carahsoft

CLOUDERA

Deploying AI in Government Without Compromising Control

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies through GSA, SEWP, ITES SW2 and a wide range of other contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Cloudera, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.

To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com

Explore More Resources:
carah.io/cloudera-resources

Join Events & Webinars:
carah.io/clouderaevents

Discover Technology Solutions:
carah.io/cloudera_solutions

Learn About Procurement:
carah.io/Cloudera-contracts

Connect With Our Team:
Cloudera@carahsoft.com
(571) 590-7100

CLOUDERA

SOLUTION BRIEF

Deploying AI in Government Without Compromising Control



The Challenge

Most AI delivery models are not designed for environments that require strict control over data, access, and operations. Many platforms depend on external APIs for model access, which means organizational data is sent outside the environment for processing. At the same time, organizations often have limited visibility into how outputs are generated and face unpredictable, consumption-based pricing models.

For organizations handling sensitive or regulated data, these limitations create immediate operational and compliance concerns. Teams may lose visibility into where data is processed, struggle to fully audit model behavior and decisions, or become exposed to external outages and policy changes beyond their control. Aligning these approaches with internal security and compliance requirements can become difficult, particularly in highly regulated or mission-critical environments.

Because of these risks, many organizations restrict AI to low-risk use cases or avoid adoption altogether.

The Cloudera Approach

Cloudera takes a different approach by bringing AI to your data rather than sending data out.

Your models of choice are developed, managed, and deployed entirely within private infrastructure, including on-premises environments, private clouds, hybrid cloud architectures, and air-gapped environments with no external connectivity.

All inference can also happen locally. Prompts, inputs, and outputs remain within the environment. Access to models is provided through secure internal endpoints, where identity, policy, and usage controls are enforced.

The result is an approach where AI operates within the same architecture, security model, and governance framework already in place.

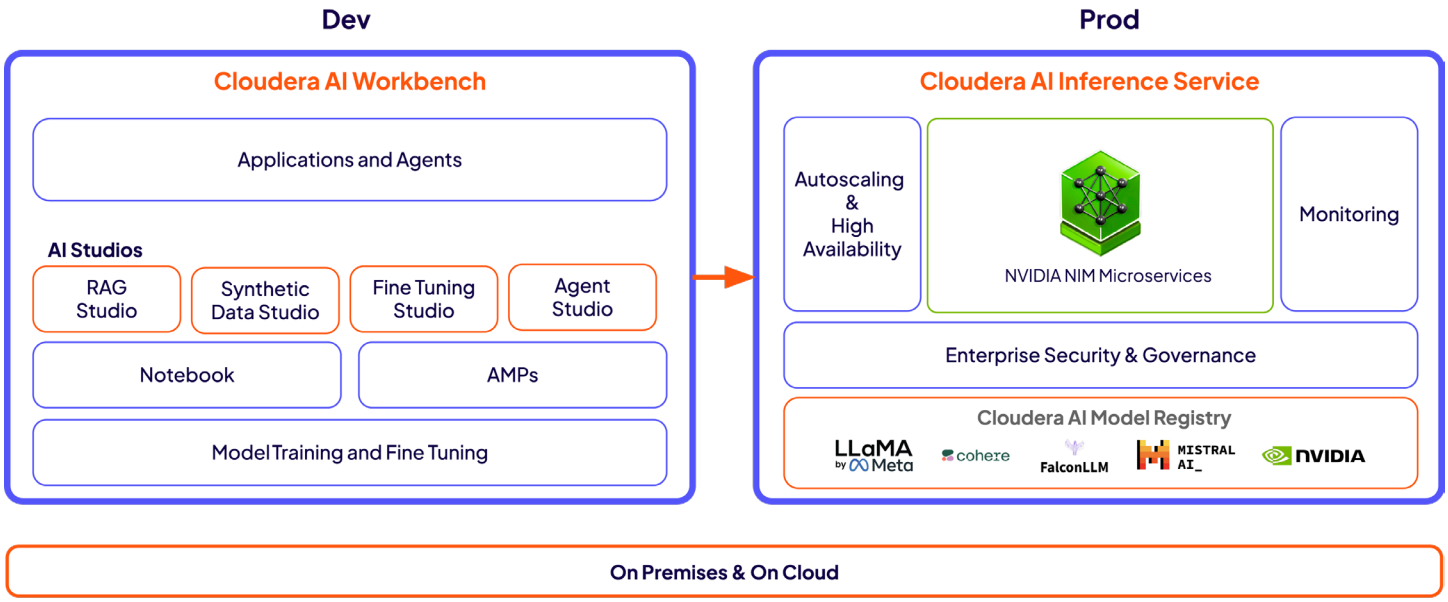
At a Glance

Most AI solutions require sending your data outside your environment. For organizations handling sensitive or regulated data, that is not acceptable.

Cloudera enables organizations to develop, train, deploy, and operate AI entirely within their own environments, enabling them to leverage AI capabilities without exposing sensitive data or introducing external dependencies.

- Models run inside your environment
- Data never leaves your boundary
- You gain full visibility into all AI activity
- The platform integrates with existing security and governance frameworks
- Costs are predictable and based on infrastructure, not usage
- Deploy in any environment, including on-premises data centers and multiple clouds

Cloudera AI: Build and Run Private AI



Cloudera provides an end-to-end platform for AI workflows that keeps development, deployment, and usage inside your environment.

The figure above shows how AI operates as part of the environment rather than outside it.

- Models are developed and prepared in Cloudera AI Workbench
- Models are managed and versioned through a centralized registry
- Inference runs locally through the AI inference service
- Access is controlled through a secure internal gateway
- Monitoring and scaling are built into the platform
- Applications and agents interact through internal endpoints only



Maintain Control of Your Data and AI

Data stays inside your environment. All data remains inside the organization at every stage. There is no external transmission of prompts, inputs, or enterprise data, and there is no sharing of outputs with third-party providers. This keeps the organization fully aligned with data residency and data handling requirements.

You can see and trace everything. Every AI interaction can be tracked and reconstructed. User identity is tied to each request, the model version and execution context are captured, and inputs and outputs are logged based on policy. Usage and system performance can also be continuously monitored. This supports compliance, oversight, and operational visibility.

Usage is controlled, not open-ended. AI usage is governed the same way as data and application access is. Access is based on existing identity and group structures, and model availability is restricted by policy. Requests are routed through controlled endpoints, and usage can be managed and constrained across teams. Organizations define who uses AI, how it is used, and where it is applied.

It scales and stays available. AI operates reliably under real workloads. Autoscaling supports changing demand, high availability ensures consistent access, and there is no dependency on public APIs or external uptime. This also aligns AI operations with continuity and recovery strategies.

Costs stay predictable. AI runs on infrastructure you control. There is no per-request or token-based billing, costs scale with planned compute resources, and there is no exposure to external pricing changes. Usage can also be governed to prevent uncontrolled consumption.

You are not locked into one model. Organizations maintain flexibility over which models are deployed. This includes support for open and commercial models, a centralized model registry for version control, and the flexibility to adapt as technology evolves.

From Capability to Application

Running a model is not the goal. Using AI to improve outcomes is. Cloudera helps organizations move from AI experimentation to production use with enterprise AI capabilities that are secure, governed, and connected to the data where it already lives.

Through Cloudera's AI Studios, organizations can apply AI to real business and mission use cases:

RAG Studio enables secure knowledge assistants using internal documents and datasets. This allows users to ask questions and retrieve answers from trusted enterprise information while keeping organizational knowledge protected and governed.

Fine Tuning Studio supports domain-specific model refinement. Organizations can improve model accuracy and relevance by adapting models to their specific language, data, workflows, and mission needs.

Synthetic Data Studio enables synthetic data generation for testing, development, and training. This helps teams build and validate AI applications without exposing sensitive, private, or regulated data.

Agentic AI Studio supports controlled AI-driven workflows for multi-step processes. It helps organizations coordinate tasks, tools, and decisions in a governed way, allowing AI to support more complex business and mission operations.

Together, these capabilities help organizations move beyond running models in isolation and apply AI securely, responsibly, and effectively across enterprise data.

Agentic AI Studio supports controlled AI-driven workflows for multi-step processes, allowing organizations to coordinate tasks, tools, and decisions in a governed way.



Integration with the Data Platform

Cloudera integrates AI directly into the enterprise data platform, enabling models to operate on trusted, governed data rather than disconnected data silos. Teams gain unified access to both structured and unstructured data while maintaining consistent governance, lineage, and policy enforcement across systems. AI workflows can also integrate with existing pipelines and operational processes, reducing disruption and accelerating adoption.

This helps ensure that AI outputs are reliable, traceable, and usable within the systems organizations already depend on.

Confidently Apply AI to Mission Priorities

The concern around AI is valid. Most approaches require giving up control in ways that do not align with how systems are designed to operate.

Cloudera's approach addresses that challenge.

AI runs inside your environment. It operates on your data. It is governed using the same controls you already have in place. There is no external exposure, no loss of visibility, no dependency on outside services, and no unpredictable usage-based cost model.

At that point, the question is no longer whether AI is too risky or difficult to control. It becomes where it should be applied to drive the most impact.

Organizations can move forward with AI in a way that is controlled, accountable, and aligned with how their systems already operate.

Ready to See What's Possible?

Schedule a secure strategy session or mission-focused demo to explore how Cloudera can bring AI to your agency's most sensitive data.

Email: Cloudera@Carahsoft.com

FedRAMP GovRAMP TX-RAMP CERTIFIED

CLOUDERA

Cloudera, Inc. | 5470 Great America Pkwy, Santa Clara, CA 95054 USA | cloudera.com

About Cloudera

Cloudera is the only hybrid data and AI platform company that large organizations trust to bring AI to their data anywhere it lives. Unlike other providers, Cloudera delivers a consistent cloud experience that converges public clouds, on-prem data centers, and the edge, leveraging a proven open-source foundation. As the pioneer in big data, Cloudera empowers businesses to apply AI and assert control over 100% of their data, in all forms, improving security, governance, and real-time and predictive insights. The world's largest brands across all industries rely on Cloudera to transform decision-making and ultimately boost bottom lines, safeguard against threats, and save lives.

To learn more, visit Cloudera.com and follow us on LinkedIn, X, and YouTube.

© 2026 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice. DS-GBR 01 June 19, 2026