Prophecy

White Paper

Managing the AI Data Governance Gap

Core principles to solve compliance friction and practical questions every data, security, and compliance leader must ask.



Executive Summary

Enterprises have rushed to adopt AI, often leaving governance as an afterthought. McKinsey's 2026 report on the state of AI trust concluded that "responsible AI maturity continues to improve, yet strategy, governance, and agentic AI controls lag behind."

For data teams, implementing AI governance starts with architecture: where data lives during processing, who reviews AI before execution, and what data is retained afterward.

This white paper provides guidance on assessing and managing the governance gap with:



- Four governance principles that hold up under the scrutiny of an audit
- Prophecy's security and AI risk ownership approach
- A six-question framework any data or security leader can use to evaluate the governance of an AI system

The real question data and security leaders need to answer is how to govern AI sustainably. **The principles below offer a starting framework.**

PRINCIPLE 1

Data That Never Has to Leave Is Data That's Never at Risk

The most common approach to enterprise AI still involves moving data somewhere else, such as into a vendor's environment, a training pipeline, or a third-party model's context window, and then attempting to govern it after the fact with policies, contracts, and monitoring. Every one of those controls is necessary, but none of them are as effective as eliminating data copies in the first place.

AI is now both a threat vector and a security tool, and the fundamentals that reduce breach cost and frequency are data discovery, classification, access control, and encryption — the same fundamentals that apply whether or not AI is involved. The lowest-risk architecture is the one with the least data to lose.

The Prophecy Way:

Prophecy does not send your data to third-party AI model providers. Its AI features work from structural metadata rather than raw data, so your underlying data stays in your environment, where it belongs.

Prophecy connects directly to your existing cloud data platforms, such as Databricks and Snowflake, and works with data in place.

Prophecy also retains only descriptive metadata, such as table and column names, rather than the underlying data itself. In sample row previews, the data is fetched temporarily and rendered in the browser over an encrypted connection. Access continues to flow through your existing identity provider and platform permissions, so users can only access data they are already authorized to see. And, there's no parallel access model to manage. Prophecy leverages the governance layer in your cloud data platform, such as Unity Catalog, rather than creating a separate set of permissions that can drift from the platform of record.

As such, your existing security policies and controls remain the foundation for how data is accessed and governed. The less data an AI system needs to move or retain, the smaller the attack surface and compliance burden on your organization.

PRINCIPLE 2

Human Oversight Must Be a Structural Control

Human-in-the-loop policies are easy to write, but hard to enforce for products that are not built for it. The distinction matters to regulators as much as it does to security teams. Regulators increasingly distinguish between AI systems that assist human decision-making and systems that act autonomously. For enterprises, that makes demonstrable human oversight a product requirement in addition to a policy requirement.

Companies that McKinsey classifies as "AI high performers" are substantially more likely to have well-defined processes for determining when model outputs require human validation before use. Oversight that lives in a process document outside of AI tooling misses the mark.

The Prophecy Way:

Prophecy's AI assists with developing and documenting data pipeline logic, but it does not independently change data, make business decisions, or act outside of defined boundaries.

Every AI-proposed change can be reviewed through the visual workflow, documentation, and underlying code before it is used, with testing in a development environment before anything reaches production. Guardrails constrain the agent to the task at hand: off-topic requests cannot affect your data workflows, and destructive actions, such as deleting a table, are refused outright. Users can halt or undo an agent at any point, and administrators can disable AI functionality entirely across their data teams.

Prophecy also applies controls to common AI failure modes. Inputs are validated and outputs are checked to help defend against prompt injection, while automated validation, human review, and testing help catch incorrect AI-generated work before it reaches production.

This is the same principle reflected in ISO/IEC 4200: effective AI governance depends on demonstrable processes and controls around how AI is used, not simply on evaluating the model itself. Prophecy makes those controls part of the development process.

PRINCIPLE 3

Govern the Full Data Lifecycle

Governance does not end when a data workflow finishes running. What a system retains after the fact — chat histories, logs, cached context — is often where the exposure sits, and it is the part of the AI stack likely to be underspecified in vendor conversations.

Cisco's 2026 Data and Privacy Benchmark Study found that clear communication about how data is used is the single most effective way to build customer confidence, ahead of any technical control. Transparency about retention is one of the more effective trust-building levers available.

The Prophecy Way:

Prophecy does not use your data to train or fine-tune AI models, and it contracts with its model providers to enforce that commitment and minimize retention on their side.

Conversation history is subject to defined retention periods and deletion controls, while data teams that require full custody can store it in their own cloud storage and retain control over access and deletion. You can also choose how much data AI can access: sample data sharing can be disabled entirely, or you can leverage de-identified or synthetic data when real data is not required.

Prophecy applies vendor oversight to the AI model providers it uses, evaluating their data access, use, and retention practices and maintaining a current subprocessor list so data teams can understand where their data may be processed.

For enterprises, AI governance is not just about what happens while an AI system is running. It is about knowing where data goes, what is retained, who can access it, and when it is deleted throughout its entire lifecycle.

PRINCIPLE 4

Establish Owners of AI Risk

Principles only matter if you are empowered to enforce them. ISO/IEC 42001 does not certify an AI model; it certifies an AI management system. The question is not "Is your AI good?" but rather "Can you show the process by which you decide, review, and correct how AI is used?"

The Prophecy Way:

Prophecy treats AI governance as an ongoing operational responsibility, with clear ownership, review processes, controls, and evidence.

- **Ownership:** AI risk follows a shared-responsibility model. Prophecy is accountable for the AI functionality, controls, and safeguards within its platform, while you remain responsible for risks associated with models or AI providers you independently select or integrate.
- **Review before release:** Changes to prompts, guardrails, or permitted model providers are treated as changes to controls, and reviewed accordingly, rather than as routine product changes.
- **Provider selection and re-review:** Model providers are assessed before adoption against data-use, retention, and security criteria, and re-reviewed on an ongoing basis. The same oversight applies to relevant subprocessors.
- **An off switch that is tested:** Users can halt or undo an agent mid-task, while administrators can disable AI functionality across their teams. These controls provide a defined way to stop AI activity when needed.
- **Evidence by default:** AI-related activity is captured through the platform's audit logging and stored in your isolated log space, providing an auditable record of AI activity and supporting ongoing review.

6 Governance Questions to Ask About Your AI System

The principles and best practices above translate into six concrete questions that any data or security leader can put to a vendor.

1

Where does my data live while the AI works on it?

Does the vendor copy data into its own environment, or process it in place? Processing in place eliminates entire categories of risk.

2

What can the AI do without human approval?

Look for structural controls that define what the AI can and cannot do, including reviewable outputs, boundaries on actions, testing before production, and the ability for users or administrators to stop AI activity.

3

Is my data used to train the model, and how is that enforced?

Ask to see the contract language that prohibits training and limits retention to move beyond marketing claims into the details.

4

What is retained after the task is done, for how long, and who can delete it?

This should be a specific, documented retention schedule, not a general statement about “minimal” data collection.

5

Are the vendor’s controls independently verified against recognized standards?

SOC 2 Type II, ISO/IEC 27001, alignment with the NIST AI RMF, and progress toward ISO/IEC 42001 are the current reference points. Self-attestation is a lower bar than an independent audit.

6

Whose models are they, and can I bring my own?

Ask which models and third-party providers the vendor relies on, whether you can use an approved model endpoint, and what happens to your workflows if your model or data platform changes. The strongest architectures preserve your ability to change providers, platforms, and technologies without rebuilding the work.

Conclusion

Data teams that successfully navigate the next phase of enterprise AI will be those that build governance directly into the tool architecture.

Prophecy is built on that premise. AI can do more of the data development work while enterprise teams retain control over their data, their workflows, and what ultimately reaches production. By minimizing data exposure, making human oversight structural, governing the full data lifecycle, and establishing clear ownership of AI risk, Prophecy gives data teams a way to adopt AI without compromising the standards and controls their data environments require.

The Standard Prophecy Holds Itself To:

AI that works within enterprise governance, rather than asking governance to work around AI.



Explore Prophecy's trust and security approach

Learn more at prophecy.ai/trust

Prophecy

© 2026 SimpleDataLabs DBA Prophecy. All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form without the prior written permission of the publisher. Prophecy and all other third-party trademarks mentioned herein are the property of their respective owners.