

HCLTech |  Red Hat

AI Factory powered by Red Hat

Engineering enterprise AI with
HCLTech and Red Hat® AI



From proof of concept to production scale

As AI adoption accelerates, enterprises are shifting their focus from developing individual models to establishing an AI operating foundation that delivers governance, performance, scalability and cost efficiency across the enterprise.

To realize long-term value, organizations need a consistent way to build, deploy, manage, run and continuously evolve AI workloads across hybrid environments while supporting emerging technologies such as Agentic and Physical AI.



HCLTech AI Factory powered by Red Hat® AI



HCLTech AI Factory combines HCLTech's engineering-led AI Factory model with Red Hat® AI's open platform to help enterprises operationalize AI at scale.

It brings together the infrastructure, platform, data and services needed to support the complete AI lifecycle. The solution enables organizations to build, deploy, manage and scale AI across hyperscalers, neo-cloud, on-prem, hybrid, edge and sovereign environments.

The solution simplifies AI lifecycle management, accelerates model deployment, enables efficient inference and integrates private enterprise data to support trusted AI outcomes. With an open, scalable architecture and built-in governance, it helps enterprises operationalize predictive AI, GenAI and Agentic AI while reducing operational complexity and optimizing infrastructure utilization.

Services under AI Factory

What we deliver

HCLTech AI Factory combines AI-ready infrastructure, platforms, models and operations to help enterprises build, operate and scale AI with confidence. Instead of isolated AI initiatives, it delivers a repeatable, industrialized approach to enterprise AI.

Service 1: AI data center design and build	Service 2: AI data center support	Service 3: Cognitive Infrastructure Services
Scope: End to end facility design and build covering power, cooling and space up to the raise floor	Scope: Physical management of racks and IT devices	Scope: An IT infrastructure and platform lifecycle management
Services: • DC consulting and advisory • DC engineering, design and procurement • MEPF and control systems build and project management • Testing, commissioning and ORAT • Digital twin services	Services: • DC facility operations • DC sustainability management • Technical/operational audits • Rack assembly and integration • Break-fix and repair services • Fiber/optics and OSP support • 24*7 onsite support • Inventory management/RMA • Robotics and physical AI	Services: • AI Factory deployment across private, public and cloud • Edge AI • Sovereign AI • AI lab as a service • Developer workspace as a service • Managed AI platform as a service • FinOps for AI • Inference as a service • Model as a service



How does Red Hat® AI help scale AI?

Red Hat® AI enables organizations to manage and monitor predictive and generative AI models across environments, from single-server deployments to large-scale distributed platforms. It also provides a framework-agnostic platform to securely operationalize AI agents with built-in governance, security, observability and tracing.

Powered by open-source technologies and a strong partner ecosystem, Red Hat® AI bridges development and production while delivering enterprise-grade performance, stability and GPU support.

Red Hat® AI helps organizations:

- Manage AI development, training and tuning for predictive and GenAI
- Deploy, manage and scale AI inference, agentic workflows and AI applications
- Build AI applications on a centralized Red Hat OpenShift-based Kubernetes platform
- Deliver end-to-end security and governance across the AI lifecycle
- Optimize resource utilization for model training and inference

AI Factory-enabled use cases:

AI Factory	Representative use cases
Physical and Industrial AI	Digital twins, robotics, vision AI, edge AI, smart assets, fleet and infrastructure management
Industry-specific AI Factories	Manufacturing, energy, logistics, retail, healthcare, financial services, regulated and sovereign environments
AI Force	Application modernization, code engineering, developer productivity, AIOps, data lifecycle management, code optimization, engineering lifecycle automation
AIoT	Predictive maintenance, quality inspection, safety compliance, intrusion detection, performance optimization
GenAI	Knowledge assistants, copilots, intelligent automation, content creation, code generation, customer experience, synthetic data generation
Agentic AI	Digital assistants, customer support, IT Ops and DevOps agents, supply chain automation, HR and finance workflows, enterprise copilots
HPC	Simulation, scientific computing, engineering, rendering and visualization

Why enterprises choose HCLTech AI Factory

1. Engineering-led foundation

Built on three decades of infrastructure engineering expertise, HCLTech combines its engineering-first approach with Red Hat® AI to deliver secure lifecycle management, governed AI operations and optimized resource utilization across hybrid cloud environments.

2. Blueprint-led architectures

Validated reference architectures and standardized design patterns optimize performance, cost, energy efficiency, scalability and ROI from the outset.

3. Full-stack AI lifecycle coverage

HCLTech delivers end-to-end AI lifecycle capabilities spanning AI data center design, operations, infrastructure, cloud, data, platforms, applications and AI models through its IP-led solutions.

4. Built for sustainable scale

Designed for high-density GPU environments with integrated governance, compliance, energy efficiency, liquid cooling readiness, sovereignty and Responsible AI practices.

5. Comprehensive AI ecosystem

Deep partnerships across research institutions, OEMs, infrastructure providers, hyperscalers, platform vendors and AI model providers accelerate AI adoption while optimizing outcomes.

6. Production-grade AI operations

Industrialized AI operations include proactive monitoring, GPU orchestration, observability, security, FinOps and rapid remediation to support always-on production AI.



7. Proven enterprise delivery

Backed by global AI and Cloud Native Labs, a large pool of AI, cloud and infrastructure experts and consistent recognition from Gartner, IDC, ISG, Everest Group, Avasant and others, HCLTech helps enterprises design, operate and scale AI with confidence.

Your AI Factory journey starts here!

Whether you are:

Defining your AI Factory strategy

Scaling proven AI use cases across the enterprise

Preparing for the next generation of Agentic and Physical AI

HCLTech is ready to help you:

Assess AI Factory readiness

Build and operate AI at scale with confidence

Design your enterprise AI operating model

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

HCLTech | Supercharging
Progress™

hcltech.com