Article

Do it yourself, or do it right? How to get your AI projects off the ground

HPE Private Cloud AI provides a solution to help you achieve your AI ambitions.

There's no question generative AI is changing our understanding of what machines can do and how businesses will operate in the 21st century. But enterprises now face the quandary of how best to deploy AI: Should they develop their own AI practices or customize solutions built by others? And what is the optimal cloud environment for these workloads?

Today, half of enterprises are running some or all their AI deployments in the public cloud.[1] The primary driver here is speed. If you want to get an AI proof of concept (POC) up and running quickly, or need to run a batch of small-scale experiments to identify promising use cases, using public infrastructure can be faster and requires a smaller initial investment than building out your own.

But when it comes to deploying AI across an organization, the benefits of public cloud can quickly fall away. While up-front costs may be relatively small, relying entirely on public infrastructure may be more expensive over the long run. In fact, over a three-year period, total cost of ownership can be up to 60% lower for organizations using private cloud AI.[2]

"There are instances where testing in the public cloud makes sense," says Michael Corrado, a senior worldwide manager for AI and private cloud at Hewlett Packard Enterprise. "But as you move into production, you need a reliable platform you can build on — one whose costs aren't going to explode as you scale."

Managing your own infrastructure helps keep operating expenses from spiraling out of control as your needs grow. And having these systems highly accessible makes it easier to fine-tune performance and optimize for specific use cases.

Perhaps most important, operating in a private cloud allows you to keep your company's proprietary data close to home. Unlike with the public cloud, you're not sharing infrastructure with anyone else. The security of your data is assured, not assumed, offering you greater peace of mind while also avoiding regulatory issues around data sovereignty.

DIY vs. getting things done

For most AI use cases, private cloud allows for greater control over cost, data, and performance. But that still leaves the question of whether to build your own bespoke AI environment from scratch or acquire a turnkey solution and customize it to your needs.

For decades, IT shops have wrestled with the build vs. buy question, and the adoption of GenAI is no different. If you've already got AI infrastructure up and running, have in-house expertise, and rely on specific tools or are pursuing unique use cases, building on top of what you already have makes a lot of sense.

So far, however, most enterprise AI efforts have fallen short. In fact, nearly 90% of AI experiments and POCs never make it into production.[3]

There are several possible causes for that, notes Corrado. The IT teams running the experiments may have been unable to keep up with the furious pace of AI innovation. The use cases may not have demonstrated enough business value. Or the organization lacked sufficient data to scale the POCs to enterprise-wide production.

But perhaps the number one reason is that the organization did not have the resources or expertise to bring its AI projects to fruition.

"A lot of enterprises are stuck in the proof-of-concept stage," says Corrado. "They might have hundreds of AI projects in the works, but for various reasons they don't fully launch. Potentially great projects never get off the ground, and organizations never find out which ones might provide real value."

An added benefit of a private cloud solution: Total cost of ownership for a turnkey AI system can be up to 50% lower than a DIY approach, thanks to reduced engineering, operational, and licensing costs.[4]

[1] "2025 State of the Cloud Report," Flexera, March 19, 2025.

[2] "The Economic Benefits of HPE Private Cloud AI With NVIDIA AI Computing by HPE," Enterprise Strategy Group, March 2025.

[3] "88% of AI pilots fail to reach production — but that's not all on IT," CIO.com, March 25, 2025.

[4] "The Economic Benefits of HPE Private Cloud AI With NVIDIA AI Computing by HPE," Enterprise Strategy Group, March 2025.

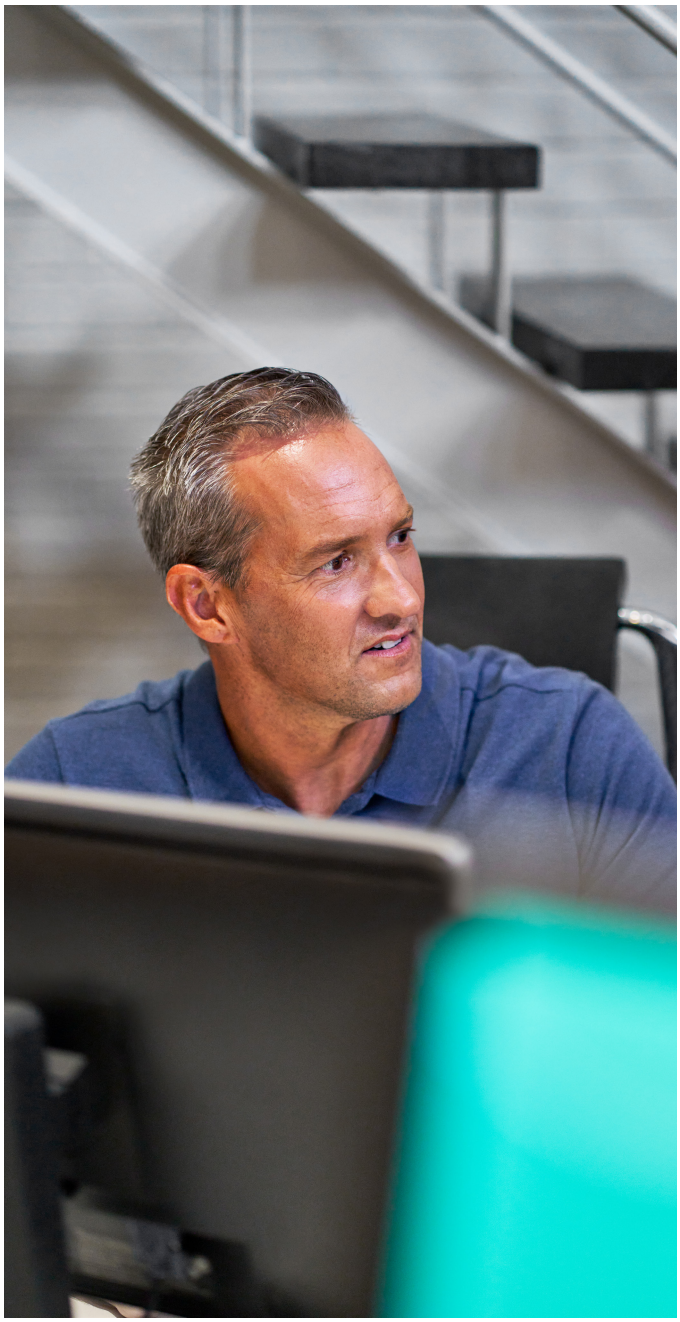

Turnkey AI: Add data and stir

For most enterprises, buying and fine-tuning a pre-built solution can not only save enormous amounts of time but reduce complexity. This is why HPE and NVIDIA have teamed up to provide integrated hardware and software solutions that allow enterprises to accelerate the launch of their AI projects.

"These solutions allow you to get started quickly and cost-effectively, by providing the technology and the expertise you need to make it work," Corrado notes.

HPE Private Cloud AI is a turnkey solution designed to help enterprises achieve their AI ambitions. The rich ecosystem of proprietary and open-source tools simplifies infrastructure configuration and management, allowing IT teams to rapidly deploy AI workloads and scale projects while maintaining the privacy and security of enterprise data.

Built on HPE ProLiant servers, the pre-built solutions come in configurations that range in size depending on whether you need just inferencing or intend to run inferencing, retrieval augmented generation, and fine-tuning.

HPE AI Essentials Software offers a comprehensive suite for AI development and deployment, while NVIDIA® AI Enterprise Software provides AI models and frameworks for the leading AI use cases. This collection of curated self-service tools allows developers to build AI applications up to 90% faster than through traditional methods.[5] Organizations can deploy technologies like GenAI chatbots with a single click.

Top-level advisory services are available to smooth out any post-deployment wrinkles, adds Corrado. And because both HPE and NVIDIA are constantly making the latest technology available to customers, enterprises need not fear rapid obsolescence or the inability to quickly scale.

"You know these solutions are going to work because HPE and NVIDIA have tested and validated them," Corrado says. "And because both the hardware and software are constantly updated, you never need to worry about falling behind. You simply add more infrastructure to get the newest stuff."

[5] "HPE Private Cloud AI with NVIDIA AI Computing by HPE: Essential to Accelerating GenAI Industrial Transformation," The Futurum Group, January 2025.

The clock is ticking

As businesses race to take advantage of the insights and innovations AI can bring, time to market becomes a key differentiator. Enterprises don't have the luxury of waiting for the technology to settle down, laboriously cleansing their data, or pondering which use cases to experiment with, says Corrado.

The time to get started is now.

"We recommend organizations start small but dig in quickly, fine-tuning things as they go," Corrado says. "You can get real value quickly out of deploying new customer service bots or eliminating inefficiencies in your sales process, but you'll never know until you start doing it."

Visit HPE.com

Learn more at

HPE.com/ai

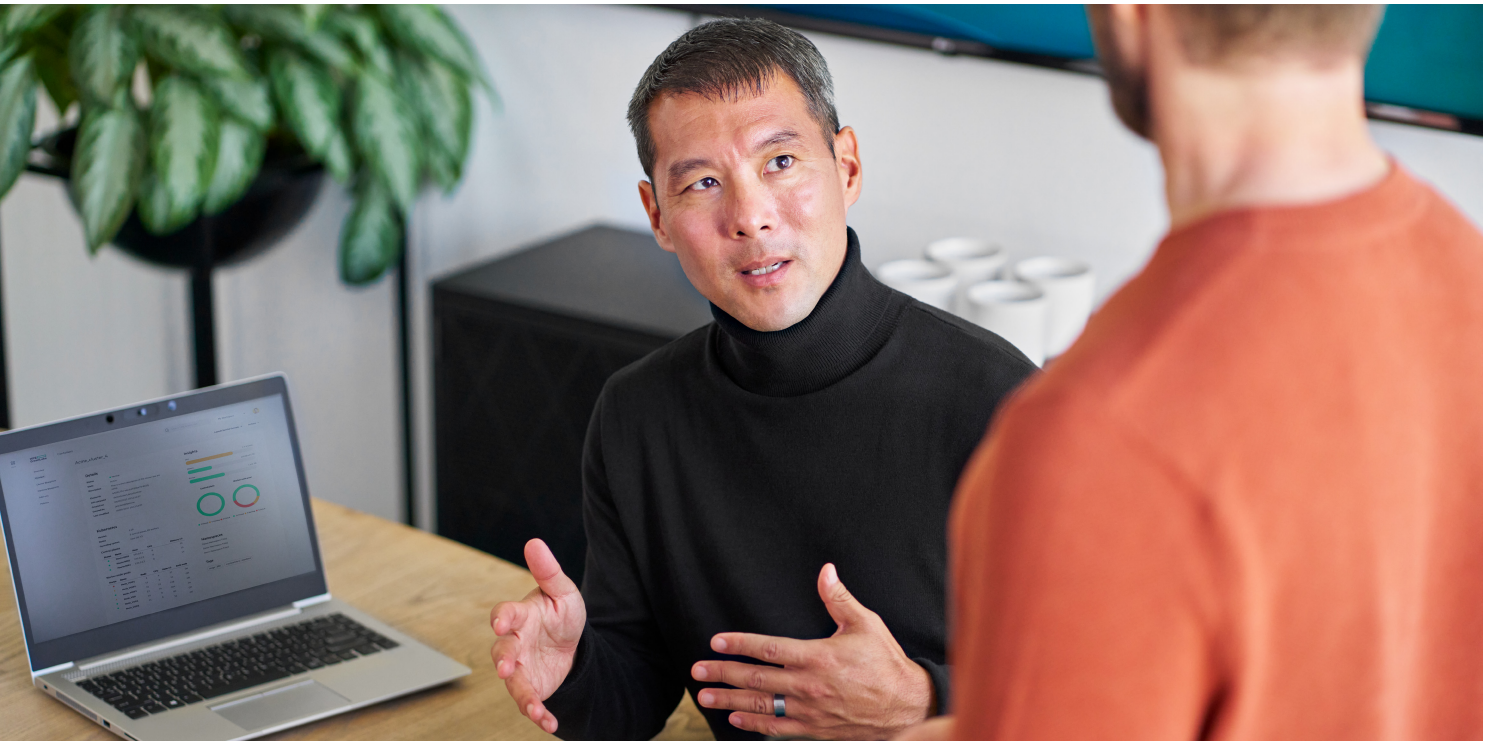

Chat now

a50013432ENW

HEWLETT PACKARD ENTERPRISE

hpe.com