MIS

EXECUTIVE BRIEF · AI POLICY

AI Usage Policy *vs.* Governance Policy

Two policies, two purposes, one program.

How they differ, why you need both, and where to start.

A short executive briefing for SMB leadership, IT, and compliance teams deciding how to structure AI policy.

DOCUMENT CONTROL

Document Type	Executive Brief / Reference
Audience	Executives, IT/Security leaders, Compliance, HR, AI program owners
Companion Document	MIS AI Governance Policy (SMB Compliance-Aware Template)
Version	1.0
Prepared	April 29, 2026

Why this brief exists

SMBs often write a single “AI policy” that tries to do two jobs at once — setting strategy and dictating user behavior. The result is a document that’s too operational for executives and too abstract for users. This brief explains the difference between the two policies and shows how they reinforce each other.

MIS

1. The Bottom Line

Two distinct documents do two distinct jobs. An AI Usage Policy tells employees how they are allowed to interact with AI on a day-to-day basis. An AI Governance Policy tells the organization how it decides which AI gets used, who owns the risk, and how AI activity is overseen. Each is necessary; neither is sufficient.

> *Governance sets the rules of the game. Usage defines how people play within those rules.*

In small and mid-sized businesses, a single "AI policy" often tries to cover both — and ends up doing neither well. Executives find it too tactical, frontline staff find it too abstract, and auditors find that critical control language (vendor approval, risk tiering, oversight) is missing. Splitting the two clarifies authority, sharpens the audience for each document, and produces evidence that maps cleanly to frameworks such as NIST AI RMF, ISO 42001, HIPAA, CMMC, PCI DSS, and customer security questionnaires.

The rest of this brief unpacks what belongs in each policy, walks through real scenarios, flags the most common pitfalls, and finishes with a quick self-assessment so you can see where your program stands today.

2. At a Glance

Use this comparison as a reference card. It expands the typical "usage vs. governance" table by adding the dimensions that matter most when you are scoping, approving, or auditing each document.

DIMENSION	AI USAGE POLICY	AI GOVERNANCE POLICY
Primary purpose	Define acceptable and unacceptable behavior with AI tools and data.	Define how the organization controls, oversees, and approves AI activity.
Audience	Employees, contractors, end users, vendors with workforce access.	Executives, AI Governance Committee, IT/Security, Privacy/Compliance, Legal.
Level	Tactical / operational — day-to-day behavior.	Strategic / oversight — risk, accountability, decision rights.
Owner	HR, IT, or department lead, with input from Security.	Executive Sponsor and AI Governance Committee, with Legal and Compliance.
Approval authority	Manager / IT for individual tool use within rules.	Executive / committee for risk tiers, regulated data, and high-impact use.

DIMENSION	AI USAGE POLICY	AI GOVERNANCE POLICY
Goal	Prevent misuse, leakage, and unsafe output.	Manage risk, alignment, and regulatory and contractual exposure.
Typical content	Acceptable use, prohibited actions, data handling, incident reporting, consequences.	Roles, risk tiering, approval workflow, vendor due diligence, monitoring, exceptions.
Change frequency	Updated as tools, threats, and workflows shift — often quarterly.	Reviewed at least annually and after material AI, regulatory, or business changes.
Evidence produced	Training records, acknowledgments, ticket/incident logs, DLP alerts.	AI inventory, risk assessments, approval records, vendor reviews, BAAs/DPAs.
Enforcement	HR / IT — access removal, retraining, disciplinary action.	Executive / Compliance — contract remedies, regulator notification, program changes.
Failure mode	Shadow AI, regulated data leakage, hallucinated customer-facing content.	Unapproved vendors, no audit trail, unmanaged agents, compliance gaps.

3. AI Usage Policy — The Operational Layer

An AI Usage Policy is the rulebook your workforce reads and acknowledges. It sits alongside your acceptable use, data classification, and information security policies and translates organizational risk decisions into specific instructions a person can follow at their keyboard.

3.1 What belongs inside

Acceptable use. Which AI tools are approved, for what kind of work, and with what data classifications.

Prohibited actions. hat must never happen — entering regulated data into public tools, using personal accounts for Company work, generating unverified customer or legal content.

Data handling expectations. Redaction, minimization, and what to do before pasting customer, employee, or financial information into a prompt.

Personal device and account rules. Whether AI may be used from BYOD, personal accounts, or browser extensions, and the configuration required if so.

Output verification. The user's responsibility to fact-check, review citations, and review AI-generated code or communications before use.

Incident reporting. Where and how to report suspected leakage, prompt injection, suspicious outputs, or unapproved tool discovery.

Consequences. What happens when the rules are broken — retraining, access removal, disciplinary action, contract impact.

3.2 Who owns it

Typically, HR, IT, or a department lead, with input from security and compliance. The owner's job is to keep the document readable, current, and tied to a workable training and acknowledgment process.

3.3 In practice

SCENARIO · Marketing wants to use a public AI tool for blog drafts

Situation: *A marketing manager has been pasting product copy and customer testimonials into a free public AI chatbot to speed up draft writing.*

What each policy does: The Usage Policy provides the rule — customer information may not be entered into unapproved public tools, and only listed enterprise AI tools are approved for marketing content. The policy points the manager to the approved tool, names what may be pasted (public materials, draft copy with no PII), and tells them to retain a human review step before publishing.

SCENARIO · An engineer pastes source code into a freemium coding assistant

Situation: *An engineer enables a free coding assistant browser extension that uploads code snippets to an unknown vendor for completions.*

What each policy does: The Usage Policy specifies that source code may only be processed by approved enterprise tools with the right contract, retention, and training-disabled settings. The policy tells the engineer to disable the extension, report the activity, and use the approved tool. It also defines the consequence path if the rule is repeatedly ignored.

3.4 Signs your usage policy is weak

- It tells people to "be careful" without listing approved tools and prohibited data types.
- It does not address personal accounts, browser extensions, or AI features embedded inside SaaS products already in use.
- There is no clear path to report an AI incident, and no examples of what counts as one.
- It has not been updated since a major new tool, vendor, or AI feature was introduced.
- Workforce members have not signed an acknowledgment within the last twelve months.

4. AI Governance Policy — The Oversight Layer

An AI Governance Policy is the document leadership, auditors, and customers will ask for. It defines how the organization decides whether to adopt an AI capability, how it tiers risk, who approves what, and how AI activity is monitored over time. It produces the evidence trail that maps to recognized frameworks and customer security questionnaires.

4.1 What belongs inside

Roles and decision rights. Executive Sponsor, AI Governance Committee, AI Owners, IT/Security, Privacy/Compliance, Legal, HR, and the workforce.

Risk tiering and approval. How AI use cases are classified (low through critical/regulated) and what approvals are required at each tier.

AI inventory and classification. An auditable list of approved, pending, retired, and denied AI use cases, with data scope and owners.

Vendor due diligence. Required contracts (BAA, DPA, AOC), security attestations, data-use restrictions, and AI-specific safeguards.

Compliance overlay. How AI is treated under HIPAA, CMMC, PCI DSS, GLBA, GDPR, state privacy laws, SOC 2, ISO 42001, and customer commitments.

Monitoring and evidence. Logging, access reviews, performance monitoring, exception tracking, and management review cadence.

Incident response and escalation. AI-specific incident categories, regulatory escalation, and post-incident learning.

Continuous improvement. Annual review, change triggers, and metrics reported to leadership.

4.2 Who owns it

The Executive Sponsor, with the AI Governance Committee. Privacy/Compliance, IT/Security, and Legal are formal contributors. Sign-off authority for high-risk and regulated-data AI sits with the Executive Sponsor and the committee — not with individual managers.

4.3 In practice

> **SCENARIO · Sales wants to deploy an AI agent that drafts customer emails**
>
> **Situation:** *A sales leader requests an AI agent that reads CRM data, drafts outbound emails, and — if approved — sends them automatically. The agent would touch customer contracts, pricing, and contact data.*

What each policy does: The Governance Policy classifies this as a high-tier use case (customer-facing, customer confidential data, autonomous action). The intake form runs through Legal, IT/Security, and Privacy/Compliance. Vendor terms are reviewed for data training, retention, and breach commitments. Approval requires human-in-the-loop on send, a kill switch, logging of every drafted and sent message, and a quarterly review. The Usage Policy then translates the approval into clear rules for the sales team using it.

SCENARIO · A SaaS vendor turns on an embedded AI feature without notice

Situation: *An existing helpdesk SaaS application enables a new generative AI summarization feature by default. It begins reading ticket bodies that contain customer PII and, in some cases, ePHI.*

What each policy does: The Governance Policy is the reason this is a controlled event rather than an undetected leak. Vendor change-notice expectations, embedded-feature handling, and data classification rules trigger a review. The committee can disable the feature, require contractual amendments, and update the AI inventory. The Usage Policy then tells staff what to do with the feature in the meantime.

4.4 Signs your governance policy is weak

- There is no AI inventory — nobody can produce a current list of approved AI tools and use cases.
- Risk tiering is informal or undocumented; the same approval path is used for low-risk drafting and customer-facing automation.
- Vendor reviews do not cover AI-specific risks (training on customer data, prompt and output retention, agent permissions).
- AI incidents have no defined category, owner, or escalation path distinct from generic IT incidents.
- Customer or auditor requests for evidence of AI oversight require ad-hoc scrambling rather than producing existing records.

5. How the Two Policies Work Together

Governance and usage are not parallel documents — they are layered. Governance decides what is permitted and under what conditions. Usage takes those decisions and turns them into concrete behavior the workforce can follow. Training, monitoring, and incident response operate across both.

Layer	What it does	Primary artifacts
Governance (top)	Sets risk appetite; classifies use cases; approves tools, vendors, and high-impact AI; defines roles and oversight.	AI Governance Policy, AI inventory, risk assessments, vendor reviews, committee minutes.
Usage (middle)	Translates governance decisions into operational rules people follow with approved tools and data.	AI Usage Policy, approved-tool list, quick rules, training and acknowledgment records.
Behavior (bottom)	What actually happens at the keyboard when a user is drafting, summarizing, integrating, or automating.	Tickets, prompts, outputs, DLP alerts, audit logs, customer-facing artifacts.

When the layers are aligned, a customer questionnaire, regulator inquiry, or internal audit can be answered with the same evidence chain that runs the program day-to-day. When they are misaligned, the documents contradict reality and the program has to be reconstructed from memory each time someone asks.

6. Where to Start

A common SMB instinct is to write the Usage Policy first because it is the most visible to staff. We recommend the opposite sequence. Usage rules that are not anchored in governance decisions tend to be guesses dressed up as policy.

Step	What to do	Output
1. Frame governance	Adopt an AI Governance Policy. Name an Executive Sponsor and committee. Define risk tiers and approval paths. Stand up an AI inventory and intake process.	Approved Governance Policy; named owners; baseline AI inventory; risk-tier definitions.
2. Translate to usage	Write the AI Usage Policy from the governance decisions. List approved tools, prohibited actions, data rules, and the incident path. Keep it short and readable.	Approved Usage Policy; workforce quick-rules card; acknowledgment workflow.
3. Train and operationalize	Roll out training, configure SSO/MFA/logging on approved tools, communicate the incident path, and begin the review cadence.	Training completion records; tool configuration evidence; incident playbook in use.
4. Review and improve	Run the quarterly and annual reviews defined in the Governance Policy. Update both documents when tools, vendors, regulations, or business priorities change.	Updated inventory and policies; metrics report; lessons-learned records.

7. Common Mistakes to Avoid

One policy for both jobs. A single “AI policy” that mixes strategy and behavior is hard to maintain and difficult to use as audit evidence. Split them, and let each speak to its real audience.

Usage rules without an inventory. If the Usage Policy says “only approved AI tools,” there must be a published list. Without an inventory, the rule is unenforceable and indistinguishable from no rule at all.

Governance without operational reach. A polished Governance Policy with no Usage Policy or training behind it produces shelfware. Decisions made by the committee never reach the people doing the work.

Treating embedded AI features as new tools every time. Many AI risks now appear as features inside existing SaaS. Both policies must address embedded AI explicitly — default-on summarization, copilots, and agents — not just standalone AI products.

Skipping vendor AI-specific review. A standard SOC 2 or vendor questionnaire does not cover model training on customer data, prompt and output retention, agent permissions, or model-update notification. Governance must add these explicitly.

No clear incident definition for AI. Generic IT incident categories miss the most common AI events — regulated data pasted into a public tool, prompt-injection-driven exfiltration, an agent taking an unauthorized action. Define them, or they will be missed.

8. Quick Self-Assessment

A short check for executive and program owners. If you cannot confidently answer “yes” to most of these, it is time to revisit the relevant policy.

Governance

☐ We have a written AI Governance Policy approved at the executive level.

☐ We can produce a current AI inventory listing approved, pending, and denied use cases with owners and risk tiers.

☐ Risk tiers, approval paths, and the standard for regulated-data use cases are documented and applied consistently.

☐ AI vendors are reviewed against AI-specific criteria — training on our data, prompt/output retention, agent permissions, model-update notice.

☐ We have an AI incident category, an escalation path, and at least one practiced response.

Usage

- ☐ We have a written AI Usage Policy that names approved tools and lists prohibited data types.
- ☐ Workforce members have completed AI acceptable-use training and signed an acknowledgment in the last twelve months.
- ☐ The Usage Policy addresses personal accounts, browser extensions, and AI features embedded in existing SaaS.
- ☐ Reporting an AI concern — regulated data exposure, suspicious output, unapproved tool — is a documented, well-known path.
- ☐ When a new AI tool or feature is approved through governance, the Usage Policy or quick rules are updated within a defined window.

Alignment

- ☐ Both policies reference each other and use the same definitions for data classification, risk tier, and AI use case.
- ☐ Approval decisions made by the AI Governance Committee are visibly reflected in the workforce-facing Usage Policy or its supporting materials.
- ☐ Audit and customer evidence requests can be answered using existing records rather than ad-hoc effort.

9. Companion Document

This brief explains the distinction between AI Usage and AI Governance Policies. The companion document, the MIS AI Governance Policy (SMB Compliance-Aware Template), provides a full governance policy ready to be tailored to your organization — covering scope, definitions, governance principles, roles and RACI, AI inventory and classification, risk tiering and approval, acceptable and prohibited uses, data protection, compliance overlays for HIPAA / CMMC / PCI DSS / GDPR, vendor due diligence, AI development and agent controls, output review, monitoring and evidence, incident response, training, and continuous improvement, plus appendices for intake, risk assessment, vendor questions, workforce quick rules, compliance mapping, and a 30/60/90-day implementation roadmap.

Recommended next step

Treat this brief as the framing document for an executive conversation. Pair it with the MIS AI Governance Policy template to anchor the program, then derive a focused, plain-language AI Usage Policy from the approved governance decisions. Adopt, train, and review on the cadence the governance document defines.