look past the "AI" label itself. The potential beneficiaries here are rarely the companies branding themselves as AI businesses. They are the companies solving a real execution problem, wherever that problem happens to sit.

DISCLAIMER

This document has been prepared exclusively for informational and educational purposes. It is intended to illustrate certain technological and economic developments relating to artificial intelligence and associated value chains.

References to specific companies, sectors, technologies or market segments are provided solely as examples to support the analysis and discussion contained herein. Such references do not constitute investment research, investment advice, a recommendation or a solicitation to buy, sell or hold any financial instrument. References to listed companies are provided solely to illustrate the structure of the AI value chain and should not be construed as views regarding the value, suitability or expected performance of any security.

Any views expressed in this document reflect opinions as of the date of publication only and may change without notice. No representation or warranty is made as to the accuracy, completeness or future validity of the information contained herein.

Any discussion of market developments, industry trends or business models should not be interpreted as an indication of future performance. Past performance is not a reliable indicator of future results.

Investors should conduct their own analysis and seek independent professional advice before making any investment decision.

Norman K. does not undertake to update this document and accepts no liability for any investment decision made on the basis of the information contained herein.

LOCATIONS

