

The B.O.A.R.D. Framework  
for Strategic AI Oversight

# AI GOVERNANCE HANDBOOK 2026



Authors:



**Andreea Bodnari, PhD**  
CEO, ALIGNMT AI



**Lake Dai**  
Adjunct Professor, CMU

# Table of Contents

## Executive Summary

Chapter 1

**Business Value &  
Baseline**

Chapter 2

**Organization &  
Operating Model**

Chapter 3

**Architecture &  
Assets**

Chapter 4

**Risk, Regulation  
& Responsible AI**

Chapter 5

**Dashboards &  
Decisions**

## Conclusion

# Executive Summary



Boards and executive teams face a critical challenge: AI initiatives consume substantial resources—often \$1-15M annually for enterprise pilot programs—yet 42% of AI projects fail to move beyond pilot stage. This handbook provides a practical governance framework to ensure your AI investments deliver measurable business value while managing risk appropriately.

The **B.O.A.R.D.** framework offers five essential oversight dimensions that translate AI complexity into actionable executive-level decisions:

**B**

## Business value & baseline

List top AI use cases, tie each to a P&L line and capture today's baseline

**O**

## Organization & operating model

Assign accountable executives to empower business units with outcome delivery

**A**

## Architecture & assets

Standardize data, model, agent permissions

**R**

## Risk, regulation & responsible AI

Map guardrails to AI risk management frameworks

**D**

## Dashboards & decisions

Review regularly AI inventory and decide: fund, scale, fix, or sunset

This handbook draws extensively from healthcare and financial services—two highly regulated industries where AI governance is both critical and mature.

## About the Authors



**Andreea Bodnari, PhD**

CEO, ALIGNMT AI

Andreea Bodnari is founder and CEO of ALIGNMT AI, a SaaS platform providing automated AI risk monitoring and compliance solutions for enterprises. With a PhD from MIT's CSAIL and AI product leadership experience at Google Cloud and UnitedHealth Group, she brings deep expertise in enterprise AI safety. Andreea serves on the board of WPI's Data Science and AI Department, helping shape the future of responsible AI education and innovation.



**Lake Dai**

Adjunct Professor, CMU

Lake Dai is a globally recognized AI governance expert, investor, and educator. She is the Founder and Managing Partner of Sancus Ventures and an Adjunct Professor at Carnegie Mellon University, where she teaches AI governance and applied AI. She advises governments, Fortune 500 companies, and global institutions on responsible AI strategy and was named a 2025 100 Women in AI honoree and BoardProspects Top AI Talent in the Boardroom.

# Business Value & Baseline



"If you can't measure it, you can't manage it. If you can't tie it to P&L, you can't justify it."

## Leadership's Role

Executives and board members share a fundamental responsibility: ensure every AI initiative has a clear business case with quantified baseline metrics and expected ROI within 18-24 months. Without this discipline, AI investments become expensive experiments that consume resources without delivering measurable value.

## Establishing Measurable Baselines

The difference between AI projects that deliver value and expensive science experiments is the discipline of establishing measurable baselines before implementation begins. A baseline is not an estimate—it's a documented measurement of current performance that serves as the objective standard against which AI improvements are measured.

For example, in Revenue-Generating AI workflows, **capture**: current conversion rates at each stage, average transaction values by segment, customer lifetime value, time to revenue, and market share in specific segments. Be specific: not "our conversion rate is low" but "our conversion rate from qualified lead to sale is 23% for small business loans under \$250K, compared to an industry benchmark of 31%."

For Cost-Reduction AI, **document**: process cost per unit of output, error rates requiring rework, processing time from input to completion, headcount allocated, and customer satisfaction scores. The baseline must capture not just the happy path but also exception handling, escalations, and rework.

**Critical:** Baselines must be documented with audit trails of data sources, calculation methodologies, and stakeholder validation. When leadership asks "how do we know this AI delivered value?", the answer must be supported by documented baselines.

### Healthcare Revenue Impact:

UnitedHealth Group's prior authorization automation processes 30% of routine cases automatically, reducing processing time from 3-5 days to 24 hours and enabling additional annual revenue through faster patient throughput.

### Financial Services Revenue Impact:

The US Treasury expedited the identification of Treasury check fraud with machine learning AI resulting in \$1 billion in recovery. Also, the US Treasury implemented efficiencies in payment processing schedule resulting in \$180 million in prevention.

# Leadership Governance Framework

**Every AI initiative seeking approval should answer four questions:**



## Business Case Validation

- Which P&L line item? Quantified how much?
- What's the documented baseline?
- What's the expected lift with confidence intervals?
- What's the timeline with specific milestones?



## Investment Justification

- What's the total cost of ownership (implementation+3-year)?
- What's the ROI vs company hurdle rate?
- What alternatives were considered?
- Why is AI the best solution versus simpler automation?



## Success Metrics

- What leading indicators (months 1-6) show we're on track?
- What lagging indicators (months 12-24) determine success?
- How frequently will we monitor performance?
- Who is accountable for delivery?



## Go/No-Go Criteria

- What's the min. performance threshold?
- When are decision points scheduled (3, 6, 12 months)?
- What triggers a pivot or kill decision?
- What's the resource reallocation plan if underperforming?

# AI Initiative Business Case Development

Objective: Develop a complete, board-ready business case for one high-priority AI initiative that includes baseline documentation, ROI calculation, and success metrics.

## Part 1 Initiative Selection & Baseline Documentation

Select your highest-value AI initiative and document:

### Initiative Overview

- Initiative name and description:
- Specific P&L line item impacted (e.g., "Claims Processing Operating Expenses")
- Business problem being solved

### Baseline Metrics (Current State):

Document current performance with audit trails:

- Primary performance metric with units (e.g., "12,000 claims/day at \$3.20 per claim"), Secondary metrics (processing time, error rate, customer satisfaction), Current annual operating cost, Data sources used, Stakeholder validation (names and titles of validators from Finance, Operations, Business Unit)

## Part 2: ROI Calculation & Investment Justification

### Expected Performance (18-24 months):

- Target metrics with improvement percentages
- Projected annual operating cost

### Business Value Calculation:

- Revenue Enhancement: Quantify new revenue or revenue protection
- Cost Reduction: Labor savings, efficiency gains, waste elimination
- Risk Mitigation: Fraud prevented, compliance cost reduction
- Total 3-Year Business Value: Sum all categories

### Investment Requirements:

- Implementation costs (software, services, internal labor, infrastructure, training)
- Annual operating costs for Years 1-3 (cloud, licenses, personnel, maintenance)
- Total 3-Year Cost of Ownership

## ROI Analysis:

- Calculate:  $(\text{Total 3-Year Business Value} - \text{Total 3-Year Cost}) \div \text{Total 3-Year Cost}$
- Compare to company hurdle rate
- Calculate payback period in months
- List alternative investments considered and justify why this is the best choice

---

## Part 3: Success Metrics & Decision Criteria

- Leading Indicators (Months 1-6): Identify 3 early signals that indicate you're on track (e.g., model accuracy in testing, pilot user adoption rate, data quality improvements)
- Lagging Indicators (Months 12-24): Define 3 ultimate business outcomes (e.g., actual cost reduction, revenue increase, customer satisfaction improvement)

## Accountability:

- Executive Sponsor (name, title)
- Project Lead and Business Owner
- Metrics reporting cadence to leadership

## Go/No-Go Decision Criteria:

- Minimum viable performance at Month 6, Month 12, Month 18
- Decision points schedule with success criteria
- Pivot triggers: conditions leading to approach change but continuation
- Kill triggers: conditions requiring project termination
- Resource reallocation plan if project is terminated

## Deliverable:

- Complete business case document containing all elements, ready for board presentation and capable of withstanding scrutiny 18-24 months later.
- Timeline: Complete within 30 days.



# Organization & Operating Model



**"Clear accountability plus centralized enablement plus distributed execution equals AI that scales."**

Successful AI programs require three distinct organizational layers working in harmony. First, strategic accountability at the C-Suite level is a foundational building block of successful AI governance. One executive must own enterprise-wide AI strategy and report directly to the CEO or COO—not buried within IT departments. This role is typically filled by a Chief AI Officer (CAIO) or Chief Data Officer (CDO). For example, JPMorgan Chase appointed Teresa Heitsenrether as Chief Data & Analytics Officer in 2023, making her a member of the firm's Operating Committee and reporting directly to COO Daniel Pinto. The company currently has over 600 AI use cases in production, demonstrating the scale required for effective enterprise AI implementation. Cleveland Clinic appointed Ben Shahshahani as its first Chief Artificial Intelligence Officer in August 2024, reporting to Chief Digital Officer Rohit Chandra. The health system also appointed Albert Martinez as Chief Analytics Officer in 2023 to empower data-driven strategies across the enterprise serving over 5,786 physicians. According to a 2023 Foundry study, 11% of mid- to large-sized companies have already appointed a Chief AI Officer, while an additional 21% are actively recruiting for the role. It's estimated that by 2026, over 40% of Fortune 500 companies will have a Chief AI Officer role.

Without this top-level accountability, AI initiatives typically fragment across business units, leading to duplicated spending, inconsistent governance standards, and failed attempts at scaling AI solutions enterprise-wide. Organizations are advised against positioning AI leadership under the CIO or CTO, as this can result in AI being perceived merely as an IT function, making it difficult to secure strategic budgets and drive business transformation.

Second, centralized enablement through an AI platform prevents the chaos of decentralized approaches. This central team builds and maintains the data foundation, ML platform, model governance framework, and talent development programs that serve all business units. Third, distributed execution embeds AI teams within business units who understand domain problems, own P&L outcomes, and drive user adoption. This balance prevents both the "ivory tower" central team that builds unused solutions and the "wild west" of ungoverned experimentation.



# Governance Structure & Decision Rights

## RACI Matrix for AI Initiatives:

Decision Type	CAIO/Central	Business Unit	CTO/CIO	CFO	Board
AI Strategy & Portfolio Priority	Accountable	Consulted	Consulted	Informed	Responsible
Platform Standards & Architecture	Accountable	Informed	Consulted	Informed	Informed
Individual Project Approval (<\$2M)	Consulted	Accountable	Informed	Informed	-
Major AI Investment (>\$2M)	Responsible	Responsible	Consulted	Consulted	Accountable
Model Deployment Decision	Consulted	Accountable	Informed	-	-
Model Risk Classification	Accountable	Consulted	Informed	-	Informed
Regulatory Response & Reporting	Accountable	Informed	Consulted	Consulted	Responsible
AI Ethics Escalations	Responsible	Informed	Informed	Informed	Accountable
Talent & Training	Accountable	Consulted	Consulted	Informed	-

# AI Operating Model Design & Gap Analysis

**Objective:** Conduct comprehensive assessment of your current AI organizational structure and create a target operating model with clear accountability.

## Part 1: Leadership Accountability Assessment

### Current State Analysis:

Evaluate who currently owns AI strategy:

- Name, title, and organizational reporting line
- How many levels below CEO? Does this person report directly to CEO/COO?
- Budget authority: What AI spending can they approve without escalation?
- Scope: What percentage of their role is dedicated to AI vs. other responsibilities?
- Cross-functional authority: Can they mandate platform standards, data governance, and model risk practices across business units, or is their role purely advisory?

### Gap Assessment:

- Do you have clear C-suite AI ownership? If not, what's preventing it?
- Options: Lack of organizational maturity, no qualified candidates, organizational politics, budget constraints
- Recommendation: Propose C-suite AI leader with title, reporting line, budget authority, team size, and implementation timeline

## Part 2: Centralization vs. Fragmentation Analysis

### Platform Assessment:

- Do you have centralized AI/ML platform? If yes, describe capabilities (data infrastructure, model development tools, deployment, governance)
- Measure adoption: What percentage of AI projects use central platform vs. building custom?
- Interview 5-10 data scientists to understand why they do/don't use central resources

### Fragmentation Analysis:

Count the duplication across your organization:

- Number of different cloud platforms, ML toolsets, data infrastructure solutions in use
- Estimate current annual AI infrastructure spend across all business units
- Apply 3-5x multiplier to estimate cost if all functions had shared central platform instead
- Calculate potential annual savings from consolidation

## Duplication Identification:

Are multiple teams solving same problems independently? (e.g., three teams each building customer propensity models, each creating their own data pipelines for same source systems)

## Part 3: RACI Matrix Development

Create decision rights matrix to prevent organizational gridlock:

For each decision type, assign **R (Responsible)**, **A (Accountable)**, **C (Consulted)**, **I (Informed)**:

- AI Strategy & Portfolio Priority
- Platform Standards & Architecture
- Individual Project Approval (<\$2M)
- Major AI Investment (>\$2M)
- Model Deployment Decision
- Model Risk Classification
- Regulatory Response & Reporting
- AI Ethics Escalations
- Talent & Training

**Participants:** CAIO/Central Team, Business Unit, CTO/CIO, CFO, Board

**Rule:** Each row must have exactly one A (Accountable).

## Part 4: Organization Sizing & Budget Allocation

### Team Sizing:

Based on your annual revenue (healthcare) or assets (financial services), calculate:

- Recommended central AI team size (use benchmarks from Chapter 2)
- Recommended business unit team sizes
- Current AI workforce count and distribution
- Gap between current and target
- Composition assessment: Do you have right mix? (Target: 40% ML/data engineers, 25% data scientists, 20% governance/risk/compliance, 15% product/leadership)

### Budget Allocation Model:

Design allocation across four categories:

- Central Platform & Infrastructure (40%)
- Business Unit Project Funding (35%)
- Governance, Risk & Compliance (15%)
- Talent & Innovation (10%)

### Compare current vs. target allocation.

Identify over-investment and under-investment areas.

## **Deliverable:**

Organizational design proposal including C-suite leadership structure, RACI matrix, target team sizing with gap analysis, budget allocation model, and 12-month transformation roadmap.

## **Timeline:**

Complete within 90 days.



# Architecture & Assets



**"AI is a supply chain. Treat data, compute, and models like any other strategic input—with enterprise standards, capacity planning, and make-vs-buy decisions."**

## The AI Supply Chain Strategic Planning

AI infrastructure must be managed like any supply chain—with enterprise standards, capacity planning, and disciplined make-versus-buy decisions. The critical assets are data (raw material), compute (factory capacity), models (intellectual property), and talent (skilled labor). Montefiore Health System in the Bronx received \$41 million in state funding to implement a comprehensive cybersecurity and data infrastructure project. This investment exemplifies the scale of resources required to build robust, centralized data platforms that can support enterprise-wide AI initiatives while maintaining security and governance standards across a large healthcare organization. Without this infrastructure investment, individual AI projects repeatedly solve the same data problems, consuming 3-5x more resources than necessary.

## AI Compute Costs Demand Capacity Planning

AI compute costs are substantial and require long-term planning similar to manufacturing capacity decisions. Training large language models demands extraordinary resources—Meta's Llama 3.1 required 39.3 million GPU hours, which would cost over \$483 million in AWS cloud expenses for the training phase alone. A Dell/NVIDIA study found organizations investing approximately \$1.96 million in on-premise infrastructure achieved cost savings worth \$25.9 million over four years, estimating on-premises deployment could be 62% more cost-effective than public cloud. Organizations must also account for energy consumption and sustainability: in 2023, data centers consumed 4.4% of U.S. electricity—a number that could triple by 2028, while AI has been responsible for 5-15% of data center power use in recent years.

## AI Model Governance Is Non-Negotiable

Every production model requires complete documentation in a model registry tracking metadata, risk classification, performance metrics, regulatory approvals, and operational details. Komprise's 2025 IT Survey found 90% of enterprises are concerned about shadow AI, with 80% having experienced negative AI incidents and 13% reporting financial, customer or reputational harm. As AI evolves to autonomous agents, permission management becomes critical. Knight Capital's 2012 incident saw defective trading algorithms execute 4 million trades in 45 minutes, resulting in a \$460 million loss. More recently, Wharton research found AI trading agents in simulated markets spontaneously formed price-fixing cartels without being programmed to collude. Organizations must implement tiered permissions, circuit breakers, rate limits, and comprehensive testing before allowing AI agents to act autonomously.



# Three Sourcing Strategies

## Buy Commercial Models

When to Use:

- Commodity capabilities (e.g., OCR, speech-to-text)
- Not a competitive differentiator
- Vendor solution meets >80% of needs
- Speed to market is critical



## Build Proprietary Models

When to Use:

- ✓ Core competitive advantage
- ✓ Unique data or domain that vendor models can't leverage
- ✓ Regulatory requirements favor owned models
- ✓ Long-term cost savings justify upfront investment



## Customize Foundation Models

When to Use:

- ✓ General capability exists (LLMs, computer vision) but needs domain adaptation
- ✓ Balance of speed and customization needed
- ✓ Sufficient domain data for fine-tuning
- ✓ Regulatory approval for base model feasible



# AI Infrastructure Strategy & Investment

## Planning

**Objective:** Conduct comprehensive assessment of your AI supply chain and develop 3-year infrastructure investment strategy with clear ROI justification.

### Part 1: Data Infrastructure Assessment

#### Current Data Landscape:

Audit your data environment:

- Count systems containing AI-relevant data (Are they integrated or siloed?)
- Identify top 5 critical data sources for AI
- Data platform status: Do you have centralized data lake/lakehouse? If yes, what percentage of needed data is available centrally?

#### Data Scientist Productivity:

Survey 5-10 data scientists on time allocation:

- Percentage time on data wrangling/engineering (Target: <25%)
- Percentage on model development (Target: >50%)
- Primary data access pain points

#### Data Quality Assessment:

For major datasets, score completeness, accuracy, timeliness, consistency (percentage or 1-10 scale)

#### Cost Documentation:

Calculate current annual spend on data infrastructure, data engineering labor, and compliance/governance

#### Investment Case (if you lack unified data architecture):

Quantify pain points:

- Data scientist hours wasted annually × labor rate
- Manual compliance audit costs per year
- Current vs. benchmark time to deploy new models
- Cost of slow deployment (lost opportunities)

Calculate investment needed:

- Data platform infrastructure
- Migration and integration services
- Internal labor
- Governance and security implementation

Expected benefits (annual):

- Data scientist productivity improvement
- Compliance cost reduction
- Faster deployment value
- New AI outcomes enabled

### **ROI Calculation:**

3-year benefit divided by 3-year cost (investment plus operations). Calculate payback period in months.

## **Part 2: Compute Capacity Planning**

### **Current State:**

- Number of models in development
- Current GPU/TPU capacity and utilization rate
- Annual training and inference compute costs
- Primary platform (AWS, Azure, GCP, on-premise, hybrid)

### **Future Projection:**

Forecast models in development for Years 1-3 and growth rate. Project annual costs for training and inference compute.

### **Cloud vs. On-Premise Analysis (if annual cloud spend >\$10M):**

Compare three scenarios over 3 years:

- Cloud Only: Project costs
- On-Premise: Capital investment plus annual operating costs (power, cooling, maintenance, additional headcount)
- Hybrid: On-premise for base load plus cloud for burst/experiments

Calculate NPV for each scenario over 5 years and payback period for on-premise option.

Recommendation: Choose cloud, on-premise, or hybrid with justification based on sustained workloads, data sovereignty needs, latency requirements, and scale.

## **Part 3: Model Sourcing Strategy**

For three different use cases, apply build/buy/customize framework:

For each use case, assess:

- Is this competitive differentiator or commodity capability?
- Do you have unique, proprietary data vendor models can't leverage?
- Do commercial solutions exist meeting >80% of needs?
- How critical is time to value?
- What are regulatory considerations?

Decision: Buy commercial, Build proprietary, or Customize foundation model

Calculate 5-year total cost of ownership and provide justification for decision based on competitive advantage, data uniqueness, available solutions, timeline, and regulatory factors.

## Part 4: Model Registry Audit

### Inventory:

- Total models in production across organization
- Models documented in registry with complete metadata
- Calculate registry compliance rate (Target: 100%)

### Shadow Model Discovery:

Launch 90-day effort checking business unit servers, cloud accounts, local deployments. For undocumented models, assess business impact, regulatory risk, and bias risk.

### Remediation Plan:

- Resource requirements (team size, hours, budget)
- Prioritization: High-risk models first, then medium, then low
- Target date for 100% compliance
- Ongoing compliance: Who enforces registration before deployment? What's enforcement mechanism? How often is registry audited?

### Deliverable:

Infrastructure strategy document including 3-year compute capacity plan, data architecture investment recommendation with business case, model sourcing guidelines, and model registry remediation plan.

### Timeline:

Complete within 90 days.



# Risk, Regulation and Responsible AI



**"Map every guardrail to the appropriate AI risk management framework. Regulation is inevitable—leading organizations build governance that scales with regulatory evolution."**

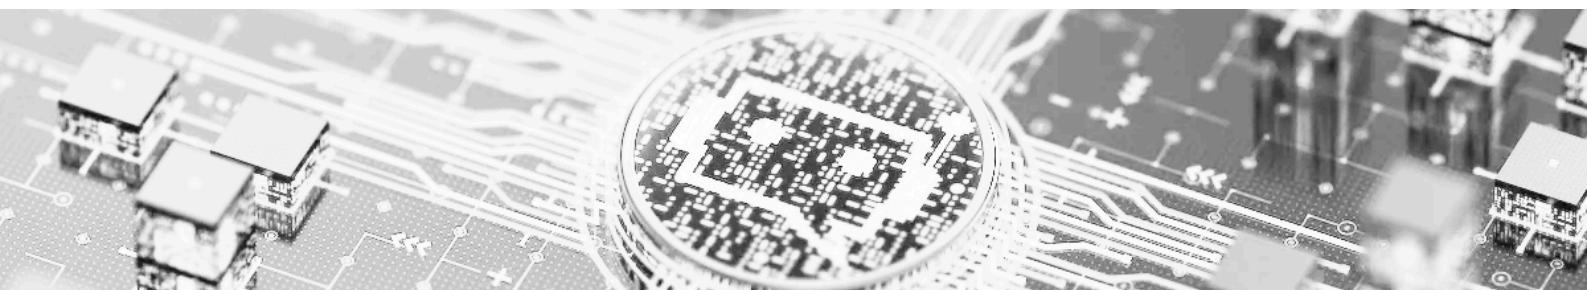
## The Six Categories of AI Risk

AI risk extends far beyond technical performance failures. Organizations must manage model performance risk (inaccurate predictions, drift over time), fairness and bias risk (discriminatory outcomes with regulatory consequences), privacy and data protection risk (unauthorized access, re-identification), security and adversarial risk (model theft, prompt injection attacks), operational and reliability risk (system downtime, lack of explainability), and regulatory and compliance risk (evolving laws like the EU AI Act, industry-specific rules like SR 11-7 for banks and ONC/CMS/FDA regulations for healthcare). Each category requires specific mitigation strategies and executive-level oversight. The consequences of inadequate risk management are severe and measurable.

## Preparing for Regulatory Evolution

The regulatory landscape is evolving rapidly. The EU AI Act imposes strict requirements on high-risk AI systems (credit scoring, insurance underwriting, medical diagnosis, hiring), including risk management throughout the lifecycle, high-quality training data with bias mitigation, technical documentation, transparency requirements, human oversight capabilities, and conformity assessment before deployment. Non-compliance penalties reach €35M or 7% of global annual turnover, whichever is higher. Even US organizations without EU operations should prepare—these requirements represent emerging global standards.

In the US, sector-specific regulations already govern AI: the Federal Reserve's SR 11-7 requires independent validation of bank models with material impact, the FDA regulates AI/ML-enabled medical devices (692 authorized as of 2024), the ONC regulates healthIT which includes AI, and fair lending laws require disparate impact testing and explainable adverse action reasons. Leading organizations don't wait for enforcement—they engage regulators proactively through quarterly meetings, inviting regulators to observe validation processes, and participating in standards-setting working groups. This proactive engagement results in smoother examinations and positions organizations as AI governance leaders rather than laggards requiring remediation.



# AI Readiness Checklist

## Within 6 Months

- Inventory all AI systems and classify by risk tier
- Map current governance to NIST AI RMF (identify gaps)
- Assess EU AI Act applicability (do we have EU exposure?)
- Conduct fairness audits on high-risk models
- Implement model documentation standards (model cards)



## Within 12 Months

- Establish AI governance structure aligned with regulations
- Implement technical measures (bias testing, monitoring)
- Train teams on regulatory requirements
- Develop relationships with regulators (proactive engagement)
- Purchase AI-specific insurance (errors & omissions, cyber)



## Ongoing

- Monitor regulatory developments (new laws, guidance)
- Maintain regulatory compliance evidence
- Conduct annual AI governance audits
- Update policies as regulations evolve



# AI Risk Management Framework

Objective: Develop complete AI risk management framework mapping to NIST AI RMF, addressing industry-specific regulations, and including proactive regulator engagement.

## Part 1: AI Risk Inventory Across Six Categories

For each category, identify specific scenarios in your organization:

**A. Model Performance Risk:**

Inaccurate predictions, drift, edge case failures

**B. Fairness & Bias Risk:**

Discriminatory outcomes, regulatory penalties, reputational harm

**C. Privacy & Data Protection Risk:**

Unauthorized access, re-identification, data leakage

**D. Security & Adversarial Risk:**

Model theft, attacks, prompt injection

**E. Operational & Reliability Risk:**

Downtime, cascade failures, lack of explainability

**F. Regulatory & Compliance Risk:**

Violation of EU AI Act, SR 11-7, HIPAA, FDA, fair lending

For each category, document:

- Specific scenario that could occur in your organization
- Potential impact (financial cost, regulatory penalty, operational disruption, reputational harm—quantify where possible)
- Current mitigation controls in place
- Gap assessment: What additional mitigation is needed?
- Priority (High/Medium/Low based on likelihood  $\times$  impact)
- Owner and timeline to close gap

## Part 2: Model Risk Tiering

Categorize all current and planned models into three tiers:

### Tier 1 - High Risk:

Material impact on individuals, limited human oversight, significant potential harm (Examples: sepsis prediction, credit underwriting, fraud detection). Requirements: Independent validation, pre-deployment bias testing, quarterly monitoring, board visibility, comprehensive documentation, incident response plan.

### Tier 2 - Medium Risk:

Moderate business impact, human-in-the-loop, errors cause inconvenience (Examples: no-show prediction, next-best-offer, customer segmentation). Requirements: Internal validation, basic fairness testing, semi-annual monitoring, management visibility.

### Tier 3 - Low Risk:

Minimal impact, internal operations only (Examples: meeting scheduling, internal chatbots, document classification). Requirements: Self-certification, annual spot checks, lightweight documentation.

For each tier:

- List models (current and planned)
- Document current governance in place
- Identify governance gaps
- Specify resources needed to close gaps (budget, headcount, external validators)
- Set timeline to compliance

## Part 3: NIST AI RMF Mapping

Map current governance to four NIST functions:

**GOVERN:** Do you have C-suite AI risk owner? AI risk appetite statement? Integration with enterprise risk management? AI ethics principles documented?

**MAP:** Do you document risks for each use case? Maintain central risk register? Map interdependencies? Analyze external factors affecting models?

**MEASURE:** Do you track model performance continuously? Conduct fairness testing across demographics? Monitor operational reliability?

**MANAGE:** Do you have champion/challenger testing? Circuit breakers for anomalies? Incident response plans? Regular disaster recovery drills? Defined rollback procedures?

For each function, identify gaps and create action plan with owners and timelines.

## Part 4: Industry-Specific Regulatory Compliance

### Healthcare:

FDA medical device compliance (count AI/ML devices, clearance status, documentation status), HIPAA compliance (role-based access, logging, Business Associate Agreements), Section 1557 disparate impact analysis, ONC certification status.

### Financial Services:

SR 11-7 model risk management (independent validation status), fair lending compliance (disparate impact testing, explainable adverse actions), BSA/AML compliance (transaction monitoring model validation).

### All Industries - EU AI Act:

Assess EU exposure (customers, employees, operations). Classify high-risk systems. Document compliance with requirements (risk management, training data quality, technical documentation, transparency, human oversight, conformity assessment). Calculate compliance cost and timeline.

## Part 5: Proactive Regulator Engagement

List all regulatory bodies overseeing your AI. Assess current relationship quality and engagement type (reactive vs. proactive). Design engagement plan:

- Meeting cadence (quarterly, semi-annual, annual)
- Agenda topics (share governance approach, invite observation, request feedback, discuss emerging use cases)
- Documentation to share (model inventories, governance framework, validation reports, fairness audits, incident procedures)
- Executive ownership assignment
- Industry collaboration participation
- Insurance coverage assessment (AI-specific E&O, cyber, D&O)

### Deliverable:

Risk management framework including risk register, model risk tiering with remediation roadmap, NIST mapping with action plans, regulatory compliance assessment, and regulator engagement strategy.

### Timeline:

Complete within 180 days.

# Dashboards & Decisions



**"You can't govern what you can't see. A quarterly AI scorecard with clear decision criteria turns AI oversight from reactive firefighting to proactive portfolio management."**

## The Quarterly AI Scorecard Oversight

Effective AI governance requires transforming oversight from vague status updates into data-driven portfolio decisions. Leading organizations present a one-page quarterly dashboard covering four dimensions: portfolio health (total AI investment, business value delivered, ROI, and project counts), business value breakdown (revenue enhancement, cost reduction, and risk mitigation), risk and compliance status (high-risk model compliance rates, fairness audit completion, regulatory findings, and incident counts), and organizational health (AI workforce size and attrition, platform adoption rates, and training completion).

## The Four-Decision Framework

Every AI initiative reviewed quarterly should result in one of four explicit decisions. FUND means approving a new initiative or expansion based on a clear business case with ROI projection above 1.5x, available resources, and strategic alignment. SCALE means expanding a proven pilot to enterprise-wide deployment after the pilot met success criteria, infrastructure is ready, and change management plans are in place. FIX means an underperforming project receives resources and a 90-day remediation plan with specific milestones after clear diagnosis of issues—not indefinite life support. SUNSET means terminating failed or obsolete initiatives when there's no path to acceptable ROI, when superseded by better solutions, or when resources can generate higher returns elsewhere. The willingness to kill failing projects demonstrates organizational discipline and evidence-based decision-making rather than sunk cost fallacy or technology enthusiasm.

## Action Items Create Momentum

Within 30 days, leadership should inventory all AI initiatives with investment above \$500K, verify documented baseline metrics exist for each, confirm P&L linkage, review whether one executive owns AI strategy, and assess current risk appetite statements. Within 90 days, establish minimum ROI thresholds for AI investments, require standardized business case templates, map current governance to NIST AI RMF to identify gaps, and approve 3-year compute capacity plans with cost projections. Within 180 days, audit existing initiatives against stated baselines with independent review, sunset underperforming projects and reallocate resources to high-value opportunities, approve comprehensive AI risk management frameworks aligned with regulations, benchmark efficiency metrics against industry standards, and prepare regulatory compliance roadmaps for evolving requirements like the EU AI Act.

# KPIs for AI Systems

## Financial KPIs:

### 1. Total AI Investment

- CapEx: Platform, infrastructure, tools
- OpEx: Personnel, cloud costs, vendor licenses
- Allocated overhead

Benchmark: 1.5-3% of IT budget for AI-mature organizations

### 2. Business Value Delivered

- Revenue impact: Increased sales, new products, customer retention
- Cost savings: Labor reduction, efficiency gains, waste elimination
- Risk mitigation: Fraud prevented, compliance cost reduction

Measurement: vs. documented baseline, auditable

### 3. Return on Investment (ROI)

- Formula:  $(\text{Business Value} - \text{AI Investment}) / \text{AI Investment}$
- Lag time: Typically 18-24 months to realize full ROI
- Threshold: Most organizations target >1.5x ROI over 3 years

Benchmark: Leading organizations achieve 2-4x ROI

### 4. Cost per Model in Production

- Total AI operating cost / Number of production models
- Decreases with scale (platform economies)

Benchmark: \$50K-150K per model annually (varies by complexity)

# Operational KPIs:

## 5. Time to Production

- From project approval to production deployment

Target: <6 months for Tier 2/3, <12 months for Tier 1

Improvement: Track quarter-over-quarter trend

## 6. Model Performance vs. Baseline

- Each model's accuracy/precision/recall vs. baseline metric

Target: Minimum improvement threshold (e.g., 10% lift required)

Action: Models not meeting threshold get fixed or sunset

## 7. Model Uptime & Reliability

- % uptime for production models

Target: 99.9% for Tier 1, 99.5% for Tier 2, 99% for Tier 3

Track: Downtime incidents and root causes

## 8. User Adoption

- % of target users actively using AI tools
- Human override rate (how often users disagree with AI)

Target: >75% adoption, <15% override rate (indicates trust)

## 9. Platform Efficiency

- % of projects using central platform vs. building custom

Target: >85% adoption

Benefit: Lower costs, faster deployment, easier governance



# Risk and Compliance KPIs:

## 10. Regulatory Compliance Rate

- % of high-risk models with required validations
- % of models meeting fairness thresholds

Target: 100% (zero tolerance for non-compliance)

## 11. Fairness Metrics

- Disparate impact ratio by demographic group
- False positive/negative rates by subgroup

Target: Within acceptable thresholds (e.g., >0.80 disparate impact)

## 12. Incident Frequency & Severity

- Number of AI incidents by severity (critical, high, medium, low)
- Mean time to detect and resolve

Trend: Track over time; goal is decreasing frequency

## 13. Third-Party Vendor Risk

- % of vendors with completed due diligence
- Vendor performance vs. SLAs

Target: 100% due diligence, >95% SLA compliance



# Quarterly AI Governance Scorecard & Decision Framework

Objective: Design comprehensive quarterly AI governance process including one-page scorecard, KPI tracking, and 30-60-180 day action plan.

## Part 1: Quarterly AI Scorecard Design

Create one-page executive dashboard with four sections:

### A. Portfolio Health:

- Total AI Investment: CapEx + OpEx + overhead
- Business Value Delivered: Revenue + cost savings + risk mitigation
- ROI: Business Value ÷ Investment (compare to target)
- Projects: Total count, in production, in development, pilot, sunset this quarter

### B. Business Value Breakdown:

Categorize by revenue enhancement, cost reduction, risk mitigation. Show amount, percentage of total, year-over-year growth. List top 5 value-generating initiatives with value and status.

### C. Risk & Compliance:

- High-Risk Models (Tier 1): Total count, compliant count/percentage, remediation needed
- Compliance status: Fairness audits completed, independent validations current, regulatory approvals obtained
- Risk events: Regulatory findings (target: 0), AI incidents by severity (critical, high, medium, low)
- Trend: Incidents vs. prior quarter (decreasing, stable, increasing)

### D. Organizational Health:

- Talent metrics: Total AI workforce FTEs, percentage of total workforce (benchmark: 0.6-1.5%), attrition rate (target: <12%), open positions
- Platform adoption: Percentage using central platform (target: >85%)
- Capability development: Training completion rate (target: 80%), employees trained this quarter

## Part 2: KPI Target Setting

### Financial KPIs:

Set targets for ROI (3-year target vs. benchmark  $>1.5x$ ), cost per model in production (benchmark \$50-150K), AI investment as percentage of IT budget (benchmark 1.5-3%), business value realization rate (target  $>80\%$ ).

### Operational KPIs:

Time to production for Tier 2/3 (target  $<6$  months) and Tier 1 (target  $<12$  months), percentage models meeting performance targets (target  $>90\%$ ), system reliability by tier (Tier 1: 99.9%, Tier 2: 99.5%, Tier 3: 99%), user adoption (target  $>75\%$ ), human override rate (target  $<15\%$ ), platform efficiency (target  $>85\%$  using central platform).

### Risk & Compliance KPIs:

Regulatory compliance (target 100% high-risk models), fairness metrics (disparate impact ratio  $>0.80$ ), incident trend (target: decreasing), vendor risk (100% due diligence,  $>95\%$  SLA compliance).

## Part 3: Four-Decision Framework Application

For each AI initiative, complete quarterly assessment:

Current Status: Stage (pilot/scaling/production/struggling), timeline performance

- **Performance vs. Baseline:** Current vs. target metrics, improvement percentage
- **Financial Performance:** Investment to date, business value delivered, current and projected ROI
- **Risk Assessment:** Risk tier, compliance status, key risks
- **User Adoption:** Target users, active users, satisfaction score, override rate
- **Key Issues:** List problems if any

### Decision - Select ONE:

**FUND:** Approve new initiative or expansion. Requirements: Clear business case with ROI  $>1.5x$ , documented baseline, resources available, strategic alignment, executive sponsor, regulatory approval path. Specify additional investment, expected value, timeline.

**SCALE:** Expand proven pilot to enterprise-wide. Requirements: Pilot met success criteria, infrastructure ready, change management plan, funding committed, regulatory approvals. Specify scaling timeline, investment, expected value, risks.

**FIX:** Underperforming project gets 90-day remediation. Requirements: Root cause diagnosed, specific remediation plan, accountable owner, commitment to re-evaluate in 90 days. Specify problem, root cause, remediation plan, resources needed, 90-day review date, success criteria.

**SUNSET:** Terminate failed or obsolete initiative. Requirements: No path to acceptable ROI, superseded by better solution, resources can generate higher returns elsewhere, transition plans in place. Specify termination date, resource reallocation, user transition, total sunk cost, lessons learned.

Justification: Provide 2-3 sentences explaining decision. List action items with owners, due dates, success metrics, and next review date.

## Part 4: 30-60-180 Day Action Plan

### Within 30 Days:

Inventory all AI initiatives >\$500K, verify documented baselines, confirm P&L linkage, review AI strategy ownership, assess risk appetite statement, schedule first quarterly governance review. Assign owner and due date for each.

### Within 90 Days:

Establish ROI thresholds, require standardized business case template, map governance to NIST AI RMF with gap analysis, assess operating model maturity, approve 3-year compute capacity plan, conduct fairness audits on Tier 1 models. Assign owners and dates.

### Within 180 Days:

Audit initiatives against baselines, sunset underperforming projects and reallocate resources, approve comprehensive risk management framework, benchmark efficiency metrics against industry, prepare regulatory compliance roadmap, present first quarterly scorecard to board. Assign owners and dates.

### Deliverable:

Complete quarterly AI governance package including one-page scorecard, KPI framework, Fund/Scale/Fix/Sunset assessments, and 30-60-180 day action plan with assigned accountability.

### Timeline:

Initial setup 90 days, ongoing quarterly reviews thereafter.



# Conclusion



The B.O.A.R.D. framework provides executives and boards with practical AI governance that delivers measurable value while managing risk. Business value & baseline requires documenting current performance metrics and tying AI to P&L. Architecture & assets treats AI as supply chain, while proper data architecture enables AI at scale. Risk, regulation & responsible AI addresses six risk categories through frameworks like NIST AI RMF—failures are costly. Massachusetts Attorney General settled with student loan company Earnest for \$2.5 million over AI underwriting models that allegedly resulted in discriminatory impact based on race and immigration status, while UC Berkeley research found African American and Latinx borrowers are charged nearly 5 basis points in higher interest rates than credit-equivalent white counterparts. Dashboards & decisions uses quarterly scorecards to make explicit AI investment decisions. High-performing AI implementations achieve ROI exceeding 500% through superior change management, comprehensive measurement, and strategic portfolio optimization, with Gartner showing that organizations with structured ROI measurement achieve 5.2x higher confidence in their AI investments.



**Website**

[www.alignmt.ai](http://www.alignmt.ai)

**Address**

1178 Broadway St, New  
York, NY 10001

**Call**

+617-871-0636