

Category

Data & AI

# **Responsible AI & Governance Playbook (Risk, Ethics, Compliance)**

A practical, risk-tiered governance model that lets AI initiatives move fast on low-risk use cases while properly controlling high-risk ones.

- Executive
- Intermediate
- Template Included

- Mithun A. Sridharan
- 11 April 2024
- 8 minute(s)

## Overview

Stand up a risk-tiered AI governance model, aligned to frameworks like the NIST AI RMF and EU AI Act, that speeds up low-risk AI work instead of slowing everything down equally.

How is this different from general enterprise risk management?

It's a specialized layer that plugs into existing enterprise risk and data governance functions rather than replacing them, adding AI-specific concerns — model bias, explainability, hallucination risk, autonomous decision-making — that general risk frameworks weren't originally built to assess.

Do we need to fully comply with the EU AI Act if we don't operate in the EU?

If you have no EU operations, customers, or data subjects, direct compliance may not apply, but the Act's risk-tier structure is widely used as a de facto industry standard and is a reasonable, defensible baseline to align to regardless — many U.S. and global regulators are converging toward similar risk-based frameworks.

How do we stop the high-risk review board from becoming a bottleneck?

Keep the four-tier system honest — the review board should only see genuinely high-risk use cases, which in most enterprises are a small fraction of the total inventory once tiering is applied correctly. If the board is reviewing dozens of cases a month, the tiering criteria are likely too conservative.

Who should have final authority to block a high-risk use case from launching?

The governance function needs real authority chartered by an executive sponsor (typically the CRO or General Counsel), not just an advisory voice — a governance board that can be overruled by a business unit VP with no escalation path isn't governance, it's a suggestion box.

What's the very first artifact we should build if we're starting from nothing?

The model and tool inventory. Every other part of this playbook — tiering, review processes, incident response — depends on first knowing what AI is actually in production, and most organizations are surprised by how incomplete their starting picture is.

```
{ "@context": "https://schema.org", "@type": "FAQPage", "mainEntity": [ { "@type": "Question", "name": "How is this different from general enterprise risk management?", "acceptedAnswer": { "@type": "Answer", "text": "It's a specialized layer that plugs into existing enterprise risk and data governance functions rather than replacing them, adding AI-specific concerns — model bias, explainability, hallucination risk, autonomous decision-making — that general risk frameworks weren't originally built to assess." } }, { "@type": "Question", "name": "Do we need to fully comply with the EU AI Act if we don't operate in the EU?", "acceptedAnswer": { "@type": "Answer", "text": "If you have no EU operations, customers, or data subjects, direct compliance may not apply, but the Act's risk-tier structure is widely used as a de facto industry standard and is a reasonable, defensible baseline to align to regardless — many U.S. and global regulators are converging toward similar risk-based frameworks." } }, { "@type": "Question", "name": "How do we stop the high-risk review board from becoming a bottleneck?", "acceptedAnswer": { "@type": "Answer", "text": "Keep the four-tier system honest — the review board should only see genuinely high-risk use cases, which in most enterprises are a small fraction of the total inventory once tiering is applied correctly. If the board is reviewing dozens of cases a month, the tiering criteria are likely too conservative." } }, { "@type": "Question", "name": "Who should have final authority to block a high-risk use case from launching?",
```

```
"acceptedAnswer": { "@type": "Answer", "text": "The governance function needs real authority chartered by an executive sponsor (typically the CRO or General Counsel), not just an advisory voice — a governance board that can be overruled by a business unit VP with no escalation path isn't governance, it's a suggestion box." } }, { "@type": "Question", "name": "What's the very first artifact we should build if we're starting from nothing?", "acceptedAnswer": { "@type": "Answer", "text": "The model and tool inventory. Every other part of this playbook — tiering, review processes, incident response — depends on first knowing what AI is actually in production, and most organizations are surprised by how incomplete their starting picture is." } } ] }
```

Subscriber access

## Unlock this playbook

This playbook — including every framework, template, and step-by-step section — is available free to Think Insights subscribers. Enter your email to unlock it instantly and get our weekly insights newsletter. No account needed, and access is remembered on this device.

First Name

Last Name

Email

Country

I agree to data processing in accordance with GDPR

Leave this field blank

## References

## Related playbooks

## Data & AI

### **Multimodal AI for Document, Image, Audio & Video Workflows Playbook**

A framework for using multimodal AI across document, image, audio, and video workflows — building genuinely integrated processing that handles multiple content

- Practitioner

- Intermediate
- Template Included

[Read playbook](#)

Data & AI

## **AI for Sustainability, Climate Risk & Transition Planning Playbook**

A framework for using AI in sustainability, climate risk, and transition planning — modeling transition pathway scenarios and climate risk exposure with genuine

- Executive
- Advanced
- Template Included

[Read playbook](#)

Data & AI

## **AI Storytelling, Internal Communication & Adoption Campaigns Playbook**

A framework for AI storytelling, internal communication, and adoption campaigns — using specific, concrete narratives of genuine AI-driven outcomes rather than

- Practitioner
- Beginner
- Template Included

[Read playbook](#)

Talk to our AI governance team

Author

Mithun A. Sridharan

I'm Mithun A. Sridharan, Founder of this website - Think Insights - on Strategy, Management Consulting, Leadership, Digital Transformation, and Data Literacy. Follow me on social media or connect with me on LinkedIn for updates.