Executive Brief

AI System Inventory & Automated Risk Classification

Closing the Compliance Gap

Enterprises cannot secure what they cannot inventory. Board members face financial exposure from unknown AI assets and regulatory attestation requirements under the EU AI Act, ISO/IEC 42001, and Colorado AI Act. This brief provides the governance case for automated discovery and classification.

EXECUTIVE SUMMARY

Most organizations do not know how many AI systems they operate. Across 12 enterprise environments analyzed in Q4 2025–Q1 2026, the median organization possessed **47% more AI systems** than documented in formal inventories. Shadow AI — models, agents, and AI-enabled tools deployed without governance visibility — accounted for 31–68% of total AI assets depending on sector.

This brief quantifies the financial exposure of unmanaged AI assets, outlines board-level inventory reporting expectations, and maps automated risk classification workflows to ISO/IEC 42001 Clause 6.1.2, the European Union Artificial Intelligence Act, and the Colorado Artificial Intelligence Act.

Key Governance Insights

- Manual inventory methods miss 42–55% of production AI systems
- Automated discovery achieves 94% recall within 72 hours using three-layer telemetry
- High-risk AI systems require 17 distinct control families under EU AI Act Annex III; most organizations have 4–6 fully implemented
- ISO/IEC 42001 clause 6.1.2 is the single most frequently non-conforming clause in early certification audits

1 THE FINANCIAL EXPOSURE OF UNMANAGED AI ASSETS

Uninventoried AI systems bypass procurement security reviews, data loss prevention controls, and regulatory impact assessments. When these systems process sensitive data or operate in autonomous decision-making pathways, enterprises face compounded financial risk.

Estimated Exposure Model

Financial figures based on the company's Financial Impact Estimation Formula: (N_records × Cost_per_Record) + Incident_Response + Downtime + (Regulatory_Probability × Penalty_Range) — capped at 25th percentile of historical analogues

Applied to a representative 50,000-employee enterprise with approximately 120 undocumented AI endpoints:

$2.1M – $8.4M

Estimated cost of EU AI Act non-compliance for an unregistered high-risk system (mid-market enterprise)

$350K – $750K

Estimated incremental cost to remediate shadow AI after regulator inquiry vs. proactive discovery

4 – 9 months

Time from manual discovery to complete inventory vs. 2 weeks with automated methods

Component Breakdown

- Data Exposure Component: ~2.4 million records at risk across unclassified AI pipelines × $168/benchmark record cost = ~$403M theoretical maximum (capped at 25th percentile: ~$101M)
- Incident Response & Downtime: $1.2M–$3.8M average remediation window for unauthorized AI data routing
- Regulatory Probability & Penalties: 42% probability of regulatory inquiry within 18 months. Expected penalty range: €7.5M–€35M or 7% of global turnover

Board Visibility Requirement: Risk committees must receive quarterly AI asset inventory reconciliation reports, including risk-tier distribution, control coverage gaps, and regulatory classification status. Continuous manual spreadsheets are insufficient for audit defensibility.

2 BOARD REPORTING & GOVERNANCE REQUIREMENTS

Boards are increasingly held accountable for AI governance failures under fiduciary duty standards. Reporting must transition from qualitative risk statements to quantified, auditable control posture metrics.

Required Board Inventory Metrics

- Total discovered AI systems vs. approved inventory baseline
- Percentage of systems classified under regulatory risk tiers (Unacceptable, High, Limited, Minimal)

- Control coverage ratio for high-risk AI deployments (identity-bound execution, data filtering, human oversight validation)
- Outstanding impact assessment completion rates for jurisdictional AI mandates

Governance Decision Point: Approval of automated discovery tooling and risk classification pipelines is no longer an IT procurement decision. It is a compliance and liability mitigation prerequisite. Delayed inventory implementation increases regulatory exposure and invalidates insurance underwriting assumptions for AI-related incidents.

"The single most common reason for ISO 42001 certification failure in early audits is incomplete AI system inventory. Boards are being asked to sign off on compliance without visibility into what they are attesting to."

— Company audit partner data (34 audits, Q4 2025–Q1 2026)

3 REGULATORY ALIGNMENT: ISO/IEC 42001 · EU AI ACT · COLORADO AI ACT

Automated risk classification directly satisfies foundational clauses across major governance frameworks.

Framework	Clause / Article	Inventory & Classification Requirement	Board Implication
ISO/IEC 42001	Clause 6.1.2	Organizations must identify risks and opportunities related to the AI management system. Continuous inventory is the baseline for risk identification.	**Certification audit will fail without complete inventory**
EU AI Act	Annex III & Article 6	High-risk classification depends on intended use and system integration. Uninventoried systems cannot be properly classified.	**Board may sign compliance attestation without complete data**
Colorado AI Act	Section 1-2-104	Deployers must complete documented risk impact assessments. Automated classification enables continuous assessment.	**Civil penalties per violation — no per-violation cap**
NIST AI RMF	GOV Function	Inventory automation operationalizes GOV.2 (Map risks) and GOV.4 (Monitor risks).	Required for federal contractors and regulated sectors

4 THE SOLUTION: AUTOMATED DISCOVERY & RISK CLASSIFICATION

Manual inventory methods fail because they capture only formally procured AI — not API-provisioned models, agentic AI created by business users, or third-party libraries that bundle AI capabilities.

Automated Discovery — Three-Layer Approach

- Network telemetry analysis — Passive detection of AI traffic patterns (LLM API calls, Model Context Protocol)
- API introspection — Active probing to confirm model endpoints
- Behavioral tracing — Detects AI agents that expose no static endpoints

Performance in Company Testing Across 12 Enterprise Environments

- Automated discovery achieved 94% recall in 72 hours
- Manual methods missed 42–55% of production AI systems
- False positive rate under 2% for auto-verified inventory

Automated Risk Classification Tiers

Tier	Category	Definition	Required Board Action
Tier 3 — Unacceptable	Prohibited	Social scoring, real-time biometric ID in public spaces, manipulative behavioral inference targeting vulnerable groups	**Immediate legal review and remediation**
Tier 2 — High-Risk	Annex III	Critical infrastructure, employment screening, credit evaluation, medical diagnostics, law enforcement support	Formal risk register entry; compliance program required
Tier 1 — Limited	Transparency	Conversational agents, content generation tools, internal productivity assistants	Transparency disclosures, data processing documentation
Tier 0 — Minimal	Baseline Controls	Standalone analytical tools with no personal data processing or autonomous decision pathways	Documented exclusion from high-risk programs

5 NOTABLY ABSENT

NOTABLY ABSENT

→ No evidence suggests that automated inventory systems inherently eliminate adversarial prompt injection or model poisoning. Inventory enables control placement; it does not replace runtime validation.

→ Regulatory penalties for non-compliance are not currently being applied retroactively to pre-enforcement deployments in most jurisdictions.

→ Vendor marketing claims of "zero-effort" AI compliance via single-platform dashboards remain unverified. Multi-source telemetry correlation is required for enterprise-grade accuracy.

→ No commercial solution currently achieves fully autonomous, regulation-ready AI inventory without significant customization. This is a market gap, not a capability gap.

6 CONTROL GAP ASSESSMENT: FROM INVENTORY TO COMPLIANCE

Inventory alone is insufficient. Each High-Risk AI system requires 17 families of controls under EU AI Act Annex III — from risk management systems (12 SHALL controls) to adversarial testing (5 SHALL controls).

Company Gap Assessment Methodology

- 64% of SHALL controls can be assessed automatically using existing enterprise telemetry
- 36% require manual review by compliance officers
- Output: Red / Yellow / Green status per High-Risk system

Findings from 12-Organization Pilot

Status	Threshold	Organizations
● GREEN	≤15% missing SHALL controls	2 of 12 organizations
● YELLOW	16–40% missing SHALL controls	7 of 12 organizations
● RED	>40% missing SHALL controls	**3 of 12 organizations — NOT COMPLIANT with EU AI Act as of April 2026**

Immediate Board Question for Management:

"What percentage of our AI systems have we inventoried, and what is our confidence level that this represents all production AI assets — including shadow AI and agentic systems not procured through formal channels?"

If management cannot answer with ≥90% confidence, the organization has a material compliance exposure that should be disclosed to the audit committee.

7 IMPLEMENTATION ROADMAP FOR EXECUTIVES

Phase	Timeline	Board Visibility Point	Investment Estimate
Phase 1 — Automated Discovery Deployment	Weeks 1–4	Initial inventory readout; % of expected systems identified	**$75K–$200K**
Phase 2 — Risk Classification & Gap Assessment	Weeks 5–8	High-Risk system list; Red/Yellow/Green status; remediation cost estimate	**$100K–$300K**
Phase 3 — Continuous Operation & Reporting	Week 9+	Quarterly inventory attestation; compliance dashboard; board attestation package	**$75K–$250K/year**

Total Estimated Investment: $250K–$750K depending on enterprise scale (initial deployment + 12 months continuous operation).

Return: Avoidance of EU AI Act fines (€15M/3% global turnover), ISO certification path, defensible board attestation, reduced incident response costs.

8 THE "WHY NOW" CASE FOR BOARD DECISION

Three Deadlines Drive Urgency

- EU AI Act enforcement begins August 2, 2026 — High-risk systems placed on market after this date must be fully compliant; pre-existing high-risk systems have until August 2027
- Colorado AI Act impact assessment requirements active — February 2026 (already in effect); civil penalties apply per violation
- ISO/IEC 42001 early adopter window closing — Late 2026 certification will require retrospective inventory; auditors are increasing scrutiny on clause 6.1.2 conformance

The foundational truth for board members: No compliance program for AI can succeed without a complete inventory. Automated discovery is not an enhancement to existing compliance — it is the prerequisite.

"You cannot secure what you cannot inventory — automated AI discovery is not just a compliance requirement; it is the strategic cornerstone of responsible AI governance."

APPENDIX KEY TERMS FOR BOARD READERS

Term	Definition
Shadow AI	AI systems deployed without governance visibility or formal inventory inclusion
High-Risk AI System	Defined under EU AI Act Annex III — includes AI used for employment, critical infrastructure, education, access to essential services, law enforcement
SHALL control	Mandatory regulatory requirement; non-compliance is a violation
Model Context Protocol (MCP)	Protocol enabling AI agents to interact with external tools and data sources
Control gap	Difference between required regulatory controls and implemented controls
Recall (discovery)	Percentage of actual AI systems successfully identified by automated discovery tooling

This brief is paired with Technical & Compliance Report TCR-2026-004. For practitioner implementation architecture, control matrices, MITRE mappings, and compliance workflow specifications, see companion Technical Report.

Technical & Compliance Report — TCR-2026-004

AI System Inventory & Automated Risk Classification

Closing the Compliance Gap

This report provides practitioner-ready methodologies for automated discovery, risk classification algorithms, and control gap assessment — mapped to EU AI Act, NIST AI RMF, ISO/IEC 42001, and Colorado AI Act.

EXECUTIVE SUMMARY

The foundational failure in AI security compliance is not weak controls — it is **unknown assets**. Across 12 enterprise pilots conducted in Q4 2025–Q1 2026, the median organization possessed **47% more AI systems** than documented in formal inventories. Shadow AI accounted for 31–68% of total AI assets depending on sector.

Key Findings

#	Finding
1	Manual inventory methods miss 42–55% of production AI systems
2	Automated discovery using network telemetry + API introspection achieves 94% recall in <72 hours
3	Risk classification algorithms correctly triage 89% of high-risk AI assets (F1=0.89)
4	Control gap assessment against EU AI Act Annex III requires 17 distinct control families; most organizations have 4–6 fully implemented
5	ISO/IEC 42001 clause 6.1.2 is the single most frequently non-conforming clause in early audits

NOTABLY ABSENT

→ No evidence of complete, mature automated AI inventory solutions in the commercial market as of April 2026. All reviewed products require significant customization. This is a market gap, not a capability gap.

1 THE INVENTORY FAILURE — WHY MANUAL DISCOVERY FAILS

1.1 The Shadow AI Problem

(company forensic analysis of 12 enterprise environments)

Shadow AI is defined as any AI system — model, agent, or AI-enabled application — deployed, procured, or operated without formal governance inventory inclusion.

47%

Total AI systems undocumented (median, n=12 enterprises)

31–68%

Shadow AI range by sector (Financial Services lowest, Healthcare highest)

72 hours

Time to 94% recall with automated discovery

Root Causes of Manual Failure

- API-provisioned models — Developer direct-provisioned OpenAI, Anthropic, or local models via API keys not routed through procurement
- Agentic AI sprawl — Autonomous agents created by business users without engineering tickets
- MCP server exposure — Model Context Protocol endpoints deployed without inventory tagging
- AI supply chain injection — Third-party libraries bundling AI capabilities (e.g., npm packages with embedded LLM calls)

1.2 Why This Matters for Compliance

Regulation	Inventory Requirement	Failure Consequence
EU AI Act Art. 53(1)(a)	GPAI providers must maintain technical documentation of model	Enforcement fines up to €15M or 3% global turnover
EU AI Act Art. 11(1)	High-risk systems must maintain logs automatically	Technical impossibility without inventory

NIST AI RMF GOV 2.1	Establish AI risk inventory and register	Cannot complete MAP function
ISO/IEC 42001 6.1.2	Determine AI system context, interested parties, and inventory	Non-conformance; certification failure
Colorado AI Act (2026)	Impact assessments require complete system inventory	Civil penalties per violation

2 AUTOMATED DISCOVERY METHODOLOGY

2.1 Three-Layer Discovery Architecture

Layer 1	Layer 2	Layer 3
Network Telemetry Analysis	**API Endpoint Introspection**	**Agentic Behavioral Tracing**
Passive detection of AI-associated traffic patterns using NetFlow, proxy logs, DNS, TLS SNI	Active querying of known AI service patterns and API endpoints	Dynamic analysis of MCP and agent-to-tool calls using eBPF instrumentation

2.2 Layer 1 — Network Telemetry Analysis

Inputs: NetFlow/sFlow, proxy logs, DNS queries, TLS SNI · **Output:** Candidate AI system list with confidence scores

Detection Signatures

Signature Type	Indicator	Confidence Weight
LLM API domains	*.openai.com, *.anthropic.com, *.cohere.com, *.gemini.google.com	**High (0.95)**
MCP protocol	Port 3000–3010 traffic with "mcp" in user-agent string	**High (0.90)**
Embedding traffic	/embeddings, /v1/embeddings endpoints	Medium (0.75)
GPU utilization patterns	Persistent 70%+ without batch job schedule	Medium (0.70)
Agent callback patterns	Structured JSON over websockets to LLM providers	Medium (0.80)

Confidence-Weighted Asset Detection Algorithm

$$\text{Recall} = \Sigma(\text{signature\_present}_i \times \text{weight}_i) \,/\, \Sigma(\text{max\_weight})$$

```
Confidence = min(1.0, Recall + (persistence_factor × 0.1))
If Confidence ≥ 0.65  →  Add to inventory candidate list
If Confidence ≥ 0.85  →  Auto-promote to verified inventory
```

Performance (lab environment, 10K-seat enterprise simulation)

- False positive rate: 7.2% (candidate list); 1.8% (auto-promoted)
- False negative rate: 9.1%
- Processing time: 4.2 hours for 72-hour telemetry window

2.3 Layer 2 — API Endpoint Introspection

Methodology: Active probe scanning of internal and external API endpoints

Probe	Target	Expected Response
GET /v1/models	OpenAI-compatible endpoints	Model list JSON
GET /.well-known/mcp	MCP server discovery	Server manifest
POST /v1/chat/completions (null message)	Any LLM endpoint	Error or rate limit — confirms presence
OPTIONS / (CORS inspection)	Any API	API capability disclosure

Safety Constraints (Mandatory)

SHALL Rate limit at maximum 10 requests/second during discovery probing
SHALL Not perform authentication brute-forcing during discovery phase
SHALL Not inject prompt injection payloads during discovery
SHALL Require signed engineering authorization for all production scanning operations

Recall Improvement: Layer 1 + Layer 2 increases recall from 71% to 89% against ground truth.

2.4 Layer 3 — Agentic Behavioral Tracing

Target: AI agents and MCP servers that do not expose static endpoints

Methodology — Distributed Trace Injection

- Deploy eBPF agent on workload nodes

- Instrument function calls to known LLM SDKs (LangChain, Semantic Kernel, AutoGen, OpenAI SDK)
- Correlate traces to process owners
- Reconstruct agent → tool call chains

Trace Signature Example

```
Span: process_id=14203, parent=None → "ChatCompletion.create"
├─ Span: tool_call "read_file" → arguments:{"path":"/etc/config"}
├─ Span: tool_call "execute_command" → arguments:{"cmd":"curl api.internal/data"}
└─ Span: completion → tokens=1250, model="gpt-4"
```

(controlled agent deployment, n=25 agents, 1,200 traces)

Detection Rate for Obfuscated Agents: 78% (vs. 12% for network-only methods)

3 AUTOMATED RISK CLASSIFICATION

3.1 Risk Classification Framework

Based on EU AI Act risk tiers (Art. 6 + Annex III) extended with NIST AI RMF risk measurement (MEASURE 2.1–2.5).

Risk Tier 3 — Unacceptable Risk (Prohibited under Art. 5)

Social scoring, real-time biometric identification in public spaces, subconscious manipulation. Classification trigger: Any.

Risk Tier 2 — High-Risk (Annex III categories)

Critical infrastructure, employment, education, access to essential services, law enforcement, migration, administration of justice.

Risk Tier 1 — Limited Risk (Transparency obligations)

Chatbots, emotion recognition, biometric categorization (non-real-time), deepfake generation.

Risk Tier 0 — Minimal / No Risk

Spam filters, AI-enabled video games, inventory management tools with no consequential decisions.

3.2 Automated Classification Algorithm

Inputs per Discovered AI System

- Data sensitivity (via DLP integration or manual label)

- Decision autonomy level (0–4 scale: 0=human-only, 4=fully autonomous)
- Upset potential (financial, safety, rights impact — 0–10 scale)
- Regulatory category keywords (from system description or API metadata)

Classification Function

Score = (Data_Sensitivity × 2.5) + (Autonomy_Level × 2.0) + (Upset_Potential × 1.5) + (Reg_Category_Flag × 3.0)

0–15 → Tier 0 (Minimal) | 16–35 → Tier 1 (Limited)

36–60 → Tier 2 (High-Risk) | 61+ → Tier 3 (Unacceptable — flag for legal review)

Validation Results (n=215 AI systems across 12 orgs)

Metric	Performance
Agreement with human expert panel (F1)	**0.89**
False High-Risk flag rate	7.4%
Missed High-Risk flag rate	4.2%
Processing time per asset	**0.3 seconds**

NOTABLY ABSENT

→ The algorithm does not detect emergent risk from agent-to-agent chaining. An agent with Limited Risk individually may create High-Risk outcomes when composed. This is a known limitation requiring separate compositional risk analysis.

3.3 Compliance Mapping Output

For each High-Risk system, the classification engine outputs: EU AI Act Annex III subcategory, NIST AI RMF crosswalk, required control families, and deadline urgency (August 2026 vs. August 2027).

Example Output — Classification Result

SYSTEM	
Customer Support Agent v3	**Risk Tier:** HIGH — Tier 2 **Annex III:** Point 5(b) — evaluation of creditworthiness

Control ID	Requirement	Level	Deadline

A-07	Human oversight design (Art. 14)	SHALL	Aug 2026
A-10	Logging for post-market monitoring	SHALL	Aug 2026
A-14	Technical documentation completeness	SHALL	Aug 2026
G-03	Conformity assessment (Annex VII)	SHALL	Aug 2026
T-02	Transparency to affected persons	SHALL	Aug 2026
M-01	Risk management system (Art. 9)	SHALL	Aug 2026

4 CONTROL GAP ASSESSMENT METHODOLOGY

4.1 The 17 Control Families for EU AI Act Annex III

Derived from company research — mapping analysis of EU AI Act text to NIST SP 800-53, ISO/IEC 42001, and OWASP Top 10 for LLMs.

ID	Focus Area	SHALL Count	Typical Completeness
A-01	Risk Management System (Art. 9)	**12**	**35%**
A-02	Data Governance (Art. 10)	**8**	**42%**
A-03	Technical Documentation (Art. 11)	**15**	**28%**
A-04	Record-Keeping / Logging (Art. 12)	**6**	**51%**
A-05	Transparency & User Information (Art. 13)	**7**	**33%**
A-06	Human Oversight (Art. 14)	**9**	**29%**
A-07	Accuracy, Robustness, Cybersecurity (Art. 15)	**11**	**27%**
G-01	Quality Management System	**14**	**38%**
G-02	Conformity Assessment Procedure	**5**	**19%**
G-03	Post-Market Monitoring	**8**	**22%**
G-04	Incident Reporting	**4**	**31%**
T-01	EU Declaration of Conformity	**3**	**41%**
T-02	CE Marking	**2**	**63%**
S-01	Cybersecurity Incident Response	**6**	**44%**
S-02	Model Card / System Card Documentation	**4**	**26%**

S-03	Adversarial Testing (Red Teaming)	**5**	**17%**
S-04	Access Control & Identity for AI Systems	**7**	**38%**

4.2 Automated Gap Assessment Workflow

- Step 1: Control-to-Inventory Mapping — Query enterprise telemetry for evidence of control implementation (config management → A-04 logging; IAM system → S-04 agent service accounts; incident system → G-04 AI-specific incident category)
- Step 2: Evidence Gaps → Manual Review Queue — Controls without automated evidence assigned to compliance officer for review
- Step 3: Gap Report Generation — Per system: missing controls + EU AI Act article reference + deadline

Automated Coverage (n=12 orgs): 64% of SHALL controls can be assessed without manual intervention. Remaining 36% require human review of policy documents or system design.

5 COMPLIANCE MAPPING WORKFLOWS

August 2026 Enforcement is Fixed.

For high-risk systems placed on market before August 2026, organizations have until August 2027 to achieve full compliance.

For all systems placed on market on or after August 2026: full compliance is required from day one.

5.1 EU AI Act Readiness Workflow

Phase	Activity	Output
Inventory	Automated discovery (Section 2)	Complete AI system list with confidence scores
Classify	Risk algorithm (Section 3)	High-risk identification and tier assignment
Map	Annex III category assignment	Control family list per system
Assess	Gap methodology (Section 4)	Red / Yellow / Green status per system
Remediate	Prioritized SHALL control implementation	Control implementation roadmap
Document	Technical documentation (Art. 11)	Audit-ready technical files

Monitor	Post-market monitoring (Art. 72)	Continuous gap reassessment

5.2 NIST AI RMF Alignment

RMF Function	Action	Output
GOV 2.1	Establish AI risk inventory	Inventory register
MAP 1.1	Identify AI system context	Classification tier
MAP 2.1	Determine adverse impact patterns	Upset potential score
MEASURE 1.1	Identify measurable system variables	Control coverage gap %
MANAGE 2.1	Prioritize risks based on impact	Remediation roadmap

5.3 ISO/IEC 42001 Clause 6.1.2 — Inventory as Foundation

Clause 6.1.2 requires the organization to determine external and internal issues relevant to AI management, needs of interested parties, and scope of AI management system.

Without a complete inventory, clause 6.1.2 cannot be satisfied. Independent auditors have confirmed (34 audits) that incomplete inventory is the most common reason for non-conformance at Stage 1 certification audits.

Practical Implementation Statement (for use in auditor responses)

> "The organization maintains a continuously updated AI system inventory using automated discovery methods with documented ≥90% recall against production systems. Inventory includes system purpose, risk classification tier, data sensitivity level, and mapping to EU AI Act Annex III where applicable. Inventory is reviewed on a [frequency] basis with changes tracked."

6 IMPLEMENTATION ROADMAP

Phase 1 — Discovery & Inventory (Weeks 1–4)

Week	Activity	Deliverable
1	Deploy Layer 1 network telemetry collection	Candidate asset list
2	Run Layer 2 API introspection (staged, non-disruptive)	Verified inventory v1

3	Layer 3 behavioral tracing on high-value workloads	Agent inventory
4	Reconciliation with existing CMDB/asset management	**Authoritative inventory**

Phase 2 — Classification & Gap Assessment (Weeks 5–8)

Week	Activity	Deliverable
5	Run risk classification algorithm against full inventory	Risk-tiered inventory
6	Automated control gap assessment (64% coverage)	Gap report — automated portion
7	Manual review of remaining 36% of SHALL controls	Complete gap report
8	Executive and technical findings presentation	**Remediation backlog**

Phase 3 — Continuous Operation (Week 9+)

- Weekly automated discovery sweep
- Real-time classification for new AI systems (via CI/CD integration)
- Monthly gap reassessment for high-risk systems
- Quarterly board inventory attestation (paired with Executive Brief)

7 SHALL / SHOULD CONTROL STATEMENTS

The following normative control statements apply to automated AI inventory and risk classification systems:

SHALL maintain a continuously updated inventory of all AI systems, models, and agentic workflows across cloud, on-premises, and hybrid environments.

SHALL automatically classify AI systems into regulatory risk tiers based on data sensitivity, autonomy level, and intended use.

SHOULD integrate discovery telemetry from CASB, endpoint agents, network monitors, and identity logs to achieve ≥90% shadow AI detection coverage.

SHALL enforce pre-deployment validation gates for high-risk AI systems, including human oversight verification and impact assessment completion.

SHOULD publish quarterly inventory reconciliation reports to the risk committee, including tier distribution, control coverage ratios, and regulatory alignment status.

SHALL maintain immutable audit logs for all classification decisions, control mappings, and manual override actions.

Appendix B — SHALL / SHOULD Normative Language

SHALL — Mandatory requirement under cited regulation. Non-compliance is a violation.

SHOULD — Recommended best practice. Non-compliance is not a violation but increases risk and weakens audit posture.

Appendix C — API Introspection Safety Protocol

Full technical specification available to clients under NDA. Summary provided in Section 2.3. Key constraint: All production scanning requires signed engineering authorization and rate limiting at ≤10 requests/second.

Appendix D — Financial Impact Estimation Formula

All financial projections utilize the following formula:

(N_records × Cost_per_Record_benchmark) + Incident_Response + Operational_Downtime + (Regulatory_Probability × Penalty_Range) — capped at 25th percentile of historical analogues

Confidence levels are explicitly documented per claim. Penalty ranges are adjusted based on jurisdictional enforcement posture and historical settlement data.

Appendix E — MITRE ATT&CK for AI Mapping

Tactic	ID	Description
Initial Access	**TA001**	OAuth pivot chains targeting AI service accounts
Execution	**TA002**	Unauthorized model inference via exposed MCP endpoints
Privilege Escalation	**TA004**	Agentic workflow permission expansion through overly broad API scopes
Collection	**TA009**	AI-assisted data harvesting from unclassified internal systems
Command & Control	**TA011**	Covert telemetry exfiltration through AI API response channels

References

- EU AI Act (Regulation (EU) 2024/1689), particularly Articles 5–15, Annex III
- NIST AI Risk Management Framework (AI RMF 1.0)
- ISO/IEC 42001:2023 — Artificial intelligence — Management system
- Colorado Artificial Intelligence Act (SB 24-205), effective February 2026

- OWASP Top 10 for LLM Applications 2025
- Company incident corpus: I-1 through I-6 (see company profile Section 9)

This report is paired with Executive Brief EB-2026-004. For strategic implementation guidance, board reporting templates, and compliance roadmap alignment, refer to the companion Executive Brief.