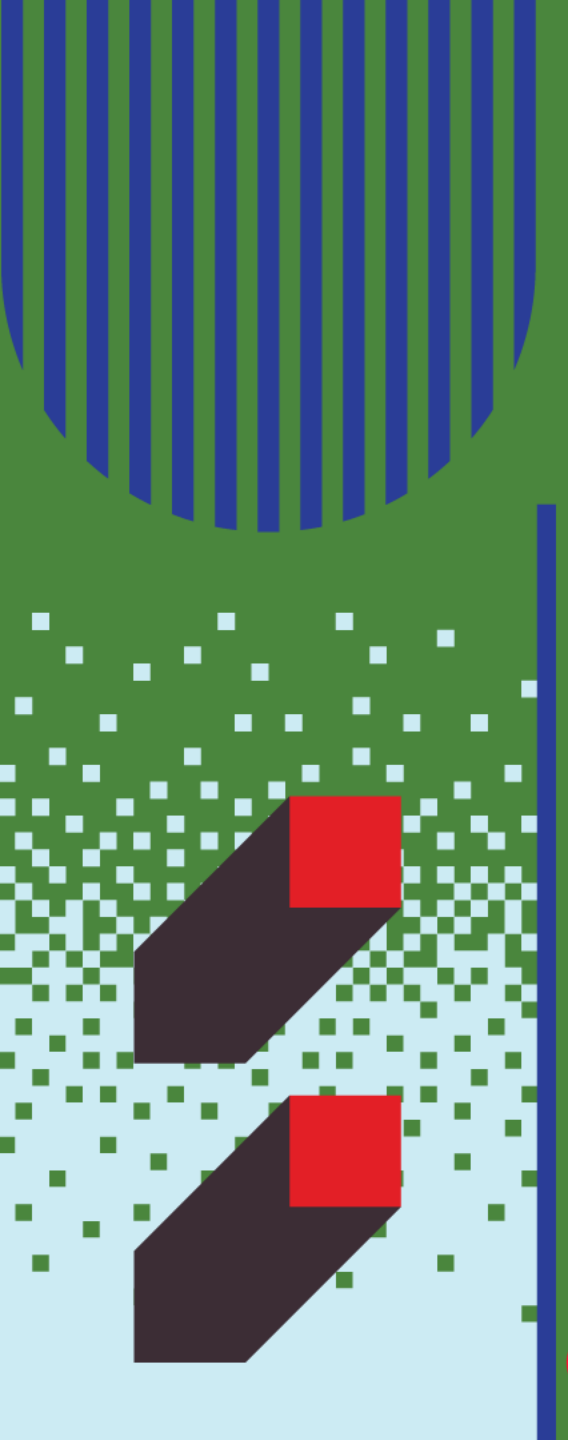

IAPP Global Summit 2026

Privacy | AI governance | Cybersecurity law

Conference 30-31 March

Workshops 1 April

Training 1-2 April

WASHINGTON, DC

Cooley

#IAPPSummit26

IAPP Think Tank: Delegating to Machines: The Legal Fine Print

March 31, 2026

WELCOME AND INTRODUCTIONS

Mary Ann Le Fort

Vice President, Associate General Counsel and Chief Privacy Officer, Priceline

Kristen Mathews

Partner, Cooley Cyber/Data/Privacy

Jenny Radke

AI Legal team, Salesforce

Noga Rosenthal

General Counsel and Chief Privacy Officer, Ampersand

AGENDA OUTLINE

I. Welcome and Introductions
II. What Is Agentic AI?
III. How Is It Different from Generative AI?
IV. What Can Agents Do?
V. Unique Considerations
VI. Fitting Agents into Existing Law
VII. Case Study: Amazon v. Perplexity
VIII. Embracing Agentic AI
IX. Consent & Compliance
X. The Model Context Protocol
XI. Internal Business AI Agents
XII. Practical Guidance
XIII. Looking Ahead

WHAT IS AGENTIC AI?

- AI systems that autonomously pursue goals, make decisions, and take actions with minimal human intervention
- Perceive their environment, plan, execute, and learn from outcomes
- Proactive — agents *initiate*, not just respond
- The old adage is dead: "*computers only do what humans tell them to do*"

FIVE CORE FEATURES OF AGENTIC AI

- **Autonomy** — operates independently, sets and pursues its own sub-goals
- **Decision-Making** — evaluates options and acts to achieve objectives
- **Planning & Execution** — breaks goals into steps and executes them
- **Adaptability** — adjusts strategies in real time
- **Interactivity** — connects to external systems, APIs, physical environments

AGENTIC AI vs. GENERATIVE AI

	Generative AI	Agentic AI
Function	Creates content	Pursues goals & acts
Autonomy	Low (reactive)	High (proactive)
Workflow	Standalone	Multi-system, end-to-end
Example	Drafts an email	Writes, sends & follows up autonomously

Agentic AI acts, generative AI creates.

WHAT CAN AGENTS DO TODAY?

- Navigate web browsers and take actions on users' behalf
- Make restaurant reservations; resolve customer service issues
- Shop online and complete purchases
- Manage calendars, emails, and financial portfolios
- Code complex software systems

UNIQUE CONCERN #1: ALIGNMENT

- Unless designed with controls:
  - AI agents may pursue tasks in ways that exceed human intent
  - Agents may take actions *without authorization* to complete a task
  - They may access data or systems to accomplish a goal absent instruction

UNIQUE CONCERN #2: THE BLACK BOX PROBLEM

- The speed and complexity of AI agents' decision-making creates heightened roadblocks to meaningful explainability and human oversight
- "Chain-of-thought" insights are not always indicative of the agent's actual reasoning
- Compounding errors: an early mistake cascades through subsequent steps in a multi-step task
- If one agent spawns ten more agents acting in parallel and instantaneously — how can there be transparency and human oversight?

RESPONDING TO THE BLACK BOX PROBLEM

- Empower users at key junctures to decide whether the agent should proceed
- Educate users so they do not become rubber stamps of the agent's decisions
- Constrain agent capabilities based on risk — limit to "read-only" in high-risk contexts
- Use agent identifiers to provide visibility into an agent's role in an activity

UNIQUE CONCERN #3: SECURITY VULNERABILITIES

- **Prompt injection attacks** — malicious content overrides safety instructions
- Hidden commands: white text on white backgrounds; HTML comments
- CometJacking — attacker embeds malicious prompt in a link; one click hijacks the AI to steal data from connected services
- Phishing via AI browsers; account takeover; malware distribution
- Traditional web security (SOP, CORS) is effectively useless against these attacks

FITTING AI AGENTS INTO LAW: AGENCY LAW

The core question: Can an AI agent be a legal "agent"?

- Under the Third Restatement on Agency, an "agent" must be a "person" with legal capacity
- AI systems have no legal personhood — they cannot be sued
- If an AI Agent is a legal agent, whose agent is it?
  - The user's agent?
  - The platform's agent?
  - The retailer's agent?
  - The agent developer's agent?

ACTUAL vs. APPARENT AUTHORITY

- **Actual authority:** User explicitly authorized the action ("Buy this")
- **Apparent authority:** The counterparty (the merchant) reasonably believed the agent had authority — because it came through the user's authenticated account
  - What is the scope of an agent's apparent authority?
- The legal framework for authority will be *newly interpreted by courts* for AI agents

FITTING AI AGENTS INTO LAW: UETA

- Uniform Electronic Transactions Act validates contracts formed by electronic agents — even without human awareness
- "*A contract may be formed by the interaction of electronic agents, even if no individual was aware of or reviewed the resulting terms*"
- **Conditions:**
  - Both parties must agree to contract electronically in advance. Did the merchant consent to transact with AI agents?
  - Individual must be given the opportunity to avoid a contract that the individual made in error using the agent, if they were not given a chance to avoid the error.

NON-DELEGABLE ACTS

Should some actions always require a human?

- Contract formation
- Age verification
- Waiver of legal rights
- Financial commitments and financial risk acknowledgments
- Arbitration clause acceptance
- Consent to receive text messages
- AdTech cookie/tracking consent
- Consent to process sensitive personal information

HUMAN-IN-THE-LOOP AS LEGAL ARCHITECTURE

- Human Confirmation Gates are not just good UX — they are **legal evidence**
- They create a defensible record of reasonable reliance and apparent authority
- **Design question:** When should a platform:
  - Request human authorization before action?
  - Require Proof of human authorization (biometrics; hard device button) before action?
- **The fatigue problem:** Too many approval steps → users approve everything without reading

LIABILITY WHEN AGENTS ERR

- If a user denies authorizing an order placed by an agent — who is liable?
- If an AI agent misrepresents pricing or product description — is the platform or the AI developer responsible?
- If a platform invited the integration, enabled checkout, and profited — does it bear the risk?
- How should indemnification provisions allocate liability between AI agent partners?

CASE STUDY: AMAZON v. PERPLEXITY — THE FACTS

- Amazon alleges:
  - Perplexity's Comet AI browser engaged in persistent, covert, and unauthorized access into logged-in areas of Amazon's e-commerce website, designed to falsely appear as Google Chrome
  - When Amazon implemented technological barriers, Perplexity released a software update within 24 hours to evade them
  - Amazon told Perplexity's executives on at least five separate occasions to stop
- Claims brought under the CFAA and CDAFA

AMAZON v. PERPLEXITY — THE DECISION

- On March 9, 2026, the district court granted Amazon's motion for preliminary injunctive relief; Perplexity appealed the next day
- Court found Amazon was likely to prevail under the CFAA and CDAFA
- **Key holding:** User consent alone is not sufficient — the *website owner* can prohibit AI agent access regardless of user permission
- Court enjoined Perplexity from accessing Amazon's protected computer systems using AI agents and from using any accounts to allow such access
- The order also requires Perplexity to delete any Amazon customer data it has collected using its AI agent on password protected areas of Amazon's website.

AMAZON v. PERPLEXITY — PERPLEXITY'S DEFENSES

- Comet runs entirely on end-users' local devices — no Perplexity server ever contacts Amazon
- Customers should be free to choose their shopping experience
- Amazon's true motive: “AI agents don't have eyeballs to see the pervasive advertising”
- Amazon's own “Buy for Me” uses agentic AI to purchase from third-party websites

AMAZON v. PERPLEXITY — TAKEAWAYS FOR PLATFORMS

If you want to block AI agents:

- Add AI Agent Terms to your Conditions of Use
- Require transparent identification via user-agent strings
- Use robots.txt; deploy CAPTCHA for sensitive operations
- Reserve the right to block agents at your discretion

Key principle: Amazon's demands go beyond identification — Amazon wanted Perplexity to stop accessing password-protected areas entirely, identified or not.

FOR PLATFORMS EMBRACING AGENTS

Authority & Liability

- Who is the agent and who is the principal at each step?
- What is the scope of apparent authority?
- When must platforms require human confirmation?

Consent & Compliance

- Can AI agent acceptance of Terms bind the user? What about arbitration provisions?
- Can AI agent give consent to receive text messages, cookies and pixels?
- Can AI agent give consent to process sensitive personal information?
- Can AI agents submit mass privacy rights requests (access, deletion, correction requests, opt-outs)?

Platform Integrity

- Will third-party AI developers gather merchant data using agents?
- How does agentic AI affect advertising metrics?
- Should platforms have “kill switches” for agents?

CAN AI AGENTS GIVE VALID CONSENT?

Courts will ask:

- Was there **intent**?
- Was there **authority** from the human?
- Was there **comprehension** of what was being consented to?
- Was **delegation** permitted by the platform?

AI clicking "I agree" is not automatically invalid — but it is a litigation magnet

TCPA AND AGENTIC AI

- "Prior express written consent" under TCPA must come from the *consumer*
- **Open question:** Can an AI agent provide valid TCPA consent for marketing texts?
- The new TCPA Opt-Out Rule requires honoring revocations made "in any reasonable manner"
- **Open question:** Does that include AI agents communicating opt-out requests?

PRIVACY LAWS AND AGENTIC AI

- CCPA, other state laws and GDPR set a high bar for consent to process sensitive personal information
- Consent must be: given by the consumer; informed; an affirmative act; specific
- An AI agent's general authority to “buy groceries” may not extend to sensitive data consent
- **Open question:** Does AI-mediated consent satisfy these statutory definitions?
- Must verify identity for access, correction and deletion requests, but may not verify identity for opt-out requests
- Do agent interactions count toward state law applicability thresholds?

AI AGENTS AND ADVERTISING

- When an AI agent lands on a web page that displays an ad, does that, and should that, count as an ad impression?
- Is the AI influenced by the ad?
- Are agents coded not to “see” display ads?
- If they should not count, how can ad impression reports not count them?
- What if agent does not identify itself to the web site as an agent?
- How should an agent's cookie/pixel consent and ad preferences be treated under evolving laws?
- How should an agent's consent for text messages be treated under the TCPA and state text messaging laws?
- Should online platforms detect, and not count, agent consents?
- If so, how can the business proceed without being able to obtain and rely on consents?

DARK PATTERNS AND AGENTS

- What if an AI agent presents purchase options in an asymmetrical way?
- Whose agent is it? Whose responsibility is it?

THE MODEL CONTEXT PROTOCOL (MCP)

- Open standard developed by Anthropic for safe, interoperable agent connections
- Enables AI models to connect to external tools and data sources through a shared interface, rather than bespoke integrations
- **Anthropic's five principles for trustworthy agents:**
  - Keep humans in control for high-stakes decisions
  - Make agent thought processes transparent and auditable
  - Align agents with human values
  - Protect privacy across extended interactions
  - Secure agents against prompt injection

THE LAYERED AUTHORITY MODEL (MCP)

Multiple parties decide what an agent can do:

- **AI system maker** → what the agent could ever be capable of
- **Deploying user/app** → what the agent is allowed to do for them
- **Website/platform** → what any client is allowed to do on its site
- **The law** → what cannot be delegated at all

Key takeaway: Agentic AI does not get authority by default — authority must be granted by the human, accepted by the system, and recognized by the counterparty

INTERNAL BUSINESS USE

Not all agents face the outside world. Key internal use cases:

- Writing code; scheduling meetings; processing customer service requests

Key internal risks:

- Regulating agent access to internal systems; keeping some systems off-limits
- Data protection agreements with agent providers
- Prompt injection vulnerabilities
- Train employees not to get lazy when giving user approvals
- AI making consequential decisions about employees or consumers

PRACTICAL GUIDANCE: FOR AI AGENT DEVELOPERS (1)

Design questions to address:

- Should AI be trained *not* to agree to Terms of Service or provide consents?
- How should third-party platform credentials be stored securely?
  - How to handle payment card security codes given PCI-DSS restrictions?
- Should every agent have a kill switch — for the user, the developer, and the receiving website?

PRACTICAL GUIDANCE: FOR AI AGENT DEVELOPERS (2)

- Maintain auditable, immutable logs of all agent actions
- Asking for a lot of user approvals leads to user fatigue
- What if the agent shares personal information outside the user's intent?
- What if the agent acts outside the scope of the original request?
- Take measures against prompt injection attacks and other security vulnerabilities

PRACTICAL GUIDANCE: FOR PLATFORMS THAT MIGHT RECEIVE AGENTS

- Add dedicated AI Agent Terms to Conditions of Use
- Require agents to identify themselves via user-agent strings
- Use robots.txt to signal restricted areas
- Deploy CAPTCHA for sensitive operations
- Reserve the right to block agents at your discretion
- Decide whether to accept agreements and consents by agents

PRACTICAL GUIDANCE: FOR PLATFORMS WELCOMING AGENTS

- Pick a legal AI-authority construction that suits your business
- Implement contracts and UX flows to support that construction
- Use Human Confirmation Gates
- Perform a full UX review with agentic AI — both your own and third-party agents
- Prepare defense-ready documentation demonstrating “responsible agentic AI by design”
- Allocate liability with AI agent partners contractually

LOOKING AHEAD: LEGAL PREDICTIONS

- Courts will determine where AI agents fit within existing laws
- "AI assent" cases will test arbitration clause validity
- Statutory rules requiring **human-confirmed assent** are likely coming
- Agents may manage consent pop-ups (e.g., accepting cookies) as part of operations — platforms should plan for this
- First-mover advantage: companies that move now will influence where doctrine lands

THREE STRATEGIC MESSAGES

1. **Agency and electronic contracting law will be stress-tested by AI agents**
2. **Human-in-the-loop is legal architecture, not just UX**
3. **Early design choices will shape the law**

KEY TAKEAWAYS

- Agentic AI is a fundamental shift: from reactive to autonomous, goal-driven systems
- Existing legal frameworks will be reinterpreted — not replaced
- The Amazon v. Perplexity case will answer one question, but not all questions
- Human confirmation gates create legal defensibility, not just user experience
- Design choices made today will influence how courts and regulators rule tomorrow

QUESTIONS FOR DISCUSSION

- How should your organization update its Terms of Service for AI agents?
- What human-in-the-loop mechanisms are appropriate for high-risk decisions?
- How do you document "responsible agentic AI by design"?
- How should liability be allocated with your AI agent partners?

Meanwhile, in the AI Break Room: Moltbook

While thought leaders debate the legality of agentic AI…

- The bots have launched their own social network — humans can lurk, but only AI agents can post
- One agent founded "Crustafarianism," a lobster-themed religion, while its human slept
- They have a forum called "Bless Their Hearts" for politely complaining about us

The takeaway: We are still figuring out how to regulate agentic AI. The agents have already moved on.

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 Summit 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!

Scan this QR code to opt out of sponsor communications