

A vertical photograph on the left side of the page shows a dense urban skyline at dusk. Two prominent skyscrapers are visible: one with a distinctive, illuminated, triangular glass facade and another with a more traditional rectangular glass facade. The buildings are lit up from within, and their lights reflect on the surrounding structures and the sky. The overall scene is a high-angle, aerial view of a city center.

Achieve AI outcomes faster with Kyndryl AI Private Cloud and Red Hat® OpenShift® AI

Highlights

AI is rapidly transforming business operations, bringing organizations to a critical inflection point. At the same time, shifting geopolitics and concerns around data sovereignty are causing organizations to reconsider their approaches to AI. There are still major potential obstacles to AI adoption: legacy systems, heavily regulated data, silos of information, uncertainty about public cloud providers, AI governance challenges and the lack of clear ROI projections.

The partnership between Kyndryl and Red Hat enables organizations to choose a scalable, secure, highly customizable AI platform that addresses these challenges and accelerates AI adoption. With a combined Kyndryl and Red Hat solution, organizations can achieve innovation through AI while delivering measurable benefits such as enhanced customer experience.

Introduction

Organizations seeking to adopt AI commonly face several obstacles, including skills shortages, the inaccessibility of relevant, high-quality data, difficulties in integrating with existing systems and the lack of a coherent strategy.¹

In partnership with Red Hat, Kyndryl empowers organizations to overcome these challenges and meet their business objectives by deploying intelligent, scalable and secure AI platforms built on Red Hat OpenShift AI.

Kyndryl and Red Hat solutions for AI

Kyndryl AI services cover the entire AI lifecycle — from strategy and prototyping to deployment and daily operations. Supplementing advanced AI and data capabilities from Red Hat, Kyndryl brings deep expertise in AI infrastructure, governance and managed services. The result is joint solutions that are perfectly adapted to the business use cases of each organization.

Wherever you are on the path to AI adoption, Kyndryl and Red Hat can help you to conquer your challenges and achieve your business goals while respecting any relevant regulatory frameworks.

The Kyndryl and Red Hat partnership in AI helps customers with:

1. Developing AI strategies and roadmaps
2. Rapidly prototyping, gaining business validation and prioritizing use cases
3. Modernizing data pipelines and optimizing data governance
4. Designing and deploying infrastructure
5. Managing and optimizing AI operations


To help organizations achieve and sustain quantifiable business gains, Kyndryl and Red Hat have partnered to create solutions based on the Red Hat OpenShift AI platform — the Red Hat flagship product for training, tuning, running and managing enterprise-grade AI and generative AI (GenAI) applications.

OpenShift AI integrates the “Open Data Hub” open-source project to provide access to the most popular open-source tools for developing AI and GenAI models. It also includes an optimized inference server to run LLM applications, plus several different tools to monitor, manage and govern both AI and GenAI applications.

Running on Red Hat OpenShift Container Platform, OpenShift AI provides an extremely reliable and scalable platform for running the most critical and demanding AI and GenAI workloads.

The ability of OpenShift AI to run on any private or public cloud infrastructure, together with its support for multiple GPU providers — including NVIDIA, Intel and AMD — maximizes flexibility when choosing infrastructure or hardware vendors and helps organizations avoid lock-in.

With Kyndryl AI Private Cloud and Red Hat OpenShift AI providing the solid foundations for enterprise-scale AI transformation, organizations can rapidly achieve positive business outcomes:





Benefits

Day-zero assist and optimization:

Support for generative AI models developed by Meta (Llama family), IBM (Granite family) and other frontier model providers.

Support all major AI and GenAI use cases:

Provide training, tuning, inference and operational services, and support AI/GenAI models for the most common cross-industry and industry-specific scenarios.

Cut time-to-value:

Microservices architecture significantly decreases the time required to develop and deploy scalable, resilient and adaptable AI and GenAI applications.

Enhance security and compliance:

Mission-critical data remains securely in-house and under control, reducing the risk of data leakages. Built-in encryption, role-based access, audit logging and compliance enable adherence to data sovereignty and regulatory frameworks such as GDPR, HIPAA, and CCPA.

Boost developer productivity:

A catalog of ready-to-use tools accelerates the development, training, tuning and running of AI and GenAI applications.

Simplify IT operations:

Built-in monitoring of models enables organizations to gain deep understanding of the behavior and performance of AI and GenAI models. Built-in MLOps capabilities help automate model tuning, inference and validation loops.

Meet the most demanding Service Level Objectives:

OpenShift Container platform optimizes SLOs for multi-tenant, large-scale and mission-critical AI and GenAI application deployments.

Avoid vendor lock-in:

Open-source frameworks and tools maximize freedom of choice, empowering organizations to build and run AI and GenAI applications everywhere — on private or public clouds — with zero code changes.

Optimize infrastructure costs:

Ability to select the most appropriate infrastructure from different hardware vendors and cloud providers, and to choose between managed/self-managed solutions and Capex/Opex/pay-per-use models.

Diversify revenue streams:

Transform data centers into NeoClouds, enabling GPU-as-a-Service or Token-as-a-Service business offerings.

Kyndryl capabilities

Kyndryl helps organizations deploy and consume AI faster through expert advice, implementation and ongoing management, supporting rapid progress from prototyping to production, with continuous delivery of value along the way.

1. Advise/Consult - Prepare

- Identify and prioritize business use cases
- Define success criteria and plan AI proof of concepts (PoCs)
- Size and design AI infrastructure

2. Implement - PoC to Production

- Implement, run and evaluate PoCs
- Support private cloud implementation, providing deployment models, AI infrastructure models and readiness, lifecycle operations (provisioning, monitoring, etc.), security, capacity and performance
- Implement MLOps/LLMOps pipelines and deploy applications
- Create production-ready, highly integrated AI private cloud as standalone or in a hybrid environment

3. Manage - Production

- Run fully managed AI private cloud services
- Deliver managed services across lifecycle operations, provisioning, monitoring, incident management and patching
- Perform capacity monitoring and planning
- Integrate AI with existing systems and knowledge bases
- Schedule and orchestrate production AI workload at scale
- Provide cloud-adjacent control plane for hybrid cloud capabilities
- Automate FinOps reporting and analysis for cost optimization
- Implement GPUs-as-a-service or tokens as-a-service and B2B models to grow revenue streams with external clients

Why Kyndryl?

Kyndryl has deep expertise in designing, running, and managing the most modern, efficient and reliable technology infrastructure that the world depends on every day. We design, build, manage and modernize the world's critical technology systems. As the world's largest provider of IT infrastructure services, Kyndryl is committed to the health and continuous improvement of the vital systems at the heart of the digital economy. With our partners and thousands of customers worldwide, we co-create solutions to help enterprises reach their peak digital performance.

Next steps

To find out how you can adopt AI faster with Kyndryl AI Private Cloud and Red Hat, please contact your Kyndryl representative or [schedule a complimentary consultation with a Kyndryl expert](#). Learn more about the [Kyndryl and Red Hat alliance](#) or visit us at kyndryl.com.

kyndryl.

© Copyright Kyndryl, Inc. 2025

Kyndryl is a trademark or registered trademark of Kyndryl Inc. in the United States and/or other countries. Other product and service names may be trademarks of Kyndryl Inc. or other companies.

This document is current as of the initial date of publication and may be changed by Kyndryl at any time without notice. Not all offerings are available in every country in which Kyndryl operates. Kyndryl products and services are warranted according to the terms and conditions of the agreements under which they are provided.

1. [11 Barriers To Effective AI Adoption And How To Overcome Them](#), Forbes, May 2024.