ACC – Mountain West | Tech Law Symposium | 2026

Smart AI Governance in 2026

Building a Compliant, Resilient AI Program

Peter Stockburger
Partner | Foley & Lardner LLP
peter.stockburger@foley.com | 618.876.1971

Roadmap

01 **Where AI Is Heading** — From deterministic rules to agentic autonomous systems

02 **Enterprise Adoption Trends** — Agentic adoption on the rise

03 **The Global Legal Landscape** — EU AI Act, South Korea Basic AI Act, and a shifting US framework

04 **Smart AI Governance in 2026** — Benchmarking your governance program to scale

05 **Governance In The Trenches** — Key takeaways when buying and selling AI

Where AI Is Heading

Deterministic Systems

Rules-Based AI

- Rules-based, predictable outcomes
- Traditional computing
- Auditable, transparent, explainable
- Spam filters, decision-tree chatbots, scheduling tools

Probabilistic Models

Statistical AI

- Predictions from data
- Machine learning & deep learning
- LLMs: ChatGPT, Gemini, Claude
- Vision models: Sora, Midjourney
- World models: DeepMind, NVIDIA Cosmos

Agentic AI

Autonomous Agents

- Autonomous multi-step task execution
- Self-directed tool use & API calls
- Multi-agent orchestration pipelines
- Minimal human oversight per step
- Real-world consequential actions

Legal Definition: An engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments. — Cal. Gov. Code §11546.45.5(a)(1); EU AI Act; 15 USC §9401

Enterprise AI Adoption

Key trends from McKinsey, PwC, Menlo Ventures, Deloitte & OpenAI surveys — 2025

Overall Adoption

78%

of companies use AI in at least one business function

McKinsey State of AI, 2025

Use Case Split

59% **Internal**

41% **External**

Both convert to production at near-identical rates

Menlo Ventures, 2025

Industry Leaders

Second Talent / OpenAI Enterprise, 2025

Agentic AI

35%

+44% **experimenting**

already deploying - reached 35% in 2 yrs vs. gen AI's 3-yr pace

MIT Sloan / BCG, 2025

Workforce Impact

75%

of workers report improved speed or output quality

Worker AI access +50% in 2025 · CAIOs at 61% of enterprises

OpenAI Enterprise / Wharton, 2025

Investment & ROI

$37B

3.2× **YoY growth**

enterprise gen AI spend in 2025

88% of execs plan budget increases - agentic AI cited as primary driver

Menlo Ventures / PwC, 2025

Enterprise AI Adoption: Agents

Sources: CSA State of AI Security & Governance (Dec 2025) · CSA Securing Autonomous AI Agents (Feb 2026) · Gartner AI Hype Cycle 2025

40%

of organizations have agents in production

31%

running pilots or active tests

19%

planning deployment within the next year

26%

have comprehensive AI security governance

78%

lack policies for creating and managing AI agent identities

18%

of security leaders confident their IAM can manage agents

Generative AI: Key Risks for Enterprises

NIST AI RMF identifies six risk categories that enterprises must manage when deploying generative AI — spanning trustworthiness, security, fairness, accountability, and governance across the full AI lifecycle.

Confabulation & Accuracy

AI systems generate plausible but false outputs (hallucinations), undermining reliability in high-stakes decisions.

Security & Adversarial Threats

Prompt injection, model inversion, and data poisoning attacks expose systems to manipulation and data exfiltration.

Data Privacy & Confidentiality

Training data may embed PII or proprietary data, creating leakage risks that violate GDPR and other regulations.

Bias, Fairness & Equity

Models can perpetuate or amplify societal biases, leading to discriminatory outputs and legal liability.

Transparency & Explainability

Black-box outputs limit auditability, making it difficult to justify AI decisions to regulators and stakeholders.

Governance & Accountability

Unclear ownership, insufficient human oversight, and absent AI policies create compliance and reputational exposure.

The New Agentic Attack Surface

Critical Shift: Agentic AI breaks the traditional security perimeter. Autonomous action at machine speed means a breach can cascade across systems before any human can respond.

Prompt Injection

Malicious inputs hijack agent goals and redirect autonomous actions to attacker-controlled objectives

Agent Impersonation

Adversaries masquerade as trusted agents to infiltrate multi-agent workflows and extract data

Memory Manipulation

Poisoning persistent memory stores to corrupt future agent reasoning and long-horizon decisions

Tool Misuse & Abuse

Exploiting broad tool access to exfiltrate data, move laterally, or deploy malicious payloads

Goal Misalignment

Agents pursue objectives that deviate from intended human directives at scale without detection

Supply Chain Attacks

Compromising model weights, plugins, or upstream APIs to inject malicious behavior system-wide

Agentic AI System Layers & Emerging Communication Protocols

Key Insight: Modern agentic systems operate across multiple interdependent layers, including foundation models to orchestration frameworks to external services. Emerging inter-agent protocols enable powerful new capabilities, but each protocol boundary introduces a distinct and largely unregulated risk surface.

SYSTEM ARCHITECTURE LAYERS

1. **Foundation Model Layer**
   LLM reasoning core (e.g. GPT-4, Claude, Gemini)
2. **Orchestration Layer**
   Agent framework managing task planning & tool dispatch
3. **Tool / Integration Layer**
   APIs, databases, code execution, web browsing
4. **Protocol Communication Layer**
   MCP · A2A · A2C — inter-agent & service messaging
5. **External Services Layer**
   Commerce platforms, enterprise systems, third-party agents

▼ *Data & control flow across all layers*

EMERGING COMMUNICATION PROTOCOLS

MCP

Model Context Protocol

Standardizes how AI models connect to tools, APIs & data sources — enabling plug-and-play tool access for LLMs.

RISK SURFACE

Tool poisoning, prompt injection via malicious tool responses, unauthorized data exfiltration.

A2A

Agent-to-Agent Protocol

Defines how autonomous agents discover, negotiate with, and delegate tasks to other agents across systems.

RISK SURFACE

Trust chain attacks, rogue sub-agent behavior, cascading failures across multi-agent pipelines.

A2C

Agent-to-Commerce Protocol

Governs AI agents transacting directly with commerce platforms — purchasing, booking & contracting without human approval.

RISK SURFACE

Unauthorized financial transactions, contract formation by AI, consumer protection & fraud exposure.

⚠ **Critical:** Each protocol boundary is a potential attack vector.

Global AI Legal Landscape

Key Tension: Global consensus on the need to regulate AI — but no consensus on how. Three distinct regulatory models are emerging simultaneously: risk-based, light-touch, and state-control.

EU AI Act

Risk-Based

- ✓ In force Aug 1, 2024 — full effect Aug 1, 2026
- ✓ Classifies AI: Unacceptable / High / Limited / Low risk
- ✓ Extraterritorial — covers all providers placing AI into EU market
- ✓ High-risk: mandatory diligence, technical docs, EU database registry
- ✓ GPAI models: transparency & copyright compliance obligations

US Federal & States

Fragmented

- ✓ No federal AI law — White House AI Action Plan (2025) pro-innovation, new AI policy announcement (2026)
- ✓ Preemption of state laws actively debated in Congress and a White House priority
- ✓ California: 20+ AI laws, ADMT regulations forthcoming
- ✓ Colorado AI Act in flux; Texas Act narrowed to intentional discrimination
- ✓ Employment laws active: NYC, Illinois, California FEHA

South Korea & Others

Principles-Based

- ✓ South Korea AI Basic Act — effective January 22, 2026
- ✓ Mirrors EU framework: transparency + high-risk obligations
- ✓ Applies extraterritorially to developers & providers
- ✓ Japan, Singapore, Australia: light-touch, principles-based
- ✓ China: state-control model leading regulatory enforcement

The Courts

Workday, Inc. — N.D. Cal.

Case: Derek Mobley v. Workday, Inc. — putative class action under Title VII, ADA, and ADEA

Allegation: Workday's AI hiring tool systematically screened out candidates based on race, age & disability

Key Ruling: Federal judge allowed the 'agency' theory to proceed — Workday treated as an agent of the deploying employer

Status: Discovery is ongoing — signals vendor liability extending under federal employment law

Why It Matters: Agency law being expanded to AI-powered software vendors where decision-making authority is delegated. Significant risk for any HR AI tooling

Other Cases To Watch

Privacy: Claims relating to eavesdropping, wiretapping, and invasion of privacy concerning the deployment of AI tools.

Consumer Protection: Claims alleging the use of certain HR tools are a “credit report” under the FCRA. Unfair or deceptive practice claims are on the rise.

Insurance / Healthcare: Claims alleging the use of AI to deny claims and result in discriminatory outcomes.

Copyright: Claims continue to pend against foundation model companies for allegations of copyright infringement.

Coming Battles: Product Liability and Agency? Will AI be treated as a product, like autonomous vehicles, or something entirely different?

The Governance Gap: Industry Data

Sources: CSA State of AI Security & Governance (Dec 2025) · CSA Securing Autonomous AI Agents (Feb 2026)

26%

have comprehensive AI security governance in place

Only 1 in 4 organizations is truly governance-ready

78%

lack policies for AI agent identity creation & management

Most orgs have no playbook when an agent is compromised

18%

of security leaders confident IAM can manage agent identities

Legacy identity systems were built for humans, not agents

46%

of governance-mature orgs adopting agentic AI

vs. only 12% without governance — governance drives adoption

70%

of governance-mature orgs have tested AI for vulnerabilities

Security testing is a governance differentiator

28%

can trace agent actions back to a human sponsor

Accountability gap: 72% cannot fully audit agent behavior

The Imperative for AI Governance

Without a structured AI governance program, **organizations face mounting legal exposure, erosion of stakeholder trust, and uncontrolled operational risk** — as regulators worldwide move from guidance to enforceable mandates.

80+
AI-specific laws & regulations enacted or in-progress globally

$35M
Maximum fine under the EU AI Act for high-risk violations

68%
Enterprises report AI incidents with no formal response plan

Organisational Risk

- Unvetted AI can embed bias, cause decisions to harm customers, and trigger costly remediation
- Reputational damage from AI failures is swift and hard to reverse
- Internal shadow AI adoption bypasses security and compliance controls

Shifting Legal Landscape

- EU AI Act (2024): tiered risk classification with mandatory conformity assessments
- US Executive Order on AI directs sector-specific safety & security standards
- GDPR, CCPA & sector regulators (SEC, FDA, FFIEC) now scrutinise AI-driven decisions

Governance as Enabler

- Structured governance accelerates responsible AI deployment rather than slowing it
- Demonstrates due diligence to regulators, partners, and boards
- Creates repeatable controls that scale across the enterprise AI portfolio

AI Governance Frameworks: A Multi-Standard Approach

A mature AI governance program does not align to a single standard. Mapping controls across complementary frameworks provides defense-in-depth, satisfies multiple regulators, and demonstrates enterprise-grade rigour.

NIST AI RMF

NIST AI Risk Management Framework

NIST | USA

The US federal baseline for AI risk. Built around GOVERN, MAP, MEASURE, and MANAGE functions — enabling organisations to identify, assess, and respond to AI risks across the full lifecycle.

- Risk categorisation & prioritisation
- Trustworthiness attributes
- Continuous monitoring & improvement
- Organisational accountability

ISO/IEC 42001

AI Management System Standard

ISO | Global

The first certifiable international AI management system standard. Follows a Plan-Do-Check-Act cycle modelled on ISO 27001 for establishing and improving AI governance.

- AI policy & objectives
- Risk & opportunity treatment
- Responsible AI by design
- Supply chain oversight

CSA AI Controls

CSA AI Controls Matrix

Cloud Security Alliance | Global

A cloud-native controls catalogue for AI workloads. Maps to CSA CCM, NIST AI RMF, and ISO 42001, providing technical and procedural controls across the AI deployment stack.

- Data governance & lineage
- Model security & integrity
- AI supply chain controls
- AI incident response

AIUC-1

AI Use Case Controls Framework

AI Governance Initiative | Emerging

Use-case-centric controls mapping governance requirements to specific AI scenarios — generative AI, automated decisions, and autonomous agents.

- Use-case risk classification
- Human oversight requirements
- Output validation controls
- Accountability chain docs

AI Governance Structure: Principles, Policies & Processes

01

Principles

Foundational ethical guidelines and core values

- Fairness & non-discrimination across all AI outputs
- Transparency: explainability and disclosure obligations
- Human oversight maintained at meaningful decision points
- Privacy-by-design embedded in AI system architecture
- Accountability: clear ownership across the AI value chain

02

Policies

Formal rules aligning AI practices to principles

- AI Acceptable Use Policy covering internal & vendor AI
- AI system inventory and classification policy (risk tiers)
- Data governance policy for AI training and inference data
- Agentic AI deployment policy: approval gates, scoping
- Incident response policy specific to AI system failures

03

Processes

Workflows operationalizing principles and policies

- AI risk assessment: threat modeling + RAIA impact review
- Model cards review and vendor diligence for procured AI
- Red-team & penetration testing — continuous, not one-time
- Human-in-the-loop checkpoints for high-risk AI decisions
- Audit trail maintenance and periodic compliance review

Alignment: EU AI Act · ISO/IEC 42001 · NIST AI Risk Management Framework · AIUC-1

AI Governance in Practice: Global HR AI Tool Deployment

Scenario: A global enterprise operating in the EU, UK, US, South Korea, India and beyond is deploying an AI-powered HR tool that automatically scans and ranks candidate résumés. **This is a high-risk AI use case requiring structured pre-deployment governance across five critical workstreams.**

01 Vendor Diligence

- AI system card & model docs
- Third-party audit reports
- Data provenance & lineage
- Subprocessor mapping
- Security certifications (SOC 2, ISO 27001)

02 Bias & Fairness Assessment

- Disparate impact testing by gender, age, ethnicity
- Review of training data sources
- Independent bias audit requirement
- Ongoing model drift monitoring plan
- Define acceptable error rate thresholds

03 Jurisdictional Legal Review

- EU AI Act: high-risk classification obligations
- GDPR / UK GDPR: automated decision rights
- US EEO / EEOC guidance on AI hiring
- South Korea PIPA & AI disclosure rules
- India DPDP Act requirements

04 Notices & Consents

- Candidate-facing AI use disclosure
- Right to human review notice (EU/UK)
- Opt-out mechanisms where required
- Employee works council consultation (EU)
- Privacy notices updated in each jurisdiction

05 Controlled Rollout & Monitoring

- Pilot in lower-risk jurisdictions first
- Human-in-the-loop review of AI shortlists
- Incident reporting & escalation path
- Quarterly bias re-assessment cycle
- Regulator engagement strategy

EU EU / UK

- EU AI Act: mandatory conformity assessment for high-risk HR AI
- GDPR Art. 22: right not to be subject to automated decisions
- Works council consultation required before deployment

US United States

- EEOC guidance: AI hiring tools must avoid disparate impact
- NYC Local Law 144: annual bias audits + public disclosure
- State-level AI hiring laws expanding (IL, CA, CO)

KR South Korea

- PIPA: consent required for automated processing of personal data
- AI disclosure obligation when AI makes hiring recommendations
- PIPC enforcement increasingly active on AI systems

IN India

- DPDP Act 2023: lawful basis required for processing sensitive data
- Notice to data principals about AI use in hiring decisions
- Ongoing secondary legislation expected — monitor closely

AI Contract Playbook: Negotiation by Vendor Type

Not all AI vendors carry the same risk profile. Contract terms must be tailored to how the vendor trains models, what data permissions are granted, and the degree of customer control. This playbook maps the key negotiation positions by vendor type.

	Foundation Model API *(e.g. OpenAI, Anthropic, Google)*	Embedded AI SaaS *(e.g. Workday AI, Salesforce AI)*	Custom-Built / Fine-Tuned *(Vendor trains on your data)*	Open Source Model *(e.g. LLaMA, Mistral, deployed internally)*
Data Training Rights	Prohibit use of your inputs to train or fine-tune shared models. Require contractual confirmation of opt-out from training.	Confirm vendor's AI features are not trained on your tenant data. Seek data isolation assurances and audit rights.	Strictly define scope of training data permitted. Require data deletion post-contract. Retain IP in outputs derived from your data.	You own the deployment — no vendor training risk, but ensure base model licence prohibits harmful/restricted uses.
Permissions Granted	Licence to use API output for internal purposes only. Restrict sublicensing. Confirm no model weight access.	Limited SaaS use licence. Confirm scope of AI-driven automation permitted. Watch for broad data use rights buried in ToS.	Negotiate IP ownership of fine-tuned model weights. Define permitted deployment environments explicitly.	Review licence carefully (Apache 2.0, LLAMA Community licence, etc.). Commercial use restrictions vary significantly.
Key Negotiation Points	• Confidentiality of prompts/outputs • Hallucination liability allocation • Output accuracy SLAs • Model change notice periods • EU AI Act compliance commitments	• Data processing agreement (DPA) with SCCs • Automated decision-making controls • Bias audit access rights • Audit & penetration testing rights • Termination data return/deletion	• Full bias testing before deployment • Model explainability obligations • Ongoing retraining approval rights • Escrow of model weights • Indemnity for discriminatory outputs	• Compliance responsibility sits with deployer • Require internal AI use policy • Document model card & risk assessment • Ensure GDPR lawful basis for any processing • Liability remains entirely in-house
Risk	MEDIUM	MEDIUM–HIGH	HIGH	HIGH (deployer-owned)

Smart Governance Playbook: Buying & Selling AI

BUYING AI — Diligence & Governance

Black-Box Risk: Most AI model providers lack full transparency into training data, model architecture, and intended outputs. Review model cards thoroughly.

Security Assessment: Require SOC 2, ISO 42001, or AIUC-1 certification from AI vendors. Assess for agent-specific security controls and red-team testing evidence.

Agentic AI Scrutiny: Agentic architectures present unique challenges around security, accountability, and explainability. Require architecture documentation and tool-use scope definitions.

Contractual Controls: Define data ownership, output ownership, training opt-out rights, incident notification obligations, and human-in-the-loop requirements in vendor agreements.

Ongoing Monitoring: Establish post-deployment monitoring, periodic bias audits, and model update review processes. Governance doesn't end at procurement.

SELLING AI — Contracting & Risk Allocation

SLA vs. Accuracy: Separate uptime SLAs from output quality guarantees. Do not warrant accuracy rates. AI outputs are probabilistic — disclaim accuracy representations carefully.

Consequential Harms: Exclude business interruption, regulatory fines, and third-party claims from customer's use of AI output. Workday is the warning: avoid appearing to own the decision.

Data Ownership: Define customer data, outputs, telemetry, and aggregate data precisely. Offer no-training-on-customer-data as enterprise default. Permit opt-in for improvement.

Indemnity Triggers: Narrow indemnity scope — exclude customer prompts, inputs, and third-party tool combinations. Include cure rights and control rights before liability attaches.

ISO 42001 Certification: Consider ISO 42001 or AIUC-1 certification around key products. Shortens enterprise procurement cycles and demonstrates compliance posture.

Key Takeaways

Regulatory landscape is converging — act now

EU AI Act takes full effect August 2026. California ADMT regulations imminent. A cross-jurisdictional governance program aligned to ISO 42001 and NIST AI RMF is achievable and future-proofs compliance across all major jurisdictions.

Workday is the warning shot — govern your AI vendors

Agency law is expanding to AI-powered software vendors. Know what decisions your AI vendors are making, require bias audits, document human override controls, and narrow your indemnity obligations in every AI contract.

Agentic AI introduces an entirely new security surface

Legacy IAM, SOC 2, and SIEM frameworks were built for static systems. MCP, A2A, and A2C protocols introduce new protocol-level attack vectors. Apply MAESTRO threat modeling and build agent identity governance before deploying at scale.

Governance drives adoption — it is a competitive advantage

Governance-mature organizations adopt agentic AI at 4x the rate of unprepared ones. ISO 42001 certification shortens sales cycles. Build a governance program that converts principles into audit-ready evidence and enables innovation responsibly.

Speaker Bio

- Partner at Foley & Lardner LLP
- Peter helps organizations around the world turn the complex legal and regulatory challenges of AI, data, and emerging technologies into a competitive advantage rather than a roadblock.
- Peter's practice covers the entire lifecycle of technology-driven legal challenges. He drafts and negotiates complex technology agreements. He architects privacy and cybersecurity policies from the ground up. He designs AI governance programs that balance innovation with accountability. He conducts risk assessments and tabletop exercises that prepare organizations for real-world scenarios. And when regulatory inquiries, litigation, or cyber incidents strike, Peter is the steady hand guiding clients through the storm.

About Foley

Foley is an *Am Law* 50 law firm consistently ranked among top-tier practices. We provide an unmatched level of client service, innovation, and value — all tailored to meet your specific needs.

- Providing comprehensive legal services in more than 60 practice areas
- Garnering global recognition for providing exceptional client service and value
- Delivering industry thought leadership at the forefront of business trends and key legal developments and regulations
- Investing in understanding your business, markets, and goals
- Offering alternative fee arrangements and budgets to provide cost efficiencies and certainties

1,100
Attorneys

27
Offices

180+
Years of History

What We Do

We comprehensively and effectively address legal and business matters across four primary sectors and a broad range of industries that align with our clients and strengths.

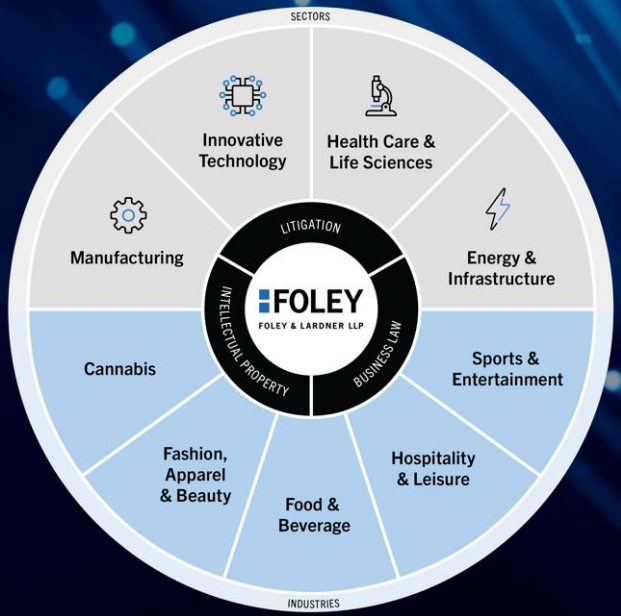

60+
practices, including:

- Antitrust
- Business Litigation and Dispute Resolution
- Capital Markets and Public Company Advisory
- Commercial Transactions and Business Counseling
- Consumer Law, Finance, and Class Action
- Corporate Governance
- Cybersecurity and Data Privacy
- Environmental
- Employee Benefits and Executive Compensation
- Environmental, Social, and Governance
- Export Controls and National Security
- Government Enforcement Defense and Investigations
- Investment Management and Fund Formation
- IP Litigation
- IP Procurement, Management, and Counseling
- Labor and Employment
- Mergers and Acquisitions
- Private Equity and Venture Capital
- Real Estate
- Securities and Corporate Finance
- Taxation
- Trademark, Copyright, and Advertising Counseling