

Update on Agentic AI

Agentic AI will impact nearly every industry over the next 4 to 5 years – there are significant opportunities for portfolio companies to enhance productivity, unlock new revenue streams and build lasting competitive advantages.

November 2025

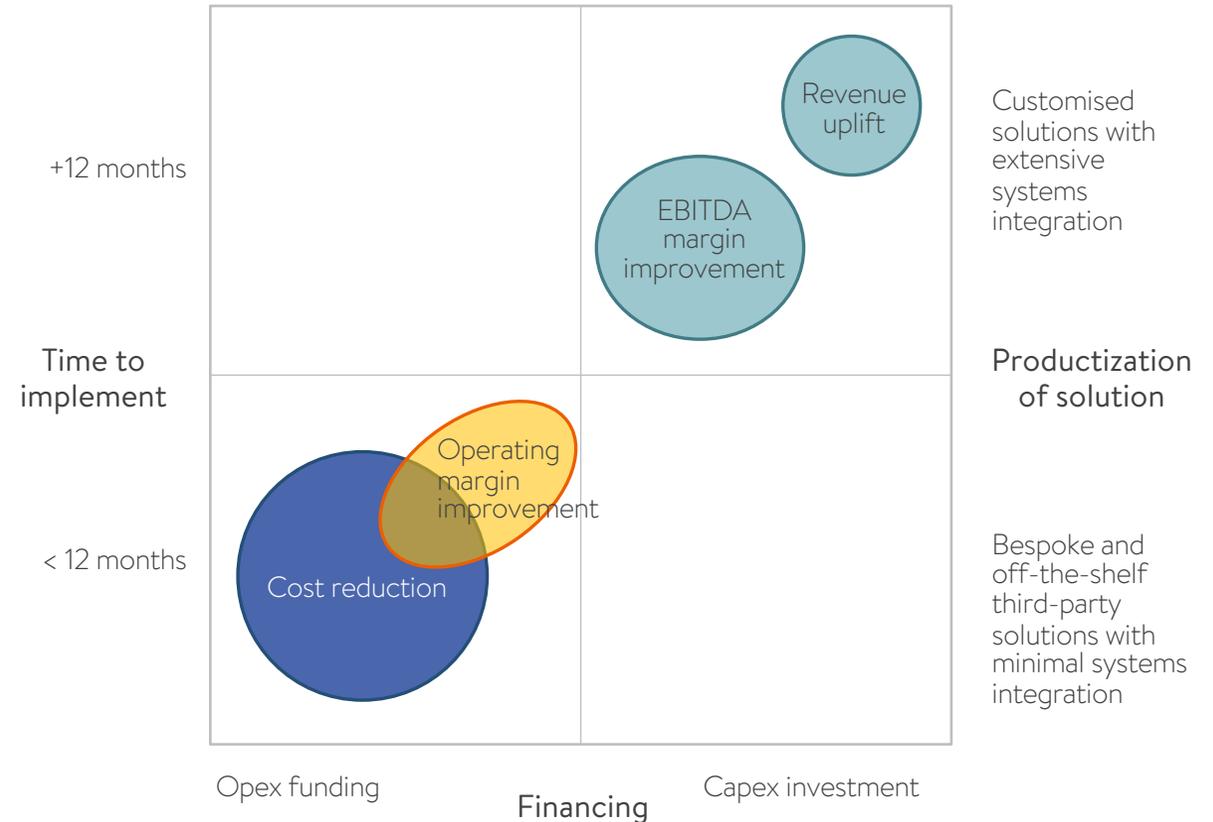
Two types of Agentic AI for investors

The most impactful use of Agentic AI for investors is in operational processes of portfolio companies that rely on repetitive human decision-making. In this context there are two types of Agentic AI:

- **Enterprise-wide:** a solution focused on EBITDA margin improvement and/or revenue uplift, addressing a business issue spanning multiple operational functions (for example, a B2B services cross-selling agent that draws from sales, marketing, customer services, finance and legal to build the scope and set the timing of the next sales pitch to an existing client, including designing a tailored pricing model to maximise Customer Lifetime Value). This type of solution requires proprietary enterprise-wide data and best-in-class enterprise data management, architecture and governance.
- **Single operational discipline:** a solution focused on cost reduction that is contained within one or two operational functions and addressing a narrow operational scope, such as the initial screening of job applicants within the HR remit. This type of Agentic AI requires proprietary data only of the operational function (which is combined with the third-party data embedded within the reasoning LLM), and can utilise data infrastructure provided by third-parties without the need to transform enterprise-wide systems.

Note that these two types of Agentic AI solutions are not mutually exclusive. Nor is one better than the other (because the value of the Agentic AI solution is the business outcome, not the scope of technical implementation) but enterprise-wide solutions are more complex, time consuming and capital intensive to implement than operational function solutions.

Landscape of Agentic AI solutions for investors



Key

- Single operational discipline.
- Multiple operational disciplines (but not enterprise-wide solution).
- Enterprise-wide Agentic AI solutions.

About the technology

Agentic AI today

Agentic AI is AI-powered software that uses sets of AI tools (ML, LLM models etc.).

The purpose of Agentic AI is to do a specific task that involves making decisions, such as adjusting software settings to deploy IT systems, choosing a book to buy for your birthday which you will actually want to read, creating customised pricing models in real-time to maximise the chance of a sale, or scheduling a £mm maintenance shutdown of a chemical processing plant.

To make decisions autonomously, Agentic AI uses natural language models (LLMs or SLMs) to reason – they are not simple ‘yes or no’ decisions that an ML model could make within well-defined boundaries, but ‘creative’ decisions involving uncertainty and interpretation.

Agentic AI a natural evolution of the technology (not a new type of AI) so organisations don’t need to retrain their data scientists or software developers to create or deploy it. But they do need a deeper understanding of operations processes and the role of automation and technical governance.

The technology is now very good in operational processes that involve decision making using a range of input data (it is accurate, reliable and fast). As a result, it is driving a revolution in process automation – you could say that the previous ‘operational version’ of Agentic AI was actually Robotic Process Automation (RPA).

Why Agentic AI is now developing so quickly

The critical aspect of Agentic AI is the ability to make the correct decision using the reasoning capability of natural language models. In the last 18 months there has been \$600bn+ investment (conservative estimate) in language models and associated infrastructure in Europe and North America, resulting in fast and reliable LLMs that have less hallucination and provide multi-modal capabilities (read and create images, video and audio alongside text). In addition, the model makers such as OpenAI, Google and Mistral, and the service providers around them (notably LangChain) have driven standardisation of processing techniques, data privacy, model safety and APIs.

This maturing of the underlying technology means that it is now easier and cheaper to build Agentic AI applications that are so fast and so reliable they can be used in many commercial use cases without human oversight.

Adoption & impact

On the demand side, there are two factors driving adoption:

- **Low barriers to adoption:** Many portfolio companies have sufficiently tech-enabled operations to implement Agentic AI. Of those that don't, many are currently investing in enterprise-wide data infrastructure for analysis, process improvement and reporting. This infrastructure is the platform for Agentic AI that spans multiple operational functions within an organisation. In parallel, many vendors are creating Agentic AI solutions that focus on single operational functions which can be implemented without an expensive transformation of enterprise-wide data infrastructure.
- **Materially significant returns for investors:** Implementing Agentic AI in operational processes, from using heavy machinery in the field to checking expense claims in a company, is now proven. It can deliver such large increases in revenue or such large reductions in cost that it actually shifts the EBITDA multiple.

Limits to adoption of Agentic AI in Europe

Under GDPR, the UK and EU have a *de facto* ban on automated decision-making in processes that handle personally identifiable information, or processes that impact human outcomes. However there are no reasons to stop Agentic AI adoption in purely internal processes of both private and public sector organisations.

Scale of impact – why the hype is not hype

Two (of many) proven use cases of Agentic AI:

T Mobile

Customer services (all activities including managing customer's phone, resetting account passwords, contracts...)

Pre agentic AI: Annual service FTE cost of \$3bn.

With agentic AI: <\$5m investment, 60% reduction in human effort, annual service FTE cost now \$500m (83% reduction).



Upstream oil exploration + downstream refining

Pre agentic AI: 2024 revenues of \$22bn.

With agentic AI: \$10m investment over 2.5 years. Immediate gain of +\$500m revenues in 12 months. On track for £1.5bn increase in revenues in 2.5 years (+6.8% revenue increase).

Case study: increasing revenue



Occidental Petroleum Corp

Oxy is an US oil business exporting in the United States and the Middle East, as well as a petrochemical manufacturer in the United States, Canada, and Chile.

In 2024 Oxy hired Patrick Bangert (the former global head of AI of Samsung Data Services and a current Advisor with eXcentius) to transform their physical operations. This wasn't tinkering at the edges: the purpose was to deploy AI at every stage of the value chain to have a *coordinated impact on revenues*.

Patrick built a team of diverse skillsets – alongside data scientists, he brought in experts in oil exploration, operating drilling machines, managing refining processes, maintaining refineries, trading oil,... he built a business and process team, not an IT team. The single objective of the team is to produce more barrels of oil per day (the most important driver of revenues for Oxy).

In an oil business there are many operational processes that impact the barrels per day. Upstream, for example, the time taken to drive a new well is a key factor, so how an operator directs a drilling machine is critical. A drilling machine generates survey data as it drives into the ground as well as using field survey data from other sources, and the role of the operator is to process both sets of information to make course corrections to avoid dead ends. Any miscalculation delays the time to bring the well on stream, even in an already proven reserve. So the team at Oxy designed an Agentic AI application that interprets survey data and drilling machine systems data in real-time, adjusting the direction and other parameters of the drilling machine as it is drilling, without the input of the human machine operator.

Downstream, the largest influence on barrels per day is the throughput capacity of the oil refinery. This is driven by the condition of the processing equipment (which wears out with use) and by the availability of complete systems (like most processing plants, a refinery will have several systems doing the same job, allowing one to be taken offline for maintenance without shutting down the entire production). Maintenance is an expensive business: replacing a single \$100k pump can easily cost \$10m per day in lost revenue.

The team built an Agentic AI system to schedule plant maintenance without human intervention. It reads plant equipment data and process sensor data in real-time, interprets the entire refinery throughput (not just the individual system), understands availability of maintenance resources (people, equipment) and predicts the net impact on revenue.

Implementing Agentic AI across the enterprise to achieve a single objective takes time: there is resistance to change, and a need to reengineer the company's IT systems into an enterprise-wide, edge-capable data infrastructure with common data architecture, low latency, high availability and high governance.

In the first 12 months (including building the team from scratch and with only limited IT transformation) the team delivered an increase in revenues of \$500m for less than \$3m investment. With a very high degree of certainty, the team are on track to increase revenues by \$1.5bn within 2.5 years, at a total cost of \$10m.

Case study: transforming EBITDA in an operational process



iLink Digital

iLink Digital are a global data and analytics IT services company, building hyper-automation infrastructure and applications for private and public sector operations.

By 2025 the central recruiting function of a Dubai based conglomerate was interviewing over 40,000 candidates per year for blue collar positions across the group.

Despite having over 50 recruitment team members, interviewing was slow and prone to errors. It was also expensive (because of the team size) so the conglomerate brought in iLink Digital to tech-enable the end-to-end interview process. iLink discovered that the initial stage of the application process was the most time consuming and costly.

In the initial stage, the first step is to screen in-bound curricula vitae (written in a myriad of visual formats and writing styles, with varying levels of detail and using a wide range of key phrases to indicate experience). The second step is to email the applicant about their success, or otherwise, of passing the initial screening.

For candidates passing the initial screening, a staff member arranges a time for a remote interview and then sets up the call in Microsoft Teams. During that interview a staff member takes notes to grade the applicant on key metrics, most of which are subjective and thus open to interpretation. In the final step, the team processes the interviews: reviewing grades, deciding who has passed, communicating with all interviewees, and writing up the notes so that successful applicants can be moved to the next phase of the application process.

To improve this process, iLink designed an Agentic AI application that reads all in-bound CVs sent into any email or any website across the group. The agent works out which job the candidate is applying for and

then screens the CV. It authors and sends rejection emails without any human intervention. For the candidates it thinks likely to succeed, the agent then communicates via email to arrange a time for a remote interview. It authors emails autonomously, including responding to applicant queries and requests for different interview times. Once the interview date is agreed, the agent schedules the video call in Microsoft Teams and sends the details to the candidate.

The first interview is held by the Agentic AI, not by a human. The agent operates a bot, which looks like a bot and tells the candidate it's a bot. It asks a set of questions of the candidate and reacts as a human would to the candidate, for example nodding in agreement to convey empathetic understanding. In real-time during the interview, the Agentic AI processes the words of the candidate's response and reads the behaviour of the candidate, for example analysing where the candidate looks at the screen or how quickly they respond to questions.

The Agentic AI then autonomously decides which candidates have passed or failed the interview, communicates with the candidates, and sets up the next stage of the application process with the recruitment team and successful candidates.

iLink's Agentic AI application is delivering the same recruitment outcomes in materially less time (for example, it can hold 250 interviews in parallel). As a result, the central HR group has reduced headcount by 50%, with an overall cost reduction of 45%.

Recommended investor actions by investment timing & investment type

	Existing portfolio company		Intended investments	
	Exit < 12 months	Exit in 2-3 years	Stand-alone target	Part of M&A / bolt-on strategy
Objectives	Reduce deal drag. Protect the EBITDA multiple.	Increase EBITDA. Build a 12month+ history of higher operating margins.	Identify where Agentic AI should be deployed to maintain and/or increase Enterprise Value at exit. Size the value opportunity and required technology investment to inform price negotiations.	Identify opportunities to deploy Agentic AI across the existing and intended group of businesses to increase group Enterprise Value at exit. Map minimum operations technology requirements for potential bolt-ons.
Actions	<ul style="list-style-type: none"> • Author an Agentic AI strategy focused on single operational disciplines (articulate ROI and the high-level implementation plan with costs, timeline and resources). • Design (not implement) a systems architecture that will support enterprise-wide Agentic AI, showing how the current IT systems can be naturally evolved into the required state, outlining data architecture and approach to data governance. 	<ul style="list-style-type: none"> • Implement Agentic AI focused on a single operational discipline. • Author an enterprise-wide Agentic AI strategy: articulate focus areas, approach to enterprise data infrastructure and systems integration, ROI and high-level implementation phases with costs, timelines and resources. 	<ul style="list-style-type: none"> • Conduct an AI opportunity and readiness assessment – addressing both enterprise-wide and single operational disciplines. • Integrate Agentic AI strategy into the technology operations value creation strategy developed pre deal. 	<p>Pre deal:</p> <ul style="list-style-type: none"> • Conduct an AI opportunity and readiness assessment (enterprise-wide and single operational disciplines) for the initial acquisition. • Integrate Agentic AI strategy into the value creation strategy for the target. <p>Post deal:</p> <ul style="list-style-type: none"> • Develop a group Agentic AI strategy based on the intended range of group business activities.
Investor KPIs	Investor clarity: articulation of how the portfolio company can adopt Agentic AI (both single operational discipline and enterprise-wide). Management capability: proof that Agentic AI is understood by the team to be backed by the next investor.	Sustained improvement in operating margin. Successful implementation.	Agentic AI strategy as part of the investment thesis. Clarity for the Investment Committee on technology investment, ROI and value protection.	Group technology strategy for the M&A investment strategy. Agentic AI strategy as part of the investment thesis for the immediate target. Clarity for the Investment Committee.
Implementation timeframe	4-6 weeks.	6 months.	2-3 weeks	2-3 weeks pre deal, plus 2 weeks post deal.



©2025 SPV Insights Limited, United Kingdom.

For further information please contact
Paul Oliver at paul.oliver@excentius.com.