




ACTIVANT RESEARCH

Open Source Generative AI

From OpenAI to open source AI: how it unbundles the AI stack



Q3 2025

Jono Vickery

Introduction

Will the best Generative AI models be open or closed source 10 years from now? We believe that is the wrong question. Rather, how do you build the best AI *system*, and which models (plural) does it make use of?

In today's AI paradigm, closed source models built by research labs with \$10bn+ budgets dominate benchmarks and leaderboards. However, open source models are rarely far behind, typically matching performance within 6 months. This performance parity matters because history shows that demand for cost, performance and customizability will push the balance toward open source, but only if these models can keep up.

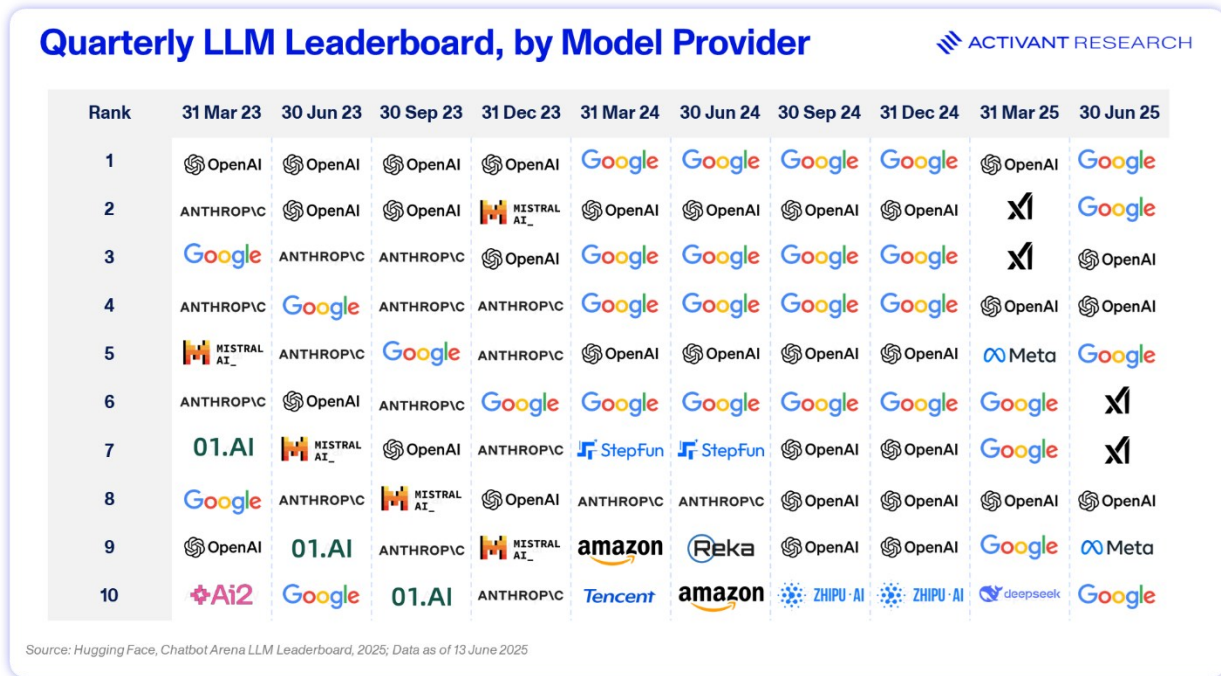
With techniques like fine-tuning and quantization, open source models can outperform massive closed source models on cost, latency and task-specific performance, all while supporting data security and regulatory compliance. Developers want to work with open source, and they don't need to compromise. Combining multiple fine-tuned, open source models into compound AI systems can beat monolithic closed source models at every dimension.

We believe this suggests that an increasing share of the Generative AI market will be captured not by massive models, but by the infrastructure layer that makes these systems possible.


The Status Quo: Dominance of Closed Source Models

When Generative AI is discussed, the names that likely come to mind are OpenAI, (which kicked off the Generative AI revolution with ChatGPT), Anthropic's Claude, and Google's Gemini. These companies secured this leadership position in societal mindshare by pairing state-of-the-art model research with a consumer interface, the chatbot, which we have all become so accustomed to. The chat interface turned out to be the perfect way to get this highly sophisticated technology into the hands of end users, including those with limited technical skills. Today, ChatGPT reports ~500 million weekly active users, and OpenAI has crossed \$10bn in run-rate revenues.^{1,2} By embedding Generative AI into search, Google's AI Overviews now reach 1.5bn users.³

However, good UX and product engineering alone are not what earned these companies their position. In fact, Google might be exceptionally weak in the [product](#) department, relying heavily on its technical strength. AI users are fickle, and leadership is backed up by **leaderboard dominance**. Consider LMArena, a leaderboard often referred to as the "[industry's obsession](#)" and even referenced in official [Google conferences](#). Since its inception, the top spots have been almost entirely held by Google, Anthropic, and OpenAI.



Not only do these providers offer state-of-the-art models, but they also deliver fully integrated stacks: they deploy the underlying GPUs/TPUs, optimize the infrastructure stack, and wrap it all up so that developers can add AI to their products with a simple API call. Unsurprisingly, there is strong crossover to leading in API inference—OpenAI clearly leads developer mindshare, with 63% of developers using its API.⁴



A lot of noise has been made about open source AI, but one thing is overwhelmingly clear today: **the dominant providers are all closed source**. Training state-of-the-art (SOTA) Generative AI models is expensive. Its reported to have cost OpenAI ~\$100 million to train GPT-4, a model now far removed from SOTA.⁵ Top labs are building clusters of 100,000 GPUs, with capital costs reported to be \$4 billion.⁶ And infrastructure costs are only 50% - 70% of the known costs, with staff making up the rest.⁷ Those costs are also intensifying, with some AI researchers now being offered pay packages worth “nine figures”.⁸ Scale is an extreme advantage: modern AI is not a world where some principled developers collaborating in forums can build a competing product, and it certainly looks like closed source providers may continue to dominate.

However, the dominance of a closed source provider in an infrastructure-level technology has little historical precedent. Developers value control, customizability, and the ability to mitigate vendor lock-in—advantages that open ecosystems can facilitate. In machine learning, open source frameworks like TensorFlow and PyTorch dominate; in databases, Postgres and MySQL hold the lead. Open source frequently wins at the infrastructure level.

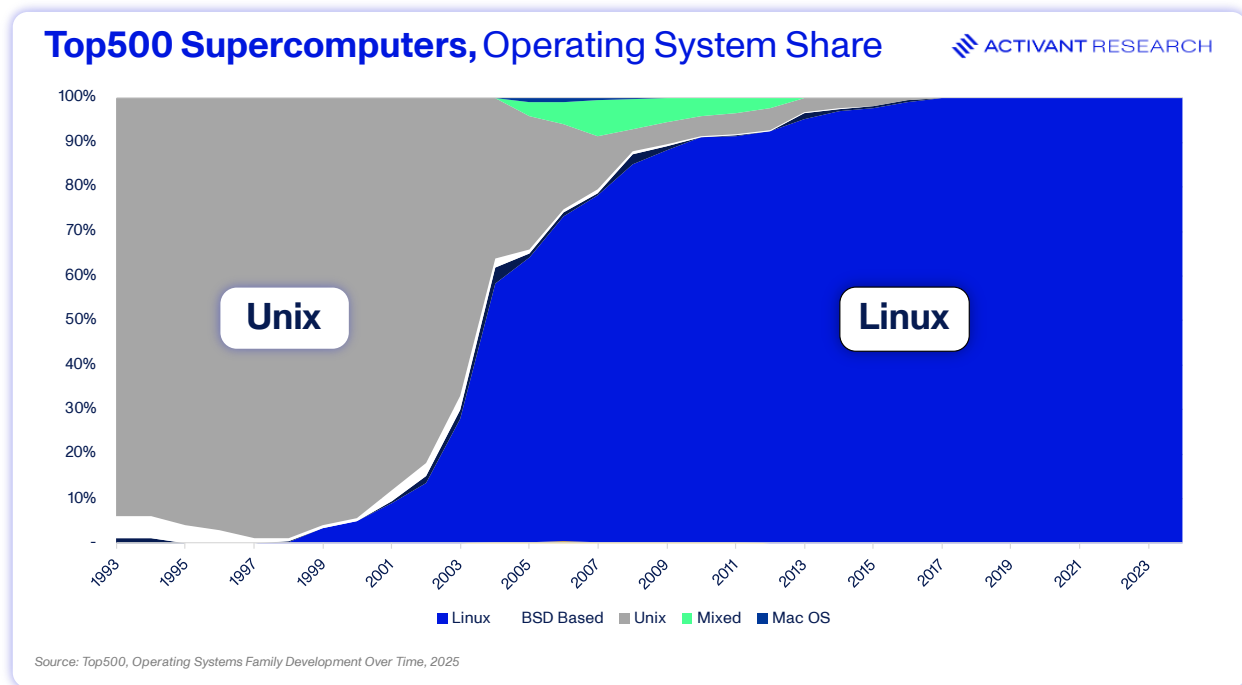
The real question is whether this will carry over to the market for foundation models, and we analyze past technological shifts for clues in the battle between open and closed models.

Lessons From History

Looking at past technological waves like operating systems, databases, developer tools, SaaS apps, and web browsers is significantly revealing as it relates to when and why technologies tend to open or closed. Customization, control, and cost-sensitivity tend to flip the scales in open source's favor, while ease of use and product quality can be a boon for closed source.

Tailoring Tech: Customization and Control Matters

Windows dominates the PC operating system market with 70% share, while Linux lags at only 4%.⁹ But turn to cloud infrastructure, and Linux starts to flip the script, holding >60% of the server OS market.¹⁰ Linux's open source and modular design allows developers to install only the essential OS components and maintain granular control over the kernel, services, and libraries, creating highly optimized, lean, and secure server environments. **When users need deep customization and control, tech tilts toward open source.** To drive this point home, consider the Top500 supercomputers, where every single floating-point operation per second matters: Linux claims 100% market share.



This consideration is also why machine learning libraries, where customization empowers researchers to build new systems, are dominated by open source PyTorch (over 70% of research implementations use PyTorch).¹¹ That's up from less than 10% in 2017, when TensorFlow, another

open source library, was dominant. PyTorch outpaced TensorFlow thanks to *deeper* customizability (with [dynamic graphs](#)), ease of use through Python, and a strong academic following, where network effects reinforced its status as the standard for AI papers. **Ease of use proved especially important, helping fuel broader community adoption.**

But What Does it Cost?

For the buyer of a Windows PC, the OS is a once-off and invisible cost. For server builders, using a proprietary OS can become extremely costly at scale. An open source server OS, even with a higher upfront cost, offers attractive cost leverage as servers scale across a datacenter. In the PyTorch example, a costly proprietary library would have been a non-starter for research use cases with uncertain payoffs. **When cost matters, tech tilts open.**

The Technical User Factor

For developers and engineers, open source often shines. Consider the dominance of Docker in containerization (over 60% share), or VS-Code in IDEs (used by 73% of Developers).^{12,13} Customizability enables differentiated product experiences. However, the extensibility of the tools needs to match the technical skills of their users.

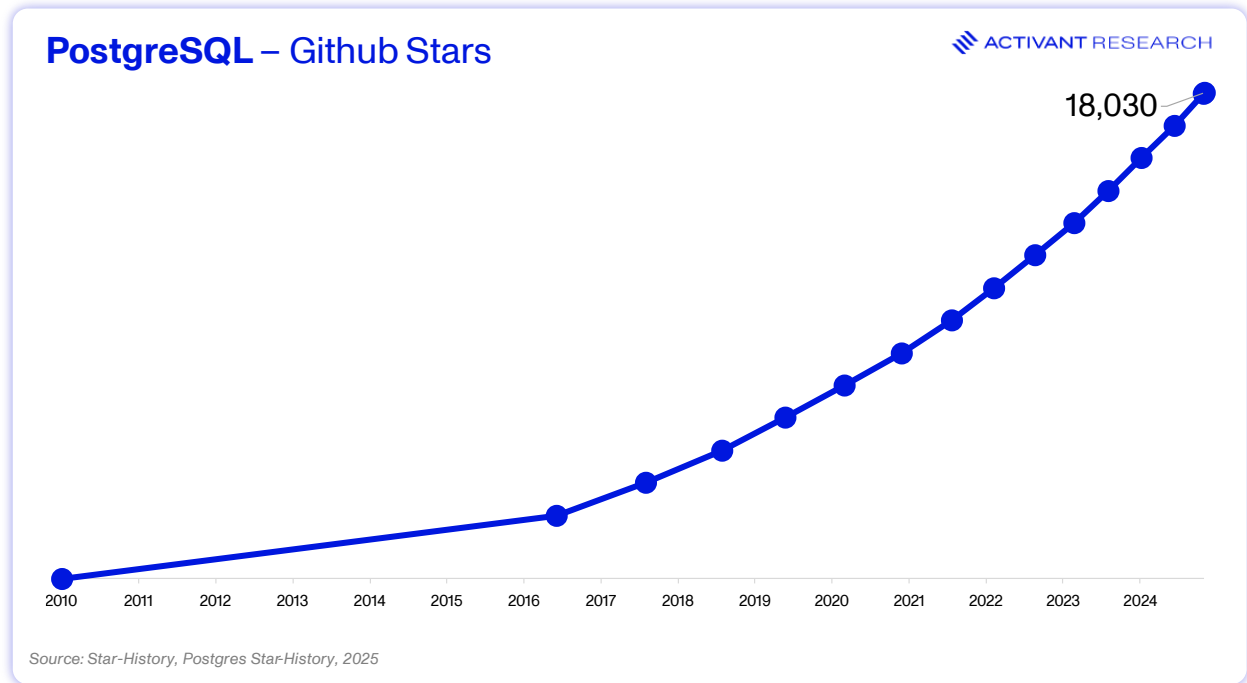
In ERP, CRM and other common SaaS apps, companies like SAP, Salesforce and Workday thrive with closed source products, even though these systems often require extensive customization. Business users require intuitive, easy-to-use, visual UI. Closed source systems thrive in these situations because the providers invest heavily in R&D to build and optimize the system end-to-end and ongoing maintenance to ensure reliability. Of course, they handle the customization issue for you, but often at the cost of [billions](#) in consulting and professional services fees.

For open source to succeed, the end user needs to be technical enough to leverage that increased customizability and control, and savvy enough to save your company the extreme costs that come with closed source providers.

Building Blocks: Infrastructure

Infrastructure tools are most likely to have technical users, with high customizability requirements and a sensitivity to costs as they scale. Tying together the points above, **infrastructure tech tends toward openness**. It's also key to note that infra tech is often about creating standards, where ease of adoption is key, and providing a tool for free is a great way to bootstrap adoption. From Apache Kafka (used by >80% of the Fortune 100), to PostgreSQL, the list is long.¹⁴ The case of Postgres also leaves us with one more lesson: **the tide can turn**. For a long time, the database market was dominated by the early closed source innovators in Oracle and MySQL. However, as the tech

commoditized over time, the ability to get 80% of the features at no cost became extremely compelling, catapulting Postgres into the position of one of the most popular relational databases.



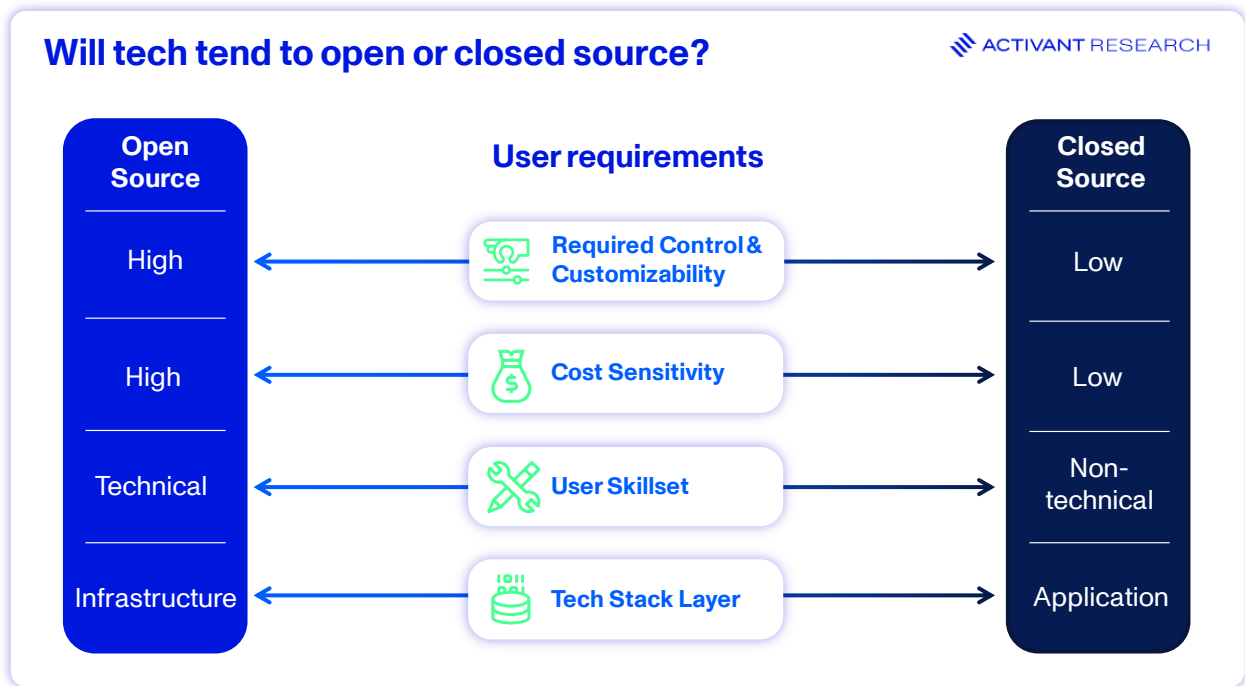
But Sometimes, It Doesn't Work That Way

Microsoft's closed source Windows dominated PCs, both because the end user was less technical but also because Microsoft created a tightly knit ecosystem where the prevailing apps that PC users wanted worked best on Windows. Closed source iOS has a ~25% market share of the mobile operating system market because Apple bundles it into an integrated product that they control and optimize end-to-end.¹⁵

One can easily see such a future with OpenAI: while their models are seeing success independently today, they are also making significant moves at the app layer with ChatGPT and Codium. ChatGPT has already largely "won" the consumer battle, and that consumer mindshare could easily be the defining factor that makes OpenAI a long-term market-share holder in the market for Generative AI models.

And sometimes, openness can just be too complicated. OpenStack, an open source cloud platform launched just a few years after AWS, struggled to gain traction and is ultimately considered to be a failed experiment. First, it was excessively technical, even for technical users, ease of use is a key consideration. Second, while it promised to be cheaper than the cloud, it failed

to account for the fact that the cloud's pay-as-you-go pricing model, with no upfront infrastructure costs, was part of its allure for many early adopters.

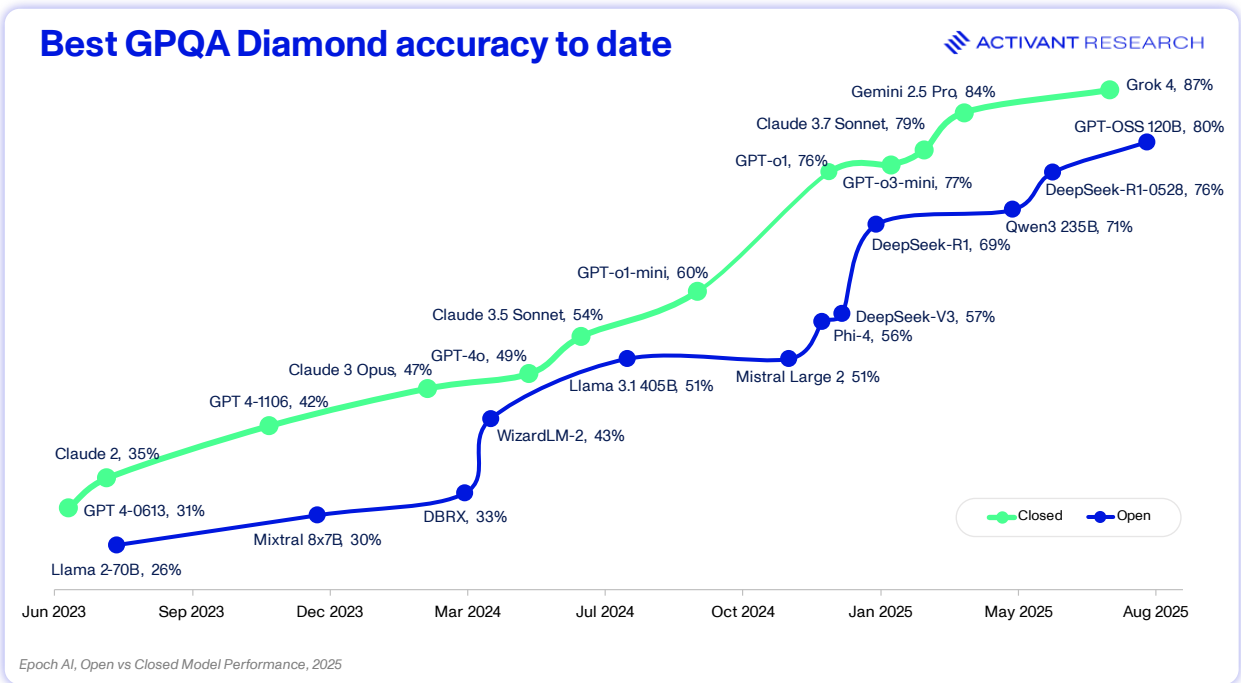


In Generative AI, **models serve as infrastructure** for technical users: AI developers. They require extensive customization to optimize for hardware configurations, cost and latency, and to fine-tune models for domain-specific use cases. History suggests that, despite the current status quo, Generative AI models are likely to tend toward open source over time, and we know that early leads can be overturned.

These fundamentals will matter less if the capital costs to build SOTA models exceed \$4bn. But in January 2025, a small and obscure Chinese hedge fund launched a reasoning model, DeepSeek R1. It ranked close to the leading models at the time, with a reported training cost of just \$6 million.^{16,17} The amounts that large companies are willing to spend make headlines, but look under the surface, and you'll see deep innovation is possible on limited budgets.

Technical Innovations Driving Open Source AI


Consider the chart below: it's clear that closed source has the advantage, outperforming open source on the Graduate-Level Google-Proof Q&A (GPQA) at every period since the launch of GPT-4.¹⁸ What's interesting, however, is that the extent of this lead never grows: open source consistently catches up after about six months.



The closed source community leads by scaling, and the open source community maintains (*lagged*) performance parity through algorithmic innovations and hardware optimization. For example, DeepSeek’s breakthrough was achieved through [Multi-head Latent Attention \(MLA\)](#), [DeepSeek Mixture of Experts](#) (algorithmic enhancements) and [Mixed-Precision training](#) (Hardware Optimization). In the spirit of open source, these breakthroughs are available to the entire community and allow for the improved cost, performance and control over LLMs. In addition to the algorithmic enhancements that we’re seeing, there are three key technical innovations driving the competitiveness, flexibility and cost-effectiveness of open source.

1. Fine tuning

Fine tuning is the process whereby a pre-trained model is retrained on a much smaller but task-specific dataset to specialize its behavior. Fine tuning a model on a company’s past customer service conversations can ensure that its new answers are more specific, accurate and better at understanding company-specific jargon. One of today’s SOTA coding models is [Qwen3-coder](#), a mid-sized (480B) open source model fine-tuned on coding data, indicating that fine-tuning can be a method for driving frontier performance.¹⁹ **It can also keep costs down. A fine-tuned version of Llama 2-7B, capable of running on most laptops, outperforms GPT-4 (a model that runs on 8 x Nvidia H100 GPUs) on some tasks.**



Fine-tuning is clearly effective, but running a new dataset through the entirety of the model weights can be slow (~24 hours per experiment) and can become costly with multiple experiments. [LoRA](#), widely adopted since 2023, essentially adds a trainable parameter to models, which allows for fine-tuning by training only a small fraction of the model's total parameters. Using this method, researchers were able to fine-tune a 65B parameter model on a single 48GB GPU, making the method widely accessible and cost-effective.²⁰ Critically, LoRAs are small files which could technically be swapped in/out and shared, meaning that open source models can function as a platform base, with LoRAs acting as flexible plug-ins for domain specificity.

2. Model Distillation

Distillation is the method where a large model's outputs (or input/output pairs) are used to train a much smaller model but achieve similar performance. For example, Llama 4's Maverick and Scout models are trained using distillation from their largest model, Behemoth. Maverick's GPQA score of 70% is only 390bps lower than that of Behemoth, but the model has one-fifth the parameter count, massively reducing the costs of both training and inference. **Distillation means that its possible to shrink models with minimal impact on performance.**

This creates highly performant models with cheaper, faster inference. For the open source community, distillation turns large, open source teacher models into a platform for experimentation, potentially giving rise to entirely new approaches in the future. For enterprises,

smaller student models can be simpler to deploy, lowering the bar to adoption and enabling new use cases, such as offline, on-device assistants in industries with sensitive data. Distillation can also help place guardrails around the training data of the distilled model, creating models aligned with a company's culture, vision or compliance requirements.

3. Quantization:

With quantization it has become common to run large language models in 8-bit or 4-bit precision without noticeable performance loss vs. traditional 16- or 32-bit precision. While this compression method can result in performance loss if not done correctly, the open source [GPTQ](#) algorithm allows one-shot conversion of full models to 4-bit weights. The bottom line is more open innovation that allows AI developers to solve for the cost and latency that they desire.

And it's not just technical innovations that make open source a key consideration.

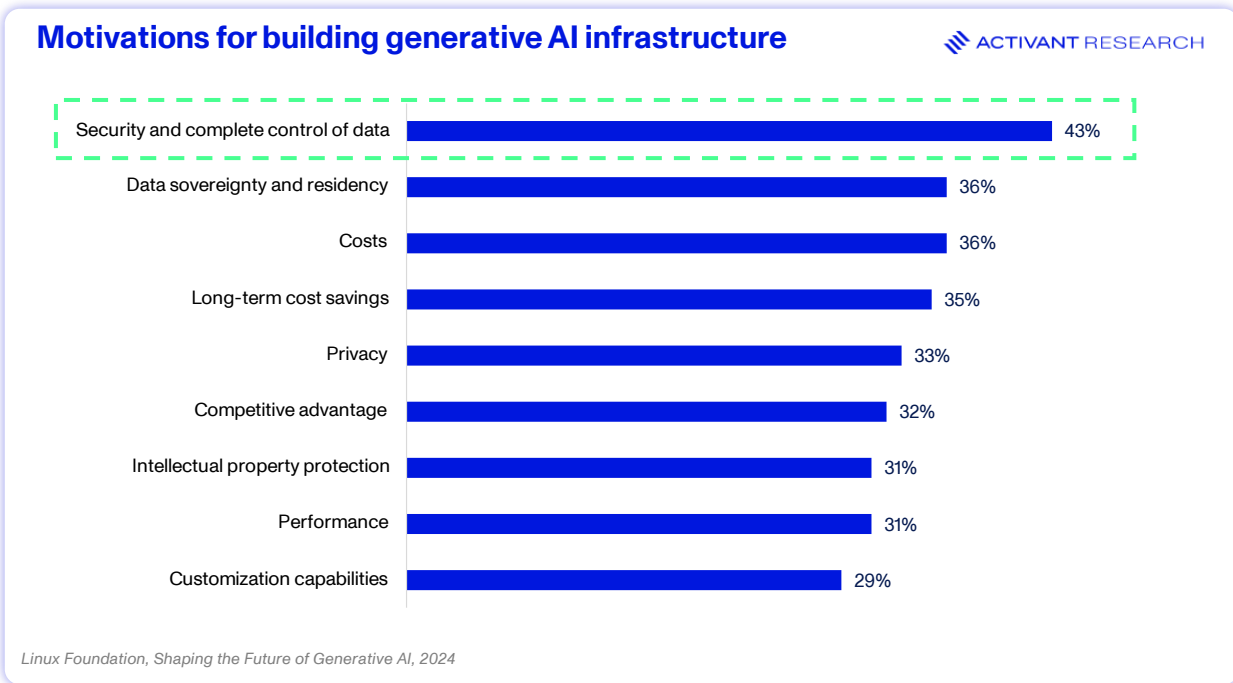
Data Gravity

In the broader IT world, 45% of workloads remain off the public cloud (being either on-premises or in virtual private cloud). Progress has been stubbornly slow, decreasing only 5 percentage points from the 50% that were still outside the public cloud in 2022.^{21,22} In fact, some 8% – 9% companies plan to fully repatriate their cloud workloads back to on-premises.²³ There are a few key recurring themes, notably the desire for greater cost control, performance optimization, operational control, and, critically, **data security and compliance**.

As we wrote in our research on [Data Security](#) last year, data security is a mess for many organizations, with more than half having suffered a material data breach in the past 12 months.²⁴ Chief among the drivers of these breaches are fragmented data estates with inconsistent governance and security policies, so it's no surprise that many CISOs feel better consolidating their sensitive data on-premises. Further, the evolving regulatory landscape, particularly in Europe, prevents many companies from sending sensitive data across borders. Data is the gravity that keeps many workloads away from the cloud.

The relevance to open source AI couldn't be greater. Building a business-critical AI application with closed source models accessed via API means sending vast amounts of your sensitive data to that AI provider's datacenter, likely in near real-time. For those who felt that the public cloud was not appropriate for their non-AI workloads, this setup will likely be a non-starter.

That's why data privacy and security are the most cited barriers to Generative AI adoption, the top challenge when implementing RAG, and the primary motivation for building an entity's own Generative AI infrastructure.^{25,26,27}

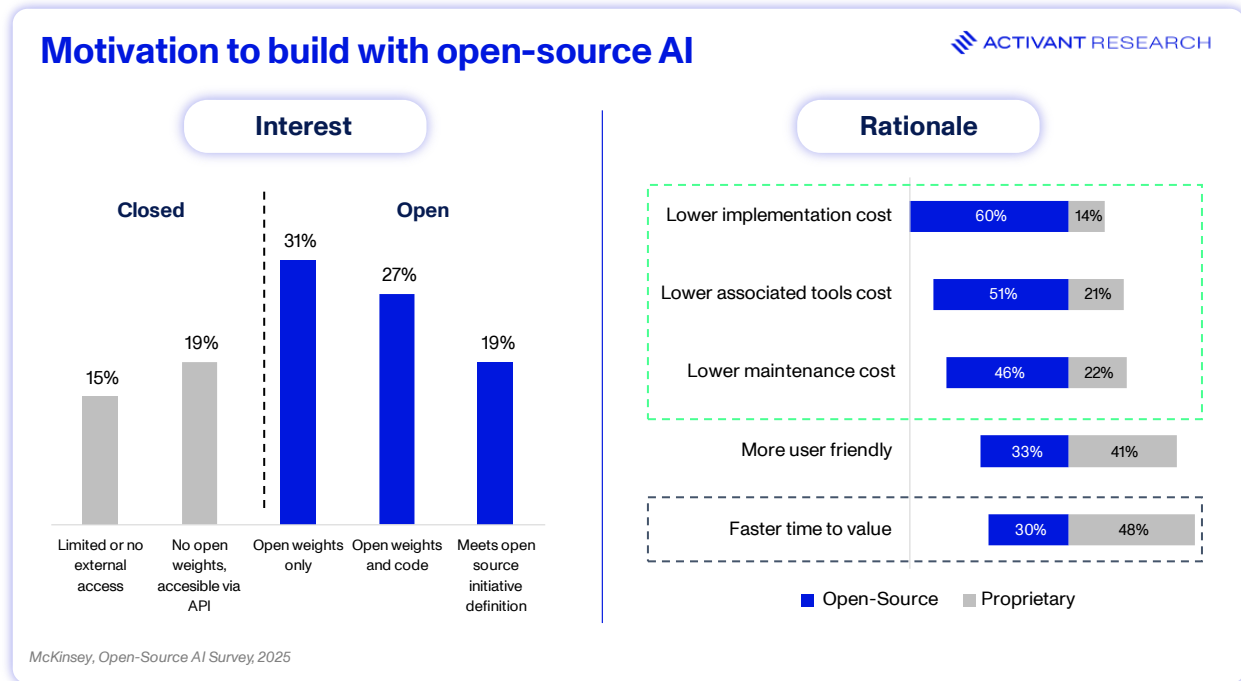


These issues couldn't be more acute in highly regulated industries like banking, where [Goldman Sachs](#) is using open source models in its own environment to ensure that its data is not used for model training and remains within its controlled environment.²⁸ Or consider the case of Australian bank [ANZ](#), which used OpenAI for experimentation but deployed fine-tuned versions of Llama in production to ensure stability and data sovereignty.²⁹

Just like the cloud, we can't expect 100% of AI workloads to be directed outside a company's private environment, and that means using open source.

The Compound AI Systems Hypothesis

While closed model APIs might hold the lead in the enterprise today, developers *want* open source AI. 83% of organizations believe that AI needs to be increasingly open and 73% of organizations expect to increase their use of open source generative tools.³⁰ In fact, it might be the case that the main reason closed source is winning today is faster time to value.

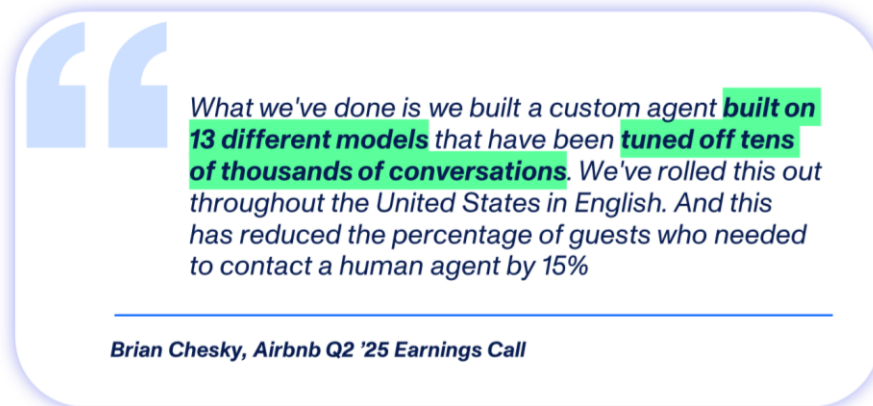


Whether it's due to the need to optimize task-specific performance, reduce costs, control the stack with more granularity, or ensure that sensitive data is secure and compliant, open source AI models are going to play a *much* greater role than they do today. Only ~25% of companies have deployed to production, which helps explain the current dominance of closed model providers – these were the easy pickings, the simple deployments.³¹

We are still early. As generative AI deployments expand to more complex, regulated entities with greater needs to optimize cost and performance, they will increasingly rely on open source solutions.

In fact, these complex deployments will not simply use one model. **Mature AI systems will consist of multiple models, orchestrated to deliver a given service.** In compound AI systems, the required balance of latency, cost and performance is achieved through a blend of large models, small fine-tuned models, distilled or quantized models and, in some cases, custom models trained by the deploying organization. Complex workflows are broken up into atomic units where the optimal model is used for each step and it is open source models, with their flexibility and customizability, that make this possible.

This is not just academic. ChatGPT itself is a system that leverages OpenAI's largest foundation model along with the Dall-E image generator and a code interpreter. FactSet's text-to-financial-formula tool hit a quality ceiling until they implemented a compound AI system: a fine-tuned open source model handles user queries, embeddings are created for formula retrieval, second model re-ranks the formulas, and a large model generates the final output.³² Similarly, Airbnb's new customer service tool employs 13 different models:

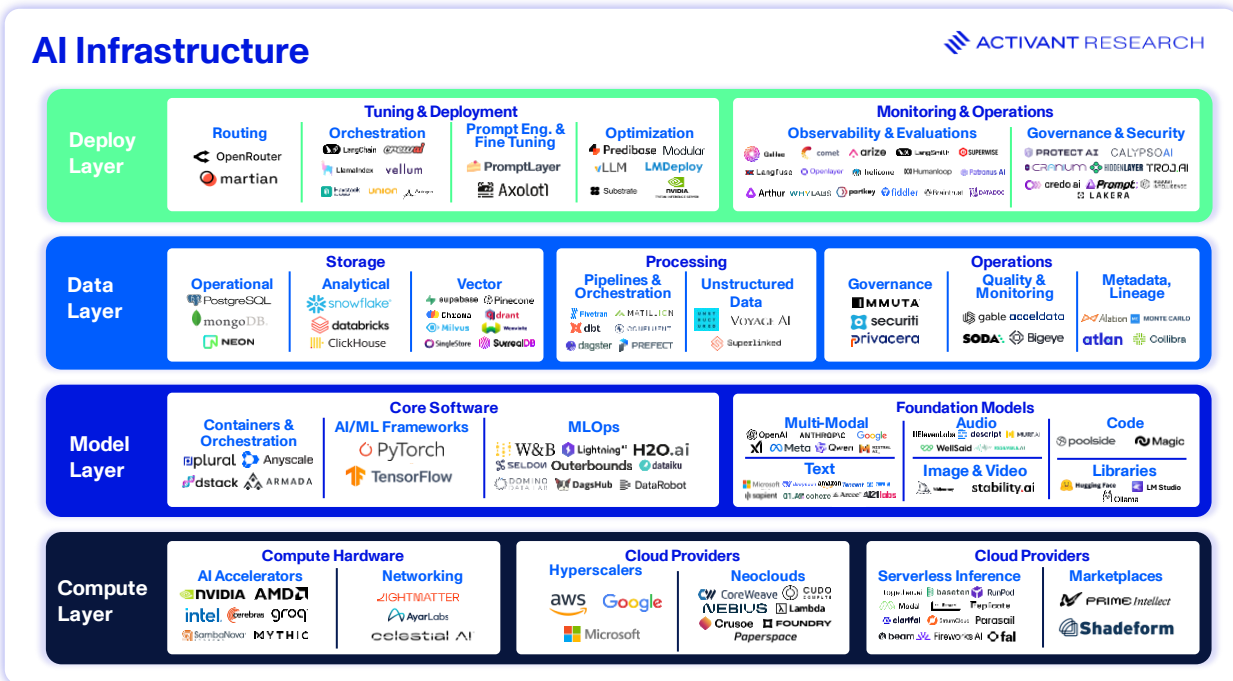


The emergence of compound AI systems tells us that not only can open source be cheaper, more secure, and assist with regulatory compliance, but that when built into the correct systems, it's just outright better.

We believe that the writing is on the wall: AI is going to shift away from monolithic systems built on closed models, to modular and compound systems with open source as the backbone. The reason that this is particularly interesting is what it means for the start-up ecosystem.

Unbundling the Infrastructure Tech Stack

In a world where a few closed source model providers dominate, the so-called “AI Infrastructure stack” is mostly abstracted away, and AI developers only interact with the big three cloud providers and a few API services (OpenAI, Google Vertex, Anthropic). **But, building compound AI systems, complete with open and closed models, shifts the value from individual models to the software that makes these systems possible.** Open source AI unbundles the AI stack.



In this new AI stack, **Neoclouds** like [Coreweave](#) and [Nebius](#) counterposition the Hyperscale Clouds by providing priority access to raw Nvidia GPUs, without the complexity of the broader Hyperscaler ecosystem (AWS offers 200+ services³³). They excel in availability of top-of-the-range GPUs. Unlike Google, Microsoft and Amazon, these Neoclouds have no internal uses for their GPUs and gain priority access to new chips—a setup that suits Nvidia’s competitive positioning.

Serverless Inference allows developers to get access to any model as a service, without having to provision or think about GPU infrastructure. [Replicate](#) allows developers to instantly deploy any of 50,000 open source models, as well as run custom fine-tuning. Serverless inference can be 2 – 3 times faster than deploying on traditional cloud, as multi-step cloud workflows are abstracted into a few lines of code and [cold-start](#) times are reduced to sub-second by companies like [Modal Labs](#).³⁴

MLOps platforms like [Outerbounds](#) are key to model development, providing model versioning, dependency management, and workflow orchestration. **Model Libraries** like [Ollama](#) allow developers to run open source models on their own hardware as simply as calling an API. This removes one of the critical bottlenecks to open source’s success, the significant complexity required to get started, further tilting the argument in favor of open source.

Data is, of course, one of the key inputs to AI training. In enterprise AI infrastructure, the **Data Layer** is critical at inference time - Retrieval Augmented Generation (RAG) can boost performance even

more than fine-tuning or prompt engineering.³⁵ As we've noted [before](#), RAG enables enterprises to tap the 80%+ of unstructured data that was previously ignored, giving rise to a new unstructured data pipeline. [Unstructured.io](#) automates the retrieval, transformation, and staging of unstructured data for LLMs, while [Voyage AI's](#) embedding and reranking models can improve RAG accuracy and reduce the cost of vector storage 3x.³⁶

Further, AI and agentic AI increase the need for [DataOps](#). Metadata gives AI contextual understanding of the data it is acting on, lineage ensures explainability of AI actions for compliance, and data quality ensures that AI is acting on robust, up-to-date data. Simply put, the data layer is a force multiplier for the performance of these AI systems. [Soda](#) provide out-of-the-box quality checks and converts natural language SQL for generative AI quality checks, while [Atlas](#)'s active metadata platform crawls the entire data estate to automate tasks across the DataOps lifecycle.

Once the model and its source data are finalized, there is still huge leverage in the infrastructure used at the deployment layer. **Routing** platforms like [Martian](#) can outperform SOTA models by directing queries to the model best suited for each task. Some users cut inference costs [100x](#) with [OpenRouter](#)—a significant nod to the compound AI systems hypothesis. To **optimize deployments**, [Modular](#) provides a high-performance programming environment for deploying models on distributed computing environments, scaling performance to over 10,000 tokens per second (compared to a human average of five tokens per second). [LangChain](#) is establishing itself as a standard in **Orchestration**, offering a software library of modular components to chain together models and vector databases into the complex workflows required by compound AI systems.

Finally, **Observability and Evaluations** platforms like [Langfuse](#), [Galileo](#), and [Comet](#) can show a detailed flow of requests and responses for debugging and understanding model behavior. They also track KPIs such as latency, cost, and quality scores, and allow users to give feedback or evaluations on model outputs.

An Open Future

Look past the OpenAI API, and there is a thriving ecosystem of tools for deploying AI systems, replete with open source models autonomously routed or chained together in predefined workflows and optimized all the way down to the silicon. The innovation in open source models makes these systems possible, and we believe that they represent the next step in the AI deployment wave—faster, cheaper, better. As focus shifts from which model beats some arbitrary benchmark to which systems drive tangible business outcomes, value moves to the software layer that enables these systems.

-
- ¹ [Forbes, ChatGPT Fuels \\$300 Billion Valuation, Waymo Taps Uber, AI Wins SXSW, 2025](#)
 - ² [TechCrunch, OpenAI claims to have hit \\$10B in annual revenue, 2025](#)
 - ³ [Google, I/O 2025 Keynote, 2025](#)
 - ⁴ [Vellum, State of AI, 2025](#); *Public survey of 1,285 people in December 2024*
 - ⁵ [Epoch AI, How Much Does It Cost to Train Frontier AI Models, 2025](#)
 - ⁶ [Semianalysis, 100,000 H100 Clusters, 2024](#)
 - ⁷ [Epoch AI, How Much Does It Cost to Train Frontier AI Models, 2025](#)
 - ⁸ [Axios, Meta launching AI superintelligence lab with nine-figure pay push, reports say, 2025](#)
 - ⁹ [StatCounter, Desktop Operating System Market Share Worldwide, 2025](#)
 - ¹⁰ [Fortune Business Insights, Server Operating System Market, 2025](#)
 - ¹¹ [PyTorch, 2024 Year in Review, 2024](#)
 - ¹² [Datadog, 10 Insights on Real-World Container Use, 2023](#)
 - ¹³ [Stack Overflow, 2024 Developer Survey, 2024](#)
 - ¹⁴ [Apache Kafka, Powered by, Accessed 2024](#)
 - ¹⁵ [StatCounter, Mobile Operating System Market Share Worldwide](#)
 - ¹⁶ [DeepSeek-AI, DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning, 2025](#)
 - ¹⁷ *It is noted that the \$6 million figure may be misleading, as it refers to the pure compute cost of the final training run only, ignoring infrastructure costs and the time spent on experimentation*
 - ¹⁸ *GPQA is a multiple-choice exam with answers considered resistant to simple google searches, and is used to benchmark reasoning capabilities of AI models*
 - ¹⁹ [WebDevArena Leaderboard, as of 30 July 2025](#)
 - ²⁰ [Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, Luke Zettlemoyer, QLoRA: Efficient Finetuning of Quantized LLMs, 2023](#)
 - ²¹ [Flexera, State of the Cloud, 2025](#)
 - ²² [Flexera, State of the Cloud, 2022](#)
 - ²³ [IDC, Storm Clouds Ahead: Missed Expectations in Cloud Computing, 2024](#)
 - ²⁴ [Rubrik, The State of Data Security, Statistics reflect 5000+ Rubrik customers across 67 countries, 2023](#)
 - ²⁵ [Bain, Survey: Generative AI's Uptake Is Unprecedented Despite Roadblocks, 2024](#)

²⁶ [Anaconda, The State of Enterprise Open-Source AI, 2025](#)

²⁷ [Linux Foundation, Shaping the Future of Generative AI, 2024](#)

²⁸ [Odd Lots, Interview with Goldman Sachs CIO, 2024](#)

²⁹ [Rico Zhang \(ANZ Head of Technology\) via LinkedIn, Leveraging Gen-AI for Quality Engineering: Our Journey from OpenAI's GPT-3.5 to In-House Llama2, 2023](#)

³⁰ [Linux Foundation, Shaping the Future of Generative AI, 2024](#)

³¹ [Vellum, State of AI, 2025; Public survey of 1,285 people in December 2024](#)

³² [Databricks, FactSet: Implementing an Enterprise GenAI Platform with Databricks, 2024](#)

³³ [Amazon, About AWS, 2025](#)

³⁴ Activant ecosystem feedback

³⁵ [Linux Foundation, Shaping the Future of Generative AI, 2024](#)

³⁶ [Voyage AI, Home Page, 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/ 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.