



Professional Services, Software & Solutions Group

GUIDE

The AI Readiness and Governance Guide

How to adopt AI responsibly: an assessment, a policy template, a risk method and a governance operating model for organizations of every size.

PS3G Inc.

1201 N Market Street, Suite 111, Wilmington, DE 19801 · Email hello@ps3g.com

Contents

- **01** What AI governance is, and is not
- **02** The AI readiness assessment
- **03** Shadow AI: finding and governing the tools people already use
- **04** An AI acceptable-use policy template
- **05** The AI risk assessment method
- **06** Data governance prerequisites for AI
- **07** A governance operating model for small, mid-size and enterprise organizations
- **08** Controls for generative AI systems
- **09** A 30-60-90 day rollout plan and a maturity checklist

Who this is for

The people who have to answer "are we allowed to use AI for this?" and then make the answer stick: business owners, founders, department heads, compliance and privacy leads, and IT and engineering leaders. It applies equally to a twelve-person professional services firm whose staff paste client documents into a chatbot, a regional healthcare group evaluating an ambient documentation tool, and an enterprise with a dozen copilots in production. The mechanics scale; the principles do not change.

How to use this document

Read Sections 01 and 02 first; Section 02 is a scored assessment you can complete in one working session. Sections 03 to 08 each stand alone: shadow-AI discovery, a one-page acceptable-use policy you can adopt as-is, a risk classification method with a worked example, the data prerequisites most AI projects trip over, three operating-model variants sized to your organization, and a control set for generative AI. Section 09 turns it all into a 30-60-90 day plan with a maturity checklist.

"AI" here covers predictive models, generative models and LLMs, retrieval-augmented (RAG) applications, copilots embedded in business software, and agents that take actions on a user's behalf. Laws and frameworks are cited factually, as a map of what exists; nothing here is legal advice. Where we give a number it is a rule of thumb from PS3G's own engagements, not a statistic.

01 What AI governance is, and is not

AI governance is the set of decisions, roles, records and controls that let an organization say, with evidence, which AI systems it uses, what they may do, who is accountable, and how it knows they work as intended.

Governance is not an ethics statement, a one-time training video, or a committee that approves nothing. It is also not a ban: organizations that prohibit AI outright end up with the most unmanaged usage, because the work moves to personal accounts and phones. Governance is an inventory, a decision record per system, named owners, impact classification, proportionate controls, a policy people have read, a review cadence, and evidence you could show a customer, auditor or regulator. It became a business requirement when generative AI entered daily work, customers started asking about it in contracts, and the frameworks below gave it a vocabulary.

Reference	What it is, factually	Why it matters to you
NIST AI RMF 1.0	Voluntary US framework with four functions: Govern, Map, Measure, Manage.	The common language for AI risk in US contracts.
ISO/IEC 42001	Certifiable international management-system standard for AI, structured like ISO 27001.	The route to a certificate if customers or regulators ask for one.
EU AI Act	Risk-tiered regulation: prohibited practices, high-risk obligations, transparency duties, general-purpose model rules; being phased in.	Applies if your AI systems or output reach the EU market.
Colorado AI Act (SB 24-205)	State law now taking effect; duties on developers and deployers of "high-risk" AI used for consequential decisions (employment, lending, housing, healthcare, insurance, education).	Requires reasonable care, impact assessments, notices and a risk-management program.
Sector rules	HIPAA (PHI and business associate agreements), financial model-risk guidance, FTC consumer-protection authority, employment law.	Already apply to AI: a tool touching PHI needs a BAA; a hiring model is subject to employment law.

Not legal advice. The table is a map, not an opinion on your obligations; confirm specifics with counsel.

02 The AI readiness assessment

The assessment scores five dimensions separately, because the weakest dimension sets the pace. Score each row 0 (not in place) to 4 (in place, documented, reviewed in the last year).

Dimension	Statement to score	Score 0–4	Evidence you would point to
People	A named person or group owns AI decisions and can say no.		Charter, meeting notes
	Staff have been trained on what they may put into AI tools.		Training record, policy acknowledgment
Process	There is a written path from "I want to use this tool" to a decision.		Intake form, decision log
	AI vendors go through the same procurement and security review as other software.		Vendor review checklist
Data	Data is classified (at minimum: public, internal, confidential, regulated).		Classification standard
	We can say which systems hold personal data, PHI or payment data.		Data inventory or map
Technology	Single sign-on and MFA cover the tools people use for work.		Identity provider coverage report
	We can see which SaaS and AI tools are in use from identity, network or expense data.		SaaS discovery report
Risk	A risk register exists and the laws and contracts constraining AI use are mapped.		Risk register, obligations list
	We know which decisions about people (hiring, pricing, eligibility, care) could involve AI.		List of consequential decisions

Reading the result

Add each dimension (maximum 8); the pattern matters more than the total. A dimension at 0–2 will block any serious AI project and belongs in the first 30 days of the Section 09 plan. At 3–5 it is workable: proceed with low-impact use cases while closing gaps. At 6–8 it is an asset. Most organizations we meet score highest on Technology and lowest on Process and Risk.

03 Shadow AI: find it, then govern it

Every organization already uses AI; the only question is whether leadership knows. "Shadow AI" is the use of AI tools, features and accounts that were never reviewed: a free chatbot on a personal email, a browser extension that summarizes pages, an AI feature a SaaS vendor switched on in an update. Nobody is acting in bad faith; they are trying to get work done faster.

How to find it

- 1 Ask, without penalty.** A short anonymous survey ("which AI tools have you used for work this month, and for what?") finds more than any scan, if people believe the answer will not be used against them.
- 2 Look at identity, network and expenses.** Identity-provider OAuth consents, DNS and proxy logs, and expense reports each reveal AI services in use.
- 3 Read your vendors' release notes.** The AI you did not choose is often inside tools you already pay for: CRM, office suite, ticketing, video conferencing, EHR. Check whether those features are on and where the data goes.

How to govern it

Triage into three buckets. **Sanction** tools that are fit for purpose: move people to a business account with single sign-on, a data-processing agreement, training opt-out and retention controls, and announce it as the approved route.

Substitute tools that are useful but cannot meet your terms, by offering an approved equivalent before switching the old one off. **Stop** only the genuinely unacceptable cases, and explain what to use instead. Repeat the pass quarterly.

Finding (illustrative)	Bucket	Action
Marketing team drafting with a consumer chatbot on personal accounts	Sanction	Business seats with SSO and no-training terms; migrate within 30 days
Finance staff pasting invoices into a free summarizer extension	Substitute	Enable the equivalent feature in the approved office suite; remove the extension
Clinical staff using an unreviewed transcription app on patient encounters	Stop, then substitute	Stop now (PHI, no BAA); fast-track review of a compliant ambient documentation vendor

04 An AI acceptable-use policy template

Replace the bracketed items and adopt. Pair it with an Approved AI Tools list (tool, approved data level, approved uses, owner); the list changes monthly, the policy should not.

[Organization] – AI Acceptable Use Policy (v1.0, [date])

1. Purpose. How staff, contractors and partners of [Organization] may use AI tools (chat assistants, copilots, code assistants, AI features in business software, agents) while protecting our clients, data and obligations.

2. Approved tools. Use only tools on the Approved AI Tools list, through your work account; personal accounts, free tiers and unlisted tools must not be used for work data. Request new tools via the AI tool request form; [owner] responds within [10] days.

3. Data you may use. Public and internal data may be used in any approved tool; confidential data only in tools marked "Confidential-approved"; regulated data (personal data, PHI, payment card data, client-confidential material) only in tools marked "Regulated-approved" for that type.

4. Your responsibility for output. AI output is a draft, not a decision. You are responsible for anything you send, publish, submit or act on: verify facts, figures, citations and code, and do not present AI output as professional judgment where a client, patient or regulator expects it.

5. Decisions about people. AI must not be the sole basis for a decision that materially affects a person's employment, pay, credit, housing, insurance, education, healthcare or access to services. Such uses require a documented risk assessment first, and a human must make and explain the final decision.

6. Agents and automation. Any AI that can take actions (send messages, change records, execute transactions, run code) must pass the risk assessment, run with least access, and have human approval for irreversible actions or those above [threshold].

7. Security and reporting. Do not disable logging, safety or data-protection settings. Report harmful, discriminatory, confidential or clearly wrong output, or data entered contrary to this policy, to [contact] within one working day. Good-faith reports are never penalized.

8. Review. Owned by [owner]; reviewed at least [annually] or on a material change.

Acknowledgment. I have read and understood this policy. Name: _____ Date: _____

05 The AI risk assessment method

Four lightweight steps that work for a small firm reviewing its first copilot and for an enterprise with hundreds of systems, because the work scales with the impact of the system, not the size of the organization.




Figure 1 – The four-step AI risk assessment loop; Step 4 feeds back to Step 1 on any material change.

Step 1 – Inventory

One row per AI system, feature or agent: name, vendor, owner, purpose, users, data read and written, actions taken, and effect on decisions about people.

Step 2 – Classify by impact

Class	Typical characteristics	Illustrative examples
Minimal	Internal productivity; a human reviews every output; no regulated data; no actions; no effect on decisions about people.	Drafting internal memos; summarizing public research; code suggestions reviewed in pull requests.
Limited	Confidential or customer data; output reaches customers; or low-consequence, reversible actions; human in the loop.	Customer-service draft replies; RAG over internal documents; meeting transcription.
High	Influences consequential decisions about people; regulated data at scale; autonomous action with irreversible effect; or a regulated category under law.	Resume screening; clinical decision support; eligibility or pricing; agents that execute payments or change production systems.

SECTION 05 (CONTINUED)

Controls and sign-off

Step 3 – Apply controls proportionate to the class

Control	Minimal	Limited	High
Vendor terms	Reviewed once per tool	Plus data-processing agreement or BAA where applicable	Plus contractual audit and incident-notification rights
Human review of output	User's own judgment	Defined reviewer before external use	Documented human decision-maker
Access and logging	SSO, MFA; vendor default logs	Role-based access to source data; prompts and outputs retained per schedule	Least privilege; full audit trail with model version, inputs, outputs, reviewer
Testing, notice, cadence	None; annual review	Spot-check accuracy before launch; disclose AI assistance; review on change	Pre-launch evaluation for accuracy and bias; monitoring; impact assessment; notice to affected people; quarterly review

Step 4 – Sign-off

Every system gets a one-page decision record: classification, controls, residual risks accepted, who accepted them, the date, and the re-review trigger. For Minimal systems this is a line in a register the AI owner signs. For High systems the accountable business executive signs, not IT. Where a legal high-risk definition applies (EU AI Act, Colorado AI Act), it is authoritative and this table is the floor.

Worked mini-example (illustrative). A sixty-person insurance brokerage wants a vendor assistant to draft replies to policyholder emails. Inventory: vendor-hosted LLM, reads the email and policy record, takes no action. Class: *Limited* – customer-facing output, confidential data, a human sends every reply. Controls: no-training terms and a data-processing agreement; replies reviewed by a licensed agent; prompts and outputs logged; thirty drafts spot-checked. Sign-off: head of customer operations; re-review in twelve months. Had the tool decided claim eligibility, the class would be *High*, sign-off would move to the chief operating officer, and an impact assessment would come first.

06 Data governance prerequisites for AI

Most AI projects that stall do so for data reasons, not model reasons. A RAG assistant is only as trustworthy as the documents it retrieves; a copilot with access to a shared drive exposes whatever the drive exposes; a model trained on a table with no lineage cannot be defended. Five prerequisites are the minimum; each can start as a spreadsheet.

1. Classification

Four levels suffice: public, internal, confidential, regulated. Label repositories, not just documents, and connect an AI tool only to repositories whose level it is approved for.

2. Retention

AI creates new copies: prompts, outputs, embeddings, cached context, vendor logs. Decide before launch how long each is kept and who can delete it, and confirm the vendor's retention matches.

3. Lineage

For every dataset an AI system reads or learns from, record source, last refresh, owner and transformations. For RAG, know which document version was answered from; for models, be able to reproduce the training set.

4. Access

If a copilot can read every file a user can, over-shared folders become a search engine for secrets. Review sharing before connecting an assistant, and give it a least-privilege identity of its own.

5. Personal data, PII and PHI handling

- ☐ Is there a data-processing agreement, and for PHI a business associate agreement, with the vendor and each sub-processor in the chain?
- ☐ Is customer data excluded from model training by contract, enforced at the tenant level rather than left to each user?
- ☐ Can you honor a deletion or access request when some data sits in vendor logs, embeddings or caches?

A simple test. If you cannot answer "what data did this assistant have access to on the day it gave that answer?" within an hour, the prerequisites are not in place. The fastest fix is narrow: govern the two or three sources the first use case needs, then extend.

07 A governance operating model for small, mid-size and enterprise organizations

The same four questions need answering everywhere: who decides, who reviews, how often, and where is the record. Pick the column that matches you today; move right only when volume makes the current model a bottleneck.

Element	Small (roughly under 50 people)	Mid-size (roughly 50–1,000)	Enterprise (1,000+, multi-entity or regulated)
Who decides	The owner or managing partner, advised by whoever runs IT.	A named AI owner (often head of IT or compliance) with a working group: IT, legal, one business lead.	An AI governance committee under an executive sponsor: legal, privacy, security, data, risk, business units.
Who reviews	The same person; outside help for anything High.	Working group reviews Limited and High; AI owner approves Minimal.	Business unit approves Minimal; AI office reviews Limited; committee approves High.
Intake and record	One inbox; one spreadsheet as the register.	A service-desk ticket type; register with decision records attached.	GRC workflow; system of record linked to asset inventory, vendor management and risk register.
Cadence	Register reviewed quarterly in an existing meeting; policy annually.	Working group monthly; register and discovery quarterly; policy annually.	Committee quarterly; AI office weekly; continuous monitoring of High systems; internal audit.
Training and alignment	One-hour session for everyone; NIST AI RMF as vocabulary.	Role-based training; NIST AI RMF profile mapped to existing SOC 2 or ISO 27001 controls.	Tracked curriculum; ISO/IEC 42001 where expected; regulatory mapping (EU AI Act, state and sector rules).

Three rules that hold at every size

Business owns the decision; IT owns the control. Whoever benefits from the system and bears the consequence of its output signs the risk acceptance; IT, security and data teams make the controls real. **One register, one policy, one approved-tools list.** Multiple lists means no list. **Approval has a service level.** If a Minimal request takes a month, people stop asking; fast approval of low-impact tools buys the authority to say no to high-impact ones.

08 Controls for generative AI systems

LLM applications, RAG assistants, copilots and agents have failure modes conventional controls do not cover. The OWASP Top 10 for LLM Applications is the standard catalog and a good build-or-buy checklist.




Figure 2 – Control points in a generative AI application; logging and evaluation span every stage.

Risk (OWASP LLM reference)	Controls
Prompt injection (LLM01): hidden instructions in a document, page or email make the assistant ignore its rules, leak data or act.	Treat retrieved content as untrusted; separate system instructions from user content; restrict callable tools; human approval for consequential actions; test injection patterns.
Sensitive information disclosure (LLM02): the assistant reveals data the user should not see, or confidential data leaves in prompts.	Scope retrieval to the user's own permissions; redact identifiers before the prompt; no-training and retention terms; data-loss prevention on AI endpoints.
Improper output handling and excessive agency (LLM05, LLM06): output executed unchecked; an agent with more power than its task needs.	Validate output before it reaches a browser, database, shell or API; least-privilege identity per agent; tool allow-list; human approval for irreversible steps; kill switch.
Supply chain and model change (LLM03): the vendor swaps the model or changes terms and behavior shifts silently.	Pin model versions where the API allows; re-run the evaluation set on any change; contractual notice of material changes; record the version in every log entry.

Two practices make the rest verifiable: an **evaluation set** run before launch and after every change, and **logging** of every interaction on Limited and High systems.

09 Rollout plan and maturity checklist

The plan assumes the common starting point: some AI in use, no register, no policy. A small firm does each item with one or two people, an enterprise with a working group; the sequence is the same. The aim is to make the yes to low-impact uses faster and the no to high-impact uses rarer.

Phase	What gets done
Days 1–30 See and stabilize	Name the AI owner and, if larger, the working group (07). Run the readiness assessment (02) and the shadow-AI discovery pass (03). Build the inventory (05). Adopt the policy and publish an approved-tools list (04). Stop clearly unacceptable uses, offering a substitute.
Days 31–60 Classify and control	Classify every inventory row (05). Review vendor terms; obtain DPAs and BAAs. Move sanctioned tools to business accounts with SSO. Fix data prerequisites for the first Limited use case (06). Sign decision records for Limited systems. Train all staff.
Days 61–90 Operate and prove	Complete impact assessments and executive sign-off for each High system, or retire it. Put generative AI controls on production systems (08). Hold the first governance review. Map the register to NIST AI RMF and applicable law. Score the maturity checklist and set a twelve-month target.

The maturity checklist

Levels are cumulative; Level 3 is a sensible twelve-month target at any size.

Level	Name	You are here when...
1	Ad hoc	AI is in use but there is no inventory, policy or owner.
2	Aware	Owner named; inventory exists; policy published and acknowledged; approved-tools list exists; unacceptable uses stopped.
3	Managed	Every system classified; proportionate controls evidenced; DPAs/BAAAs in place; decision records signed by business owners.
4	Measured	High systems have impact assessments and monitoring; generative systems have evaluation sets and logging; controls tested by audit.
5	Optimized	Controls enforced automatically by procurement, data, identity and risk systems; ISO/IEC 42001-ready; any regulator question answered in a day.



Professional Services, Software & Solutions Group

About PS3G

PS3G Inc. is an IT services and software company headquartered in Wilmington, Delaware, with a delivery center in India that lets us work around the clock. We design, build, secure and run software for organizations of every size.

Our services include websites and web applications, mobile apps, .NET and Microsoft engineering, cloud migration and DevOps, data platforms and data governance, AI consulting and development, AI governance, cybersecurity and compliance, software testing, managed services, and IT staffing with dedicated teams.

We also build and run our own products. DocVilla is a HIPAA-compliant cloud EHR, telehealth and practice-management platform used by medical practices across the US. We apply the guidance in this document to our own systems.

How we can help with AI governance

- **AI readiness assessment:** the Section 02 assessment run with your team, with a written gap report and a prioritized plan.
- **Shadow-AI discovery and register build:** identity, network and survey-based discovery, an inventory you own, and a first classification pass.
- **Policy, operating model and data governance:** acceptable-use policy, approved-tools list, decision-record templates, a governance model sized to you, and classification, access and retention for the data your first use cases depend on.
- **Secure generative AI builds and framework alignment:** RAG assistants, copilots and agents built with the Section 08 controls in, on Azure, AWS or Google Cloud; mapping to NIST AI RMF, ISO/IEC 42001 and the laws that apply to you.

Talk to us

Our engagement model is simple: email us → we reply with questions → a 30-minute video call we schedule → a short assessment → a written proposal. We work as fixed-scope projects, dedicated teams, managed services or advisory, whichever fits.

Email hello@ps3g.com.

PS3G Inc., 1201 N Market Street, Suite 111, Wilmington, DE 19801

This guide is provided for general information. It is not legal advice and does not create any professional relationship. References to laws and frameworks are factual summaries as understood by the PS3G team at the time of writing; confirm applicability and current requirements with qualified counsel.