Enterprise AI Governance Checklist

Balancing security with citizen developer empowerment

	VERSION CONTROL/TRACEABILITY	
	Action item	**Guidance**
☐	Require Git (or equivalent) for all AI-generated code, prompts, and skills	*No vibe code goes to production without a commit history*
☐	Enforce ownership tagging on every agent, skill, and connector	*Who built it, who maintains it, who approves changes*
☐	Log all prompt changes with timestamps and authors	*Treat prompts like code — they are code*
☐	Define a rollback procedure for AI-driven automations	*One broken prompt shouldn't require a helpdesk recovery operation*
	SECURITY/DATA PROTECTION	
	Action item	**Guidance**
☐	Classify data sensitivity levels and restrict what can be fed into AI tools	*PII, regulated data (HIPAA, GDPR), and trade secrets need explicit rules*
☐	Confirm LLM usage is private and excluded from training pipelines	*Verify model contracts — not all enterprise tiers are equal*
☐	Define which data sources agents and skills can connect to	*MCP filesystem, database, and API connectors need access controls*
☐	Audit data flows quarterly for shadow AI usage	*Shadow AI grows faster than official rollouts*
☐	Deploy input scanning (Presidio, Comprehend, or regex+NLP) on user-supplied text	*Detect SSNs, policy IDs, DOBs, and medical terms before they reach the LLM*
☐	Scan uploaded files for sensitive content at the point of ingestion	*Extract text from PDFs and docs and run the same classifier before the AI sees them*
☐	Replace free-form prompt fields with structured intake templates where possible	*Typed fields reduce accidental data exposure by design*
☐	Add a user attestation step before submitting inputs containing sensitive data	*A simple confirmation prompt shifts accountability and creates a paper trail*
☐	Inspect LLM outputs for sensitive data echoed back in responses	*If a full SSN appears in output, it was in the input — a reactive signal for audits*
☐	Surface a persistent data handling reminder inside all AI-enabled tools	*Keep policy top-of-mind without requiring users to recall training docs*
	AI DEVELOPMENT LIFECYCLE	
	Action item	**Guidance**
☐	Establish dev / staging / production environment separation for AI tools	*No prototyping directly in prod — this is the #1 incident cause*
☐	Require peer review or sign-off before any agent goes live	*A second set of eyes catches conflicting prompt logic*
☐	Define a change communication process for shared services	*Downstream users need to know when a shared skill changes*
☐	Create a lightweight AI change request template	*A Google Form beats no process at all*

VERIFICATION AND VALIDATION (TESTING)		
	Action item	**Guidance**
☐	Train citizen developers on non-determinism and hallucination risk	*They are excited innovators, not adversaries — education works*
☐	Require output validation steps in any AI workflow touching decisions	*Human-in-the-loop checkpoints for high-stakes outputs*
☐	Build test prompts and expected outputs into every skill or agent	*Regression-test AI like you would test code*
☐	Document known failure modes for each deployed agent	*What does it get wrong, and how often?*
INSTRUMENTATION		
	Action item	**Guidance**
☐	Instrument token consumption per team, project, and use case	*Runaway token spend can invert ROI fast*
☐	Set spend alerts and hard limits on API and connector usage	*Budget guardrails — not kill switches*
☐	Track storage and compute costs tied to AI workflows	*Cloud costs compound — log them from day one*
ENABLEMENT		
	Action item	**Guidance**
☐	Publish a curated library of approved, reusable skills and connectors	*Speed up builders; reduce reinvention and shadow risk*
☐	Provide AI basics + development practices training (Git, testing, prompting)	*Paired with tools, not just policy docs*
☐	Create an internal AI ideas / showcase channel	*Celebrate wins — it drives adoption and surfaces best practices*
☐	Assign an AI enablement contact per business unit	*Someone to answer questions before problems escalate*
☐	Define a fast-track approval path for low-risk automations	*Governance should not be slower than the innovation it governs*

Remember: you are liable for what AI does or does not do. Governance is not a constraint on growth — it is the condition for sustainable growth.