

Unlock business potential with Secure AI for Regulated Industries

The challenges of AI adoption

While generative AI has the potential to deliver vital efficiencies and insights for organisations in highly regulated sectors like banking and government, it can also be accompanied by some significant challenges:

- **Lack of company and industry context** – Generic AI models often lack the necessary context specific to individual companies and industries. Without this context, models may generate inaccurate or irrelevant results.
- **Data privacy and security concerns** – Regulated industries must ensure end-to-end security, encryption, and access controls to protect sensitive data and comply with privacy regulations.
- **Ownership of intellectual property** – Clear intellectual property rights are essential for organizations investing in AI, fostering trust and encouraging adoption of innovative solutions.
- **Lack of talented data scientists** – Organizations face challenges adopting Generative AI due to a shortage of skilled professionals with the necessary expertise.

Gen AI Challenges - Enterprise h2oGPTE


	Use proprietary Closed Models	Exclusive LLM using Open LLM
		h2oGPTE
Data Privacy and Security	Data is shared on cloud, Leakage Risk	No data (query, prompts, response) is shared outside
AI Governance	Limited Transparency	Full visibility and transparency
Customization Cost of Scalability	Limited customization High enterprise adoption cost	Fully customizable Fixed Cost
Access and Availability	Downtime Risk	Control over Downtime
Ownership	Owned by OpenAI / Google	Owned by enterprise
Replicability	Need to start from scratch	Replicable for other groups

Figure 1. Challenges associated with embracing closed AI models from proprietary vendors



Accelerate your AI ambitions with Secure AI for Regulated Industries

H2O.ai, Dell, and Red Hat have launched Secure AI for Regulated Industries, a comprehensive solution for secure generative AI tailored for regulated sectors. Unlike typical public cloud offerings, it ensures data sovereignty within controlled environments, allowing organizations to leverage internal data aligned with their business models.

Built for on-premises deployments, it supports unlimited text generation, custom AI application development via a low-code Python stack, and robust data privacy and security.

With support from H2O.ai's expert data scientists, this platform accelerates AI initiatives through straightforward API integrations and an intuitive user interface.

Benefits of Secure AI for Regulated Industries

- **Contextual Insights from Internal Company Documents** – H2O.ai utilizes internal documents and sensitive data to enhance AI models to better understand your business context.
- **Fortified Data Privacy and Security** – Regulated industries must ensure end-to-end security, encryption, and access controls to protect sensitive data and comply with privacy regulations.
- **Open-Source Models and Ownership** – Uses open-source algorithms for transparency, enabling easier explanations of predictions, while organizations retain ownership and avoid vendor lock-in.
- **Accelerated Adoption and Expertise** – The intuitive interface facilitates collaboration between data scientists and business analysts, while embedded expert guidance speeds up AI implementation.

Benefits of DELL Technologies

Dell Technologies accelerates your journey from possible to proven by leveraging innovative technologies, a comprehensive suite of professional services, and an extensive network of partners. Dell Technologies helps customers to bring AI to their data. Dell Technologies solution approach is focused on:

- **Simplified** – Speed up time-to-results by combining strategic guidance, validated solutions and services.
- **Trusted** – Get the most value of your data with infrastructure designed for your business needs.
- **Tailored** – Get the most value of your data with infrastructure designed for your business needs.

Dell Technologies' AI-ready portfolio Deliver the best AI performance* and simplify sourcing, deployment, and management of AI infrastructure with cutting edge technology designed for the Generative AI era.



Red Hat OpenShift Benefits

Red Hat OpenShift offers customers a powerful application platform that unlocks the power of technologies like containers and Kubernetes while reducing or even hiding these technologies' complexities so customers can focus on what matters most, the building, deploying, and running of applications at scale.

Three essential attributes help our customers choose Red Hat OpenShift:

- **Trusted** – Red Hat OpenShift ensures customers' security, resiliency, and other policies are always adhered to no matter where the application is deployed.
- **Comprehensive (integral)** – Red Hat OpenShift includes and integrates all the capabilities that are needed for developers, IT/Cloud operations, and security teams to manage applications.
- **Consistent** – Red Hat OpenShift helps customers manage, maintain, and build applications in a unified platform, supporting a broad portfolio of developer tool sets across hybrid environments.

Solution architecture: overview

The joint solution combines several pieces of industry-leading technology to accelerate the power of generative AI: H2O.ai's generative AI platform; the Red Hat OpenShift Container Platform; DELL's infrastructure stack; Nvidia graphic cards and Intel's computing power. The solution offers the ultimate in flexibility, with the ability to connect to multiple data lakes across your organisation to train and generate desired outcomes across airgapped or connected environments, on premises, in the public/private cloud or at the edge.

Overview of H2O AI Platform

200+ Data Sources



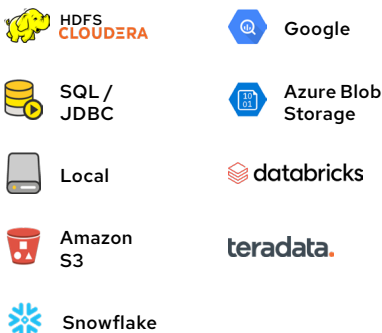
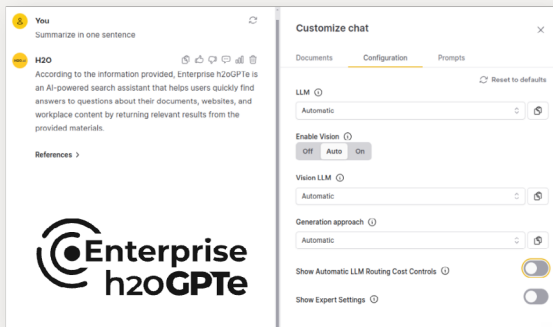


Figure 2. Overview of the Secure AI for Regulated Industries platform and key components



H2O.ai

Enterprise Multimodal RAG



Automatic Model Routing Based on Price and Performance



Enterprise h2oGPTe Key Features

H2O Agentic AI assistant leverages a code interpreter and tool calling to solve reasoning and research queries quickly. It provides data-driven, grounded insights by automatically generating code and utilizing tools such as web search, math, machine learning, data science, Python, and image generation. Agents review and refine their work until they deliver the correct solution, along with the full source code.

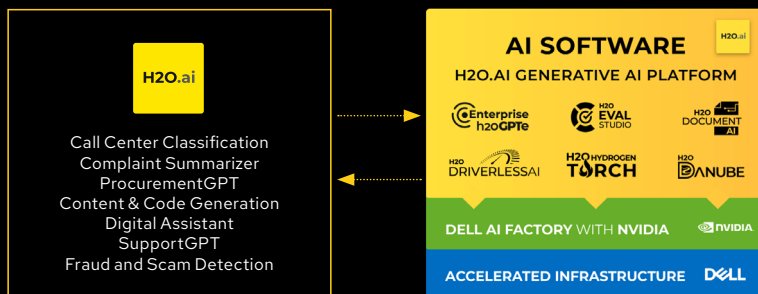
Cost/Performance based Auto-LLM Routing Automatically selects the optimal LLM for the job from all the configured models, including H2O Danube, Llama, InternVL, Mistral, OpenAI, Anthropic, Azure, Google, Amazon, Mistral, InternVL, HuggingFace, etc. Compare all real-time model performance stats graphically.

Multi-Modal (Text, Audio, Vision), Unlimited Context Support for retrieval (asking questions), extraction (guided generation of JSON adhering to your schema), and summarization of unlimited context with multi-modal capabilities, allowing Enterprise h2oGPTe to solve virtually any use case.

H2O.ai on Dell AI Factory

COST-EFFICIENT FULL STACK AI PLATFORM FOR
GENERATIVE AND PREDICTIVE AI

Leverage private, protected data to fine-tune or retrieve-augment-generate (RAG) domain-specific LLMs for multi-agent intelligent systems



On-Premise and Airgapped Generative and Predictive AI
for Enterprise and Public Sector Customers

Red Hat | Dell | H2O.ai Ecosystem Solution Brief

Accelerate Time to Production for AI Use Cases and Reduce TCO

compared to public cloud. Reduce the engineering time required to integrate technologies needed to enable data scientists and app developers to test and deploy AI use cases at scale.

On Premise AI

83% of enterprises plan to move workloads away from public clouds to private clouds and on-premises.

Secure Your Data In-House

Data privacy is non-negotiable.

Built-in Enterprise Multimodal RAG


Skip the complexity and risk of building your own RAG. Get robust Multimodal RAG with Enterprise h2oGPTe.

Cut Cloud Costs

Combat uncontrolled cloud costs and unpredictable expenses.

Cross-industry success

Secure AI for Regulated Industries has helped a wide range of organisations to be more productive, reduce risks and financial losses, and increase operational efficiency.



Financial Services

Wholesale / Commercial Banking

- Know your customers (KYC)
- Anti-Money Laundering (AML)

Card / Payments Business

- Transaction Frauds
- Collusion fraud
- Real-time targeting
- Credit risk scoring
- In-context promotion

Card / Payments Business

- Deposit fraud
- Customer churn prediction
- Auto-loan




Healthcare & Life Sciences

- Early cancer detection
- Product recommendations
- Personalized prescription matching
- Medical claim fraud detection
- Flue season prediction
- ER & hospital management
- Remote patient monitoring
- Medical test predictions




Telecom

- Predictive maintenance
- Avoidable truck-rolls
- Customer churn prediction
- Improved customer viewing experience
- Master data management
- In-context promotions
- Intelligent as placements
- Personalized program recommendations



Marketing & Retail

- Funnel predictions
- Personalized ads
- Fraud detection
- Next best offer
- Next best action
- Customer segmentation
- Customer churn
- Customer recommendations
- Ad predictions and fraud



Manufacturing

Predictive Maintenance

- Machine failure prediction
- Pre-emptive maintenance
- Fault detection
- Control Value predictions
- Early warning systems

Defect Detection

- Wafer defect detection
- Surface crack detection
- Coke prediction

Supply Chain

- Stock level demand sensing
- Demand forecasting
- Inventory Optimization



Keen to know more?

Connect to any of these partners

H2O team

Shadus Ong | shadus.ong@h2o.ai

DELL team

Sumash Singh | sumash.singh1@dell.com


Red Hat team

Irving Yap | irving.yap@redhat.com

About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. A trusted adviser to the Fortune 500, Red Hat provides award-winning support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

 facebook.com/redhatinc

 [@RedHat](https://twitter.com/RedHat)

 linkedin.com/company/red-hat

North America
1-888-REDHAT1
www.redhat.com

Europe, Middle East, and Africa
00800 7334 2835
europa@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com