



# AT HPE, AI IS A SUPERCHARGER FOR PRODUCTIVITY

With HPE Private Cloud AI, innovation comes first and scales fast



**Customer:** Hewlett Packard Enterprise  
**Industry:** Technology  
**Country:** United States

When Hewlett Packard Enterprise set out to develop its own internal enterprise artificial intelligence (AI) solution, it embarked on a path of innovation that would enhance employee productivity and showcase the broader value of its integrated AI platform. Built on HPE Private Cloud AI—one of the first fully turnkey AI infrastructure systems for the enterprise—the solution accelerates projects from pilot to production while delivering performance at scale.

## The next everything

AI is the next frontier, the next industrial revolution, an investment opportunity, a team sport... just fill in the blank. It could be all or any of them, but one thing is certain—it will be transformative.

Dan Waibel has been working on AI projects for the past 15 years, and he's experienced more than a few "wow" moments in his work. "To me, the biggest unlock right now is the story around agentic AI. We've always done well with large language models, but with agentic, the whole equation changes," Waibel says. "We're looking at the possibility of building a digital workforce that can augment what our employees are capable of achieving."

His colleague Manmeet Dhillon, head of Enterprise AI at HPE, agrees. "Just in the last six months, there's been a tipping point in how we use these platforms. Everything we can do now, we can do even better. The things we want to do next—we can see how to get there faster."

Today, Waibel serves as the global chief data officer at HPE, overseeing data strategy, digital transformation, and AI. And he's leading HPE on a journey to develop a company-wide AI platform, leveraging its own AI private cloud environment.

## A road to AI paved with expertise

It's a trajectory HPE has been constructing for years. When the company launched its HPE GreenLake cloud, it was a bold move toward a strategy of selling everything as a service, with a focus on self-service, flexible capacity, and the ability to scale rapidly without massive up-front investment.

Now, it's doing the same thing for AI. With the ChatHPE GenAI Hub, the technology leader is developing an AI environment on its own integrated hardware platform: HPE Private Cloud AI. This turnkey AI factory is the first co-engineered solution from NVIDIA AI Computing by HPE, a collaborative effort between HPE and NVIDIA that helps organizations achieve their AI objectives.

The reasons are almost boundless. "We decided to build the product on our own systems because it delivers the data security we need, minimizes costs, and gets us up and running faster. It also helps us to improve the HPE Private Cloud AI platform as a product for our customers," Waibel relates.

## Vision

Spur innovation by developing a secure enterprise AI platform to empower the workforce

## Strategy

Leverage a cost-effective private cloud solution that speeds time to value while boosting performance at scale

## Outcomes

- Delivers innovation from day one with preintegrated hardware, software, and expertise
- Saves approximately four hours per user per week by automating repetitive and time-consuming tasks
- Builds a scalable foundation for AI without inflating costs (estimated 50% savings in three to four years)



“Because HPE Private Cloud AI is configured for GenAI workloads out of the box, you don’t need to spend months building the platform. It’s your cloud. You can focus on AI first.”

– Dan Waibel, Global Chief Data Officer, Hewlett Packard Enterprise

## A project in search of a home

ChatHPE didn’t start life in HPE’s own data center. “In reality, when we first stood it up, HPE Private Cloud AI wasn’t even a product offering. It was just launched at HPE Discover Las Vegas 2024, so we developed it in the public cloud,” Waibel recalls.

The first iterations were on OpenAI through Microsoft Azure, which was a valuable experience for Waibel and Dhillon. “Launching this in the public cloud gave us a firsthand view of the challenges so many companies face when they tackle their first AI initiatives,” Waibel says.

According to Dhillon, the cloud was a stepping stone to a more stable and scalable environment: “The cloud is great when you’re first experimenting and prototyping, but it wasn’t the right place to grow.”

## Limitations in the cloud

Scaling quickly became an issue. “It was really a cost issue for us. The pay-as-you-go model works well early on, but when you turn around and look at scaling something like ChatHPE to an enterprise with 60,000 employees, the ROI was not in our favor,” Waibel adds.

The team also realized the security profile of public cloud was limited for their purposes. “There are certain types of data you don’t want to put into the public cloud—anything around legal, intellectual property, patent data—and in those cases, you want to keep your data on-prem. As far as regulatory compliance goes, that extends to data sovereignty as well,” Dhillon says.

Yet another challenge the team faced was performance variability. “If you’re running workloads in the cloud pay-as-you-go model, you really don’t have control over where those land. If there’s capacity free somewhere halfway around the world, your workload gets moved, and your performance lags. Ultimately, performance is not predictable, and that’s an issue,” Dhillon relates.



## A platform for scale, performance, and flexibility

By the time HPE was ready to move the AI hub into its own data center, it had grown from 200 users to more than 6,000, powering more than 20 applications in under six months.

The team chose HPE Private Cloud AI to expedite the migration. “Until recently, you really had only two choices in the AI marketplace: the cloud or build your own,” Waibel explains. Now, there’s a new option for those who want their AI infrastructure on-premises: HPE Private Cloud AI.

With AI-optimized infrastructure, networking, software, and foundation models, this turnkey AI factory enables faster, more accurate, and scalable implementation of AI projects. And to accelerate the deployment of generative and agentic AI applications, NVIDIA Blueprints are integrated into the solution.

With HPE Private Cloud AI, the deployment process bypassed the complexities of hardware assembly, data pipeline structuring, and software integration. And because it’s a part of the HPE GreenLake ecosystem, setup and management are significantly streamlined. With one-click deployment and built-in dashboard functionality, Waibel and team had a head start.

After the system arrived in the HPE global IT data center, ChatHPE was onboarded just one week later. After two weeks, the GenAI was inferencing. “Because HPE Private Cloud AI is configured for GenAI workloads out of the box, you don’t need to spend months building the platform. It’s your cloud. You can focus on AI and business value first,” Waibel says.

## Supercharging the talent pool

The initiative lays the groundwork for a future where HPE employees are truly augmented and supported by AI. “Ultimately, we want to be able to support our staff with agentic AI. Imagine an autonomous agent to gather leads, an agent to compile customer research, and an agent to summarize the data into talking points—it’s how we supercharge our talent pool,” Waibel says.

It’s a transformation that could completely alter the vendor–customer relationship. “Who we’re selling to is changing,” Waibel explains. “Instead of having conversations with the CIO or CTO, we’re going to be talking to CMOs, CEOs, and maybe other sales officers. It’s no longer about the 3-year server lifecycle. It’s about business use cases, vision, and strategy.”

Looking ahead three to four years, HPE can already predict a 50% savings by moving its AI from the public cloud to HPE Private Cloud AI. Furthermore, the more HPE relies on its own private cloud, the more it refines HPE Private Cloud AI for the marketplace. “We’ve collected early feedback and implemented considerable improvements,” Waibel says. “It’s already a better product than it was six months ago.”



**“Who we’re selling to is changing. Instead of having conversations with the CIO or CTO, we’re going to be talking to CMOs, CEOs, and maybe other sales officers. It’s no longer about the 3-year server lifecycle.”**

– Dan Waibel, Global Chief Data Officer, Hewlett Packard Enterprise

## Solution

— HPE Private Cloud AI

## Explore more

[HPE.com/Private-Cloud-AI](https://hpe.com/Private-Cloud-AI)

Visit [HPE.com](https://hpe.com)

## [Chat now](#)

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Azure and Microsoft are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All third-party marks are property of their respective owners.

a00151864ENW

HEWLETT PACKARD ENTERPRISE

[hpe.com](https://hpe.com)

