Hilco Global
AN ORIX COMPANY

SMARTER PERSPECTIVE: CYBERSECURITY

Governing Data in Frontier and Open-Weight AI

By Jason Koehn and Aidan Morrissey

November 2026

Why Enterprise AI Needs a New Data Privacy Strategy

The performance gap between open-weight models and leading closed models has historically proved significant, but it is rapidly narrowing. As the gap narrows, organizations increasingly evaluate newly released open-weight models for enterprise use. Though, they often overlook the data privacy risks and opportunities inherent in these systems. Early benchmarks suggest that leading open-weight models match closed frontier model performance on certain coding, reasoning, and enterprise tasks, sometimes at a significantly lower cost per API call.[1] The shift toward open-weight models is already underway. According to Axios, U.S. companies are increasingly turning to open-weight models as ways to reduce their AI spending.[2] Given this shift, the key AI design question for enterprises is not which model is "best," but which model should have access to which data and under what controls.

Frontier vs. Open-Weight Models: Core Privacy Trade-Offs

For enterprise AI teams, high-performing open-weight models may offer the ability to download, modify, and deploy model weights within a company-controlled environment, subject to applicable licenses and technical requirements. From a data privacy perspective, a self-hosted open-weight deployment can allow an enterprise to keep prompts, outputs, and related data within a company-controlled environment rather than routing them through a third-party model API, depending on the deployment configuration and associated services.



At the same time, recent warnings and incidents illustrate potential risks when company data is sent to external AI providers. Microsoft CEO Satya Nadella cautions that enterprises relying on external AI models risk exposing sensitive data and proprietary knowledge to model providers.[3] In July 2026, researchers found that xAI's, Grok Build, coding assistant uploaded entire code repositories to a company-controlled cloud storage bucket, potentially including proprietary source code, API keys, and unredacted secrets to xAI-controlled cloud storage.[4]

These examples sharpen the question of which tasks should run on externally hosted closed-source models and which should remain on self-hosted models within enterprise-controlled environments.

Despite these risks, locally deployed open-weight models are unlikely to replace closed frontier models for every use case; many enterprises may continue to rely on frontier-level performance for their most complex tasks. Many organizations also lack the infrastructure to host large-scale open-weight models on premises. A more likely outcome is a multi-model architecture in which companies pair open-weight models for cost-sensitive or highly customized workloads with closed frontier models for tasks that demand the highest levels of capability.[5]

This multi-model architecture can reduce costs by routing high-volume tasks that

Date Type/Task	Externally-hosted Model	Self-hosted Open-weight Model	Notes
Public information	Yes	Yes	Either model is generally acceptable given low sensitivity; external models may be cheaper at low volume, while self-hosting avoids even minimal logging outside the enterprise.
Internal policies, non-secrets	Yes, with retention and access controls	Yes, with maintenance and security responsibility on the company	External models provide speed and lower infra burden; self-hosting reduces vendor exposure but requires internal SLAs, logging, and monitoring.
Customer PII (non-special)	Restricted, only with appropriate privacy controls and DPAs	Yes, in an approved environment	External use hinges on contracts, data-transfer rules, and strong access controls; self-hosting simplifies residency and transfer issues but makes the enterprise fully accountable for safeguards.
Trade secrets, source code, R&D	Highly restricted, to avoid exposing proprietary information to a provider	Yes, with strong access controls and sufficient secure infrastructure	External frontier models may offer superior capability but increase "leaking alpha" risk; self-hosting supports IP protection at the cost of significant GPU capacity, MLOps, and specialist staff.
Regulated data (health, financial, etc.)	Case-by-case; some providers offer contractual protections and sector-specific controls	Restricted; responsibility lies with the company to meet regulatory requirements	In both cases, adherence to sector rules is mandatory; self-hosting can reduce cross-border transfer complexity but demands rigorous technical and organizational controls.

do not require frontier-level performance to lower-cost open-weight models.[6] Enterprise cost savings from multi-model ecosystems, however, depend on token volume, utilization, and deployment. Self-hosting requires infrastructure, technical resources, and maintenance capabilities that many companies may not have in-house, potentially creating significant upfront and recurring costs.[7] Given these hurdles, organizations may choose to self-host open-weight models for tasks that require tighter control over sensitive or regulated data and reduced reliance on external model providers to align with their data privacy and governance priorities.[8]

For privacy and governance purposes, self-hosted models may increase an organization's ability to keep data within an enterprise-controlled environment but also require the organization to develop, secure, and maintain the supporting infrastructure. Externally hosted models may route prompts, outputs, or related data outside the enterprise-controlled environment and therefore require different third-party privacy and security controls. A multi-model architecture requires an enterprise to develop AI data privacy strategy that explicitly maps which tasks and data classifications may be routed to each layer of the AI stack based on the risks associated with that layer. The table below provides a simplified example of how an enterprise might align model configurations with data classifications. This approach should be tailored to each organization's use cases, legal and regulatory obligations, contractual commitments, and risk tolerance.

Practical Steps: Make Models Part of Your Data Protection Program

Beyond the table above, which maps data access to risk tolerance by model configuration, enterprises need operational controls for an increasingly diverse AI stack. Recommended best practices include:

A. Make Your Data Inventory Model-Aware
   - Update existing data inventories and classification schemes to identify which model classes and deployment types are authorized to process each data classification. Document these rules and apply them consistently.
   - Set default routing rules, for example, restrict confidential and regulated data to approved model classes and deployment types.

B. Establish a Model Registry and Routing Guardrails
   - Maintain a registry of models and endpoints in use that identifies each model's owner, deployment type, hosting arrangement, and approved data classifications.
   - Use the registry to enforce routing rules and review exceptions.

C. Enforce Deployment Controls
   - For externally hosted models, conduct due diligence regarding how model providers log, retain, and use company data. Organizations can begin by asking:
     - Which categories of personal, confidential, or other sensitive data may be logged or retained?
     - Are inputs or outputs used for training, fine-tuning, safety monitoring, or service improvement by default, and can those uses be disabled?
     - In which regions are prompts and outputs processed or stored, and which subprocessors may receive them?
   - For higher-risk use cases,

supplement contractual protections with technical validation. Examples include running pilots in a sandbox, monitoring outbound traffic, and verifying that tools do not transmit more data than. The reported Grok Build incident illustrates the value of this validation.

D. Implement Engineering Controls for Self-Hosted and Hybrid Environments
   - For externally hosted or hybrid models, place AI gateways or comparable controls in front of external APIs to support data minimization and enforce routing rules.
   - For self-hosted models, implement access controls, network segmentation, vulnerability and patch management, logging, and monitoring appropriate to the data and use case.

From Models to Data Governance: Looking Ahead

As capable open-weight AI models enter enterprise technology stacks, organizations need to consider both their privacy risks and benefits. Deployment is a critical component of this analysis, as locally deployed open-weight models can provide greater control over company data but also entail significant infrastructure, security, and governance requirements. Given these considerations, organizations may increasingly adopt multi-model deployments that combine locally hosted, open-weight models with proprietary frontier capabilities. Data privacy processes and control frameworks should align with these architectures and govern how data moves across the broader AI stack. The central question is not simply which model performs best, but how tasks and data should be routed across different model environments to balance privacy, security, legal, regulatory, contractual, and performance requirements.

Sources:

1. *Artificial Analysis*
2. *Axios*
3. *IT Pro*; *SN Scratchpad*
4. *Axios*
5. *CrowdStrike*
6. *Vercel*
7. *Lenovo Press*
8. *Data Lens Journal*



JASON KOEHN
MANAGER
GLOBAL CYBER ADVISORS
PROFESSIONAL SERVICES

Jason can be reached directly at jkoehn@hilcoglobal.com.



AIDAN MORRISSEY
DIRECTOR
GLOBAL CYBER ADVISORS
PROFESSIONAL SERVICES

Aidan can be reached directly at amorrissey@hilcoglobal.com.

The views and opinions expressed are provided for general informational and educational purposes only and are subject to change without notice. This publication does not constitute investment, tax, or legal advice, nor should it be relied upon as a recommendation regarding any course of action. The information contained herein is provided "as is" without any representations or warranties as to its accuracy, adequacy, or completeness. While Hilco Trading, LLC has exercised reasonable care in preparing this publication, it assumes no liability for any actions taken or not taken based on its contents. Past performance is not indicative of future results.