Whitepaper

AI Governance 101

A Practical Framework for Accountability, Transparency, and Control

Amanda Ray, Associate General Counsel
Nicole Reineke, Chief AI Officer

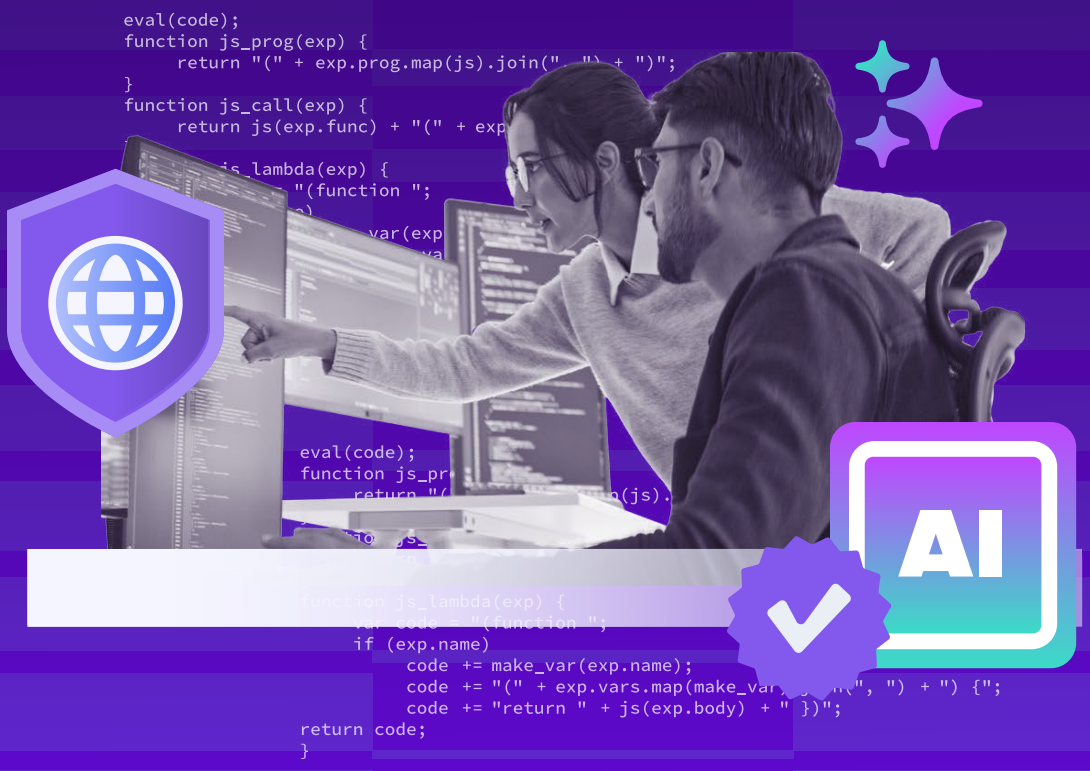

N-ABLE™

Table of Contents

Why AI Governance Matters Now

Artificial intelligence (AI) is already part of the operational systems that businesses rely on every day. AI tools summarize tickets, draft communications, analyze telemetry, and increasingly automate actions across IT environments.

As these capabilities expand, the role of AI is shifting from simple assistance to influencing or initiating operational decisions.

That shift creates a new responsibility.

Organizations must be able to explain:

- **Where AI is being used**
- **The types of data AI is processing**
- **What those systems can access**
- **What actions they can take**
- **How their behavior is monitored**

Without this visibility, AI adoption can introduce avoidable risk. Teams may connect AI tools directly to production systems, deploy automation without defined oversight, or allow AI features in third-party tools to process data in ways that are not clearly understood.

AI governance exists to prevent those situations.

Good governance does not slow innovation. It enables it. When organizations understand how AI systems operate, document their behavior, and apply appropriate controls, they can confidently adopt new capabilities while protecting their customers and their environments.

This paper provides a practical framework to help organizations do exactly that. It outlines the legal reality, including transparency and documentation, and the four controls that allow organizations to move forward with AI while maintaining accountability, transparency, and operational control.

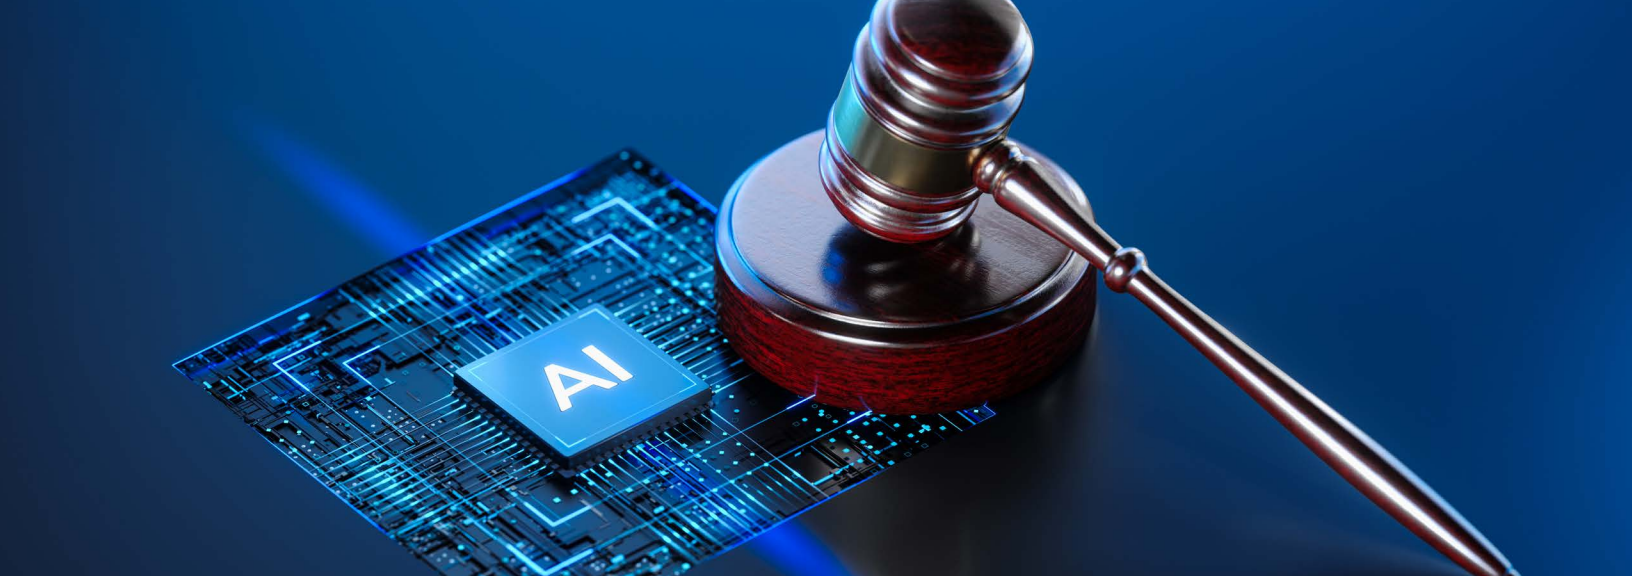

Legal Reality: AI Governance Starts with Transparency and Documentation

When teams talk about AI risk, the conversation often turns to model behavior, hallucinations, or technical safeguards. Those issues matter. But from a legal perspective, the foundation is much simpler.

The most important protection an organization has when adopting AI is transparency and documentation.

Regulators and customers rarely ask whether a company perfectly predicted every possible AI outcome. What they want to understand is whether the organization acted responsibly when introducing the technology.

In practice, that responsibility comes down to three very practical questions:

1. Did the organization obtain consent based on a clear disclosure of how AI is used?
2. Did the organization document how decisions about the system were made?
3. Did the organization implement reasonable controls around the system's behavior?

If an organization can answer yes to these, it is generally in a defensible position. If not, it may struggle to explain how the technology was introduced and managed.

This does not mean companies must stop innovating until every regulation is finalized. In fact, the opposite is true: when organizations are transparent about how AI operates, maintain clear records of their decisions, and build governance around those systems, they create the conditions that allow innovation to move forward responsibly.

Three legal principles consistently guide safe and responsible AI deployment.

Transparency

Customers must understand:

- **Whether their data is processed by AI**
- **What the AI is doing**
- **Whether their data is used to train models**
- **What decisions are automated**

Companies often underestimate this obligation and assume customers already know AI is being used. This assumption creates legal exposure.

Consent and Disclosure

Customers must be informed:

- **What systems process their data**
- **Whether new AI vendors are introduced**
- **What outputs are produced from their data**

If a company changes how data is processed, customers must be notified and allowed to opt out.

For example, if a company begins analyzing customer call transcripts with a new AI system, customers must be notified of such change.

Documentation

Documentation is your best defense.

Organizations must document:

- **What the AI system does**
- **What inputs it receives**

Clearly note the data types involved, specifically:

- Health and medical data (HIPAA-type data)
- Financial data (account info, transaction history, credit data)
- Biometric data (face, voice, fingerprints)
- Children's data
- Government identifiers (SSNs, IDs)

- What outputs it produces
- What systems it interacts with
- Who is responsible for it

An evaluation of the AI system's overall risk:

- **A high-risk AI system is one that's used to influence decisions about people, such as:**
  - Credit, lending, pricing, or eligibility
  - Employment, hiring, or promotion
  - Access to services or benefits
  - Profiling behavior

This documentation is the evidence that regulators expect to see.

Legal guidance consistently returns to the same themes: transparency, documentation, and reasonable controls. But those principles only become real protection when they are translated into daily operational practices. This translation is accomplished through the four controls.

The Four Controls Every Organization Needs

Governance cannot live only in legal language or policy documents. It must be built directly into how a company is run and how AI systems are designed, deployed, and monitored.

The following four controls translate the legal principles of transparency, consent and disclosure, and documentation into practical operating standards. Together, they create the minimum structure needed to know where AI exists, who is responsible for it, how it is communicated, and whether the organization can prove what the system did.

Control 1: Clear Ownership (AI Governance Committee)

AI systems introduce new operational and legal risks. Without clearly assigned responsibility, it becomes difficult to manage those risks or respond when something goes wrong.

Every AI system must have defined ownership and oversight. This typically begins with a lightweight AI governance committee that establishes policy, reviews higher-risk use cases, and ensures that AI deployments align with company standards for security, privacy, and operational safety.

At the system level, each AI capability must have two accountable roles: a business owner, responsible for the purpose and acceptable outcomes of the system, and a technical owner, responsible for implementation, permissions, logging, and ongoing monitoring. When ownership is clearly defined, organizations can maintain accountability, manage risk, and operate AI systems responsibly.

Control 2: AI Visibility (AI Inventory)

The first step in governing AI is knowing where it exists. Many organizations underestimate how quickly AI capabilities spread across tools, workflows, and automations. AI is no longer limited to standalone applications. It is increasingly embedded in platforms that teams already use daily.

To manage risk effectively, organizations must maintain visibility into all AI systems operating within their environment. This includes AI features built into Software as a Service (SaaS) platforms, internally developed tools, AI-enabled automations, and external AI services connected to operational systems.

The most effective way to achieve this visibility is by maintaining a centralized AI inventory that tracks every AI system and agent used across the organization. This inventory becomes the foundation for governance, risk assessment, and ongoing oversight.

Control 3: Customer Transparency and Communication for AI Systems (Communication and Policies)

One of the most important legal protections when deploying AI is clear communication with customers about how AI is used and how their data is handled.

Organizations often underestimate this requirement. A common assumption is that AI is now embedded in so many tools that customers already expect it. However, relying on that assumption creates risk.

Customers still expect transparency about:

- **How their data is processed**
- **Whether AI systems analyze their information**
- **What decisions AI systems influence or make**
- **Whether their data is used to train models**

Clear communication protects both the organization and its customers. It establishes trust, sets expectations, and ensures that the organization can demonstrate responsible AI use if regulators or customers raise questions.

Control 4: Vendor Management (Vetting Questionnaires)

Many AI capabilities used by teams come from third-party tools and SaaS platforms. Even when the AI system is provided by a vendor, the organization remains responsible for how that system interacts with customer environments and data.

Before deploying any external AI capability, an organization should verify how the vendor governs the technology, protects customer data, and documents system behavior. In particular, organizations should confirm that customer data is not used to train external models and that logging and monitoring controls exist.

If a vendor cannot clearly explain how their AI system handles data, security, and oversight, that should be treated as a risk signal.

Let's explore each of these in detail.

Control 1: Minimum AI Governance Committee Structure

AI governance does not require a large bureaucracy. What it does require is clear accountability and a consistent decision process. As organizations begin deploying AI systems and agents, someone must be responsible for reviewing risk, approving deployments, and ensuring that systems behave as expected once they are in operation.

Regulations and legal frameworks around AI are still evolving, but the core expectations are already clear. Organizations must be able to explain who owns each system, how risk was evaluated before deployment, what oversight exists, how behavior is monitored over time, and whether the company can reconstruct what the system did.

A lightweight AI governance committee provides the structure needed to answer those questions. This group establishes standards, reviews higher-risk use cases, and ensures that AI systems are deployed responsibly. The goal is not to slow innovation, but to create a clear process so teams can move quickly while maintaining accountability.

The following roles represent the minimum structure needed to operate this governance model. In smaller organizations, a single person may fulfill multiple roles.

Consider the following types:

Role	Typical Position	Responsibilities
Executive Sponsor	COO, CIO, CISO, or Head of Platform Operations	Approves AI governance policies Resolves high-risk decisions Ensures the program has funding and support Sets the organization's tolerance for AI risk
AI Governance Lead	Program lead responsible for AI oversight	Maintains the AI inventory Coordinates governance reviews Tracks approvals, exceptions, and status Maintains governance standards, templates, and processes
Business Owner (Per AI System or Agent)	Leader responsible for the business outcome	Defines the use case and business purpose Accepts operational and business risk Ensures the AI system is appropriate for the business context Approves changes to system scope

Role	Typical Position	Responsibilities
Technical Owner (Per AI System or Agent)	Technical lead responsible for implementation	Manages integrations and system connections Defines and enforces permissions Implements logging and monitoring Oversees testing and validation Maintains rollback and operational safety mechanisms
AI Review Group	Cross-functional governance team	Reviews medium- and high-risk AI use cases Typically includes security, privacy/legal, platform operations, service delivery, and product or automation leadership

Control 2: AI Visibility – Establishing the AI Inventory

Effective AI governance starts with knowing where it exists. Many organizations underestimate how quickly AI capabilities spread across tools, workflows, and automations. AI is no longer limited to standalone applications. It is increasingly embedded in platforms that teams already use every day.

To manage risk effectively, organizations must maintain visibility into all AI systems operating within their environment. This includes AI features built into SaaS platforms, internally developed tools, AI-enabled automations, and external AI services connected to operational systems.

The most effective way to achieve this visibility is by maintaining a centralized AI inventory that tracks every AI system and agent used across the organization. This inventory becomes the foundation for governance, risk assessment, and ongoing oversight.

AI systems should follow a defined lifecycle before entering production.

These steps include:

- **Registering the use case**
- **Classifying the risk**
- **Reviewing controls**

STEP 1 Register the Use Case

Document:

- **The purpose of the AI system**
- **Connected systems**
- **Data involved**
- **Expected outcomes**

Role	Item	Description/Purpose
Required Inventory Fields	Agent Name	Unique name used to identify the AI system or agent
	Description of Purpose	Clear explanation of what the agent is designed to do
	Business Owner	Person accountable for the business purpose and outcomes
	Technical Owner	Person responsible for implementation, permissions, monitoring, and maintenance
	Vendor or Internal Build	Indicates whether the system is third-party or internally developed
	Model / Provider Used	AI model or service provider used (e.g., OpenAI, Bedrock, internal model)
	Connected Systems	Systems the agent interacts with (PSA, RMM, CRM, endpoint tools, etc.)
	Data Types Accessed	Categories of data the agent reads or processes
	Permissions Granted	Access levels granted to the agent
	Actions the Agent Can Take	Operational actions the system can execute
	Human Oversight Model	Approval model used (approval required, interruptible, post-review)
	Risk Classification	Low-, medium-, or high-risk designation
	Logging Location	Where operational logs and decision records are stored
	Date Approved	Date governance approval was granted
	Last Review Date	Most recent governance review
	Current Status	Pilot, production, paused, or retired

Governance Question	Why It Matters
What problem is this agent solving?	Confirms the agent has a valid business purpose
What systems can it read?	Identifies data exposure risk
What systems can it write to?	Identifies operational impact risk
Can it trigger automation?	Determines whether it can initiate cascading actions
Can it affect a customer directly?	Determines customer experience and legal risk
Can it change configurations?	Identifies infrastructure control risk
Can it expose regulated or confidential data?	Identifies privacy and compliance implications
Can its actions be reversed?	Determines operational recoverability
Can we reconstruct its decision path?	Ensures auditability and incident investigation capability

STEP 2 Classify the Risk

Determine whether the system is low, medium, or high risk.

Risk Level	Characteristics	Examples	Required Controls
Low	• Internal only • No direct customer impact • No privileged system access • Reversible outcomes • Advisory or draft support only	• Summarizing internal notes • Drafting internal documentation • Suggesting ticket categorization • Recommending next actions for technicians	• Simple registration • Standard logging • Business owner approval • Periodic review
Medium	• Touches operational systems • May influence customer work • Limited automation • Limited system writes • Moderate data sensitivity	• Drafting customer communications for approval • Proposing remediation scripts • Suggesting patch group assignments • Prioritizing alerts for human review	• Formal review • Test evidence • Defined rollback • Tighter access controls • Structured logging • Human approval or timed override

Risk Level	Characteristics	Examples	Required Controls
High	• Autonomous actions in production • Privileged access • Customer data exposure • Destructive or irreversible changes • Security, compliance, or financial implications	• Implementing auto-remediation on customer endpoints • Modifying account permissions • Deleting records • Enforcing security containment measures • Executing customer-facing decisions without review	• Human approval or timed override • Formal governance review • Executive signoff • Enhanced testing • Continuous monitoring • Strong approval boundaries • Post-incident review • May be prohibited until maturity improves

STEP 3 Review Controls

Before approving an AI system or agent for deployment, confirm that appropriate safeguards are in place. This review ensures the system has clear boundaries, oversight, and traceability.

Confirm that appropriate safeguards exist:

- Permission limits
- Human oversight
- Testing procedures
- Logging and evidence

Permission Limits

AI systems must follow the principle of least privilege. The agent should only access what is necessary for its defined purpose.

Permission Review Checklist

- [ ] The agent has a clearly defined scope of operation
- [ ] The agent only has access to systems required for its task
- [ ] Read access and write access are separated where possible
- [ ] The agent does not use shared or anonymous service accounts
- [ ] The agent does not have standing administrator privileges unless explicitly approved

- [ ] Credentials used by the agent are stored securely and rotated regularly
- [ ] Cross-tenant access is restricted or prohibited
- [ ] Access to sensitive systems (PSA, RMM, identity systems, billing platforms) is explicitly reviewed
- [ ] Access permissions are documented in the AI inventory
- [ ] Unused or unnecessary permissions have been removed

Red Flags

- One agent accessing many unrelated systems
- Unrestricted scripting privileges
- Production credentials used in testing environments
- Broad mailbox access

Human Oversight

AI systems must have defined human supervision appropriate to their risk level.

Oversight Review Checklist

- [ ] The human oversight model is clearly defined
- [ ] The responsible business owner has approved the agent's scope
- [ ] The responsible technical owner is identified and accountable
- [ ] A human approval step exists for sensitive actions
- [ ] High-risk actions cannot execute without human confirmation
- [ ] Operators have the ability to interrupt or stop agent activity
- [ ] Escalation paths are documented if the agent encounters uncertainty
- [ ] Customer-impacting actions are reviewed before execution
- [ ] The organization has a process for reviewing agent behavior patterns

Oversight Models

Oversight Level	Description	Typical Use Cases
1 – Human Approval Required	The AI agent cannot take action until a human explicitly approves the output or decision. This level is used when the consequences of an incorrect action could impact customers, security, finances, or system integrity.	• Customer communications involving commitments • Privileged administrative actions • Security response actions • Financial or billing decisions • Destructive or irreversible system changes
2 – Human Interrupt Window	The AI agent may act after a short review window unless a human intervenes to stop the action. This allows for faster automation while preserving the ability for operators to intervene before execution.	• Low-impact operational tasks • Routine internal automations • Draft-to-send communications with limited scope • Operational workflow assistance
3 – Post-Action Review	The AI agent executes actions automatically. Humans review patterns and outcomes afterward to confirm that the system behaves appropriately over time.	• Low-risk optimization tasks • Internal analytics and reporting • Triage recommendations for technicians • Nonsensitive internal routing or classification

Testing Procedures

AI systems must be tested beyond basic functionality to ensure they behave safely.

Testing Review Checklist

- [ ] Expected behavior tests have been performed
- [ ] Edge cases and abnormal inputs have been tested
- [ ] Failure scenarios have been simulated
- [ ] The system fails safely when uncertain
- [ ] The agent does not exceed its permission boundaries
- [ ] Prompt injection and manipulation attempts have been tested where applicable

- [ ] Incorrect outputs do not automatically trigger harmful actions
- [ ] The rollback process has been tested successfully
- [ ] Logging has been validated during test runs
- [ ] The testing results are documented and stored

Testing Scenarios to Simulate

- Incorrect or malicious input
- Unavailable connected systems
- Conflicting automation triggers
- Model output errors
- Permission boundary violations

Logging and Evidence

AI systems must generate logs that are sufficient to reconstruct actions and decisions. Minimum logging elements for AI and GenAI use cases include:

Log Element	Purpose
Timestamp	Records when the action occurred
Agent Identity	Identifies which AI system or agent performed the action
Version or Model Used	Captures the specific AI model or agent version used
Initiating Event	Records what triggered the action (alert, request, automation trigger)
User or System Request	Identifies the originating request or user interaction
Key Prompt or Instruction Reference	Provides traceability to the prompt or instruction set used
Connected Tools Called	Records any tools, APIs, or systems the agent invoked
Retrieved Data Source Reference	Identifies the data sources used to generate the response
Output Generated	Records the content or recommendation created by the AI
Action Executed	Documents the operational action taken
Approval Event (if required)	Records human approvals when oversight is required
Final Result	Captures the outcome of the action
Errors or Overrides	Logs failures, blocked actions, or human overrides

Logging Review Checklist

- [ ] Logging is enabled before production deployment
- [ ] Logs include the identity of the AI agent
- [ ] Logs include timestamps for all actions
- [ ] Logs record the triggering event or request
- [ ] Logs record the model or agent version used
- [ ] Logs record tools or APIs called by the agent
- [ ] Logs capture outputs generated by the AI system
- [ ] Logs record any actions executed on external systems
- [ ] Logs capture approval events if human oversight is required
- [ ] Logs capture errors, overrides, or blocked actions
- [ ] Logs are stored in a centralized logging system
- [ ] Logs are retained according to company policy
- [ ] Logs can be tied to incident or ticket IDs when applicable

Evidence Standard

If an auditor, regulator, or customer requests evidence, the organization should be able to demonstrate:

Evidence Requirement	What It Shows
What controls existed	The governance safeguards in place
When controls were reviewed	Ongoing governance oversight
What the agent did	Operational traceability
Who approved the design	Human accountability
How drift or failures would be detected	Monitoring and risk management capability

STEP 4 Approve or Reject

At this stage, the organization decides whether the AI system is ready to operate in production. This decision should not be informal or based on hallway conversations. The review group should document the decision, the rationale, and any conditions attached to the approval.

AI governance requires creating a clear record of how the decision was made. If a regulator, auditor, or customer ever asks how a system was approved, the organization should be able to show the discussion, the risk assessment, and the controls that were considered.

Possible outcomes include:

- **Approved**
  The system meets the required controls and can move to production
- **Approved with Conditions**
  The system may be deployed, but only after specific safeguards or limitations are implemented
- **Rejected or Needs Redesign**
  The system presents unacceptable risk or lacks the controls required for safe deployment
- **Pilot Only**
  The system may run in a limited environment with restricted permissions, limited scope, and defined monitoring

For every decision, the governance committee should document:

- The risk classification
- The systems and data involved
- The controls that were evaluated
- The oversight model selected
- Any limitations placed on the deployment
- Who approved the decision and when

This documentation becomes part of the AI inventory and creates the evidence needed to explain why the system was allowed to operate.

AI systems should also be reviewed regularly after deployment. Higher-risk systems require more frequent review because their actions have greater potential impact. Over time, these reviews help organizations detect drift, adjust permissions, and confirm that the system continues to operate within its intended scope.

Control 3: Customer Transparency and Communication for AI Systems

A critical legal safeguard when deploying AI is clear communication with customers about how AI is used and how their data is handled.

Organizations often underestimate this requirement. A common assumption is that AI is now embedded in so many tools that customers already expect it. However, relying on that assumption creates risk.

Customers still expect transparency about:

- How their data is processed
- Whether AI systems analyze their information
- What decisions AI systems influence or make
- Whether their data is used to train models

Clear communication protects both the organization and its customers. It establishes trust, sets expectations, and helps the organization demonstrate responsible AI use if regulators or customers raise questions.

What Organizations Should Clearly Communicate

When communicating AI usage to customers, organizations should focus on several key elements.

Communication Topic	What Customers Should Understand
Purpose of AI Usage	Why the AI system is used and what problem it solves
Data Processed	What types of customer data may be analyzed
AI Outputs	What the system produces (summaries, recommendations, automation triggers)
Model Training	Whether customer data is used to train models
Third-Party Involvement	Whether external AI vendors process the data
Customer Control	Whether customers can opt out of specific AI features

Clear explanations are far more effective than technical descriptions. Customers generally want to understand what the system does and how it affects them, not the details of how the model works.

When Organizations Should Communicate AI Usage

Communication about AI typically occurs at four points in the customer relationship.

1 At the Start

When a customer first enters into a service agreement, the organization should clearly disclose how AI may be used within the service.

This information is usually communicated through:

- **Service agreements:** Consider the N-able service agreements available at https://www.n-able.com/legal
- **Privacy disclosures:** Consider the N-able Privacy Notice: https://www.n-able.com/legal/privacy

- **Product documentation:** Refer to documentation such as https://documentation.n-able.com/N-central/userguide/Content/Configuration/ScriptSoftwareRepository/AI_Assisted_Scripting.htm
- **Onboarding materials:** Consider publishing a Pledge on the Responsible Use of Artificial Intelligence (AI), such as the one N-able has published: https://www.n-able.com/trust-center/ai-pledge

Customers should understand:

- **Whether their data may be processed by AI systems**
- **What types of analysis may occur**
- **How outputs from those systems are used**

For example, if an organization uses AI tools to analyze operational telemetry, summarize support tickets, or generate technician recommendations, this should be explained as part of the service model.

The key principle is simple: customers should understand the role AI plays in the services they receive.

2 When AI Processing Changes

If an organization changes how customer data is processed, customers should be notified.

This is particularly important when:

- **New AI vendors are introduced**
- **New types of data analysis are performed**
- **AI systems are integrated more deeply into products or workflows**

A practical example is the introduction of AI analysis for recorded interactions.

If a company begins using an AI system to analyze customer call transcripts, that change should be disclosed. Even if calls were already recorded, customers may not expect transcripts to be processed by external AI systems.

Organizations often communicate these changes by sending notifications to customers explaining:

- **What is changing**
- **What data may be processed**
- **What the AI system will do**
- **Whether customers can opt out**

This level of transparency helps maintain trust while allowing organizations to continue innovating.

3 When New AI Subprocessors Are Introduced

Many AI systems rely on external providers such as cloud platforms or model vendors.

If a new third-party system begins processing customer data, customers should be informed.

Organizations typically communicate these changes through:

- **Updates to subprocessor lists**
- **Privacy policy updates**
- **Customer notification emails**

For example, if a company begins sending customer interaction data to a new AI provider for analysis or summarization, customers should be notified that a new processor is involved.

In many cases, organizations provide customers an opportunity to raise concerns or opt out.

4 When AI Incidents or Data Issues Occur

If an issue occurs involving customer data, AI processing, or an AI system behaving incorrectly, organizations should communicate the event to customers in a timely and transparent manner.

This includes situations such as:

- **Data being processed in a way that was not intended**
- **Incorrect outputs or automated actions affecting customer environments**
- **Unauthorized access to data processed by an AI system**
- **Third-party AI vendor failures or policy changes that affect data handling**
- **AI systems acting outside of expected operational boundaries**

When incidents occur, organizations should clearly communicate:

- **What happened**
- **What systems or data were affected**
- **Whether any customer action is required**
- **What steps are being taken to resolve the issue**
- **What safeguards are being implemented to prevent recurrence**

For example, if an AI automation tool incorrectly modifies configuration settings across managed devices, customers should be informed that the issue occurred, what systems were impacted, and what remediation steps have been taken.

Similarly, if a third-party AI vendor experiences data exposure or changes their data handling policies in a way that affects customer information, organizations should promptly notify affected customers and explain the implications.

Transparent incident communication helps maintain trust and demonstrates responsible AI governance.

Communication Scenarios

Example 1: AI-Assisted Support Operations

An organization introduces an AI tool that summarizes service tickets and suggests next steps for technicians.

Appropriate communication might include:

- Explaining that AI assists in analyzing ticket data
- Clarifying that recommendations are reviewed by technicians
- Confirming that customer data is not used to train external models

Example 2: AI Analysis of Customer Call Transcripts

A company records support calls and later introduces an AI system to summarize transcripts.

Customers should be informed that:

- Recorded conversations may be processed by an AI system
- The purpose is to improve service quality and documentation
- The analysis does not change the original service commitments

Example 3: Introduction of a New AI Provider

An organization integrates a new AI platform to analyze operational telemetry.

Customers should be notified that:

- A new AI provider is being used
- The provider processes specific operational data
- The system supports service improvements, such as faster incident resolution

Strong communication practices allow organizations to adopt AI confidently.

When customers understand how their data is used and what safeguards exist, organizations gain the freedom to experiment, improve services, and deploy new capabilities.

Transparency is not a barrier to innovation. It is what makes innovation sustainable.

Organizations that clearly communicate how AI operates within their services build trust, reduce legal risk, and position themselves as responsible technology partners.

Control 4: Vendor Controls

Many AI capabilities used by organizations come from third-party platforms, SaaS tools, or embedded AI features. While these vendors may advertise strong AI capabilities, the responsibility for how those tools interact with customer environments ultimately remains with the organization.

Before deploying any third-party AI system into production environments, organizations must evaluate how the vendor handles governance, data protection, monitoring, and operational control.

A key legal insight is that vendor assurances alone are not sufficient protection. Organizations must understand how the vendor processes data and be able to explain those practices to customers.

In particular, companies should verify that customer data is not used to train external models unless that usage is clearly disclosed and approved. Many AI services include optional training features that must be explicitly disabled.

The goal of vendor evaluation is to ensure that the organization can confidently answer the following questions from customers, auditors, or regulators:

- What vendor systems process customer data?
- What safeguards exist to protect that data?
- What evidence shows that the system behaves safely?

Vendor AI Governance Review

Before deploying a third-party AI system, organizations should evaluate the vendors' AI maturity and practices. Consider asking the following.

Question	Why It Matters
Does the vendor have an AI governance framework?	Demonstrates that the vendor manages AI risk intentionally
Do they document how models are used and monitored?	Provides transparency about system behavior
Can they explain how the system is evaluated over time?	Confirms ongoing monitoring and improvement
Do they clearly document how customer data is processed?	Ensures organizations can communicate data usage to customers
Do they provide documentation supporting regulatory compliance frameworks?	Supports privacy and regulatory obligations

Strong governance documentation makes it easier for organizations to explain AI usage transparently to customers and regulators.

Security and Data Protection

Question	Why It Matters
What customer data is retained by the system?	Determines data exposure risk
Is customer data used to train models?	Protects customer intellectual property and confidential data
Can model training on customer data be disabled?	Ensures data ownership protections
How is tenant separation handled?	Prevents data leakage across customers
What credentials or permissions are required?	Determines operational risk
What audit logs are available?	Enables traceability and incident investigation

Behavior Controls

Question	Why It Matters
What guardrails exist to limit unsafe outputs?	Prevents harmful system behavior
How does the vendor mitigate prompt injection or manipulation?	Protects operational systems from malicious input
How does the vendor test the system for unsafe behavior?	Demonstrates responsible model development
Is there evidence of ongoing safety evaluation?	Ensures the system improves over time
Can the vendor explain how model behavior changes are managed?	Supports operational predictability

Operational Controls

Question	Why It Matters
Can the AI feature be limited by role or permission?	Allows organizations to control exposure
Can the feature be disabled quickly if problems occur?	Enables rapid incident response
Can logs be exported to the organization's logging systems?	Supports auditing and monitoring
Are model updates communicated to customers?	Prevents unexpected behavior changes
Can organizations configure how the system interacts with operational tools?	Protects production environments

Governance Rule

If a vendor cannot clearly answer basic questions about:

- Data handling
- System monitoring
- Model training
- Logging and audit evidence

The organization should treat that as a risk signal and reconsider deployment.

Responsible AI adoption requires that organizations understand not only what the AI system does but also how the vendor manages the system behind the scenes.

What Companies Should Stop Doing Immediately

As organizations begin adopting AI tools and agents, some common practices create unnecessary operational, security, and legal risk. Many of these behaviors arise from treating AI systems as simple productivity tools rather than systems capable of influencing or taking action in real environments.

To reduce risk while continuing to innovate, organizations should immediately stop the following practices:

Practice to Stop	Why It Creates Risk
Allowing staff to connect AI tools directly to production systems without review	Unreviewed integrations can expose sensitive data or grant unintended system access
Treating AI agents like harmless assistants	Many AI tools can generate scripts, trigger automations, or influence operational decisions
Granting broad administrative permissions to automation tools	Excessive permissions increase the risk of unintended or destructive actions
Deploying AI systems without logging and traceability	Without logs, organizations cannot reconstruct actions or investigate incidents
Relying on vendor marketing instead of verifying controls	Vendors may advertise capabilities without clearly documenting governance, data handling, or safeguards
Assuming annual reviews are sufficient	AI systems evolve quickly and require ongoing monitoring and periodic review
Allowing AI agents to interact with customers without defined controls	Customer communications and commitments require clear oversight and accountability
Leaving ownership ambiguous	Every AI system must have a named business owner and technical owner

Organizations that address these issues early establish a strong foundation for responsible AI adoption. By removing these risky practices, companies can move forward with AI deployment in ways that support innovation while maintaining operational control and customer trust.

30-Day Action Plan

Week 1: Establish Ownership and Stop the Unknowns

The first week focuses on creating accountability and preventing unmanaged AI from spreading.

1 Name the AI Governance Leadership

Identify the core governance roles:

- **Executive sponsor**
- **AI governance lead**
- **Initial review group (security, operations, service delivery, product/automation)**

The objective is to create a small group that can approve or reject AI use cases quickly.

2 Issue a Temporary Policy

Publish a short internal notice:

- **No new AI agents or automations may be deployed into production without review**
- **All existing AI tools must be registered in the AI inventory**

This prevents shadow AI and agent sprawl while the program is established.

3 Begin the AI Inventory

Create a simple registry to capture:

- **AI tools in SaaS platforms**
- **Internal AI automations**
- **AI features embedded in existing tools**
- **External AI services connected to systems**

The goal is visibility, not perfection.

4 Assign Owners to Known Systems

For each identified AI system, assign:

- **A business owner**
- **A technical owner**

If no owner exists, the system should be reviewed or paused.

Week 2: Classify Risk and Evaluate Vendors

The second week focuses on understanding what each AI system can do.

5 Classify AI Use Cases

For every system in the inventory, assign a risk level:

- **Low risk – internal advisory or productivity tools**
- **Medium risk – systems influencing operational workflows**
- **High risk – systems capable of autonomous actions or affecting customers**

This determines the level of governance required.

6 Begin Vendor Reviews

For third-party AI systems, ask:

- **What data is retained**
- **Whether customer data trains the model**
- **How tenant separation works**
- **What logging and audit evidence exists**

If vendors cannot clearly answer these questions, treat the system as a risk signal.

7 Identify High-Risk Systems

Flag systems that:

- **Modify production environments**
- **Access privileged systems**
- **Interact with customers**
- **Process sensitive customer data**

These systems require immediate review.

Week 3: Implement Control Boundaries

The third week focuses on operational safety.

8 Define Human Oversight Levels

Apply the three oversight models:

- **Level 1 – Human approval required**
- **Level 2 – Human interrupt window**
- **Level 3 – Post-action review**

Ensure that high-risk actions require stronger human oversight.

9 Reduce Excessive Permissions

Review agent permissions and remove:

- Shared service accounts
- Standing admin privileges
- Access to unrelated systems

AI agents should follow **least privilege** principles.

10 Enable Logging and Evidence

Ensure every production AI system logs:

- Triggering events
- Actions taken
- Data sources used
- Model versions
- Approvals when required

The organization must be able to reconstruct what the agent did.

11 Review Vendor AI Settings

Confirm that:

- Customer data is not used to train models
- Tenant boundaries are enforced
- Audit logs are available

These settings often require manual configuration.

Week 4: Formalize Governance and Customer Transparency

The final week operationalizes the program.

12 Publish the AI Governance Policy

Document:

- The AI approval workflow
- Ownership requirements
- Inventory requirements
- Oversight models
- Logging expectations

Ensure teams understand the process

13 Establish the AI Governance Review Process

Schedule a regular review cadence:

- Low-risk systems – quarterly review
- Medium-risk systems – quarterly or monthly review
- High-risk systems – monthly review minimum

14 Implement Customer Transparency Practices

Confirm the organization can explain:

- What AI systems process customer data
- What the systems do
- Whether data trains models
- What vendors are involved

Transparency and documentation are key legal protections

15 Define AI Incident Triggers

Define events that require investigation:

- Unauthorized AI actions
- Unexpected behavior changes
- Vendor model updates affecting behavior
- Logging failures
- Unexplained outputs

Outcome after 30 Days

By the end of the first month, the organization should have:

- Visibility into AI systems and agents
- Named owners for each system
- A working governance committee
- A vendor review process
- Logging and evidence standards
- Defined human oversight models
- A repeatable approval workflow

This creates a minimum viable AI governance program that allows organizations to innovate safely while maintaining accountability.

Appendix A:
AI Governance Review Form (Sample)

AI System / AI Agent Governance Review Form

Field:	Entry:
Agent Name:	
Description of Purpose:	
Business Owner:	
Technical Owner	
Vendor or Internal Build:	
Model/Provider Used:	
Risk Classification:	Low / Medium / High:
Connected Systems:	
Data Types Accessed:	
Actions Agent Can Take:	
Human Oversight Model:	Approval Required / Interruptible / Post-Review:
Date of Review:	
Reviewer(s):	

Control Review Checklist

Permission Limits

Check	Verified
Agent scope and purpose are clearly defined	
Access limited to required systems only	
Read and write access separated where possible	
No shared service accounts used	
No standing admin privileges without approval	

Permission Limits

Check	Verified
Credentials securely stored and rotated	
Sensitive systems access reviewed (PSA, RMM, Identity)	
Permissions documented in AI inventory	
Reviewer Notes	

Human Oversight

Check	Verified
Business owner has approved the use case	
Technical owner responsible for operation	
Oversight model documented	
High-risk actions require human approval	
Operators can interrupt or stop agent activity	
Escalation path defined if system behaves unexpectedly	
Customer-impacting actions reviewed	
Reviewer Notes	

Testing Procedures

Check	Verified
Expected behavior testing completed	
Edge cases and abnormal inputs tested	
Failure scenarios simulated	
Permission boundary testing completed	
Prompt injection / manipulation testing performed (if applicable)	
Rollback procedures tested	
Test results documented	
Reviewer Notes	

Logging and Evidence

Check	Verified
Logging enabled prior to deployment	
Agent identity captured in logs	
Timestamp recorded for all actions	
Triggering event logged	
Model/version recorded	
External tools or APIs logged	
Outputs and actions recorded	
Human approval events recorded (if required)	
Errors and overrides logged	
Logs stored in centralized logging system	
Log retention meets company policy	
Reviewer Notes	

Final Governance Decision

Decision	Seletion
Approved	
Approved with Conditions	
Pilot Only	
Needs Redesign	
Rejected	
Conditions or Required Changes	
Next Review Date	

Appendix B: Internal Policy Language Companies Can Build On

Here is starter policy language you can adapt to your needs.

AI Use Policy Starter

All AI systems, AI agents, AI-enabled automations, and third-party AI features used in company operations or client service delivery must be registered, reviewed, and approved according to the company's AI governance process before production use.

Each AI system must have a named business owner and a technical owner. No AI system may operate in production without documented scope, defined permissions, required logging, and an approved oversight model.

AI systems may only access data and systems necessary for their approved purpose. Privileged actions, customer-impacting actions, and destructive changes require explicit control boundaries and elevated approval.

The company will maintain audit evidence sufficient to reconstruct material AI actions, decisions, and connected system interactions. AI usage outside approved channels is prohibited.

Appendix C: AI Use Disclosure for Customer Communication

Here is starter language you can adapt to your needs.

Use of AI in Our Services

Some of the services we provide use artificial intelligence ("AI") technologies, including machine learning and large language models, to support and enhance service functionality. These technologies may be used for purposes such as summarizing information, translating content, improving customer support, or assisting with operational tasks.

Where AI is used, it operates subject to appropriate safeguards designed to protect customer information. These safeguards include data minimization, use limitations aligned with the intended service purpose, and technical and organizational security measures.

In some cases, our services rely on third party technology providers that support these AI enabled features. We require such providers to handle personal data in accordance with applicable data protection and privacy laws and to implement reasonable protections for customer information. We do not permit third party providers to use customer data to train public or unrelated AI models, except where expressly agreed.

Please review our privacy and data protection documentation for more information on AI enabled services.

N-able protects businesses from evolving cyberthreats. Our AI-powered cybersecurity platform delivers business resilience to more than 500,000 organizations worldwide, leveraging advanced end-to-end capabilities, simplified workflows, market-leading integrations, and flexible deployment options to improve efficiency and drive critical security outcomes. Our partner-first approach pairs our technology with experts, training, and peer-led events that empower customers to be secure, resilient, and successful. n-able.com

---