Brochure

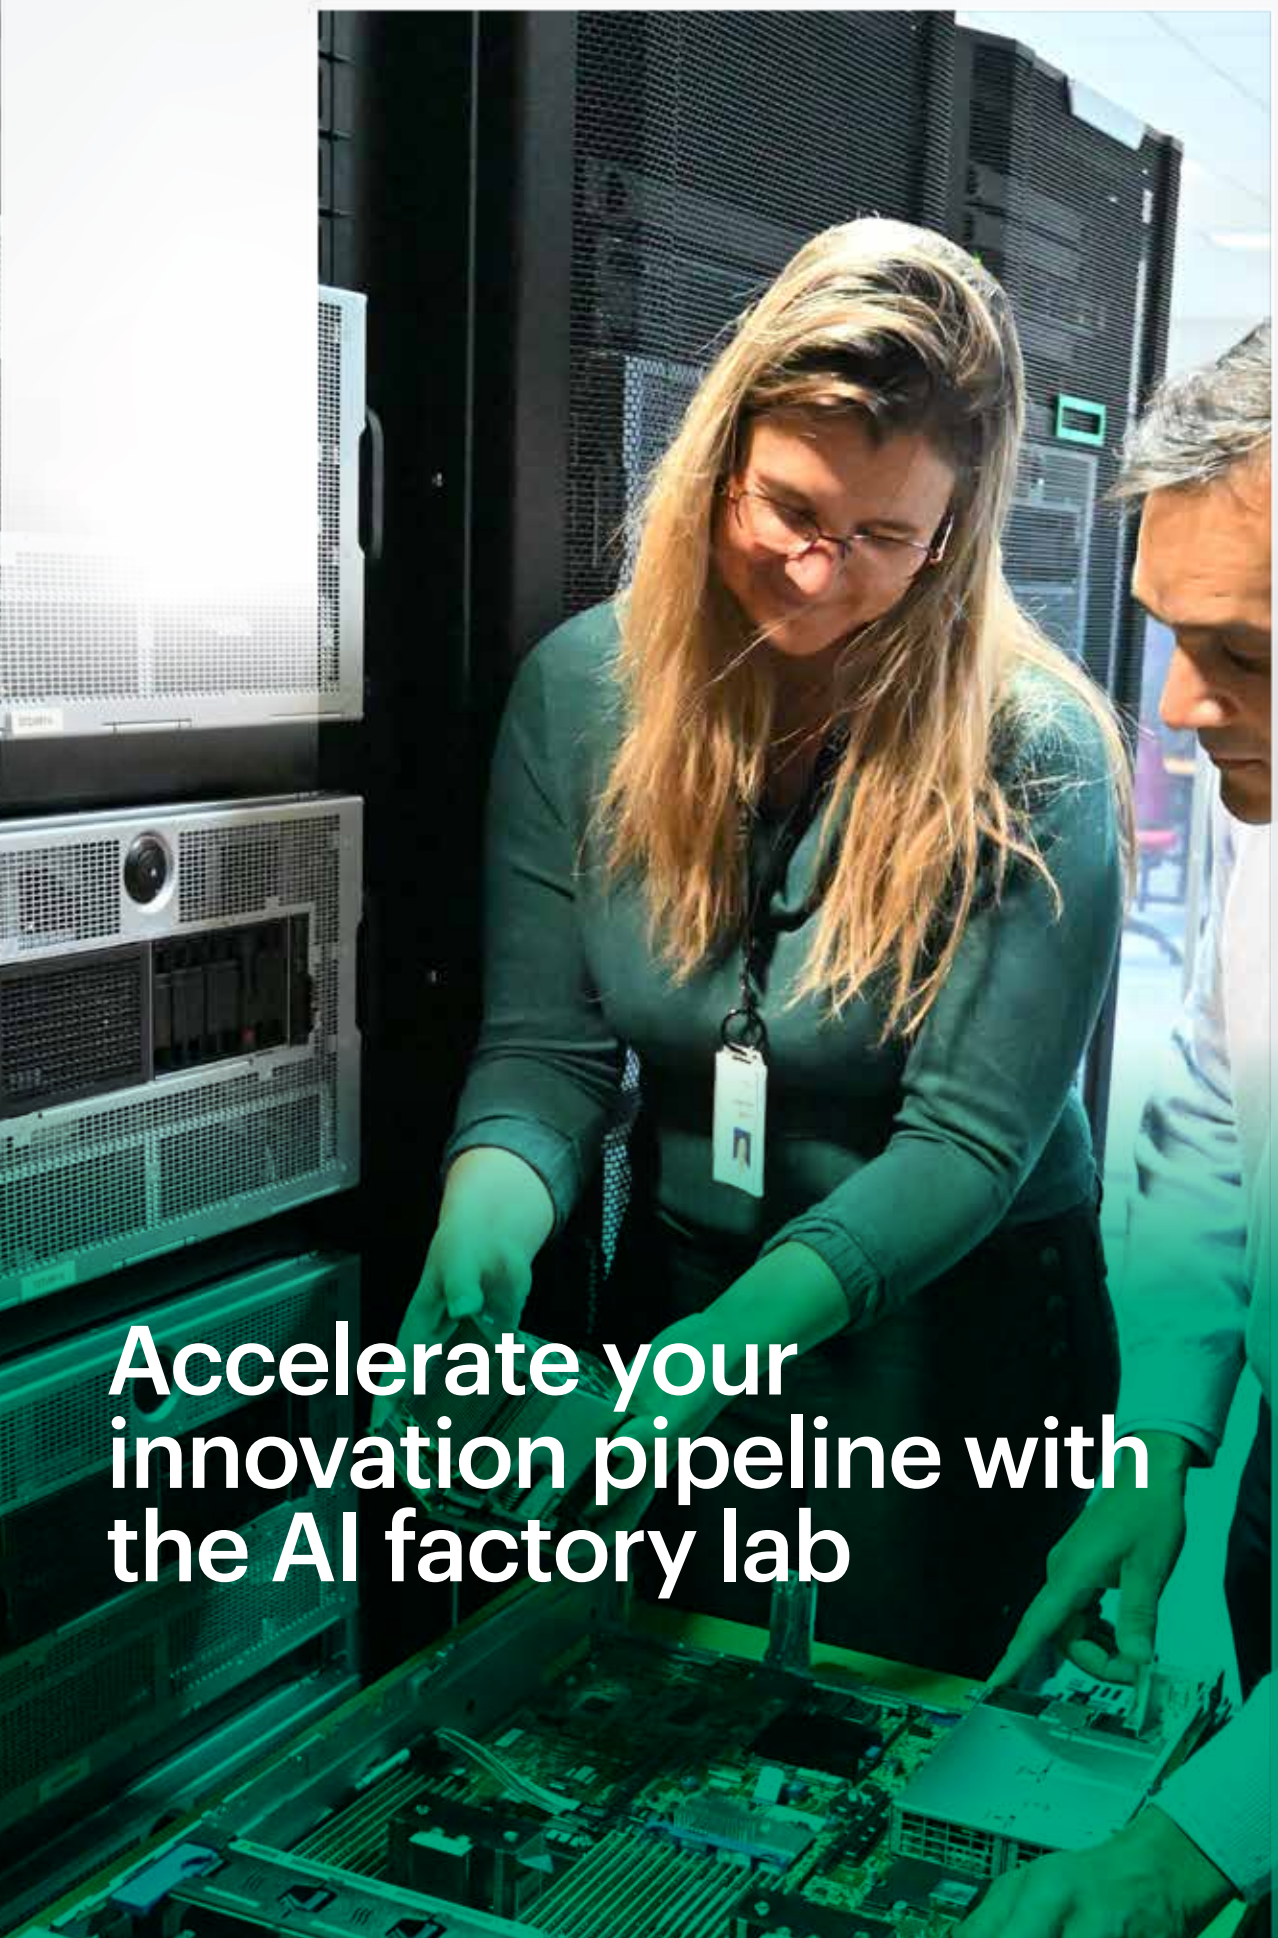

# Accelerate your innovation pipeline with the AI factory lab

# The importance of proving AI

To realize the full promise of artificial intelligence (AI), enterprises must go beyond theoretical models and confront the realities of real-world deployment. The complexities of today's AI workloads demand solutions that are tested and proven in authentic operational conditions.

Relying solely on simulations exposes organizations to performance gaps, integration hurdles and unforeseen risks that can undermine competitiveness and slow time-to-value. Validating AI solutions in environments that mirror actual production is now a business imperative. It enables organizations to fine-tune configurations, verify interoperability across the technology stack, and address scale and security challenges with confidence.

For leaders committed to driving innovation and impact, proving AI in the real world is not just best practice —it's the foundation for sustainable advantage and measurable results.

## $30-40B

**enterprise investment into GenAI.**

## 95%

**of organizations are getting zero return.**

## 5%

**of integrated AI pilots are extracting millions in value, while the vast majority remain stuck with no measurable impact.**

## 5%

**of custom enterprise AI tools reach production.**

MIT, The GenAI Divide: State of AI in Business, June 2025

# Your AI solutions, proven here

The HPE AI factory lab, sponsored by NVIDIA and located in Grenoble, France, provides the definitive environment for proving and scaling AI solutions. It is purpose-built to help European customers and partners accelerate their AI initiatives by providing a space to test applications on engineering-validated configurations.

This facility empowers organizations by providing hands-on access to a high-performance infrastructure featuring HPE Compute XD690 servers with NVIDIA B300 GPUs, HPE ProLiant DL380a servers with NVIDIA RTX PRO™ 6000 GPUs and NVIDIA AI Enterprise software. This allows for rigorous testing, training and co-innovation with HPE experts.

With dedicated support and access to world-class services, organizations can minimize risks associated with AI investments and confidently transition from concept to production. The lab serves as a collaborative hub, enabling customers to fine-tune their environments, create value with HPE AI factory solutions, and map their journey toward sovereign AI capabilities. By providing this proving ground, HPE can help enterprises transform their AI ambitions into tangible, competitive advantages with solutions validated for real-world performance and impact.

**Prove and scale AI solutions**

The lab is purpose-built to help European customers and partners accelerate AI initiatives.

**Minimize risk** associated with AI investments and confidently transition from concept to production.

**Get hands-on access** to high-performance servers with NVIDIA GPUs, networking and software, HPE storage, software and services.

**Security and compliance**

Map your journey toward sovereign AI and transform AI ambitions into competitive advantages.

# Engagement types and outcomes

The lab offers structured engagement models that map to common milestones — demonstration, proof of concept, workshops and transition to production — each with tangible deliverables designed to accelerate decision making of the AI lifecycle for maximum impact.

### Immersive demonstration

- Navigate into a full HPE and NVIDIA AI factory software stack that is purpose-built to deliver reproducible, secure and automatable AI lifecycle management at enterprise scale
- See an end-to-end AI factory scenario running entirely within a European-controlled environment

### Proof of concept

- Bring representative workloads and measure time-to-train, inference latency and throughput
- Receive a validated performance report and recommended architecture

### Workshops

- Multi-day, on-site collaboration with HPE experts
- Co-design and tune pipelines
- Deliver runbooks, automation artifacts and a production roadmap

### Transition to production

- Team trainings and hands-on experience to speed time-to-production
- Validated configuration templates, deployment automation, monitoring dashboards, cost and usage reports

## Champollion and beyond

HPE has a long heritage in Grenoble that includes early investment in AI systems, with more than 300 engineers and over 6,000 square meters of lab presence. There are also researchers in areas including energy, medicine, biotechnology, semiconductors and nanotechnology.

In May 2022, HPE announced a new AI supercomputer, "Champollion," in Grenoble, France, in collaboration with NVIDIA. Champollion is part of HPE's continued investment to HPC & AI technology in France, building on President Emmanuel Macron's national "AI for Humanity" initiative, announced in March 2018, to advance industries with AI and boost economic growth.

In June 2022, Champollion reach the #331 position in the TOP500 and #10 in Green 500 ranking the 500 fastest and energy efficient supercomputers in the world in open, efficient and scalable foundation designed for AI workloads.

## Champollion AI supercomputer

| Component | Configuration |
|---|---|
| Servers | 20x Apollo 6500 Gen10+ |
| Processors | 40x AMD EPYC™ (64 cores @ 2.45-3.5 GHz) |
| GPUs | 160x NVIDIA® A100 GPUs |
| Memory | 20TB DDR4-3200 dual rank |
| Network | 80x NVIDIA InfiniBand HDR (Connect-X6 200 GB/s) |
| Storage | 90TB NVMe local/300 TB Lustre on ClusterStor E1000 |

## AI factory at scale

In early 2026, the Grenoble lab will be enhanced with a brand-new, dedicated AI Factory at-scale system, jointly financed by HPE and NVIDIA. Purpose-built to accelerate customer immersions, the new platform will be equipped with the latest HPE and NVIDIA AI-optimized technologies — next-generation NVIDIA HGX accelerated servers, and high-bandwidth networking — to provide a production-class environment for hands-on validation, large-scale training and governance testing.

This investment expands our sovereign, Grenoble-based infrastructure and enables longer, more representative customer engagements where teams can co-design, validate and operationalize AI Factory solutions tailored to their production requirements.

| Component | Configuration |
|---|---|
| Servers | 8x HPE Compute XD690 |
| Processors | 16x Intel® Xeon® 6776P 2.3 GHz 64-core 350W processors |
| GPUs | 64x NVIDIA HGX B300 288GB |
| Memory | 16TB Dual Rank x4 DDR5-6400 |
| Network | 64x NVIDIA InfiniBand XDR/Ethernet 800GbE |
| Storage | 1.1PB NVMe local/ 3.3PB Lustre on ClusterStor E1000/ 600TB DAOS |

## Software and orchestration stack

The Grenoble AI factory at scale runs a full HPE and NVIDIA AI factory software stack that is purpose-built to deliver reproducible, secure and automatable AI lifecycle management at enterprise scale.

- Harvester provides Kubernetes®-native bare-metal provisioning and lifecycle management for GPU-enabled nodes, enabling secure multi-tenant or dedicated bare-metal clusters without sacrificing performance.
- HPE Morpheus® Enterprise software orchestrates the AI workflows and provides self-service provisioning and governance so teams can request resources, while IT can enforce quotas and maintain compliance across projects.
- Red Hat® Ansible® AWX centralizes automation and configuration management through UI and API-driven playbooks with role-based access control (RBAC) so deployments and changes are scripted, repeatable and auditable.

- Netris automates software-defined networking tasks: Infrastructure inventory and IP address management (IPAM), VLAN assignment and firewall policy automation so network topologies and tenant isolation can be tested repeatedly.
- MinIO® provides scalable, S3-compatible object storage for datasets and model artifacts, enabling realistic storage access patterns for AI training and inference.
- Observability and operational insights are provided by Prometheus® for metrics collection, Grafana® for visualization, and Thanos for long-term metric retention.
- NetBox serves as the authoritative source of infrastructure inventory and IP address management (IPAM), and Exivity® enables cost-management, usage reporting and chargeback/showback.
- Keycloak® centralizes authentication and SSO to validate enterprise identity and RBAC patterns in a sovereign context.

Together, this stack enables customers to validate not only raw model performance, but also governance workflows, access controls, monitoring and cost models that are critical for operationalizing AI at scale.

# What sets HPE apart

### World-class services

HPE delivers comprehensive services that guide organizations through every stage of their AI transformation. This includes expert consulting, deployment and integration, dedicated support, and flexible consumption options, all tailored to de-risk complex projects and accelerate outcomes.

### Edge-to-cloud expertise

By leveraging the power of HPE GreenLake and the proven Edge-to-Cloud Adoption Framework, HPE enables organizations to modernize their IT infrastructure. This holistic approach optimizes hybrid cloud environments and scales AI operations with exceptional agility and cost efficiency.

### Financing and asset management services

HPE Financial Services offers a comprehensive portfolio of flexible financing solutions designed to help organizations with predictable and easy monthly payments, as well as IT asset lifecycle services that help organizations transform and repurpose IT resources into new IT capabilities and innovations.

### Global network of customer innovation centers

Worldwide HPE Customer Innovation Centers, including the new AI Factory Lab in Grenoble, provide an immersive environment for collaboration and discovery. These centers enable customers to test solutions, engage with leading experts, and explore cutting-edge technologies in real-world scenarios—transforming bold ambitions into reality.

"Every customer brings a unique vision and need for AI. At the HPE + NVIDIA AI factory lab, we partner to turn that vision into reality."

—Nathalie Viollet, AI factory lab Manager

### NVIDIA AI Computing by HPE

To accelerate enterprise AI, HPE and NVIDIA joined forces to deliver AI factories integrating NVIDIA's accelerated computing with HPE's secure, scalable infrastructure. The NVIDIA AI Computing by HPE portfolio enables organizations to move beyond experiments and build robust, production-ready AI systems. Solutions include turnkey private cloud AI for fast starts, scalable AI factories for model builders, and sovereign AI factories for specialized data protection. These offerings feature validated technology stacks, multi-tenant support, air-gapped and federated architectures, predictable pricing, and rapid deployment. Fully integrated infrastructure and services allow IT teams to focus on developing AI use cases.

# Accelerate your innovation pipeline

AI is no longer just a technical challenge — it's a strategic imperative. The HPE AI factory lab, sponsored by NVIDIA and located in Grenoble, France, provides the definitive environment for proving and scaling AI solutions. It is purpose-built to help European customers and partners accelerate their AI initiatives by providing a space to test applications on engineering-validated configurations. With dedicated support and access to world-class services, organizations can minimize risks associated with AI investments and confidently transition from concept to production.

## Take the next step

Connect with our experts to discuss your unique challenges, schedule a consultation, or experience these solutions firsthand at the AI factory lab or at an HPE Customer Innovation Center. Let's accelerate your innovation pipeline together!

Visit HPE.com

## Contact

grenoble_ai_lab@hpe.com

Chat now

© Copyright 2026 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.

AMD is a trademark of Advanced Micro Devices, Inc. NVIDIA, NVIDIA HGX, NVIDIA RTX, and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Intel Xeon is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. All third-party marks are property of their respective owners.

a00155546enw

HEWLETT PACKARD ENTERPRISE

hpe.com