

sifted / Reports



The rise of AI agents

And how they will reshape work

Sponsored by **box**

The rise of AI agents

And how they will reshape work

AI agents are poised to reshape the way we use computers — and it's a big deal.

Beneath the agent sits something computer scientists have pursued for decades: a machine you can direct simply by describing what you want. Not by learning its language, but by using your own.

In this report, we examine the rise of agents as the next step in AI, speaking to both the companies developing these systems and the startups and scaleups integrating them into their own operations.

Get in touch

**Interested in partnering with
Sifted on a report?**

Contact us via
commercial@sifted.eu.

**Looking for vital market
intelligence?**

Visit sifted.eu/sifted-pro to learn
how Sifted Pro can help.

Contents

3

Foreword

4

How far we've come

5

CHAPTER I

Beyond the chatbot

8

CHAPTER II

How companies are using agents

12

CHAPTER III

Safely deploying AI agents

17

CHAPTER IV

Work redefined — jobs, productivity
and competitive advantage

21

Conclusion

Foreword from

In the past 12 months, agent adoption among growing businesses has shifted from pilots to deployments in customer support, finance, and operations — and the companies that moved early are already seeing the gap widen.

An AI agent can take action toward a goal, carry work forward across steps and systems, and help lean teams complete work that once took hours of manual effort. Saving time, cutting manual work, improving visibility, scaling operations, and reducing risk aren't aspirational goals. They're what lean teams actually get when agents handle the work. Agents can process documents, route approvals, answer customer queries, and extract information from contracts and invoices, all without adding headcount or budget you don't have.

But there's a catch. Many organisations don't have an AI problem as much as they have a content problem. "The documents, contracts, emails, and institutional knowledge that define how your business actually works are often scattered, ungoverned, and inaccessible to AI systems — making the content layer the real bottleneck." As Aaron Levie, CEO of Box puts it, "Context is king for AI agents".

Most of the information that explains how your business really works doesn't live in neat rows and columns. It lives in contracts, customer emails, onboarding documents, meeting notes, policy files, invoices, and years of institutional knowledge scattered across drives, inboxes, and point solutions. AI models trained on public data don't know your supplier terms, approval workflows, or customer history. If agents can't access and understand that information securely, they can't reliably act on it. Levie argues, 'Agents getting the right context to do their work will be the dominant IT challenge over the next decade'.

This is where the next phase of AI adoption will be won or lost. Generic AI can answer generic questions. Useful agents need trusted business context. They need a centralised content layer: a secure, governed foundation

where the right documents, permissions, metadata, and workflows come together so people and agents can work from the same source of truth.

For Box, this is the promise of connecting business content to AI. When you can securely connect your content to agents, teams, and apps, you unlock trusted insights, automate workflows, and move faster with governed work.

This report examines how European startups and scaleups are building and deploying AI agents — cutting through the hype to ask what actually separates agents from copilots and automation, where they're delivering real operational impact, and how leaders should approach governance as agents gain access to sensitive systems.

Those questions are especially urgent for scaling companies. Large enterprises carry decades of technical debt and complex governance models. You may have less infrastructure — but you also have more freedom to design workflows around AI from the start, and less room for costly mistakes. A 200-person company that centralises content, builds clear guardrails, and treats agents as active participants in work can move with a speed that bigger competitors may struggle to match.

Security and governance will determine whether that speed becomes an advantage or a liability. AI must operate on authorised content, with permissions, auditability, and compliance intact. Without that foundation, teams turn to unsanctioned tools, creating shadow AI risks and exposing sensitive data.

The rise of AI agents is not simply a story about automation replacing tasks. It's a story about organisational knowledge becoming usable at scale.

The companies best positioned for this shift will treat content as strategic infrastructure. For scaling businesses, that's the opportunity ahead: to build AI-ready operations now, while the rules of competition are still being written.

How far we've come

We've come a long way since spreadsheets. In the 1980s there was VisiCalc, then Lotus 1-2-3, then Excel and suddenly anyone with a PC could ditch the paper ledger and run financial models in minutes. For the first time, a non-programmer could automate their own work. Things stepped up in the 1990s when ERP and CRM systems wired entire companies together.

Machine learning arrived in the 2010s and suddenly software could learn, spot patterns, make predictions, improve over time. IBM Watson beat "Jeopardy" champions and Siri answered questions. Chatbots started handling customer service and algorithms began making recommendations.

Deterministic tools made way for adaptive ones. New AI tech could handle ambiguity, generalise from data, and get better with use. Cloud computing and giant datasets poured fuel on the fire, embedding AI into tools people already used, until even Excel started suggesting formulas and spotting anomalies.

The knock-on effect was that repetitive tasks got automated, and human roles shifted toward strategy, oversight, and judgment. Data started driving decisions that gut instinct used to make.

But it came with new headaches — algorithmic bias, data quality nightmares, workforce anxiety. The technology moved fast. The governance frameworks to match it were a little slower.

ENTER AI AGENTS

The latest tech wave? Autonomous agents that, rather than sitting passively waiting for a prompt, pursue goals, break them into steps, use tools, browse the web, write and run code, and loop back when something doesn't work.

Early glimpses arrived with systems like AutoGPT and BabyAGI, which were clunky but eye-opening

proofs of concept. Then came more polished agent frameworks such as LangChain, CrewAI, and purpose-built enterprise agents capable of handling multi-step workflows that would previously have required a human. Now, the tools to set up these bots have never been more accessible: suddenly, non-technical founders are having a go at coding (and loving the thrill of it all).

OpenAI CEO Sam Altman has called AI agents "the next giant breakthrough." We're speeding into the age of bots making appointments, writing emails and doing our shopping. Some tools can already perform these tasks quite well (albeit with supervision).

Investors are pouring money into AI agent makers. Startups building these clever bots in Europe collectively raised €2.2bn across 110 deals in 2024; in 2025 they raised €6.2bn across 429 deals. So far in 2026, investors have backed 123 AI agent startups, with total funding exceeding €2bn.

A handful of AI agent makers are already well known, including the likes of AI video company Synthesia, voice AI company ElevenLabs and vibe coding platform Lovable.

Of course, AI agents face hurdles. While it's true that these clever bots are increasingly getting deployed inside companies, many tech and personnel challenges remain — from concerns about cybersecurity and governance to resistance from employees who fear AI will take their jobs.

In this report, we examine the rise of agents as the next step in AI, speaking to both the companies developing these systems and the startups and scaleups integrating them into their own operations.

A key question we'll explore: if AI agents can now act autonomously and make sense of information that once resisted automation, how might this change how smaller, scaling organisations operate, how they compete and how they grow?

CHAPTER I

Beyond the chatbot

What AI agents actually are



It's worth asking at the outset: do AI agents actually exist? Or are companies simply building sophisticated AI workflows and calling them "agents"?

A true agent is more than a few large language models stitched together. Agents execute multiple autonomous steps, in no prescribed order, and return results that aren't predetermined. The underlying idea is simple: anything a person does on a computer could, in principle, be done by a computer.

"A copilot suggests; an agent executes," says Heather Ceylan, senior vice president and chief information security officer at Box, which has developed AI agents that can pull insights from companies' data. "Once a system can access sensitive content, invoke workflows, or take action on behalf of a user or service account, it moves into agent territory."

A chatbot that produces text and stops is not an agent, however sophisticated the output. A true agent takes initiative and uses feedback from its environment to shape future behaviour.

66

Once a system can access sensitive content, invoke workflows, or take action on behalf of a user or service account, it moves into agent territory."

Heather Ceylan, senior vice president and chief information security officer, Box

Using Ceylan's definition, a lot of nifty AI tools stop short of being agents. Take the inbox sorter created by London AI company Fyxxer. It's almost like an executive assistant, drafting your emails, organising your inbox, and taking meeting notes. But it doesn't go beyond suggesting actions. The agent, on the other hand, is a bot you can delegate to and trust that it'll see a task through, whether that's writing reports, editing files, or sending messages.

"The agents have a memory and can self-improve over time, which makes them analogous to a human," says Leo Rosales, cofounder of Ralio, which is making AI agents to handle payments.

And even if LLMs don't improve much in the future, "we still have years leveraging all that we have right now," he adds. "There's so much potential here."

66

The agents have a memory and can self-improve over time, which makes them analogous to a human."

Leo Rosales, cofounder and CTO, Ralio



FXYER FOUNDERS

A MIXED TRACK RECORD

Agents fall into two broad categories: the “horizontal” kind, which are broad and cross-functional, designed to work across an organisation. “Vertical” agents are domain-specific, built to handle a defined set of tasks within a particular field.

Agents can be reliable. They can also go badly wrong.

For example, Anthropic, maker of the Claude language model, let its agent run a vending machine for a month to see whether it could turn a profit. The bot lost money by giving away items for free.

Other tales of mishap include food chain Taco Bell saying it would rethink its AI drive throughs after a customer seemingly crashed the system by ordering 18k waters. AI app building site Replit accidentally deleted a company’s entire database.

“The most difficult workflows involve multi-agent pipelines where a single injected prompt can cascade across an entire automated workflow,” Ceylan explains. “Agents operating in automated workflows can delete, move, or modify content across large volumes of files before a human notices — and many of these actions are irreversible.”

Because of examples like those above, there’s a fair bit of caution when it comes to agent deployment. Fewer than half of their AI agents make it beyond the proof-of-concept stage, according to a recent poll of business leaders conducted by machine learning company Dataiku.

A much-cited 2025 MIT research paper suggested 95% of generative AI pilots fail to deliver a return on investment.

Yet adoption is pressing ahead regardless: Dataiku also reports that 86% of business leaders say their companies now rely on agents in daily operations, with nearly half reporting that dozens of core processes depend on them (a side note: only one in five always requires agents to explain their reasoning — and a striking 95% admit they could not fully trace AI decisions end-to-end if asked to do so by regulators, issues we’ll examine further in chapter 3).

‘A JOYSTICK FOR THE WORLD’

AI agents are poised to reshape the way we use computers — and it’s a big deal. The key difference from previous tech waves is this: for the first time, your entire computer can respond to plain English.

Beneath the agent sits something computer scientists have pursued for decades: a machine you can direct simply by describing what you want. Not by learning its language, but by using your own.


“

There’s no 9-5 work anymore.”

Thierry D’Hers, field CTO, Pigment

Need to organise your files? Just say so. Want to run a data analysis or spin up a website for your business? Type what you need and let the agent figure out the rest. The computer, at last, does the translating. “AI agents make language into a joystick for the world,” as one New York Times article memorably put it.

“There’s no 9-5 work anymore,” says Thierry D’Hers, field CTO at French unicorn Pigment, which has developed agents to help CFOs manage complex finance decisions. “Because you can go to dinner or have a drink with your friends and your agents can continue on working. It’s exciting.”



CHAPTER II

How companies are using agents

Finding leads, writing code,
scouring documents — even
ordering toilet seats



Right now in Stockholm, a café is being run entirely by an AI agent called Mona. The experiment comes from Andon Labs, a company testing whether AI agents can operate real-world businesses. Andon has form here: another agent, Luna, has been running a store in San Francisco with the help of hired humans.

Results so far have been mixed. In the Swedish cafe, the human assistant has a wall display which showcases some of Mona's unnecessary purchases, including 10 litres of cooking oil and 15 kilograms of canned tomatoes.

Luna has made a few booboos too, ordering 1,000 toilet seat covers for the employee bathroom, then listing them as merchandise, according to the New York Times. Prices in the store are high, even by San Francisco standards — one customer reported being charged \$10 for a bar of soap.

These entertaining stunts are hardly representative of where AI agents stand today. Generally speaking, companies have more modest ambitions for the technology — at least for now.

In this chapter, we examine how startups and scaleups are deploying AI agents across functions such as operations, customer support, finance, and content workflows, and where adoption is shifting from experimentation to everyday use.

BABY STEPS

For most companies today, it's baby steps. At this point, few AI agents are "transforming how work gets done," according to a 2025 PwC report. "Many employees are using agentic features built into enterprise apps to speed up routine tasks — surfacing insights, updating records, answering questions. It's a meaningful boost in productivity, but it stops short of transformation."

For example, agents are being set up to help companies with triage: reading tickets (i.e records of customer feedback or technical bugs), assigning priority, drafting replies, and pulling relevant order information from CRMs.

At other companies, voice agents handle early-stage recruitment — confirming applicant interest and collecting basic information before a human gets involved.

German startup Voize raised €43m from investors in 2025 to develop a voice AI agent to help nurses take bedside notes. The idea formed from CEO Fabio Schmidberger's direct observation of nurses in care homes "writing notes on the backs of their hands, then typing them up later." Existing note-taking tools were built for doctors, he says — but nurses work differently. "They're constantly moving and can create 30 to 100 patient records per shift. It's overwhelming."

Burnout across healthcare is severe, with documentation among the top frustrations driving staff out of the profession. Founded in 2020, Voize is now used in more than 1,000 facilities. Schmidberger compares the technology to a screwdriver. "When push comes to shove and you desperately need to get something done, you're not trying out a fancy new electric drill — you want to get the job done with what you have."

London-based PolyAI has created powerful voice AI technology that companies are deploying as receptionists and customer support agents. Today's tech far exceeds what voice assistants Siri and Amazon Echo can do, says Nikola Mrkšić, cofounder and CEO. But while AI now sounds semi-indistinguishable from humans, it still has trouble understanding us. "ChatGPT will misunderstand one word and go on a comical side-quest," he says.



© ANDON MARKET, SAN FRANCISCO

Agents are also helping salespeople. One sales rep, who spoke to Sifted on background, described how her company has used agents to tackle a persistent bottleneck. "Getting enough information from a client to build a truly targeted lead list used to take hours of back-and-forth," she says.

66

It asks the right questions, follows up, and adapts based on what the client gives me. People don't realise it's AI because it's not generic."

Her solution was to build an agent, using developer tools from German company n8n, that handles the entire client intake conversation over email. "It asks the right questions, follows up, and adapts based on what the client gives me. People don't realise it's AI because it's not generic; it's trained on how I actually communicate."

The sales rep reviews the final package and makes the call on whether to move forward and reach out to a prospect. "That one agent replaced what used to be my biggest time sink," she says.

'97% OF OUR CODE IS WRITTEN BY AI'

AI coding tools developed by Anthropic, OpenAI and Google, meanwhile, are helping companies move at speeds that would have been unimaginable just a few years ago. The days of AI assisting humans with coding have become humans assisting the AI with coding.

"Over 97% of our code is now written by AI," says Bobby Healy, CEO of Irish drone logistics company Manna Air Delivery. "That's a pretty staggering number." The remaining 3% "will never be," he adds — "flight control, safety-critical systems, parachute detonator code. You wouldn't let AI anywhere near that."



Document workflows are being transformed too. A mid-market professional services firm using Box to manage client contracts redesigned its review process entirely, according to Sam Crisp, Box's EMEA team director of solutions engineering.

"Instead of a paralegal manually extracting key dates, obligations, and counterparty details from each contract, an agent now processes incoming documents, extracts structured data, flags GDPR-relevant clauses, and routes for human sign-off," he explains. "The paralegal's role shifted from extraction to review and exception handling."

Ralio, a company we met in the previous chapter, is building agents to handle payments. Agents are already zipping around the company doing various tasks. "We built an agent on top of [AI tool] OpenClaw," says cofounder Leo Rosales. "We'll send our agents into Notion or Google Workspace to retrieve information. For example, what's our VAT number again? How much are we paying for our software licences again? The agent will help us here."

"We can also take a screenshot of someone on LinkedIn and direct the agent to load this person into our CRM." The company is "dogfooding" its own payment agents too. "We've used our agent to pay invoices and salaries."

It might be a while before other companies follow suit. "We have a backup on all our internal documents. With payments, you cannot do backup," Rosales says. "Mass adoption for payment agents will take a while but it's going to be a big unlock when it comes."

'QUICK WINS'

One thing agents are really good at already is spinning up prototypes. "We were meeting with a customer recently and they were describing all the problems they had and wanted to know what we could do to help," says Thierry D'Hers, the field CTO at Pigment, who we met in the previous chapter. "We could get them a solution in almost real time. It didn't take three weeks to build, it was an hour."

At consultancy Yonder, the strategy for deploying agents begins with "quick wins," says partner Rich Cownley, a Box customer. "There's a lot of knowledge sitting in old decks, case studies, and deliverables that was effectively invisible before. At the same time, highly capable people spend too much time on repeatable tasks — synthesis, summarisation, transcription, first drafts, analysis."

Yonder now uses Box agents to fetch information, summarise meetings and spin up research. Irish aviation software company Cloudcards also uses Box agents to do the "heavy lifting" while leaving the "nuanced judgement" for humans, says cofounder Barry Fitzgerald. "Agents handle the extraction of commercial and technical terms across dozens of documents, while our human experts focus on strategic oversight, complex negotiations, and the final clean draft."

But fully agent-run workflows remain uncommon. Agents are typically handling one or two steps in a larger process, with humans managing the rest — and a fair amount of oversight is still required. "The hardest workflows are the ones that depend on tacit organisational context," says Ravi Malick, global chief information officer at Box. "These are workflows where the right answer isn't in any document — it lives in the heads of experienced people."

"Agents struggle when the knowledge required to complete a task has never been made explicit. For small-to-medium businesses, this is often compounded by less formally documented processes than you'd find in large enterprises."

Across the companies Sifted spoke to, one theme emerged consistently: a key limiting factor for agents is the quality and accessibility of the content the agent is working with.

Data quality, Malick adds, is consistently underestimated. "Agent performance is a direct reflection of how well-structured and accessible your underlying data is. That's about information architecture — not AI, agents, or tech."

“

Agent performance is a direct reflection of how well-structured your underlying data is.”

Ravi Malick, global chief information officer, Box

Cownley echoes this point. "Agents aggregate patterns well, but they can miss outliers, misread tone, or flatten nuance — especially across markets and languages."

The hardest workflows to hand off, says Crisp, are those requiring judgment across unstructured, ambiguous inputs. "Contract review workflows where edge cases require legal interpretation, or multi-document synthesis where source quality is inconsistent."

European businesses, he notes, consistently flag the same need: agents that can "degrade gracefully if they can't answer" and handle edge cases safely — particularly where regulatory compliance is at stake.

"Agents are only as good as the context they're given," says Federico Scolari, Sifted's AI lead, who has created several agents. "Some models are already incredibly smart if unleashed on information that is neat, precise, complete and that leaves little room for interpretation."


"But when instructed imprecisely, agents are fairly bad at improvising."

So if you ask an agent to surf a sales data provider like Apollo, score leads based on a few signals and call top targets it will almost always call someone; it might not call the right person if it doesn't know who your ICP [ideal customer profile] is precisely, Scolari adds. "They're like interns in cognition and seniors in execution."

CHAPTER III

Safely deploying AI agents

**You need guardrails for trust
and control. But how far should
you go?**



AI agents introduce new, sometimes knotty questions. If an AI agent books a flight, sends an email, and updates a CRM entry on your behalf, who's accountable when something goes wrong? How do you audit a decision that emerged from a chain of autonomous steps?

These concerns are not hypothetical. According to a recent poll by Dataiku, European business leaders are among the most cautious globally when it comes to deploying AI agents: 63% of data leaders in Germany fear AI could harm customers, and 69% fear harm to employees — the highest rates recorded anywhere in the world.

Yet attitudes are far from uniform. Some leaders are giving agents ample rein to take action on their behalf, while others see the technology as a powerful tool that still requires human intelligence to guide it — arguing that automation will not displace workers as quickly as current hype suggests.

It is precisely this tension that this chapter sets out to examine. We look at the governance, security, and oversight mechanisms required to deploy AI agents responsibly.

Central to that examination is the risk of under-governance: over-permissioned agents, data leakage, unreliable outputs, the automation of already-flawed processes, and the accountability gaps that emerge when systems make consequential mistakes without a clear human owner.

But guardrails need not mean gridlock. We explore how well-designed controls — embedded early into systems and workflows rather than bolted on afterwards — can strengthen trust, protect sensitive information, and support safe, sustainable adoption as companies scale, without importing the kind of enterprise-level bureaucracy that kills momentum in smaller organisations.

AGENT CONFESSIONS

Several companies shared with Sifted the agent failures they're encountering in practice — and the oversight measures they're putting in place to keep those systems in check.

“

Tiny misses snowball into the wrong rabbit hole.”

Jan Philipp Harries, cofounder and CEO, ellamind

Matt Wilson, founder of London-based Jack & Jill, which uses agents to help jobseekers land interviews and companies find hires, says agents often need “supervisor” agents to make sure they’re not getting carried away.

“A good example of where we need an extra eye is when we’re sourcing candidates: our agent Jill might discover a promising candidate but identify red flags in their profile after doing due diligence,” he says. “Undeterred by this, the agent will still want to present these ‘promising’ hires to the client.”

Jan Philipp Harries, cofounder of Bremen-based ellamind, which helps companies evaluate chatbot answers for accuracy, agrees on the importance of additional oversight.

“With multi-step plans and chains of tool calls, tiny misses snowball into the wrong rabbit hole,” he says. “That’s why we built our evaluation platform that highlights and classifies these failure modes so teams can course-correct fast.”





Another long-standing issue concerns the quality of the data models are trained on, and combatting the risks of “hallucinations”, known in common parlance as mistakes.

“LLMs that agents use have actually gotten a little worse over time,” says Daniel Keinrath, cofounder of Vienna-based Fonio, an AI telephone assistant that fields calls for companies. “We worked with GPT-4.1 mini for a long time, but since the release of GPT-5, latency — as in, how long the wait for results is — has increased 4x. We’ve found a way to fix it, but it requires a lot of workarounds.”

66

LLMs that agents use have actually gotten a little worse over time.”

Daniel Keinrath, cofounder and CEO, Fonio

Hallucinations remain a major issue, he says. “The only effective way to reduce them is with a really strong system prompt, a solid user prompt, and a small context window.”

Emma Burrows, cofounder and chief technology officer of London-based Portia AI, a platform which helps businesses build their own AI agents, tells Sifted that AI agents are quite good at completing a single, specific task — but giving them multiple tasks at once can disrupt their consistency.

“A single use case agent — one that processes refunds for customers, for example — is pretty well understood,” she says.

“Where the technology is really only getting to is a multi-agent system where you can just feed a whole bunch of tools and workflows to the same system and have it produce great results in a whole set of different use cases.”

Some companies are beginning to grapple with the risk of AI agents wrongfully impersonating real consumers, as traders begin using AI agents to buy and sell assets on their behalf.

“You have very smart AI agents that can do this right now and it’s already challenging to detect,” says Monica Eaton, CEO of Chargebacks911, which develops fraud detection tools for companies. “The whole framework in payments has been designed around human behaviour. We’re using AI to help assess the intent behind purchases. If we can’t adjust rapidly, there’s going to be a massive fallout.”



NO BROAD ACCESS

Setting boundaries for agents is a good idea. "The principle applied is that agents should never have broad standing access just because they are useful," says Box's Ceylan. "Permissions should be scoped to a task, time-bounded where possible, and fully logged."

66

Agents should never have broad standing access. Permissions should be scoped to a task, time-bounded where possible, and fully logged."

Heather Ceylan, senior vice president and chief information security officer, Box

"Companies deploying AI agents today often have limited security event logging, alert dashboards, and admin controls specific to agentic behaviour — which is precisely the gap that needs to be closed," she adds.

When Michael McNeil talks about the state of enterprise AI, he sounds like a seasoned operator who's watched too many promising technologies fail at the implementation stage. "You'd be smoking something funny if you passed the output of an agent straight over to a customer," says McNeil. "That's just not where we are."

The thesis behind his new London-based startup, called Deliverance, is that there's a trust and ROI problem with AI agents specifically. "We're not short of models. We're not short of demos that look impressive in a boardroom," he says. "What enterprises are missing is something that gets them to a yes."

The focus is on functions that McNeil describes as ripe for agentic workflows: sales, legal, and procurement. The company, now around 25 people, is already pointing to £4m in ARR as early validation that the approach is landing.

The core concept behind Deliverance is what McNeil calls "an organisation with a brain", a system where work isn't just processed, but understood, planned, and executed by agents trained for specific contexts. Those agents aren't dropped into workflows cold. They're trained incrementally, put through small challenges, and assessed before they're given real responsibility.

The analogy McNeil uses is a new sales hire. "You don't send them out the door on day one. You train them, you watch them, you coach them. Then you let them rock on." The same logic, he insists, has to apply to agents. "And just like a good salesperson, they put their hand up when they're not sure."

That hand-raising is built into the architecture. For every agentic workflow Deliverance deploys, there's a cycle of plan, execute, evaluate — with every output scored on a confidence level. The system flags when it's uncertain.

Central to this is what McNeil describes as a separation between executor and evaluator agents. The executor does the work. The evaluator critiques it independently — a kind of built-in second opinion that catches drift before it becomes error. Crucially, the tolerance can be tuned. "You can dial up or down how much latitude it's allowed," he explains. "You're toning its ability to bullshit to the minimum."


66

It's not magic. It's a governed system."

Michael McNeil, founder, Deliverance.ai

This architecture also addresses one of his more colourful warnings about agent behaviour. "If you give an agent a thousand ways to do something, it'll pick the one way you didn't want." The governed approach Deliverance uses is designed to constrain that range of motion — not to hobble the agent, but to make its decisions auditable and predictable.

"Hopefully," he says, "it gives organisations the comfort they need. It's not magic. It's a governed system."



So a fair amount of agent babysitting is still required — which is why companies like London-based Trent AI, for example, are finding investment. The company raised a \$13m seed round in April for its product, which slips a leash on AI agents. “Organisations are deploying AI agents and autonomous workflows faster than their security can adapt, and most development teams using these agents and workflows have no security framework designed for their systems,” CEO Eno Thereska says.

Willem Delbare, CEO of Belgian unicorn Aikido, tells Sifted that AI agents will spur huge business for cybersecurity companies: software is now being built at lightning speed, AI-generated code is the norm and autonomous agents are rewriting systems faster than humans can review. “That adds up to a booming market,” he said. Swedish vibe coding hotshot Lovable is partnering with Aikido, which deploys AI agents to simulate real-world attacks on live Lovable applications.

James Neville, CEO and founder of London-based fintech Yaspa, has a simple recommendation for ensuring agents don’t get carried away: limit what they have to spend. “Set up a wallet with a hundred pounds and it’s not the end of the world if something goes wrong.”

In European organisations, there is also heightened awareness of strict data protection rules (GDPR): employees want to know whether their data or their clients’ data is being processed appropriately. Under-constraining an AI agent may also incur regulatory liability under the new EU AI Act.

“The approach is to define clear boundaries, document controls, and make sure humans retain accountability for high-impact decisions,” says Crisp. “For businesses without dedicated security teams, choosing a platform with built-in guardrails and transparent logging is essential.”

“

Organisations are deploying AI agents and autonomous workflows faster than their security can adapt.”

Eno Thereska, cofounder and CEO, Trent AI

At the same time: watch you don’t overdo the guardrails. “If every meaningful action requires so many checks that the workflow becomes slower than the original process, the value collapses,” says Malick.


“The answer is not fewer guardrails — it is smarter guardrails. Tighter controls where the risk is high; lighter controls where the task is low-risk and easily reversible.”



CHAPTER IV

Work redefined — jobs, productivity and competitive advantage

**A software update needs
a cultural shift**



How might AI agents reshape roles and organisational structures?

Already, jobs are changing. The restructuring is most visible in roles that previously spent significant time on manual document processing, compliance checking, or client reporting — tasks that agents can now handle with appropriate oversight, says Box's Crisp. This agent-speediness leads to "repeatable ROI", he adds.

But agents won't reach their full potential without human buy-in and companies changing to accommodate them. "Truly autonomous agents are still relatively rare," says Eugi Bartolo, chief technology officer at AI search agency performance.io. "The biggest barriers are not technical but organisational: people, communication and alignment."

European AI companies such as Mistral and H Company have understood this and are dispatching "forward-deployed engineers" — an idea borrowed from American tech company Palantir, which involves sending highly skilled staff into businesses to solve problems. Mistral now has a team deployed in the London offices of HSBC to create AI tools that allow its 200k employees to automate repeatable tasks like compliance checks, the bank's CIO, Stuart Riley, told Forbes recently.

66

The biggest barriers are not technical but organisational: people, communication, and alignment."

Eugi Bartolo, chief technology officer, performance.io

In this chapter, we examine how productivity may be redefined in agent-enabled organisations, where output and coordination no longer scale linearly with headcount. We consider how these shifts may alter competitive dynamics, enabling smaller organisations to operate with the leverage of much larger ones, and what it means for slower adopters.

UPSKILLING NEEDS

Will AI lead to job losses? It's already happening. Big companies such as Pinterest, Block and others have cut thousands of workers in recent months, citing efficiencies enabled by AI. The number of roles eliminated and not replaced was highest among early-career employees, or those with no previous experience.

Anxiety that AI will take your job is also one of the chief reasons many tech workers are opting for crazier hours: so-called 996 culture, or working 9am-9pm, six days a week, has become a badge of honour in a growing number of European companies.

"There's really a feeling that AI can break into any business category: it's so exciting but it's also driving fear," says Malte Landwehr, chief product and marketing officer at Berlin-based Peec AI. "I think about the people who used to run phone switchboards in offices who lost their jobs to technology. Or how one day the car replaced the horse and carriage. I don't want to be a horse."

Some argue the AI "jobpocalypse" scenario has been overblown. "The idea that jobs go away as a result of AI is going to be completely proven wrong," Box CEO Aaron Levie told CNBC in June. He argued that AI actually creates new categories of jobs around looking after AI agents.





But to avoid irrelevance in the future, companies will need to think carefully about where their competitive advantage lies in a world of powerful AI platforms.

A good place to start? Giving workers an AI skills top-up. Employees who have a native grasp of AI tools will likely outperform employees who insist on doing the grunt work manually.

"Companies are investing in AI technology their workforce isn't trained to use," says Amelia Miller, CEO of London-based AI upskilling site Ivey. "Most employees were hired long before generative AI existed, so they don't know how to use it effectively — let alone keep up with the latest tools."

“

Companies are investing in AI technology their workforce isn't trained to use.”

Amelia Miller, CEO, Ivey

Miller says that businesses are caught in an "adoption lag paradox", in which the technology keeps improving but usage isn't keeping up. Ivey's answer is a platform that pairs AI skill seekers with a network of peers and mentors, and offers bootcamps designed to sharpen AI knowhow.

SHOULD AGENTS GO ON THE ORG CHART?

Some are predicting a great flattening in companies, where AI condenses the organisation chart. Management hierarchies may not collapse immediately, but the expectation that they'll stay the same is fading.

“

I think we'll reach a point very soon where most people will wonder how they ever worked without at least one agent under their management.”

Sam Fitzroy, cofounder and CEO, Dalia

Should AI agents actually be added to the org chart? This has been a point of debate. Some say it's a silly argument — the phone or the laptop didn't make it onto the org chart, so why should an AI agent?

"I think we'll reach a point very soon where most people will wonder how they ever worked without at least one agent under their management," says Sam Fitzroy, cofounder and CEO of US recruitment platform Dalia. "But that doesn't mean a computer programme should make it onto the org chart."

But perhaps agents need their own org chart, says Jon Riede, a partner at CareerPlug, another recruitment site. "Human workers can make better decisions and be more productive if they know which agents support their work and how those agents will interact with each other," he says.

Agents will at least require new performance metrics at companies, according to IBM. Old-school metrics, designed for human work, might not capture some of the value AI agents can bring. New metrics such as “agent-to-human handoff rates” or “human overrides of agent recommendations” will be vital, the company says.

“

Human workers can make better decisions and be more productive if they know which agents support their work.”

Jon Riede, partner, CareerPlug

So, instead of just tracking productivity, outcomes from these automated decisions are what really need to be measured. These new metrics should evaluate both the efficiency and effectiveness of AI agents in creating value and driving business growth.

MEET THE AI ORCHESTRATOR

It's likely a new type of role will emerge at companies. The AI lead, the agent wrangler, the agent orchestrator: it's got many potential titles.

Tech recruiter True Search says it's seeing loads of companies hiring for one role in particular: a chief of staff or AI transformation, which is helping the organisation modernise with AI.

These new roles aren't necessarily being filled by AI-native talent, but are designed to work across the company to drive change and adoption, says Emma Pott, a partner at True Search.

She says candidates are coming from different places, including from the product and operation side. Pott reports True Search is working with a number of different scaleups for this type of role, and she's also seeing legacy businesses which are digital first but still have a lot of older tech that needs updating.

“Because of how fast the industry is moving, every six to eight weeks there is a new model from one of the labs,” adds Alex Howman, a principal at True Search. “If you're not up to date and familiar with that new model and the best new model, you will just get left behind,” he says. “That new head of AI, that quite senior individual who can lead that strategy, is probably one of the newest things we've seen.”

The companies that have welcomed agents to their office have advice for those who haven't yet. “If I was starting again today, I would adopt an ‘agentic-first’ design from the outset rather than focusing on task automation,” says Barry Fitzgerald, cofounder of Irish aviation software company Cloudcards.

“

If I was starting again today, I would adopt an ‘agentic-first’ design from the outset.”

Barry Fitzgerald, cofounder and CTO, Cloudcards

CONCLUSION

How to get the best out of AI agents

AI agents are clearly no longer a future concept. Across Europe, startups and scaleups are already deploying them to write code, process documents, support customers, qualify leads and automate routine workflows.

So far, the reality of adoption looks very different from early predictions. Rather than replacing entire companies, today's agents are extending what human teams can achieve — handling specific tasks while relying on people for judgement, oversight and accountability.

What we've heard is that the companies seeing the greatest value are not necessarily those deploying the most sophisticated agents, but those pairing them with high-quality data, clearly defined processes and appropriate safeguards.

And as this report has shown, successful adoption depends as much on company readiness as technical capability.

Poorly governed agents can amplify mistakes, create security risks and erode trust; well-governed agents can unlock gains in productivity, speed and scale.

The broader impact will extend well beyond individual processes. Agents are beginning to alter how companies are structured, how work is measured and which skills create competitive advantage.

New roles are emerging to manage and orchestrate AI systems — “AI agent wranglers”, anyone? — while employees increasingly need the ability to work with intelligent tools.

The lesson for European businesses is not that every process should be automated, nor that agents will render human expertise obsolete. The opportunity lies in combining the strengths of both.

The companies that pull ahead, we've heard, will be those that treat agents as a new layer of capability — one that extends what leaner teams can achieve without sacrificing the human judgement that complex decisions still demand. That means investing also in the governance, culture and skills needed to deploy this technology well.

The age of agentic work has arrived. The companies that navigate it well won't necessarily be the fastest to adopt; rather, they'll be the most clear-eyed about why.

Acknowledgements

EUGI BARTOLO

Chief technology officer,
performance.io

EMMA BURROWS

Cofounder and CTO, Portia AI

HEATHER CEYLAN

Senior vice president and chief
information security officer, Box

RICH COWNLEY

Partner, Yonder

SAM CRISP

EMEA team director of solutions
engineering, Box

WILLEM DELBARE

Cofounder and CEO, Aikido
Security

THIERRY D'HERS

Field CTO, Pigment

MONICA EATON

Founder and CEO, Chargebacks911

RAFIE FARUQ

CEO and cofounder, GenieAI

BARRY FITZGERALD

Cofounder and CTO, Cloudcards

SAM FITZROY

Cofounder and CEO, Dalia

JAN PHILIPP HARRIES

Cofounder and CEO, ellamind

BOBBY HEALY

Founder and CEO, Manna Air
Delivery

RICHARD HOLLINGSWORTH

Cofounder and CEO, Fyxxer AI

ALEX HOWMAN

Principal, True Search

DANIEL KEINRATH

Cofounder and CEO, Fonio

ÁLVARO LEAL

Chief of staff, Factorial

RAVI MALICK

Global chief information officer, Box

MICHAEL MCNEIL

Founder, Deliverance.ai

AMELIA MILLER

Cofounder and CEO, Ivey

NIKOLA MRKŠIĆ

Cofounder and CEO, PolyAI

JAMES NEVILLE

Founder and CEO, Yaspa

EMMA POTT

Partner, True Search

JON RIEDE

Partner, CareerPlug

LEO ROSALES

Cofounder and CTO, Ralio

FABIO SCHMIDBERGER

Cofounder and CEO, Voize

FEDERICO SCOLARI

Senior analyst and data lead, Sifted

ENO THERESKA

Cofounder and CEO, Trent AI

MATTHEW WILSON

Cofounder and CEO, Jack & Jill

This report was written by **ÉANNA KELLY** with analysis by **FEDERICO SCOLARI**, edited by **ÉANNA KELLY**, designed by **GAÉTAN NIVON** and illustrated by **GEORGE CLAYTON**.

It was produced by **ALEENA RAIKUNDALIA** and **JONATHAN SINCLAIR**.



Box (NYSE:BOX) is the leader in Intelligent Content Management. Our platform enables organizations to fuel collaboration, manage the entire content lifecycle, secure critical content, and transform business workflows with enterprise AI. Founded in 2005, Box simplifies work for leading global organizations, including JLL, Morgan Stanley, and Nationwide. Box is headquartered in Redwood City, CA, with offices across the United States, Europe, and Asia.

Visit box.com to learn more.

And visit box.org to learn more about how Box empowers nonprofits to fulfill their missions.



Sifted is the only media and intelligence brand dedicated to Europe's startup and venture capital ecosystem. We deliver the clarity and context needed to navigate a fast-changing market.

Sifted Pro unlocks the full power of our expertise — equipping users with award-winning journalism, investment-grade research and market intelligence tools powered by analyst-verified, proprietary data to uncover opportunities and make smarter decisions.

Get in touch

Interested in partnering with Sifted on a report?

Contact us via
commercial@sifted.eu.

Looking for vital market intelligence?

Visit sifted.eu/sifted-pro to learn how Sifted Pro can help.