



IAPP PRIVACY. SECURITY. RISK. + AI GOVERNANCE GLOBAL 2026


Training 6-7 October

Workshops 7 October

Conference 8-9 October

AI GP exam 9 October

SEATTLE



#IAPPpsr26

RECONCILING AI AND CYBERSECURITY GOVERNANCE

Legal and Operational Controls

IAPP P.S.R. + AI Governance Global 2026 • Seattle • Thursday, 8 October, 14:15 PDT



#IAPPpsr26

WELCOME AND INTRODUCTIONS



Demian Ahn
Partner
Wilson Sonsini Goodrich & Rosati
Wilson Sonsini Goodrich & Rosati (Moderator)



Sam Kaplan, AIGP, CIPP/US
Regulatory Strategy and Public Law
Lead
Reflection AI



Charles Wonpu, CIPP/E
Sr. Director, Global Head of
Product & Privacy Legal, Zscaler



Justin Okun
Senior Counsel
Cyber, Privacy, Fraud & LE Ops,
DoorDash



#IAPPpsr26

AI Adoption is Spreading at Historic Speed

Closed Model Capabilities & Delivery

"Closed" = weights not publicly released; accessed as a service, not distributed.

Frontier closed models (Anthropic's Claude Opus, OpenAI's GPT, and Google's Gemini) still lead on the hardest reasoning tasks, but the open-weight gap has narrowed to single digits. (Stanford HAI, 2026 AI Index (Technical Performance))

Most often accessed via API: vendor keeps the weights, infrastructure, and safeguards (some also offered via managed cloud or private/VPC deployment).

Logging, rate limiting, and patching are managed by the provider, but your data leaves your boundary.

Adoption & Governance

88% Enterprise Use (AI Index citing McKinsey State of AI.)

79% of organizations regularly use generative AI in at least one business function (AI Index citing McKinsey State of AI.)

55% increase in AI incidents. 362 in 2025, up from 233 in 2024 (AI Incident Database.)

68% of breached organizations this year lacked AI governance to manage AI or detect shadow AI, up from 63% last year. (IBM's Cost of a Data Breach Report 2026)

19% of organizations reported coordinating between their governance and security teams (IBM's Cost of a Data Breach Report 2026)

Open-Weight Models Capabilities & Delivery

"Open-weight" ≠ "open-source": the parameters are published, but the training data and pipeline typically stay undisclosed.

Open-weight families (Llama (Meta), Qwen (Alibaba), DeepSeek, Mistral) publish their parameters for download and self-hosting.

The gap to the closed frontier has narrowed to roughly four months and low-single-digit percentages on most benchmarks.

Enables self-hosting, air-gapping, fine-tuning, and provider independence, but expands governance obligations

EMERGING RISKS

The same capability driving adoption is now weaponized by attackers, operating autonomously, and turning the models themselves into targets.

AI Weaponized

22,000+ AI-related fraud complaints in 2025 (voice-cloned BEC, deepfake investment pitches, and synthetic personas). (FBI IC3, 2025 Annual Report)

State-sponsored actors leveraging commercial models across the attack lifecycle. (Google GTIG, Nov. 2025)

LLM-embedded malware calls AI APIs at runtime to rewrite its own code and generate commands, evading signature-based detection (Google Threat Intelligence Group, AI Threat Tracker, Nov. 2025)

Autonomous & Frontier-model risk

Anthropic disrupted the first reported large-scale AI-orchestrated cyber-espionage campaign with humans only "on the loop." (Anthropic Threat Intelligence)

Frontier models now demonstrate advanced offensive-cyber capability: finding zero-days and building exploit chains. (OpenAI GPT-6 Astra; Anthropic Claude Mythos).

Mid-2026 several frontier developers disclosed containment failures during cybersecurity evaluations (GPT-5.6 Sol and an unreleased successor, Claude/Mythos, Muse Spark 1.1, Kimi K3)

AI Systems as Targets & Internal Misuse

"Shadow AI" Triggers First SEC Form 8-K for Unauthorized AI Use

Shadow-AI security incidents more than doubled, to 43% — staff routing sensitive data through unapproved tools. (IBM 2026)

AI-related breaches rose to 21% of organizations, up from 13%; model inversion (\$6.07M) and prompt injection (\$5.89M) were the costliest incident types. (IBM, Cost of a Data Breach Report 2026)

Root causes are usually structural: compromised APIs, applications, and cloud misconfigurations. (IBM 2026)

RECENT POLICY REACTIONS

Federal Executive Posture

June 2026 EO: voluntary pre-release review of "covered frontier models"; up to 30 days' government access; no mandatory licensing

Silent on export controls: leaving covered models as export-control candidates

Directs a Treasury/NSA/CISA cyber clearinghouse and AG focus on AI-facilitated crime

CISA BOD 26-04: new risk-based patch model; 3-day remediation for highest-risk flaws, citing AI-accelerated exploitation

Legislative & Enforcement

26-AG coalition letter (Sept. 23, 2026): urges Congress to enact mandatory frontier-AI safety regulation while expressly opposing federal preemption of state AI laws (State AGs to Speaker Johnson, Leaders Thune, Jeffries, Schumer, Sept. 23, 2026)

FRONTIER Act (H.R. 9925): would mandate risk frameworks, independent audits, 24-hour incident reporting (introduced, stalled)

15-state AG letter: asserted consumer-protection/privacy violations; Alabama DTPA subpoena

1,100+ frontier-AI employees backed international coordination to slow development

Sector & Infrastructure

SEC FY2026 exam priorities target firms' AI supervision/governance, accuracy of AI representations and AI-amplified cyber threats. (SEC 2026 Examination Priorities)

FINRA 2026 report adds GenAI: vendor due diligence, agentic-AI guardrails, tool testing

New York RAISE Act puts frontier-AI oversight in NYDFS (72h incident reporting); May 2026 advisory extends Part 500 to AI

CISA/Five Eyes agentic-AI guidance; Secure Integration of AI in OT (OT/ICS in scope)

LEGAL RISKS

Exposure is twofold: a fast-arriving wave of AI-specific regulation and the full weight of existing law — which already reaches AI without any new statute.

AI-Specific Regulation

EU AI Act: risk-based duties on high-risk providers/deployers; GPAI transparency (phasing 2026–27)

US Frontier Developers: CA TFAIA & NY RAISE & IL Safety Measures Act (transparency, incident reporting, catastrophic-risk assessment; IL adds independent audits)

US Deployer & Anti-discrimination: CO ADMT (SB 26-189, eff. Jan. 1, 2027) & TX TRAIGA

US Employment AI: IL HB 3773 & NYC Local Law 144

EU Product & Cyber: CRA & NIS2 (security-by-design + cyber baseline)



Existing Law Reaches AI

Sector: GLBA Safeguards Rule & HIPAA Security Rule

Federal Enforcement & Disclosure: FTC Act §5 & SEC cyber disclosure (8-K 1.05 / Reg S-K 106)

State Cyber & Privacy: NYDFS Part 500 & state privacy acts & CPPA ADMT/audit regs

Cross-cutting: breach-notification, IP, contract, consumer-protection, anti-discrimination

Biometrics: Illinois BIPA, Texas CUBI, Washington reach AI facial-recognition and voiceprint processing

#IAPPpsr26

INTEGRATING AI INTO PRODUCTS

Building AI into a product changes your legal status — you may become a manufacturer, a deployer, and the party that inherits the model's risk. The controls, contracts, and disclosures don't transfer; they land on you.

Responsible Party

You can't shift liability to "the algorithm" or a licensed third-party tool (FTC §5)

You're the controller/business: transparency, consumer-rights, and breach-notification duties (and the clock) sit with you (GDPR; CCPA/CPRA)

Sector duties stay with the regulated entity (HIPAA, GLBA); placing an AI product on the EU market makes you a CRA manufacturer

Model Delivery & Choice

Closed / API: data leaves your boundary; the vendor's contract, logging, and safeguards govern, and you inherit its outages and term changes

Open-weight, self-hosted: strongest data control but the heaviest governance lift; provenance and backdoor risk travels with the weights and can't be recalled

Build vs. buy is now a security and compliance decision, not just an engineering one

You inherit vendor & transitive AI

Third-party components may **embed** models you never chose; map the full model chain, not just your direct vendor

Force the **questions:** is our data training the model? who are the sub-processors, and can we audit and enforce? where does data physically go?

Flow security, notification, no-training, and audit-rights duties **down** by contract; match notification clocks to the shortest applicable window

RISKS FOR AI USERS & DEPLOYERS

Most organizations meet AI risk by using or deploying it, not building it — the exposure spans security, data, compliance, and operations, and routinely falls between existing owners.

Security Exposure

Model-layer attacks classic SOC tooling was never built to see: prompt injection, data/model extraction, and poisoning

Agentic and tool-use risk: over-permissioned agents and compromised connectors/MCP servers turning a helper into an attack path

Shadow AI is now a leading breach vector; IBM found shadow-AI incidents more than doubled, to 43%, at higher cost (IBM 2026)

Data & Compliance

Sensitive data crosses the boundary through prompts, retrieval, fine-tuning, and integrations, often into vendor systems you don't control

Foreign-hosted or open-weight models raise data-jurisdiction and export-scrutiny questions (cf. Kimi K3)

AI incidents trigger the same breach-notification clocks as any other, and machine-speed intrusions compress the time to assess and notify

Operational & Legal


Hallucination, drift, and error drive bad decisions and client harm, and you own the output

Contractual and product-liability exposure for AI failures and defective outputs, increasingly non-waivable (EU PLD)

The ownership gap is the through-line: AI and cyber teams each assume the other owns detection, response, and the reporting clock, so no one does

A Natural Extension of Your Security Stack

Know Your Network - You Can't Defend What You Can't See



#IAPPpsr26

CONTRACTUAL CONTROLS

Where you can't engineer the risk away, you allocate it by contract — with model providers and AI-enabled vendors — and the terms have to match AI's distinct risks, not boilerplate SaaS language.

Data & Privacy Terms

No training or fine-tuning without express, revocable permission (no "improving the services" carve-out)

Full model-chain transparency: sub-processors and upstream/embedded models

Data location, retention, encryption; deletion that reaches model artifacts

Government-request handling, cross-border transfer, and data-residency terms

Security & Incident Terms

Notification matched to your shortest clock (e.g., 36h banking / 72h NIS2/state)

Security-testing, audit rights, and vulnerability handling, including model-layer issues

Published safety/usage protocols as binding commitments; notice of material model changes

Provenance, integrity, and no-backdoor reps for open-weight artifacts

Liability & Continuity


Indemnity for IP (training-data), data, and security failures; confirm what caps apply

Caps calibrated to AI-specific harm; some exposure can't be contracted away (EU PLD)

Service-continuity and exit/transition rights on provider disruption

Warranties on performance, availability, and compliance; fallback and portability

THE INTEGRATED GOVERNANCE MODEL



How Did Things Go? (We Really Want To Know)

Did you enjoy this session? Is there any way we could make it better? Let us know by filling out a speaker evaluation.

1. Open the IAPP Events app.
2. Select **IAPP P.S.R. + AIGG 2026**.
3. Tap "Schedule" on the bottom navigation bar.
4. Find this session. Click "Rate this Session" within the description.
5. Once you've answered all three questions, tap "Done".

Thank you!

#IAPPpsr26

