Creating Responsible AI Usage Guidelines for the Enterprise

Generative AI adoption isn't slowing down. Across enterprises, business teams are already using tools like ChatGPT, Copilot, and other AI-powered assistants—often without formal approval. These tools boost productivity but also pose risks.

Unvetted AI usage can lead to data leakage, regulatory exposure, and reputational damage. According to a recent report on GenAI data leakage, 8.5% of employee prompts included sensitive data posing a serious threat to a company's security, privacy, legal and regulatory compliance. And most organisations don't know it's happening until it's too late.

For digital governance leaders, the question isn't whether AI should be used—it's how to define enforceable policies aligned with business goals, legal obligations, and data standards.

This article outlines a practical, phased framework to build a responsible AI governance model.

Generative AI use is rising often unapproved creating data and compliance risks; this article offers a phased framework for responsible governance.

The Ground Reality of AI Use in GCC Enterprises

Business users are using with generative tools in ways leadership often doesn't see—until risks surface:

- A finance team uploads budget reports to a public GenAI tool.
- Customer service uses AI-generated replies with biased or legally risky language.
- Marketing integrates AI through browser plug-ins that bypass IT controls.

These are driven by productivity, but lead to regulatory scrutiny, data leaks, and confusion.

The issue isn't the tools—it's the lack of oversight. Few enterprises have a policy that defines:

- Approved AI tools
- Permissible data types
- Review responsibility
- Violation escalation

Most cybersecurity strategies still focus on systems—not user behaviour. But AI adoption starts at someone's keyboard.

Where Most Organisations Get AI Policy Wrong

Policies exist but often are not designed for how AI is used. Many respond with vague guidelines or blanket bans, which do not reduce risk but just delay exposure. Here are the most common failures we see in GCC organisations, and what they look like in practice:

Tool first, policy later

Teams often start using AI tools before confirming if they're safe, compliant, or even permitted. There's usually no vetted list of approved platforms, and no formal risk review process for free-to-use plugins or SaaS-based tools. Without a policy-first approach, governance becomes reactive and fragmented.

Data blindness

No one is tracking what types of data employees are feeding into AI tools. Sensitive numbers, customer records, even internal strategy documents—all of it can be shared with third-party models without controls or logs. If you oversee compliance or data protection, this is a quiet breach waiting to surface. And in regulated sectors, it may already count as one. What you don't see in usage patterns can still show up in audit findings.

Undefined accountability

When AI-generated content introduces risks like misleading claims, biased responses, and hallucinated figures, no one knows who is accountable for the mistake. This ambiguity is a governance risk that shows up during crises.

Siloed ownership

AI spans IT, Legal, HR, and Risk—but they often work in silos. This leads to conflicting rules and poor enforcement. Without shared ownership, policy coordination breaks down. If you're building a cybersecurity strategy that spans the enterprise, this fragmentation will derail it. A policy that isn't jointly owned won't hold. And gaps in coordination are where policy failure begins.

A Practical Framework for Responsible AI Governance

Creating an AI policy is not a legal formality but an operational directive that defines how business users innovate, protect data, and stay within regulatory bounds. We recommend a five-step governance model that balances productivity and control across departments.

Step 1 Define use cases before tools

Don't start with tools. Start with business needs and associated risks. Jumping straight to tool approval can lead to inconsistent policy enforcement. Worse, it blinds you to how AI is being used at the decision layer, where the real risk lives. Instead, start by identifying legitimate, value-generating use cases.

For example:

- Using GenAI to summarise internal policy documents may be low risk.
- Using the same tool to rewrite contract terms or automate investment briefs introduces legal exposure, data integrity concerns, and reputational risk.

By scoping use cases first, you gain two advantages:

1. You can assess risk vs. reward in context before approving technology adoption.
2. You can prioritise governance resources where they're actually needed.

This also gives business leaders ownership in shaping the policy, instead of treating it as an IT gatekeeping function.

Key considerations:

Map high-frequency tasks across departments (e.g., HR, Marketing, Finance, Operations)

Identify where AI could reduce effort without introducing business or compliance risk

Flag high-risk use cases that should be disallowed or deferred

Document these use cases as the foundation for your policy language

Pro tip: This exercise often reveals shadow usage you didn't know existed. Use it to inform your data protection, AI governance, and broader cybersecurity strategy planning.

Step 2 Draft an AI policy that prioritises simplicity and enforceability

Good policies are easy to understand and act on. If the language is too abstract, too legal, or too technical, it will be misunderstood or ignored.

Many policies also fail because they focus on permissions instead of process. For instance, stating that "financial data must not be shared with public AI tools" is useful only if employees know how to classify data, or who to ask for approval.

A weak policy language would say, "Avoid using sensitive data in AI platforms." A stronger, operational version will say, "Financial, personal, and legal data are classified as sensitive. These must not be entered into public GenAI tools like ChatGPT. For business use cases, submit a tool request via IT."

Here are the things your AI policy should clearly state:

- Approved and restricted AI tools
- Data classifications (what's allowed, grey area, and prohibited)
- Output responsibility: who owns review and sign-off
- Escalation path for violations or policy exceptions
- Penalties for misuse, framed in operational terms

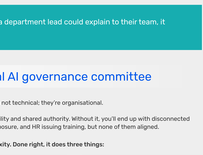

Pro tip: If your policy doesn't read like something a department lead could explain to their team, it needs work.

Step 3 Build a cross-functional AI governance committee

No single team owns AI usage. That's why most AI policy failures are not technical: they're organisational.

A functioning AI governance model needs cross-departmental visibility and shared authority. Without it, you'll end up with disconnected efforts: IT vetting tools, Legal reviewing contracts, Risk mapping exposure, and HR issuing training, but none of them aligned.

A governance committee creates a structure around that complexity. Done right, it does three things:

Tool Governance

Approve or deny new AI tools based on risk, data residency, and model behaviour.

Use Case Review

Evaluate business requests for GenAI in context—e.g., automating candidate screening, customer responses, or report generation.

Breach Oversight & Escalation

When something goes wrong, this is the team that investigates, assesses the impact, and triggers the response.

Suggested roles in the committee:

Domain	Role
IT/Security	Tool vetting, access control, and data handling oversight
Legal/Compliance	Regulatory interpretation, policy structure
HR	Policy communication, training rollout
Risk/Audit	Exposure modelling, reporting integration
Business Sponsor	Ensures the framework supports productivity, not just control

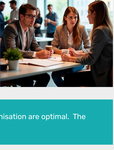

Pro tip: Keep the committee lean. Five to seven decision-makers in an organisation are optimal. The point is alignment, not bureaucracy.

Step 4 Build awareness among business users

Most AI risks do not come from infrastructure. They come from human decisions—made quickly, with incomplete context. Teams need to know what is permitted, what is not, and how their actions tie back to company policy, even if they are using AI to summarise meeting notes or write job descriptions.

To deliver business awareness as an organisation, the best practices include

- Delivering department-specific training (e.g., marketing, HR, finance)
- Using real-world examples to illustrate grey areas and violations
- Providing one central reference page for policy documents and approved tools
- Establishing a reporting channel for AI misuse or ambiguity
- Reinforcing through repeated micro-learnings, not one-time sessions.

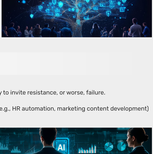

Step 5 Pilot, measure, and mature

Rolling out an AI policy across an enterprise without first testing it is a guaranteed way to invite resistance, or worse, failure.

Instead, start with a limited pilot. Choose 1-2 departments with distinct AI use cases (e.g., HR automation, marketing content development) and validate the entire governance flow: approval, usage, review, and escalation.

Use this pilot to answer the following questions:

- Can users easily follow the policy?
- Are tool requests being processed efficiently?
- Are reviewers confident in classifying acceptable output?
- Do any policy gaps appear in practice?
- Is the governance committee looped in at the right moments?

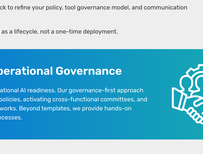

Once the pilot holds its weight under live conditions, use the feedback to refine your policy, tool governance model, and communication cadence.

Just like any cybersecurity strategy, AI governance must be treated as a lifecycle, not a one-time deployment.

How Paramount Helps You Build Operational Governance

Paramount enables enterprises move from reactive controls to operational AI readiness. Our governance-first approach supports digital leaders in defining use cases, creating enforceable policies, activating cross-functional committees, and aligning AI usage with existing cybersecurity and compliance frameworks. Beyond templates, we provide hands-on support to operationalise responsible AI across teams, tools, and processes.

Talk to Paramount to assess where your organisation stands in the responsible AI governance landscape and to turn AI intent into responsible action.

* Data Sources: From Payrolls to Patents: The Spectrum of Data Leaked into GenAI