

# Forward Deployed Lock-In

Why The Labs Are  
Moving Into Services

June 2026

AUTHOR:



Emma Rowand

Research Analyst



### About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.

On May 4th, Anthropic announced a \$1.5 billion joint venture with Blackstone, Hellman & Friedman, and Goldman Sachs.<sup>1</sup> Seven days later, OpenAI launched the OpenAI Deployment Company (DeployCo), seeded at \$4 billion, and acquired Tomoro, a 150-person consultancy, as its founding asset.<sup>2</sup> Two labs that disagree about almost everything else independently decided in the same week that the right vehicle for enterprise AI work was a separate entity with private equity in the mix. Neither expanded existing partner programs. Neither leaned harder on Accenture or Deloitte. *Why?*

The general understanding is that enterprise adoption has lagged expectations, and the wide gap between model quality and deployed value needs to be closed. As [MIT's 2025 State of AI in Business Report](#) pointed out, despite \$30-40 billion in enterprise GenAI investment, 95% of organizations are getting zero return.<sup>3</sup>

All true. But none of it explains why OpenAI guaranteed PE partners a [17.5% annual return](#), why both labs bought their way in rather than partnering (OpenAI with [Tomoro](#), Anthropic with [Fractional AI](#)), or why they landed on the same vehicle in the same week. If the goal were just faster deployments, wouldn't the cheapest path be to triple your applied AI team and white-label through system integrators?

The services gap is a real cause but not the one driving this. As we argued in [Open Source Generative AI](#), the value in this market accrues to the infrastructure and orchestration layer, not to the models themselves. The fight is over who gets embedded inside the enterprise before the models commoditize. These labs are building land-grab vehicles, and the land is *operational dependency*.

The labs are the most credible names in the room, which is what makes saying yes so easy, but that lock-in is expensive and hard to undo. Buyers should know what they're committing to, and what else is out there, before they find themselves spending 2028 ripping and replacing an entire system.

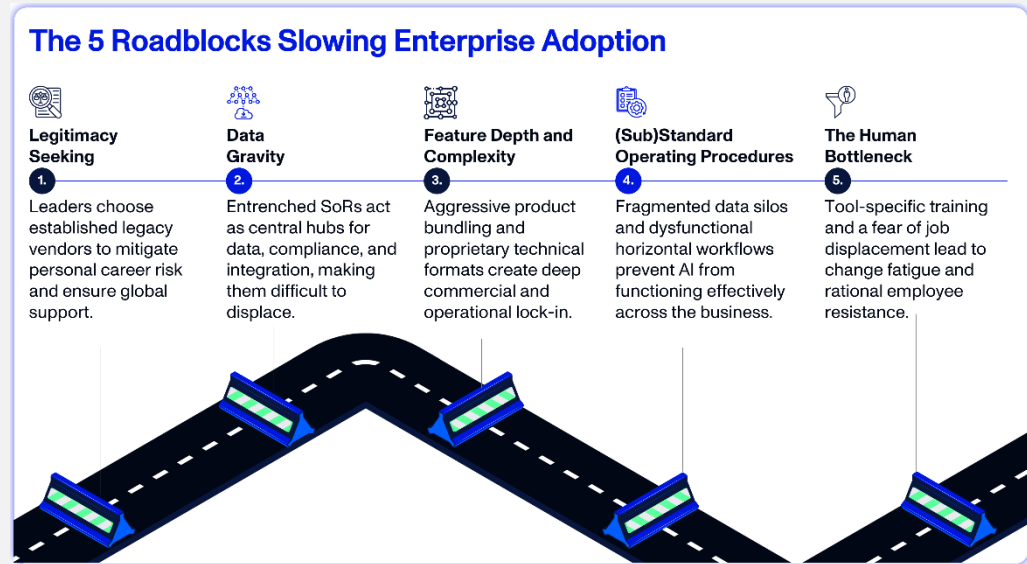
## The Palantir Comparison Is Half Right

Everyone reaching for an analogy lands on Palantir's Forward Deployed Engineer (FDE) model, and the labs aren't hiding the resemblance. But Palantir's FDEs were a product discovery mechanism: unprofitable services labor used to convert custom deployments into platform features. Foundry exists because Gotham was deployed into hundreds of environments and the patterns got productized.

DeployCo and the Anthropic JV are structurally different. The services entity consumes tokens from a platform that **already exists** upstream. The services work doesn't need to be profitable, just break-even enough to keep running while driving inference volume through and increasing the lab's compute revenue. An FDE who costs \$400K fully loaded but pulls a hypothetical \$5M+ in annual API consumption makes a ton of sense for a lab that earns 50-70% margins on the inference that FDE drives.<sup>4</sup>

## Captive Distribution by Design

Enterprise systems, once deployed, are difficult to replace. In [Selling AI-Native Service, Now](#) we identified the five primary roadblocks slowing enterprise AI adoption: legitimacy seeking, data gravity, feature depth and complexity, substandard operating procedures, and the human bottleneck.



The lab-affiliated FDE entities mitigate all five.

1. Blackstone on the cap table, adoption across a PE backer's portfolio, and Tomoro's roster that already includes Virgin Atlantic, Tesco, Mattel, and Fidelity make the buying decision a lot less lonely.
2. Anthropic's December 2025 donation of MCP to the Linux Foundation established the protocol as the emerging standard for how AI agents access data across systems, lowering the friction of connecting to siloed data.<sup>5</sup>
3. Instead of handing over an API key and walking away, FDEs go onsite and build the custom scaffolding, absorbing the integration pain themselves.
4. FDEs shadow employees and rebuild the process around the model, rather than bolting AI onto a broken workflow.
5. And because FDEs sit inside the work, they see how it really gets done, which takes quiet sabotage off the table.

**The race to deploy FDEs is a race to install operational dependencies before frontier-model capability commoditizes.** If a developer at a Fortune 500 company builds something on the Claude API, procurement can rip it out at renewal. If a forward-deployed team spends eighteen months embedding custom code into core workflows, training models on proprietary data, and redesigning the org around the system, the switching cost rises by orders of magnitude.

That race itself is a tell, because it cuts against the "SaaS apocalypse" narrative. In that story, each new model unlocks capabilities that work immediately across a company, so Anthropic or OpenAI become the horizontal operating system where

all work happens, and legacy incumbents like Salesforce, ServiceNow, and Workday get obliterated. But if raw capability were enough, the labs wouldn't need to get inside a company to understand its workflows and operations well enough to build "custom" deployments. The fact that they do is an admission that SaaS isn't dead and the true moat for incumbents is not the code itself, but the deeply entrenched data schemas, historical context, and user behavior patterns accumulated over decades of operational execution.

That need for embedding is creating a market of its own. Enterprise implementation companies are popping up everywhere. But [Ciridae](#), [Pit](#), [Foaster](#), [Tribe AI](#), and pre-acquisition Tomoro all share a single constraint: services has to be profitable, because services is the whole company. DeployCo and the Anthropic JV do not operate under that constraint. For them, services is a way to drive inference volume. If Anthropic sells Claude Code tokens to its JV at cost, and OpenAI does the same with the entity it majority-owns, the lab-affiliated shops have a cost structure no independent services firm can match.

**Which leaves one question for enterprise customers: do you want to be locked in by a vendor that profits from burning compute, or one that profits from results?**

That cost asymmetry is why the most interesting non-lab companies in this space are not really services firms at all. [Sierra](#) and [Decagon](#) look like FDE teams from the outside but they are productized businesses that own a specific workflow end-to-end. They survive because the lab does not want to be in the business of running a contact center, and the customer is buying an outcome rather than inference. The model underneath is interchangeable, which is the whole point.

The deeper mispricing is assuming operational dependency always equals durable revenue. It does for Palantir, because Foundry is hard to replicate and Palantir sells it as its own product. The same holds for the orchestration players, building their defensibility above the model layer. [Distyl AI](#) pairs its proprietary enterprise AI operating system (Distillery) with FDEs to embed itself in the workflows of telecom, healthcare, and financial services incumbents, while keeping the model layer swappable. Their moat becomes the platform that captures an enterprise's data, institutional knowledge, and domain expertise. The labs aren't building a platform product; they're absorbing the cost of a highly customized dependency on their own models.

This setup also assumes buyers are content routing operational workflows through a US-controlled frontier lab. In reality, sovereign deployment mandates, the EU AI Act, and compliance requirements in finance and defense will increasingly reject dependencies on centralized American API providers, favoring localized or isolated architectures instead. Since this layer is expensive to build either way, we believe paying a premium for a model-agnostic version that lets you swap the underlying model beats spending the same money to lock yourself into one lab.

For founders, the defensible ground has narrowed but has not disappeared. Defensibility sits at the extremes: deep, productized vertical outcomes in highly regulated sectors where labs refuse to take on compliance risk, or model-agnostic orchestration layers for sovereign and enterprise buyers who refuse to be locked in.

For investors, underwrite defensibility, not dependency. Be wary of paying platform multiples for what is really a services business riding someone else's inference margin.

And for buyers, don't trade long-term flexibility for one lab's short-term lift. As market maturity sets in, we believe the second wave of enterprise buyers will treat model-agnostic architecture as a strict procurement requirement rather than a preference.

- 
- <sup>1</sup> Wall Street Journal, [Anthropic Unveils \\$1.5 Billion Joint Venture With Wall Street Firms](#), 2026
  - <sup>2</sup> Axios, [OpenAI launches AI consulting arm valued at \\$14 billion](#), 2026
  - <sup>3</sup> MIT, [The GenAI Divide: State of AI in Business 2025](#), 2025
  - <sup>4</sup> MindStudio, [Why Anthropic's 70% Inference Margins Matter for Your API Costs — And What to Expect Next](#), 2026
  - <sup>5</sup> Anthropic, [Donating the Model Context Protocol and establishing the Agentic AI Foundation](#), 2025

**Disclaimer:** The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at [activantcapital.com/companies/](http://activantcapital.com/companies/) for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.