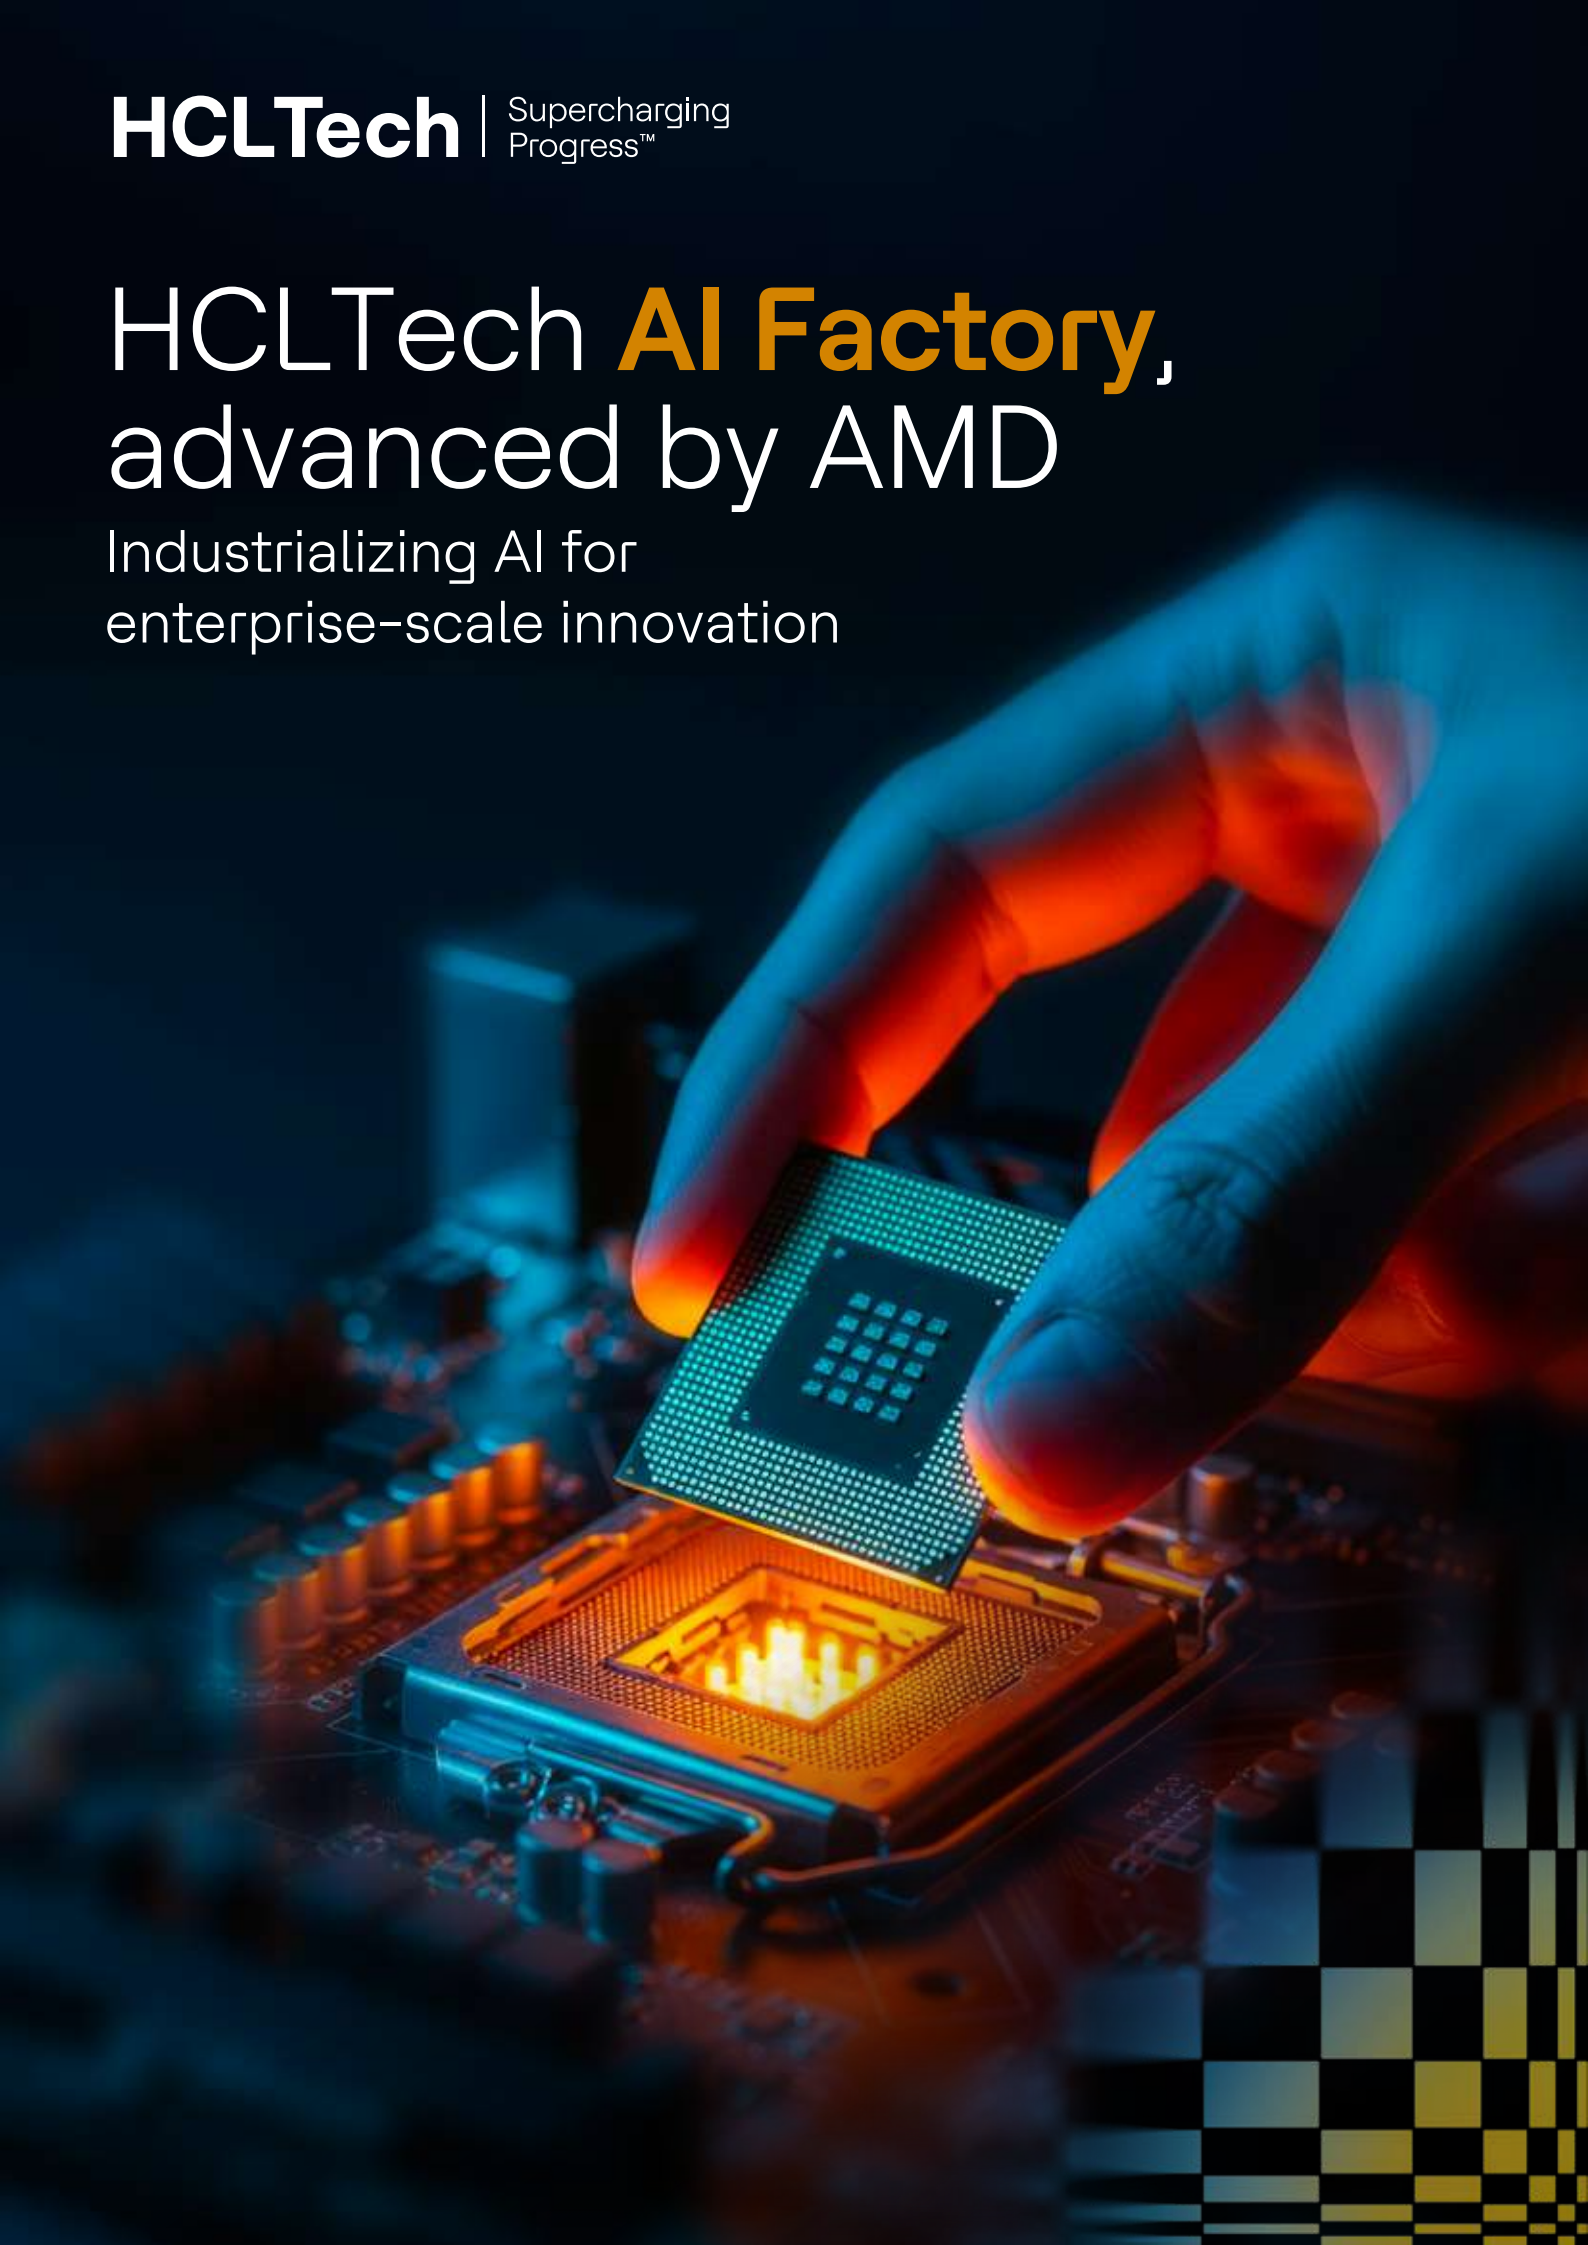
HCLTech | Supercharging Progress™

HCLTech **AI Factory**, advanced by AMD

Industrializing AI for enterprise-scale innovation



Enterprise AI is at an inflection point. Organizations across industries are accelerating investments in AI, with global AI spending expected to exceed $2.5 trillion by 2026 and AI infrastructure emerging as one of the fastest-growing technology investment areas. The AI Factory market alone is projected to reach $780 billion by 2034, reflecting a fundamental shift toward industrialized AI platforms capable of supporting enterprise-scale innovation.

Despite this momentum, the path from AI experimentation to production remains challenging. Fewer than half of AI and GenAI initiatives successfully scale beyond pilot stages, while many organizations struggle to sustain AI programs over the long term. Common barriers include fragmented technology stacks, siloed data environments, escalating infrastructure costs, operational complexity, governance challenges and increasing demands for data privacy, security and regulatory compliance. At the same time, enterprises are reassessing their AI deployment strategies, with growing emphasis on sovereign, private and hybrid AI environments that provide greater control over data, intellectual property and operational risk.

HCLTech AI Factory, advanced by AMD is designed to address these challenges by providing a scalable, open and production-ready AI foundation. Combining AMD Instinct™ GPUs, AMD EPYC™ CPUs, ROCm™ AI software, edge AI platforms and an expanding ecosystem of AI partners, AMD enables organizations to build, train, deploy and scale AI workloads across cloud, on-prem, edge and hybrid environments. Rather than introducing additional complexity, HCLTech AI Factory delivers an integrated AI infrastructure and software platform that helps enterprises accelerate AI adoption, optimize performance and cost, maintain data sovereignty and confidently transition from experimentation to enterprise-wide AI deployment.



What is HCLTech AI Factory, advanced by AMD

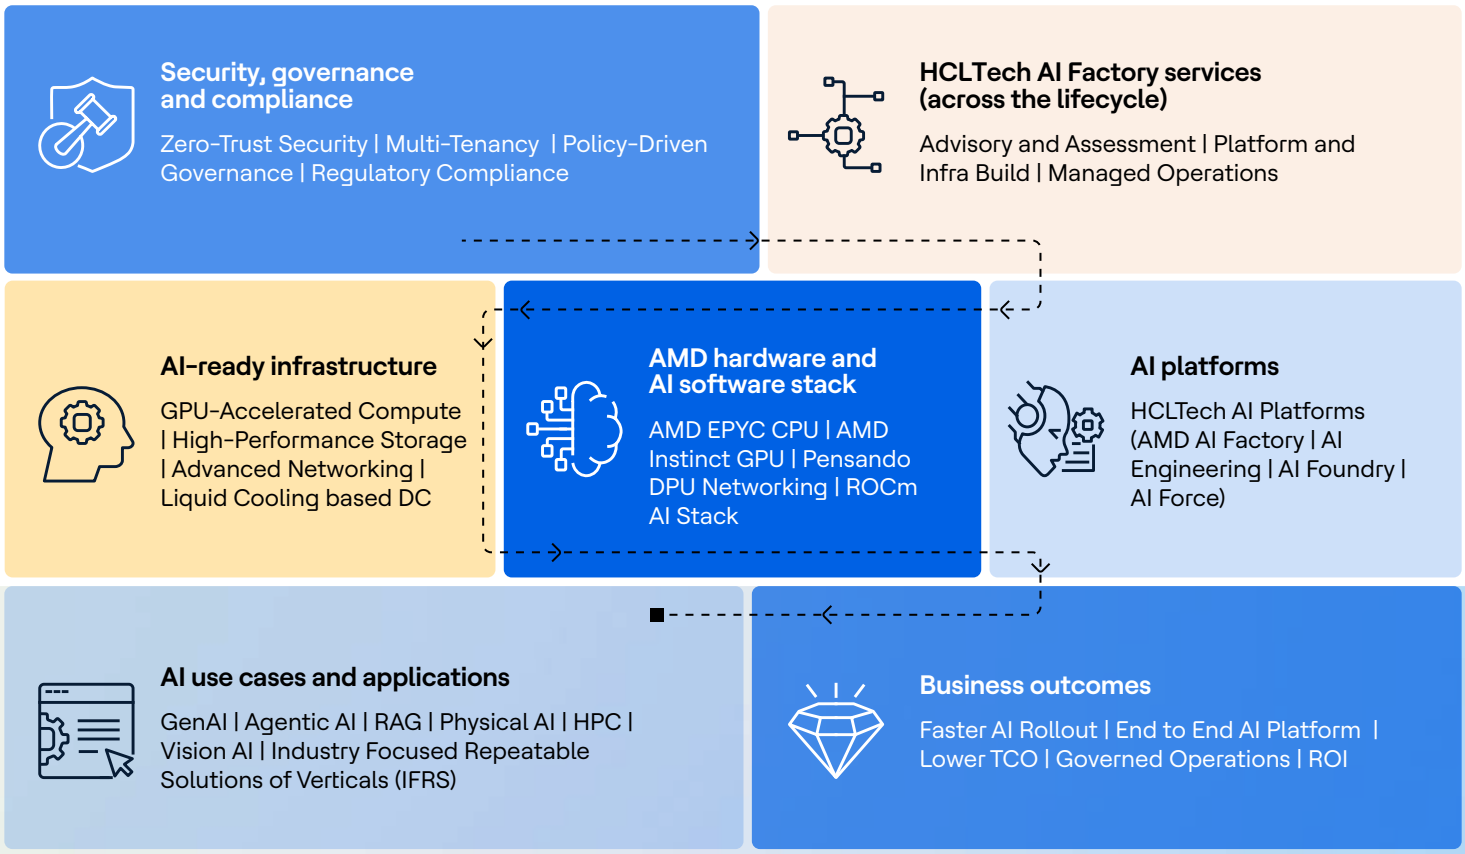
HCLTech AI Factory, advanced by AMD, is a comprehensive enterprise AI platform designed to help organizations move from AI experimentation to scalable, production-ready business outcomes.

Combining AMD's high-performance compute technologies with HCLTech's AI services, accelerators and industry expertise, the platform provides a unified foundation to build, deploy, govern and scale AI workloads with greater speed, flexibility and efficiency.

The solution brings together AMD Instinct™ GPUs, AMD EPYC™ CPUs, ROCm™ software, AI development tools, inference services and open ecosystem platforms within an integrated AI Factory architecture. It supports private, hybrid, sovereign and multicloud environments, enabling enterprises to modernize their AI infrastructure while optimizing performance, cost, security and operational control.

With open standards, flexible deployment models and industry-specific accelerators, HCLTech AI Factory, advanced by AMD empowers enterprises to industrialize AI adoption and turn innovation into measurable business impact.

Solution architecture



What we deliver

AI-ready infrastructure

Deployment of AMD-based AI infrastructure leveraging AMD Instinct GPUs, AMD EPYC CPUs, high-speed networking and enterprise storage optimized for training, inference and fine-tuning workloads.

Open AI software platform

An enterprise-grade AI platform powered by AMD ROCm, RedHat OpenShift AI, Kubeflow, MLFlow, AMD AI Workbench and AMD Inference Services that enables open and portable AI development.

Enterprise AI lifecycle services

Unified services supporting AI model development, fine-tuning, inferencing, monitoring, governance and continuous optimization.

Data and AI engineering

Integrated AI data platform supporting data preparation, vector databases, feature engineering, model pipelines and enterprise data governance.

ROCm migration and optimization services

Assessment, migration, optimization and validation services enabling enterprises to move workloads from proprietary GPU ecosystems to AMD ROCm environments with reduced risk.

Managed AI operations

End-to-end management of AI infrastructure, AI platforms, GPU clusters, model operations, observability and performance optimization.

Core platform capabilities

Capability	Description
ROCm AI Platform	Open software ecosystem supporting PyTorch, TensorFlow, ONNX, Hugging Face and enterprise AI frameworks
AMD AI Workbench	Simplified AI model development and deployment environment
AMD Inference Microservices	Optimized inferencing services for enterprise AI workloads
Resource Management	GPU scheduling, orchestration and workload optimization
Multi-Tenant AI Infrastructure	Shared AI platform with secure isolation and governance
Monitoring and Observability	Infrastructure, GPU, model and application performance visibility
Security and Governance	End-to-end AI governance and compliance management
AI Data Platform	Enterprise data engineering, vector databases and RAG enablement



AI Factory enabled use cases

AI category	Use cases
Generative AI	Enterprise copilots, knowledge assistants, content generation, code assistants
Agentic AI	Autonomous agents, IT operations assistants, customer service agents
HPC and Scientific Computing	Simulation, research computing, molecular modelling
Computer Vision	Quality inspection, medical imaging, safety analytics
Industry AI	Manufacturing, BFSI, Lifesciences and Healthcare, Retail, Telecom and Public Sector solutions

Why enterprises choose HCLTech AI Factory, advanced by AMD



Engineering-led, not tool-led. Built on decades of infrastructure expertise

Built on deep engineering heritage and three decades of experience managing complex, mission-critical IT infrastructure. HCLTech brings the operational rigor and domain depth needed to design, build and operate AI environments for real-world performance and reliability in collaboration with AMD.



Full-stack coverage across the AI lifecycle

Unlike platform-only or services-only approaches, HCLTech AI Factory delivers end-to-end AI lifecycle coverage by combining advisory and assessment, platform and infrastructure build and managed operations into a unified engagement model.



Blueprint-led, validated architectures built for ROI

Follows a blueprint-led operating model using validated reference architectures and standardized design patterns that balance performance, cost, energy efficiency and scalability. Designs are optimized upfront for ROI, not retrofitted later.



Built for private, hybrid and sovereign AI

Enterprises can deploy AI workloads in environments that align with data sovereignty, compliance, latency and security requirements while maintaining control over data and IP.



A rich portfolio of IPs, accelerators and industry frameworks

HCLTech AI Factory is powered by proprietary IPs and accelerators including AI Force, AI Foundry, AI Engineering, AI Cost Optimizer and various others enabling faster deployment and reduced time to value.



Strong ecosystem partnerships and End-to-end ownership

HCLTech works with leading technology partners across infrastructure, cloud, software and AI ecosystems to deliver production-ready AI Factory architectures.

Supports the full AI infrastructure and platform journey, from AI data center design and build to infrastructure support, platform engineering, cognitive operations and managed services.



AI platforms expertise

Dedicated adoption, optimization and migration, framework to accelerate enterprise-scale AMD AI adoption.



Global AI labs and innovation centers

Access to AMD-powered AI Labs, solution accelerators, benchmarks and proof-of-concept environments. Dedicated AI Lab Infrastructure available in HCLTech data center, Austin and shared AI Infrastructure in AMD cloud. Access to a large and growing pool of trained AI, cloud and infrastructure professionals, supported by a global network

Accelerate your AI journey with HCLTech and AMD

Whether you are modernizing AI infrastructure, migrating to open platforms or scaling AI across your organization, HCLTech and AMD are ready to help you accelerate.

- Start with an AI Factory assessment
- Build and run AI at scale with confidence
- Design your enterprise AI operating model

Write to us at **ecosystem.marketing@hcltech.com** to schedule a briefing, request a proof of concept, or explore how we can help you turn AI ambition into measurable business outcomes.

HCLTech | Supercharging Progress™

HCLTech is a global technology company, home to more than 227,000 people across 60 countries, delivering industry-leading capabilities centered around AI, digital, engineering, cloud and software, powered by a broad portfolio of technology services and products. We work with clients across all major verticals, providing industry solutions for Financial Services, Manufacturing, Life Sciences and Healthcare, High Tech, Semiconductor, Telecom and Media, Retail and CPG, Mobility and Public Services. Consolidated revenues as of 12 months ending March 2026 totaled $14.7 billion. To learn how we can supercharge progress for you, visit hcltech.com.

hcltech.com