

Executive Brief


HPE Private Cloud AI

Simplify, streamline, and secure AI projects at scale with on-prem infrastructure from Logicalis and HPE.




In a few short years, AI has become the world's most hyped technology. The promise of automation, efficiency, and productivity gains has led organizations large and small to create their own AI solutions.


Yet as enterprises move from experimentation to production, they are facing a few common challenges:



AI relies on models and mountains of data: IT teams face complex decisions about which models to use, where they’re hosted, and how to manage their lifecycles. Data management also brings concerns around sensitive IP, security, and compliance.



The tooling landscape is filled with numerous, often expensive solutions: Many of these require powerful hardware, complex integration, and continual maintenance. Even after managing data, models, and tools, organizations must solve scalability and cost predictability problems.



Infrastructure is evolving and complex: One of the most important decisions is choosing the right AI infrastructure. Because AI is constantly changing, so too is the infrastructure that supports it. Teams must have the expertise and the technology stack that can support AI initiatives at speed, at scale, and at cost.

When it comes to infrastructure, there are three approaches: the public cloud, DIY infrastructure, and private cloud AI.

2022

ChatGPT catalyzed an explosion of interest in generative AI. In just two months, it accumulated over 100 million users, ushering in the generative AI era.

2023


The wave of excitement gave rise to widespread experimentation as enterprises and developers began developing trials and proofs-of-concept.

2024

The journey hit a new milestone: the race to production. Companies around the world took their POCs into live environments, but only about one in ten POCs reached production successfully.

2025

The focus turned to agentic AI, with autonomous agents, hybrid systems, and more intuitive AI solutions.






Public Cloud AI Infrastructure

Initially, public cloud infrastructure may seem like the easiest route. It is easily accessible and scalable, with costs that appear predictable.

However, there are downsides, such as:

- Limited control over infrastructure
- Potential cloud outages
- Data sovereignty and compliance issues
- Latency and data adjacency problems
- Costs that become less predictable as projects scale
- Vendor lock-in with proprietary tools

Although the public cloud offers simplicity and can work for early projects, over the long term its limitations can restrict innovation, security, and growth.



Build Your Own Infrastructure


Building your own AI infrastructure provides the maximum flexibility, with on-premises infrastructure, data sovereignty, and increased security, among other benefits.

However, this is the most complex, costly option and suitable for very few organizations.

Challenges include:

- Finding, hiring, and retaining the right expertise
- Significant complexity for both deployment and maintenance
- Higher labor and integration costs
- Extended time-to-implementation

For most organizations, these added challenges make "DIY infrastructure" impractical or completely out of reach.





On-Premises, Private Cloud AI

The third option is the private AI cloud, which provide the best of both cloud and on-prem infrastructure.

Increasingly, organizations are opting for onsite cloud solutions, which delivers benefits such as:


- Predictable costs
- Direct control over data and security
- Modular scalability
- Easy access to the latest tools and models
- Simplified support via a single vendor

With support from managed service providers such as Logicalis, organizations can bypass these challenges and focus their efforts on what matters most: innovating with AI.


Why Private AI Infrastructure?

Infrastructure is the foundation upon which apps, automations, agents, and other innovative AI solutions are built. Like the infrastructure of a building, AI infrastructure needs to work seamlessly, without unpredictable costs, compliance issues, or complexity that slows innovation.


Private AI infrastructure systems solve for these challenges by offering:




Full-stack, turnkey,
standardized solutions



Pre-integrated and factory-
certified compute,
networking, storage,
software, and security




Designs that
streamline AI model
training, fine-tuning,
and inference
workloads



Ready-made
automations, AI
agent libraries,
blueprints, and
developer tools

Managed service providers, as mentioned, can add even more benefits by handling workload assessments, infrastructure design and deployment, ongoing management, and more. Logicalis, in partnership with HPE, is helping organizations make the most of AI with a turnkey private cloud platform that lets organizations deploy in days, not months.




HPE Private Cloud AI

HPE Private Cloud AI delivers a complete, secure AI workbench with a unified data lakehouse and rapidly deployable models, for a variety of use cases.

Part of the NVIDIA AI Computing by HPE portfolio, HPE Private Cloud AI is a turnkey platform that lets teams get started on AI projects immediately and spend more time on innovation.

NVIDIA AI Computing by HPE¹

HPE Private Cloud AI



This managed technology stack combines:

- NVIDIA accelerated computing, networking, and software
- HPE high-performance compute and storage
- HPE GreenLake Cloud services
- Curated capabilities for data pipelines, orchestration, processing, and MLOps

HPE Private Cloud AI is purpose-built for the innovators driving enterprise AI, including:

- **Administrators.** Manage infrastructure, data sources, tools, GPU resources, and user access. The platform makes infrastructure management seamless.
- **Data Scientists & ML Engineers.** Focus on experimentation, training, and model deployment, instead of tool complexity.
- **Data Engineers & Analysts.** Build data pipelines to process, transform, and visualize massive datasets while minimizing operational overhead.

1. www.lemondeinformatique.fr/actualites/lire-un-cloud-prive-pour-l-ia-pret-a-l-emploi-avec-hpe-et-nvidia-%C2%A0-94087.html

HPE Private Cloud AI

HPE Private Cloud AI delivers a complete, secure AI workbench with a unified data lakehouse and rapidly deployable models, for a variety of use cases.

Part of the NVIDIA AI Computing by HPE portfolio, HPE Private Cloud AI is a turnkey platform that lets teams get started on AI projects immediately and spend more time on innovation.

Helps Accelerate AI Initiatives²

HPE Private Cloud AI is for...

1. Administrators

Administrators to deliver AI infrastructure

Rack, stack, offer and expand the infrastructure for AI in just a few clicks, with end-to-end management. Manage data sources, tools, GPU resources, users, and access to tools and data

2. Data Scientists + ML Engineers

Data Scientists + ML Engineers to deliver AI solutions

Experiment, train, and deploy ML models and AI applications that will deliver unprecedented efficiencies, without worrying about tools or infrastructure.

3. Data Engineers + Analysts

Data Engineers + Analysts to deliver data that fuels AI

Build data pipelines, and process, transform, as well as visualize the data needed to power AI solutions while reducing the effort needed to manage data infrastructure.

From a centralized dashboard, IT operations teams can:

- Deploy and configure all core components from day one
- Gain complete visibility into hardware and infrastructure health
- Proactively monitor GPU cores, CPU capacity, and system performance
- Detect and resolve bottlenecks before they impact operations


Data scientists and engineers can use HPE AI Essentials, the built-in software platform, to gain easy access to the right resources, including:

- A pre-integrated ecosystem of tools and models, from NVIDIA AI Enterprise to open-source notebooks
- Self-service access to eliminate IT bottlenecks
- Support for both novice and expert users

With this comprehensive software stack, teams can build everything from basic models to complex agentic AI systems, in less time and with less overhead.

2. www.paths.ext.hpe.com/c/v100007030enw?x=z272Ok&cc=emea_europe&lang=en&lb-height=100&lb-width=100&lb-mode=overlay&pf_route=v100007030

HPE AI Essentials
HPE Private Cloud AI Software Stack



Build the Next Generation of Applications



Content and code generation



Customer support resolution



Detection and prediction




Employee onboarding and training




Knowledge base copilots



Self-service portals



Virtual agents and chatbots



Visual computing



Take the Next Steps with

Logicalis and HPE Private Cloud AI

Logicalis has been a trusted HPE Partner since 1993, delivering a full portfolio of solutions and services customized to our clients' needs. Empowered by a range of specializations that span infrastructure, AI, compute, storage, and more, we help businesses turn technology into a catalyst for change.

If you'd like to get started on your journey with HPE Private Cloud AI, reach out today to take the next step.

[Schedule Free Executive Consultation](#)

Why Logicalis?

- HPE Platinum Partner, trusted since 1993
- HPE solutions and services portfolio that span hybrid IT, data protection, networking, and more
- Highly certified HPE technology experts focused on hybrid data center, private cloud, and AI solutions
- Authorized HPE solution provider with global reach and consistent services in 15 countries
- Proven track record with education, government, and enterprise customers

