

Insights

WHERE DEPENDENCE SITS

WHAT THE OPENAI/MICROSOFT AND ANTHROPIC/AMAZON DEALS REVEAL ABOUT POWER, OPTIONALITY AND VALUE IN AI

Apr 23, 2026

Some of the most consequential strategic decisions in frontier AI are being made away from research labs. They are being taken by the executives negotiating collaborations between AI companies and the technology hyperscalers who have become their indispensable partners.

Two deals – OpenAI/Microsoft and Anthropic/Amazon – reward close examination, and demonstrate that these companies have taken fundamentally different routes, although both paths create significant interdependence between the AI companies and their partners.

OpenAI's collaboration with Microsoft began in 2019 as an investment with a \$1 billion headline value. Since then, it has evolved into an unprecedented partnership, involving deep technological and financial interdependence.

Microsoft has access to OpenAI's IP until 2032, including the right to integrate OpenAI models into its own products and to resell OpenAI services via Microsoft Azure. Microsoft also holds a significant equity interest in OpenAI. The commercial arrangements extend further - there are revenue sharing mechanisms across key products and platforms, and a very substantial commitment by OpenAI to purchase cloud infrastructure and related services.

Taken together, these arrangements go well beyond a conventional supplier relationship. Microsoft is not simply providing compute. It has secured a meaningful economic interest in OpenAI and its models and the ability to embed those models across its own product suite. OpenAI has secured access to the capital and infrastructure required to operate at the frontier of the industry. In return, it has accepted a high degree of interdependence with a single partner, which may in time also develop into a key competitor.

Anthropic's deal with Amazon is strikingly different. While Amazon has invested heavily in Anthropic, and Anthropic relies on Amazon for key cloud and hardware requirements, the arrangements maintain greater separation than those between OpenAI and Microsoft.

Critically, Anthropic has not granted IP rights to Amazon comparable to those agreed between OpenAI and Microsoft, and there are no widely reported revenue sharing arrangements. Claude is made available to Amazon customers through Bedrock, allowing enterprise users to build applications on top of Anthropic's models. However, this operates more as a distribution channel than as a transfer of underlying economic rights.

In essence, in addition to selling very large volumes of compute to their frontier AI partners, Microsoft has secured access to OpenAI's technology and the right to incorporate it into its own products, whereas Amazon is positioned more as a distributor and infrastructure provider without direct rights over proprietary IP.

It is also important to note how these arrangements have evolved as both the companies and the wider AI market have developed.

When OpenAI and Microsoft established their partnership in 2019, OpenAI was effectively a research lab in need of capital and compute. Since then, the value of OpenAI has increased dramatically, as have the capabilities and commercial relevance of its models. At the same time, it has become increasingly clear that OpenAI and Microsoft are, in certain respects, competitors.

The partnership has therefore been revisited and adjusted over time. That process has allowed both parties to introduce a degree of additional flexibility and to manage areas of friction, while remaining deeply intertwined.

Recent changes include an agreement that certain categories of OpenAI intellectual property, including categories of research-related IP, will be carved out from the rights granted to Microsoft prior to the contracted end date of the wider partnership, and OpenAI has secured greater flexibility in relation to its use of third party cloud providers. At the same time, key elements of the relationship, including the integration of OpenAI models within Microsoft's ecosystem and the use of Azure for core services, remain central.

Anthropic's partnership with Amazon has also deepened since it was established in 2023, with Anthropic building on its original decision to make Amazon its primary cloud provider by also naming Amazon as its primary training partner.

However, the reality is more nuanced than a simple commitment to Amazon silicon, and Anthropic has taken care not to commit entirely to a single supplier. It has expanded its use of Google TPUs for both future training capacity and inference, and has been explicit that its compute strategy spans AWS Trainium, Google TPUs and NVIDIA GPUs.

Therefore, while Anthropic does not appear to have committed – either legally or operationally – to exclusivity regarding compute, its position could be characterised as one of managed dependence.

Put another way, training frontier AI models is not a workload that can be easily moved between providers. It involves deep integration between model architecture, training frameworks, compilers and the underlying hardware. Once a model family has been developed and tuned on a particular chip architecture, switching that workload elsewhere requires significant engineering effort and may involve performance trade-offs. The same dynamic applies, albeit to a lesser extent, on the inference side.

The reality is that both OpenAI and Anthropic, through the commercial and technical decisions they have made, have created forms of interdependence with their partners.

OpenAI has, over time, negotiated greater contractual flexibility, including the ability to use multiple cloud providers for certain workloads. That is meaningful. But it sits alongside a reality in which it remains closely tied to Microsoft's ecosystem through IP rights and deep commercial integration.

Anthropic, by contrast, has preserved greater formal independence at the IP and revenue level, while still making technical choices that bind it closely to the infrastructure of its key partners, albeit in a more diversified way.

Different mechanisms, similar outcomes.

Whether through contractual integration, as in the case of OpenAI and Microsoft, or technical integration, as in the case of Anthropic and its hyperscaler partners, these companies are becoming tied to each other in ways that are not easily reversible.

The broader point is that in frontier AI, "lock-in" can take different forms. Once model development, training pipelines and deployment architectures are co-designed with a particular hyperscaler's hardware and software stack, the cost of moving away from that ecosystem becomes significant, even in the absence of formal exclusivity. Contractual rights over IP or distribution can be revisited over time, but deeply embedded technical dependencies cannot be unwound merely by negotiation.

For the hyperscalers, this is strategically powerful. Supplying compute is valuable. Shaping how that compute is used is more so. If the next generation of models is built to run most efficiently on your chips, within your cloud, then you are not just a supplier. You are part of the product.

For frontier AI companies, the trade-off is complex. Access to vast amounts of compute on workable terms is existential. But securing this in a way that embeds long-term dependence may constrain future strategic flexibility, particularly as the competitive dynamics between model providers and hyperscalers continue to evolve.

What these partnerships ultimately show is that the real negotiation and strategic decision-making is not simply over price, or even over IP. It is over where dependence sits in the stack, and how difficult it will be to change course later.

That is Where Dependence Sits, and why it matters.

*Author: **Richard Werner**, Partner, M&A and Corporate Finance*

RELATED CAPABILITIES

- M&A & Corporate Finance
- Corporate
- Technology Transactions
- Digital Transformation & Emerging Technology

MEET THE TEAM



Richard Werner

Partner, London

richard.werner@bclplaw.com

[+44 \(0\) 20 3400 2329](tel:+442034002329)

This material is not comprehensive, is for informational purposes only, and is not legal advice. Your use or receipt of this material does not create an attorney-client relationship between us. If you require legal advice, you should consult an attorney regarding your particular circumstances. The choice of a lawyer is an important decision and should not be based solely upon advertisements. This material may be “Attorney Advertising” under the ethics and professional rules of certain jurisdictions. For advertising purposes, St. Louis, Missouri, is designated BCLP’s principal office and Kathrine Dixon (kathrine.dixon@bclplaw.com) as the responsible attorney.