

Challenge

Bias and data quality at scale: Training data issues multiply across hundreds of models, making manual review infeasible.



Ownership gaps between data science, legal, and IT: No single function owns governance across the full lifecycle, leading to diffused accountability.



Continuous monitoring across the model lifecycle: Models degrade, data drifts, and requirements shift post-deployment.



Managing GenAI risk within existing governance structures: GenAI introduces hallucination, prompt injection, and data leakage risks that traditional ML frameworks miss.



Mitigation

- Automated bias detection integrated into MLOps workflows
- Data quality scoring at ingestion with threshold based rejection

- Governance committees with defined escalation paths
- CDAO or equivalent role as central governance owner

- Automated performance monitoring with alerting
- Scheduled re-evaluation cadences tied to risk classification

- Extended governance protocols for GenAI-specific risks
- Guardrail services covering cost control, safety checks, and output quality