ENTERPRISE AI GOVERNANCE HANDBOOK

AEGIS

Agentic Enterprise Governance & Intelligence Standard

A single unified framework that satisfies every applicable AI regulation — EU AI Act, NIST AI RMF, ISO 42001, Colorado AI Act, California ADMT, Illinois AIVII, NYC LL144, FTC, SEC, GDPR, and emerging federal policy — through seven interlocking operational pillars and a dedicated AI Cost Governance layer.

14+
REGULATIONS COVERED

7
GOVERNANCE PILLARS

€35M
MAX EU AI ACT FINE

60%
COMPLIANCE OVERHEAD REDUCTION

Version 1.1 · May 2026 · Confidential

Primary: US Enterprises · Scope: Agentic & Traditional AI

ARIANA.DIGITAL · PRACTICAL NEXT STEPS

Primers for better Agentic Governance

DIAGNOSE · 5 MIN

AI Transformation Readiness Diagnostic

Seven dimensions. Twenty-eight questions. Under five minutes. Benchmark your org against PwC, McKinsey, Deloitte, BCG, and Gartner 2026 research — then get a scored breakdown and priority actions you can fund this quarter.

ROI: Know which governance gap is blocking scale before you spend on pilots.

TAKE THE DIAGNOSTIC → (https://arian readiness-diagnostics.h

BUILD · SPECIALIST BENCH

Contractor Pipeline · myndQ A-Team

Nine specialist roles for regulated work that ships in production — diagnostic leads, context engineers, agentic UX, MRM, and more. Remote-first, transparent day rates, vetted through talent.myndQ.ai.

ROI: Deploy governed agentic systems without a 12-month hiring cycle.

VIEW OPEN ROLES → (https://ariana.digita digital-pipeline-jobs.

RUN · 6–12 MO PROGRAMS

AI Success Pack · Context + Experience

Forty-two percent of AI projects fail on missing context and unusable experiences. Two programs fix both: knowledge infrastructure and interaction design — the layers that determine whether AI delivers ROI in production.

ROI: Turn governance from compliance overhead into adoption and measurable value.

EXPLORE PROGRAMS → (https://ariana success-pack.

SECTION 1

Executive Overview

The enterprise AI regulation landscape has reached an inflection point. Fourteen or more active and imminent regulations now govern how organizations deploy AI systems — each with its own applicability thresholds, compliance requirements, audit obligations, and penalty structures. Running these as separate programs generates $2–8M in annual duplicated overhead and leaves dangerous gaps. **AEGIS closes that gap by unifying all regulatory obligations into a single, audit-ready governance architecture.**

14+	Jun 30	Aug 2	€35M	60%	
ACTIVE REGULATIONS	COLORADO DEADLINE 2026	EU AI ACT HIGH-RISK	MAX EU AI ACT FINE	AUDIT COST REDUCTION	EVI R FA

Urgent: Two enforcement deadlines are imminent — Colorado AI Act (June 30, 2026) and EU AI Act High-Risk AI compliance (August 2, 2026). Enterprises without a risk management program, consumer disclosures, and conformity assessments in place face live penalty exposure.

SECTION 2

The AEGIS Framework Architecture

Seven interlocking pillars, each addressing a distinct governance domain while sharing a unified evidence base, common risk language, and single audit repository. Implement once — satisfy all regulations simultaneously.

PILLAR 1

Governance Architecture

Board oversight, AI policy, CAIO accountability, vendor contracts

PILLAR 2

AI System Inventory

Registry, risk classification, EU exposure mapping, cost baseline

PILLAR 3

Risk & Impact Assessment

Pre-deploy bias audits, conformity assessment, DPIAs, red teaming

PILLAR 4

Controls & Human Oversight

Guardrails, HITL gates, agentic boundaries, cost circuit-breakers

PILLAR 5

Transparency & Rights

Consumer notices, opt-outs, explainability, anti-AI-washing

PILLAR 6

Monitoring & Response

Logging, drift alerts, cost monitoring, incident response, SEC reporting

PILLAR 7

Regulatory Intelligence

Horizon scanning, maturity growth, ISO 42001 roadmap, program evolution

CROSS-CUTTING

Cost Governance Layer

AI spend baseline, ROI gates, operational cost monitoring, vendor accountability

SECTION 3

Pillar Detail — Requirements & Accountability

PILLAR 1 — GOVERNANCE ARCHITECTURE

Organizational Accountability & Policy

Build the structures that make AI governance real and legally defensible

REGULATIONS SATISFIED

NIST GOVERN | EU AI ACT ART. 9 & 26 | ISO 42001 §5–§6 | OMB M-25-22 | CO AI ACT §4 | FTC / SEC

REQUIRED ACTIVITIES

- Establish AI Governance Board (C-suite + Legal + Tech + Finance)
- Appoint Chief AI Officer (CAIO) with defined accountability
- Publish enterprise AI Use Policy — acceptable and prohibited uses
- Create AI Ethics Charter aligned to OECD principles
- Define RACI for every AI system lifecycle stage
- Board-level AI risk reporting — quarterly minimum
- Vendor AI contract addenda: data usage, IP, audit rights, incident SLAs, cost caps
- Annual AI governance program review and maturity assessment

Agentic AI: Define escalation thresholds — which agent actions require human approval before execution. Document HITL vs. HOTL per agent use case in the AI Use Policy.

CEO / BOARD | CAIO | GENERAL COUNSEL | CISO

PILLAR 2 — AI SYSTEM INVENTORY

AI Registry, Classification & Risk Tiering

You cannot govern what you cannot see — complete visibility is non-negotiable

REGULATIONS SATISFIED

NIST MAP | EU AI ACT ANNEX III | ISO 42001 §8.4 | CO AI ACT | CA ADMT | COST BASELINE

REQUIRED ACTIVITIES

- Maintain live AI System Registry — centralized, version-controlled
- Capture: purpose, data types, affected populations, EU exposure, cost
- Risk tier classification: Prohibited / High / Limited / Minimal
- Map applicable regulations per system per jurisdiction
- Tag all consequential-decision systems (CO, CA, IL, EU)
- Include all third-party and vendor AI in registry
- Track EU AI Database registration status for high-risk systems
- Quarterly inventory refresh; shadow IT AI discovery

Agentic AI: Capture agent action scope (read/write/execute), autonomous decision boundaries, tool integrations, data access permissions, and per-agent compute cost baseline.

CAIO | ENTERPRISE ARCHITECT | BUSINESS UNITS | FINANCE

PILLAR 3 — RISK & IMPACT ASSESSMENT

Pre-Deployment Risk, Bias & Safety Evaluation

Gate deployment on evidence of safety — not assumptions. Every system earns its way into production.

REGULATIONS SATISFIED

NIST MEASURE | EU AI ACT ART. 9–11 | GDPR DPIA | ISO 42001 §8.4 | CO AI ACT | CA ADMT | NYC LL144 BIAS AUDIT | IL AIVII

REQUIRED ACTIVITIES

- Algorithmic Impact Assessment (AIA) before every production deployment
- Independent bias audit for employment AI (NYC mandatory, IL best practice)
- EU AI Act conformity self-assessment or notified body review
- Risk-adjusted ROI gate: cost-benefit sign-off before deployment approval
- Disparate impact analysis: 80% rule across all protected classes
- GDPR DPIA for any AI processing EU personal data
- Red team adversarial testing: hallucination, jailbreak, data poisoning
- All test results stored as audit evidence; re-assessment required on material model changes

Agentic AI: Test agent action chains for unintended cascades. Simulate adversarial prompts attempting to override agent boundaries. Evaluate multi-agent interaction risks and goal hijacking scenarios.

AI RISK TEAM | EXTERNAL AUDITOR | LEGAL / COMPLIANCE | FINANCE (ROI GATE)

PILLAR 4 — CONTROLS & HUMAN OVERSIGHT

Technical Safeguards, Guardrails & Human Gates

Where policy becomes enforced code — the technical implementation of governance

REGULATIONS SATISFIED

NIST MANAGE | EU AI ACT ART. 14–15 | ISO 42001 §8.5 | OMB M-25-22 AGENTIC | TX TRAIGA PROHIBITIONS | EU ART. 5 BANNED AI

REQUIRED ACTIVITIES

- Output guardrails on all production LLMs — content filtering, hallucination control
- Human approval gates for high-stakes decisions (loans, hiring, medical, legal)
- Least-privilege data access for all AI systems
- Accuracy, robustness, cybersecurity testing (EU Art. 15)
- Hard stops for EU Art. 5 prohibited practices (social scoring, manipulation)
- Model drift detection with automatic alerting thresholds
- Input sanitization and prompt injection defenses
- Override mechanisms: humans can always stop, correct, or override any AI decision

🤖 **Agentic AI:** Define explicit action allowlists and blocklists. Circuit breakers: auto-pause agent when confidence below threshold or novel situation detected. Per-agent budget caps — cost-as-a-circuit-breaker for runaway agent loops. Model selection governance: agents use least-cost model meeting accuracy requirements.

AI ENGINEERING | CISO | MLOPS

PILLAR 5 — TRANSPARENCY & RIGHTS MANAGEMENT

Disclosure, Consumer Rights & Explainability

Legal obligations to disclose, explain, and empower every person affected by AI

REGULATIONS SATISFIED

EU AI ACT ART. 13 | GDPR ART. 22 | CO AI ACT DISCLOSURES | CA ADMT OPT-OUT | CA AI TRANSPARENCY ACT | IL AIVII NOTICE | NYC LL144 AUDIT PUB. | FTC ANTI-AI-WASHING

REQUIRED ACTIVITIES

- Consumer/employee notice when AI makes or influences a consequential decision
- Plain-language disclosure statements (multi-jurisdiction templates)
- Opt-out mechanism for ADMT (California CPPA requirement)
- Human review right for significant AI decisions
- Appeal and correction process for adverse AI outcomes
- AI-generated content labeling (CA AI Transparency Act)
- Published bias audit results on company website (NYC LL144)
- No deceptive AI capability claims in marketing, investor materials, or SEC filings (FTC/SEC)

Agentic AI: Users interacting with agents must know they are talking to AI. Autonomous agent actions affecting a user's account, data, or services must be logged and disclosed upon request.

LEGAL | PRODUCT | CUSTOMER EXPERIENCE | MARKETING

PILLAR 6 — MONITORING & INCIDENT RESPONSE

Continuous Surveillance, Cost Tracking & Response

Governance doesn't end at deployment — it intensifies. Every system needs a heartbeat.

REGULATIONS SATISFIED

NIST MANAGE | EU AI ACT ART. 12 & 72–79 | GDPR BREACH NOTIFICATION | ISO 42001 §9–§10 | SEC CYBERSEC. DISCLOSURE | NIS2 DIRECTIVE

REQUIRED ACTIVITIES

- Automatic event logging on all production AI (EU Art. 12 mandatory for high-risk)
- Post-market monitoring with defined KPIs per AI system
- Bias and fairness drift alerts — statistical monitoring on demographic outcomes
- Token/API cost monitoring alongside performance and bias in unified dashboard
- Cost anomaly alerts — flag unexpected spend spikes in AI pipelines
- AI incident classification and escalation procedure
- Serious incident reporting to EU market surveillance authority
- SEC material AI incident disclosure in 8-K / 10-K filings

🤖 **Agentic AI:** Log every agent action with full context: intent → tool called → data accessed → decision → outcome. Immutable audit trails. Real-time anomalous behavior alerts. Cost-as-circuit-breaker triggers automatic agent halt at defined spend threshold.

MLOPS | CISO | AI RISK TEAM | LEGAL | CFO

PILLAR 7 — REGULATORY INTELLIGENCE & PROGRAM EVOLUTION

Continuous Improvement, Horizon Scanning & Maturity Growth

AI regulation evolves faster than any other tech domain — the framework must evolve with it or become a liability

REGULATIONS SATISFIED

ISO 42001 §10 | NIST GOVERN 6.2 | EU AI ACT POST-MARKET | CO SB 26-189 (JAN 2027) | CA ADMT (APR 2027)

REQUIRED ACTIVITIES

- Quarterly Regulatory Horizon Scan (new laws, enforcement actions, guidance)
- Annual AEGIS Framework Review — gap analysis vs. current obligations
- Track: CO SB 26-189 (Jan 2027), CA ADMT full (Apr 2027), EU products (Aug 2027)
- Monitor federal AI legislation developments in Congress
- ISO/IEC 42001 certification program (target within 18 months)
- Annual AI governance training for all AI builders and deployers
- Benchmark governance maturity level annually
- Publish annual AI Transparency Report for external stakeholders

Agentic AI: Track NIST Agentic AI RMF Profile (in development 2026), IEEE P3394 standard, and EU implementation guidance for autonomous AI agents. Update Pillars 3 and 4 as agentic AI capabilities advance.

CAIO | LEGAL / REGULATORY AFFAIRS | AI GOVERNANCE COMMITTEE

SECTION 4

Cross-Regulatory Compliance Matrix

Implement AEGIS once. Each pillar carries compliance weight across multiple regulations simultaneously — eliminating the duplication of siloed programs.

● Primary — directly satisfies ◐ Supporting — contributes – Not applicable

REGULATION	P1 GOVERN	P2 INVENTORY	P3 ASSESS	P4 COI
🇪🇺 EU & INTERNATIONAL				
EU AI Act — Prohibited (Art. 5)	●	●	◐	
EU AI Act — High-Risk AI (Art. 9–17)	●	●	●	
EU AI Act — GPAI Model Obligations	◐	●	●	
GDPR — Automated Decisions (Art. 22)	◐	◐	●	
ISO/IEC 42001 AI Management System	●	●	●	
🇺🇸 US FEDERAL				
NIST AI RMF 1.0 (GOVERN/MAP/MEASURE/MANAGE)	●	●	●	
OMB M-25-22 Federal AI Procurement	●	◐	◐	

REGULATION	P1 GOVERN	P2 INVENTORY	P3 ASSESS	P4 COI
FTC Section 5 — AI Deception & Unfairness	●	–	◐	
SEC AI Risk Disclosure (10-K / 8-K)	●	◐	◐	
🏛 US STATE LAWS				
Colorado AI Act (SB 24-205 / SB 26-189)	●	●	●	
California CPPA ADMT Regulations	◐	●	●	
California AI Transparency Act (SB 942)	◐	–	–	
Illinois AI in Employment (AIVII)	◐	●	●	
NYC Local Law 144 (Hiring Bias Audits)	–	◐	●	
Texas TRAIGA (Prohibited AI Practices)	●	◐	–	

SECTION 5

💰 AI Cost Governance Layer

Ungoverned AI spend is a board-level risk hiding in plain sight. The Cost Governance Layer threads through Pillars 1, 2, 4, and 6 — making compliance the financially optimized choice and giving the CFO, CAIO, and board full visibility into the true cost and ROI of your AI ecosystem.

$2–8M	€35M	60%	3×
Annual Savings	**Max EU AI Fine**	**Audit Cost Reduction**	**Evidence Reuse**
Versus running 14 parallel compliance programs	Prohibited AI practices — 7% of global turnover	Unified program vs. siloed compliance approach	Average regulations satisfied per audit artifact

⚠ Cost of Non-Compliance — Penalty Exposure

EU AI Act — Prohibited Practices	**€35M / 7% turnover**
EU AI Act — High-Risk Non-Compliance	**€15M / 3% turnover**
GDPR AI Processing Violations	**€20M / 4% turnover**
Illinois AIVII — Employment AI	**Uncapped damages + fees**
NYC Local Law 144	**$500–$1,500/day**
Colorado AI Act	**Civil penalties + class action**
SEC — AI Washing Enforcement	**Consent orders + fines**

✅ Cost Savings from AEGIS Implementation

- ↓ Single AI registry eliminates duplicate inventorying across Legal, Compliance, Security, IT — est. 200+ hrs/yr saved per team
- ↓ Shared evidence base: one bias audit satisfies NYC LL144 + IL AIVII + CO AI Act + EU AI Act simultaneously — audit costs cut 60–70%
- ↓ Unified vendor AI contract template replaces bespoke legal drafting — est. $50–150K/yr in outside counsel avoided
- ↓ Pre-deployment gates prevent costly post-launch remediation — AI system rebuilds average $500K–$2M per incident
- ↓ ISO 42001 certification → reduced cyber/AI insurance premiums and procurement differentiation

Five AI Cost Governance Domains

DOMAIN	AEGIS PILLAR	KEY ACTIVITIES
AI System Cost Baseline	P2 Inventory	Every AI system tagged with TCO: compute, licensing, API; cost-per-decision metric; shadow AI spend discovery
Risk-Adjusted ROI Gates	P1 + P3	No deployment without cost-benefit sign-off; expected value > compliance cost + operational cost + penalty exposure; kill switch criteria defined
Operational Spend Monitoring	P6 Monitor	Token/API cost in unified dashboard alongside bias and performance; cost anomaly alerts; model efficiency scoring (accuracy per dollar)
Vendor Cost Accountability	P1 Governance	Contract addenda: escalation caps, audit rights, exit cost limits, data portability; annual vendor cost vs. value review
Compliance Cost Allocation	P1 + P7	Governance spend mapped per system/regulation/BU; evidence reuse tracking; penalty reserve; target: compliance cost below 15% of AI budget

Agentic AI Cost Controls (Pillar 4): Per-agent budget caps with hard stops; cost-as-circuit-breaker for runaway agent loops; multi-agent cost attribution to originating business request; tasks projected to exceed defined threshold require human pre-authorization; model selection governance — agents use least-cost model meeting accuracy requirements.

SECTION 6

Agentic AI — Specialized Governance Controls

Agentic AI systems — autonomous agents that perceive, reason, decide, and act across multi-step workflows — introduce governance dimensions that traditional ML governance does not address. These eight control domains are layered onto all seven AEGIS pillars.

Objective & Boundary Definition

- Define explicit agent objectives in plain language — verifiable by humans
- Specify allowlist of permitted tools, APIs, and data sources per agent
- Define blocklist: actions the agent can NEVER take autonomously
- Set maximum action scope (read vs. write vs. execute)
- Document objective drift triggers requiring human review

Least-Privilege Access Architecture

- Agent identity and access management (IAM) separate from human users
- Scoped credentials per agent — no shared admin accounts
- Just-in-time (JIT) access escalation with automatic revocation
- Data access audited at query/record level
- Network segmentation to limit agent blast radius

Human-in-the-Loop (HITL) Gates

- Classify every agent action: autonomous, advisory, or HITL-required
- HITL required: actions affecting finances, health data, legal status, employment
- Confidence thresholds — below threshold escalates to human automatically
- Novel situation detection: unseen input patterns trigger human review
- Time-boxing: agent pauses after N actions without human checkpoint

Immutable Action Audit Trail

- Log every action: intent → tool → data accessed → decision → outcome
- Immutable, tamper-evident storage (satisfies EU AI Act Art. 12)
- Unique transaction ID per agent task chain for end-to-end traceability
- Retention policy aligned to jurisdiction requirements
- User-accessible action history for transparency obligations

Multi-Agent Orchestration Controls

- Govern agent-to-agent communication — no unchecked trust
- Define orchestrator agent accountability (legal responsibility)
- Prevent prompt injection across agent boundaries
- Test cascading failure scenarios in multi-agent pipelines
- Aggregate risk: combined agent capabilities may exceed individual tier

Agentic Safety Testing Protocol

- Adversarial prompt testing: attempts to override agent boundaries
- Goal hijacking: attacker substitutes agent objective
- Hallucination cascade: verify agent doesn't act on false retrieved facts
- Rollback capability: can every agent action be reversed?
- Chaos engineering: what happens when a dependency fails?

Agentic AI Cost Controls

- Per-agent budget caps: max compute/API spend per task — hard stop
- Cost-as-circuit-breaker: runaway loops trigger automatic halt
- Multi-agent cost attribution to originating business request
- Model selection governance: least-cost model meeting accuracy requirements
- Tasks exceeding cost threshold require human pre-authorization

Emerging Agentic AI Standards

- NIST Agentic AI RMF Profile — concept Apr 2026; monitor for final
- Model Context Protocol (MCP) governance — tool access standardization
- EU AI Act guidance for agentic systems — pending 2026 Commission notes
- IEEE P3394 Agentic AI Standard — in development
- CAIO mandate: agentic AI expected to trigger explicit CAIO roles in federal agencies

SECTION 7

Maturity Model & Compliance Timeline

AEGIS Governance Maturity Model

LEVEL 1 · REACTIVE

Ad Hoc

Governance on paper, not in practice. Compliance is reactive.

- ✓ No AI system registry
- ✓ No formal risk process
- ✓ Reactive to incidents
- ✓ No dedicated AI role
- ✓ ⚠ HIGH exposure

LEVEL 2 · DEVELOPING

Structured

Basic structures exist. Compliance documented but inconsistent.

- ✓ AI inventory exists (partial)
- ✓ AI Use Policy published
- ✓ Risk process defined
- ✓ CAIO or owner named
- ✓ Key state laws addressed

LEVEL 3 · COMPLIANT

Systematic

Full AEGIS framework implemented. All current obligations met.

- ✓ Complete live registry
- ✓ Pre-deploy AIAs for all
- ✓ All state disclosures live
- ✓ EU AI Act conformity done
- ✓ Post-market monitoring
- ✓ Board-level reporting

LEVEL 4 · LEADING

Optimized

Governance is a strategic competitive differentiator.

- ✓ ISO/IEC 42001 certified
- ✓ Automated compliance
- ✓ Real-time dashboard
- ✓ Governance in CI/CD pipeline
- ✓ Annual AI transparency report

Master Compliance Calendar

FEB 2, 2025 — ENFORCED

EU AI Act — Prohibited Practices Ban

Absolute bans on highest-risk AI (social scoring, subliminal manipulation, real-tim biometric surveillance) — enforced now. €35M / 7% turnover.

JAN 1, 2026 — LIVE

Illinois AIVII + Texas TRAIGA + CA AI Transparency Acts

Employment AI discrimination law live in Illinois. Texas categorical AI bans active. California AI transparency and frontier safety laws effective.

AUG 2, 2025 — ENFORCED

EU AI Act — GPAI Model Obligations

General-purpose AI model providers must have technical documentation, copyright policies, and systemic risk assessments in place.

⚠ JUN 30, 2026 — IMMINENT

Colorado AI Act (SB 24-205) — Enforcement Begins

Risk management programs, consumer disclosures, and algorithmic discrimination mitigation required for consequential-decision AI. Civil penalties + class action.

⚠ AUG 2, 2026 — CRITICAL

EU AI Act — High-Risk AI Full Compliance

All Annex III systems must complete conformity assessments, technical documentation, CE marking, EU database registration, and post-market monitoring. €15M / 3% turnover.

JAN 1, 2027

Colorado SB 26-189 (Revised AI Act) + ISO 42001 Target

Revised Colorado AI Act replaces SB 24-205 with broader ADMT scope. ISO/IEC 42001 certification target for Level 4 maturity enterprises.

APR 1, 2027 / AUG 2, 2027

CA ADMT Significant Decisions + EU Regulated Products

Full California ADMT enforcement including opt-outs and PIAs. EU AI Act applies to AI embedded in medical devices, machinery, and other regulated products.

SECTION 8

90-Day Implementation Roadmap

Sequenced by regulatory urgency. Phase 1 must begin immediately — Colorado enforcement is imminent.

PHASE 1 · EMERGENCY

Triage & Immediate Compliance

Now — June 30, 2026 (Colorado deadline)

- [ ] P2: Complete AI system inventory across all business units
- [ ] P2: Tag every system with cost baseline (compute, licensing, API)
- [ ] P2: Classify each system: high-risk, consequential-decision, other
- [ ] P4: Terminate any EU AI Act prohibited practices immediately
- [ ] P5: Draft consumer disclosures for CO + CA + IL systems
- [ ] P3: Run bias audit on all employment AI (IL + NYC requirements)
- [ ] P1: Assign compliance owner per jurisdiction per AI system
- [ ] P5: Publish NYC LL144 bias audit results on company website
- [ ] P3: Deploy risk management program for CO AI Act systems

PHASE 2 · HIGH PRIORITY

EU AI Act High-Risk Sprint

July 1 — August 2, 2026 (EU deadline)

- [ ] P3: Complete technical documentation (Art. 11) for Annex III systems
- [ ] P6: Implement automatic event logging (Art. 12) in production
- [ ] P4: Document human oversight procedures (Art. 14) per system
- [ ] P3: Conduct EU conformity assessment (self or notified body)
- [ ] P2: Register high-risk AI systems in EU AI Database
- [ ] P6: Launch post-market monitoring system
- [ ] P1: Brief board on AEGIS program status and EU readiness
- [ ] Cost: Formalize AI cost baseline and deploy cost monitoring

PHASE 3 · FOUNDATION BUILD

Full AEGIS Operationalization

Q3–Q4 2026

- P1: Establish formal AI Governance Board with quarterly cadence
- P4: Deploy enterprise guardrails platform (Credo AI / IBM watsonx)
- P4: Build agentic AI boundaries, HITL gates, and per-agent cost caps
- P6: Unified monitoring dashboard: bias, drift, performance, cost
- P5: Build multi-jurisdiction disclosure template library
- P1: Execute vendor AI governance addenda across top-50 AI vendors
- P7: Begin ISO/IEC 42001 gap assessment and certification planning

PHASE 4 · MATURITY

Level 4 Leadership & Ongoing

2027 and Ongoing

- P7: Achieve ISO/IEC 42001 certification
- Cost: Publish first Annual AI Cost Governance Report to board
- P7: Monitor Colorado SB 26-189 (Jan 2027) and update program
- P7: California ADMT full enforcement (Apr 2027) — update opt-out flows
- P7: EU regulated products deadline (Aug 2027) — extended conformity
- P6: Automate compliance monitoring with real-time dashboard
- P1: Publish Annual AI Transparency Report for external stakeholders
- P7: Integrate NIST Agentic AI RMF Profile when published

PUT AEGIS TO WORK

Primers for better Agentic Governance

DIAGNOSE · 5 MIN

AI Transformation Readiness Diagnostic

Score governance maturity across strategy, data, technology, talent, operating model, risk, and adoption. Use the results to sequence AEGIS pillars against your biggest regulatory and ROI gaps.

Problem solved: Stop guessing where agentic AI will stall.

TAKE THE DIAGNOSTIC → (https://ariana readiness-diagnostics.h

BUILD · SPECIALIST BENCH

Contractor Pipeline · myndQ A-Team

Financial services, healthcare, and energy engagements need specialists who have shipped under SR 11-7, HIPAA, and NERC CIP — not generalists. Pull vetted talent when Colorado and EU deadlines are non-negotiable.

Problem solved: Staff governance sprints without permanent headcount.

VIEW OPEN ROLES → (https://ariana.digita digital-pipeline-jobs.

RUN · 6–12 MO PROGRAMS

AI Success Pack · Context + Experience

Operationalize the context layer (metadata, knowledge graphs, RAG++) and experience layer (copilots people actually use) so agentic systems stay inside AEGIS guardrails after go-live.

Problem solved: Governance that survives production, not just audit day.

EXPLORE PROGRAMS → (https://ariana success-pack.