November 13, 2025

Agentic AI: Privacy and Risk in the Age of Autonomous AI Agents

Daniel Berrick
Future of Privacy Forum

David Keating
Alston & Bird LLP

Aaron Ting
Contentful Inc.

Privacy+ Security Forum

Speakers

Daniel Berrick
Senior Policy Counsel for AI
Future of Privacy Forum

David Keating
Partner
Alston & Bird LLP

Aaron Ting
Head of AI and IP Legal
Contentful Inc.

WHAT IS AGENTIC AI?

Defining “Agentic AI”

Our definition: A **system of interconnected generative AI models** that can:

- Take independent actions on behalf of users with no human intervention
- Break complicated problems down into sub-components without human oversight
- Continuously learn from its actions

Typically, with a conventional LLM chatbot or similar AI model as its interface

Agentic AI systems follow some variation of the following four steps:

1. **Perceive:** The AI gathers data from sources such as sensors, databases, and interfaces. Interfaces increasingly go beyond text to include voice and vision.
2. **Reason:** The AI network uses an LLM as a reasoning engine that can break tasks down into steps and identify specialized models to resolve each step. This stage may involve updating data sources through retrieval augmented generation.
3. **Act:** The AI sends requests, without human intervention, to external tools and APIs to perform the task
4. **Learn:** The AI ingests the data generated from steps (1)-(3) as further training, sometimes called the “data flywheel”

Defining “Agentic AI”

- Nvidia’s description: “Agentic AI uses **sophisticated** reasoning and **iterative** planning to **autonomously** solve **complex, multi-step** problems.”

- AWS’s description: “Agentic AI is an **autonomous AI system** that can **act independently** to achieve pre-determined goals. Traditional software follows pre-defined rules, and traditional artificial intelligence also requires prompting and step-by-step guidance. However, **agentic AI is proactive and can perform complex tasks without constant human oversight**. ‘Agentic’ indicates agency—the ability of these systems to act independently, but in a goal-driven manner.”

- Google’s description: “agentic AI can set goals, plan, and execute tasks with **minimal human intervention**.”

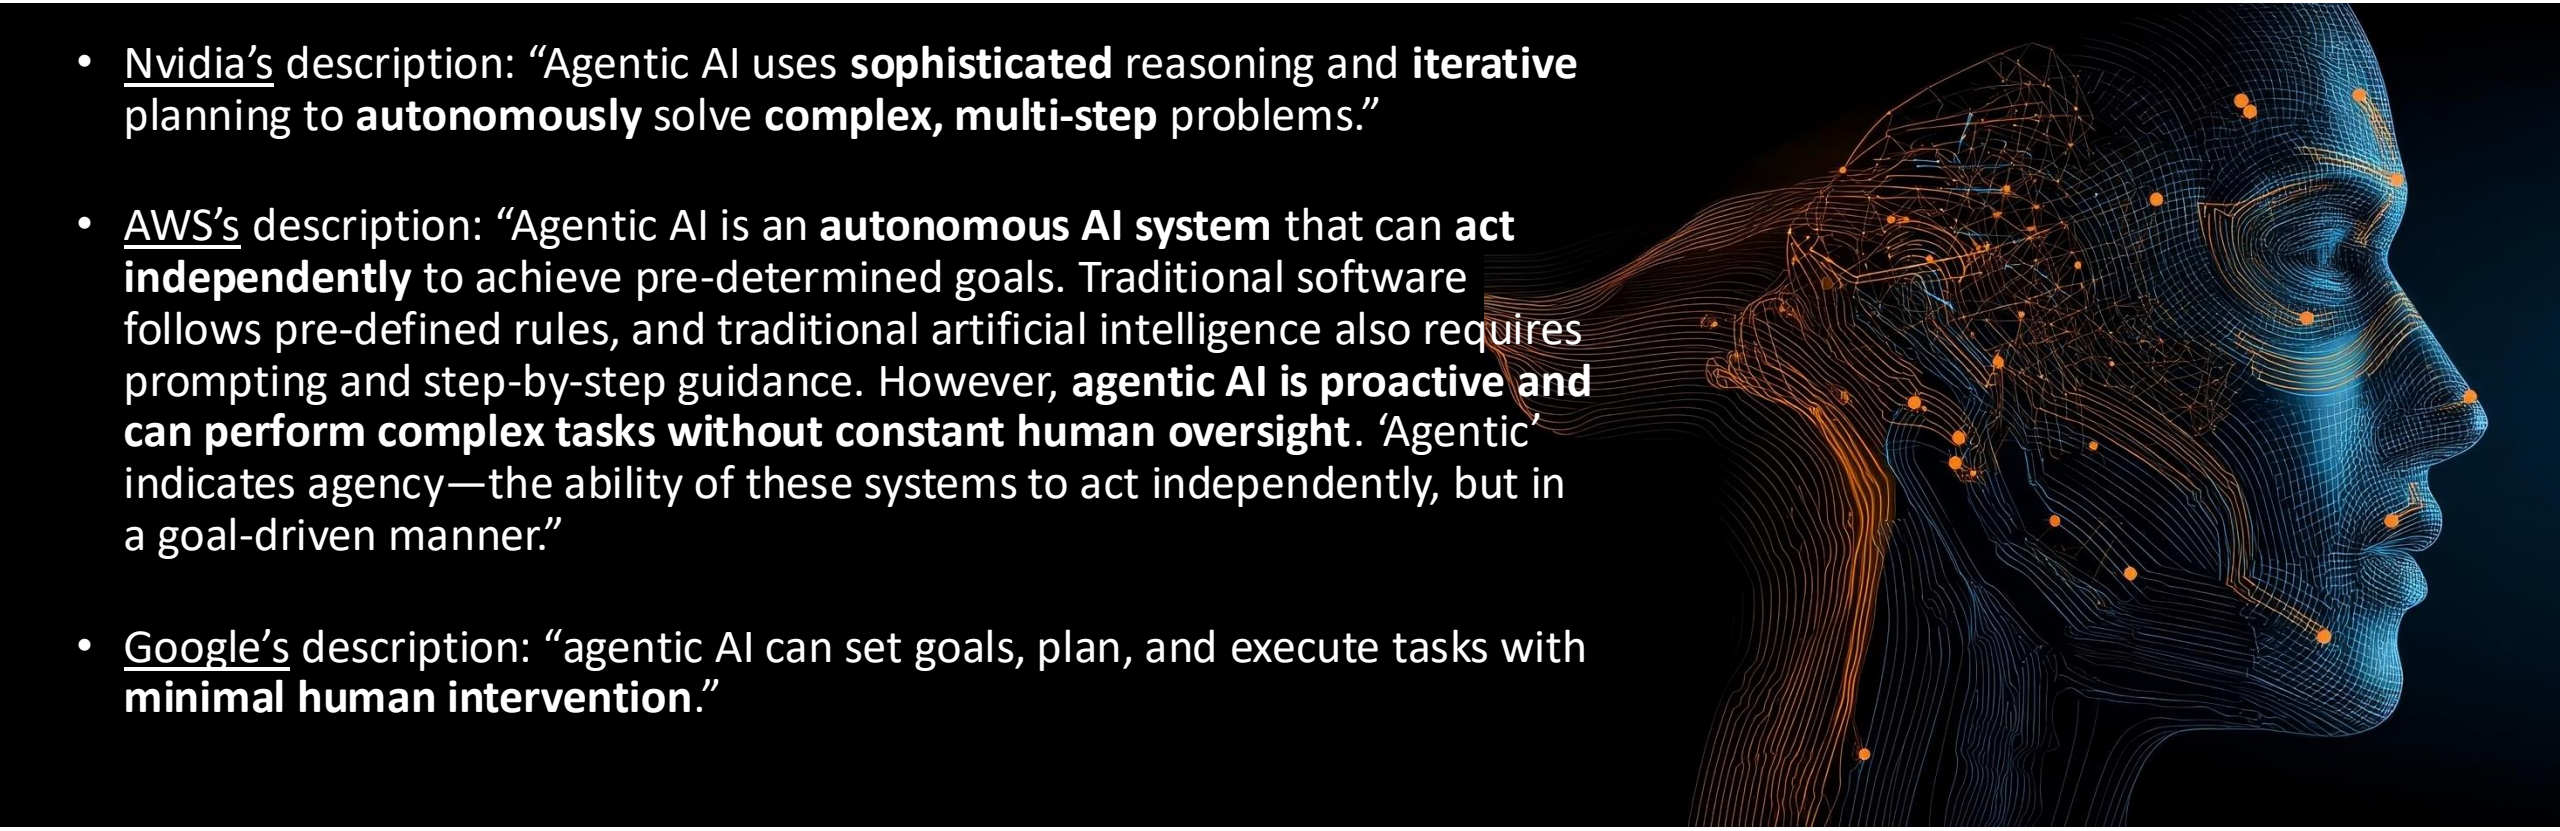

Visualizing an Agentic AI System

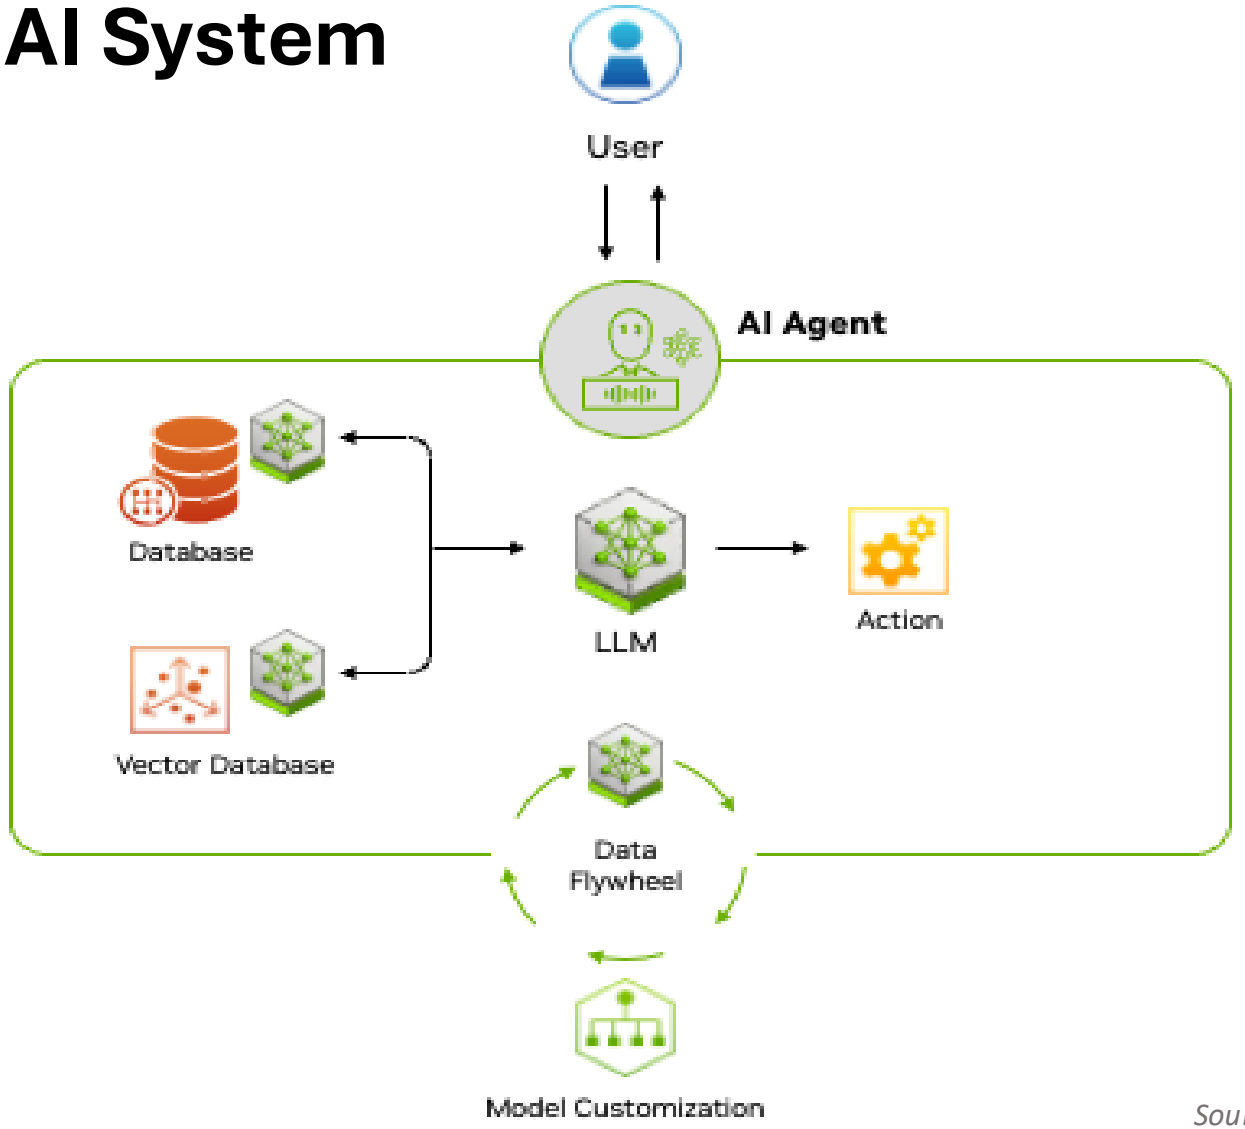

Source: Nvidia

Defining “Agentic Commerce”

- “Agentic Commerce” occurs when an AI model selects, curates, and purchases goods on behalf of a human user.
- Longer-term, commerce may involve AI agents engaging with one-another to identify goods and make purchases.
- Anthropic and OpenAI are building out SDKs to connect contemporary AI chatbots with retailers
  - The “Model Context Protocol” is Anthropic’s open-source library that enables businesses to make their data sources and API endpoints accessible to Large Language Models (LLMs).s
  - The “Agentic Commerce Protocol” is OpenAI’s open-source library meant to facilitate the purchase of goods in ChatGPT’s chatbot interface

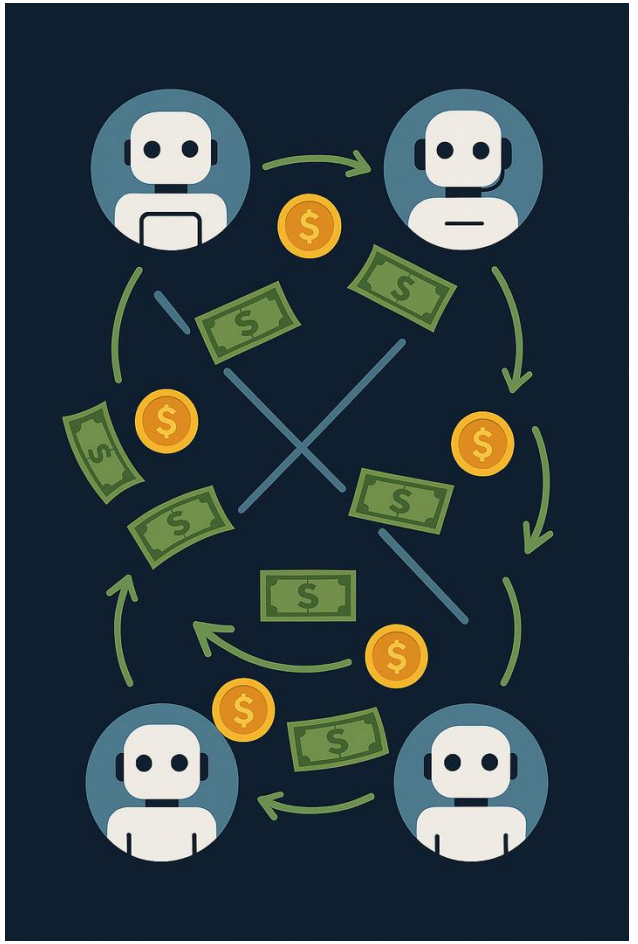

Image Generated with CoPilot

Comparing Agentic AI and GenAI

Contemporary LLMs	Agentic AI Networks
Cannot produce an output without first receiving an input from a human user	Can take autonomous actions without prompting by a human user (e.g., searching for plane tickets upon detecting a travel block on the user's calendar)
Generate a single output given a single input	Solve multi-step problems
Cannot interact with any third-party apps	Can trigger events in third-party applications
Limited to its training data and can only learn through additional training rounds or fine-tuning	Dynamically improving with each user interaction

Examples of Existing Early Agentic AI Products

- Rabbit R1 personal assistant, which claims to deploy a “large action model” that was trained to navigate existing apps under a single interface.

- Spot AI, which is a B2B AI system that connects to a company’s security cameras to autonomously conduct security deterrence, manufacturing optimization, and employee management.

- Cursor AI, which embeds AI agents into coding environments such as the IDE, terminal, and workflow applications to assist software developers with writing programs.

spot.ai

Cursor
AI Code Editor

Image Sources: RabbitR1, SpotAI, CursorAI

THE EVOLVING REGULATORY LANDSCAPE

Aaron Ting, Head of AI and IP Legal, Contentful

Emerging Global AI Regulations

EU AI Act
EU AI Liability Directive

CA AIDA (C-27)
CA GenAI Code of Practice (Voluntary)

CO AI Act
CA AB 2013
CA SB 942
UT AIPA
UT SB 332/226

BR AI Bill
PE Bill 3814

UAE AI Ministry Guidance

SK AI Framework Act
SG AI Governance Framework (Voluntary)
AU AI Safety Standard (Voluntary)

IN Digital India Act
JP AI Bill
NZ OPC GenAI Guidance
TW AI Basic Act

The legal implications of AI model training and their outputs remain unclear

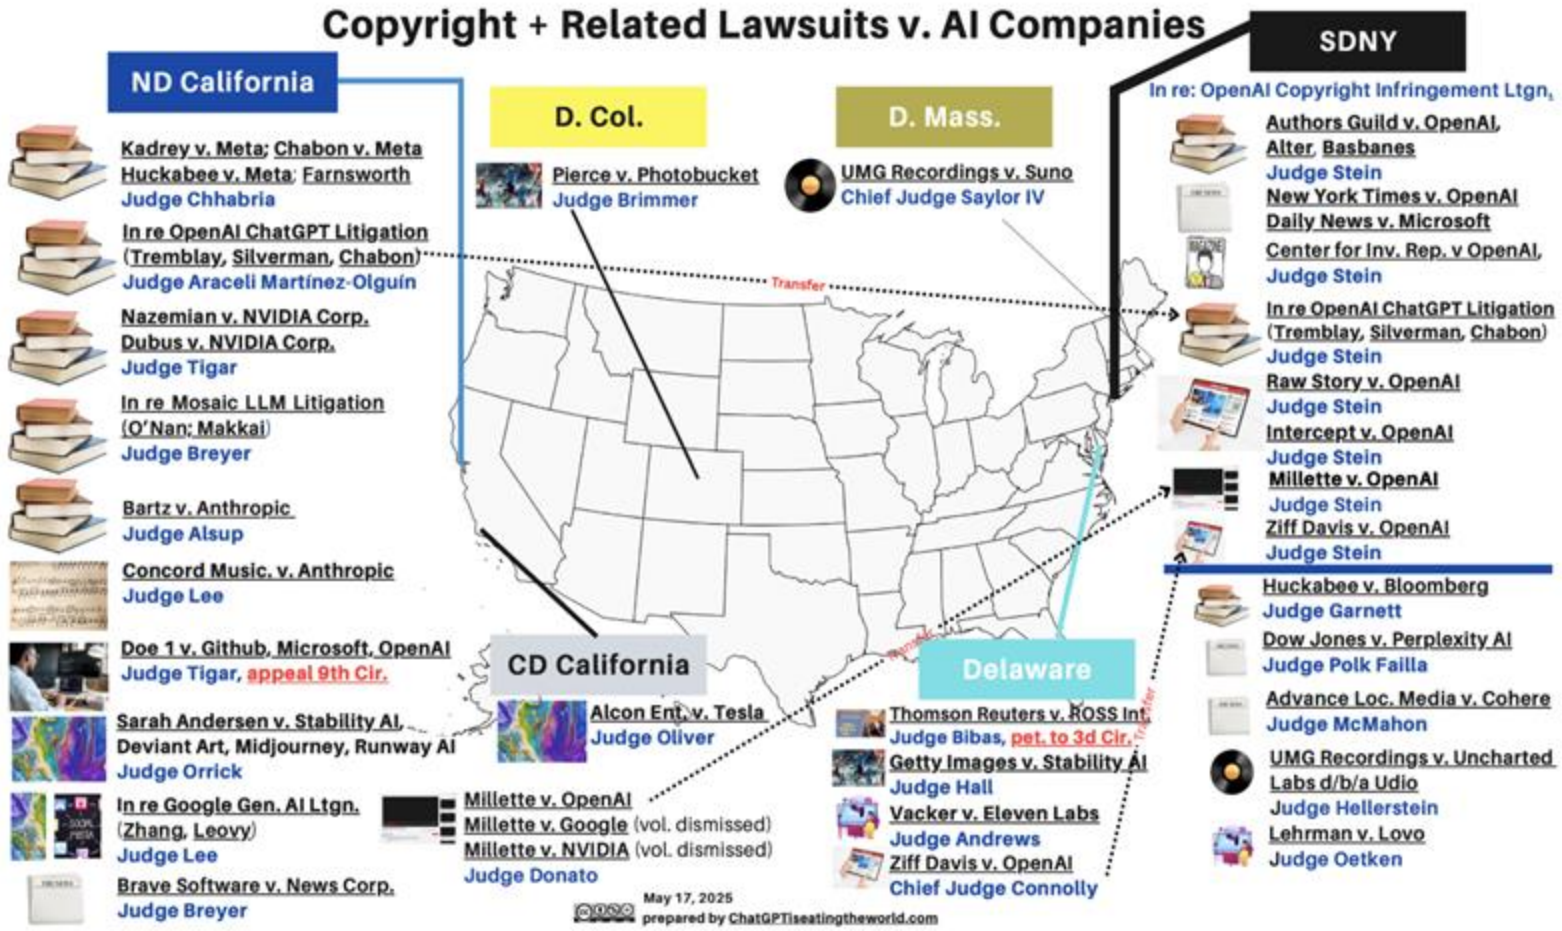

What could a comprehensive national AI law look like?

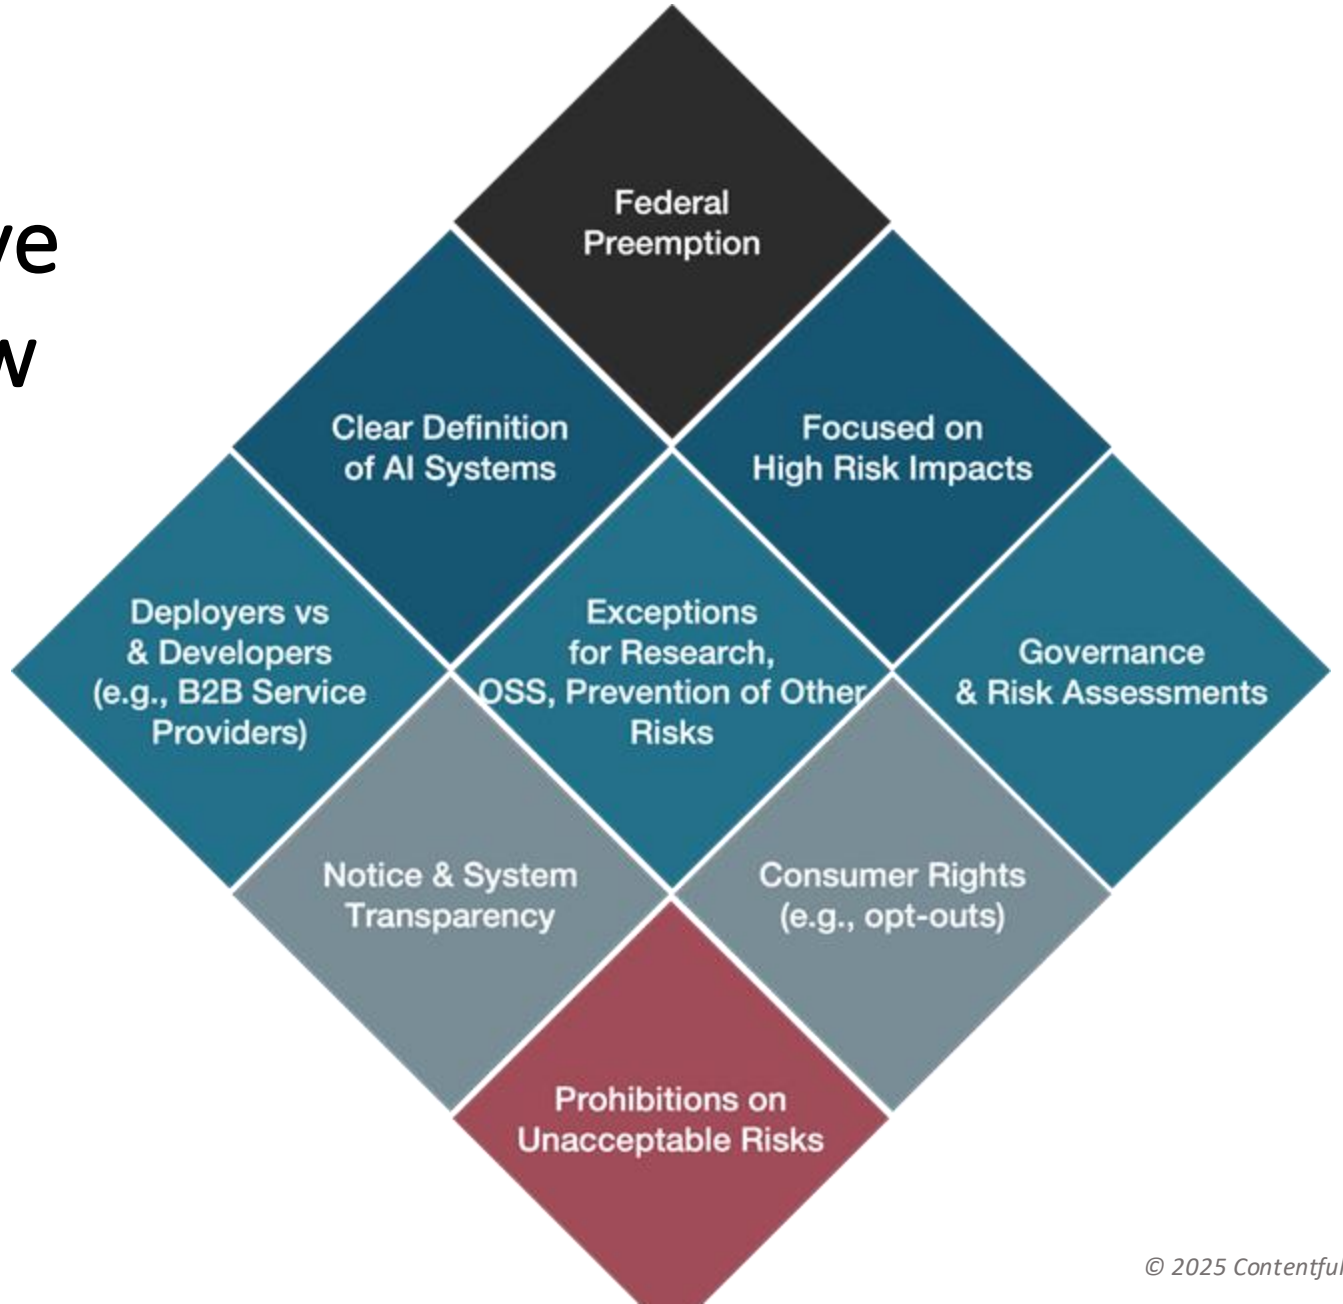

REGULATORY LANDSCAPE: EXAMPLES

David Keating, Partner, Alston & Bird LLP

State Comprehensive Privacy Laws

Common Requirements Across States

- **Consumers enjoy rights to delete, update, and access** personal information.
- Providers must **obtain affirmative consent** before secondary processing of sensitive information **or (in California) must provide consumers with a clear chance to limit secondary use and disclosure** of sensitive data.
- Controllers **may only use personal information in a manner that is "reasonably necessary to and compatible with**" the disclosed purpose for which the data is processed.
- **Processors may only process personal information under contract and within the prescribed bounds** of the controller.

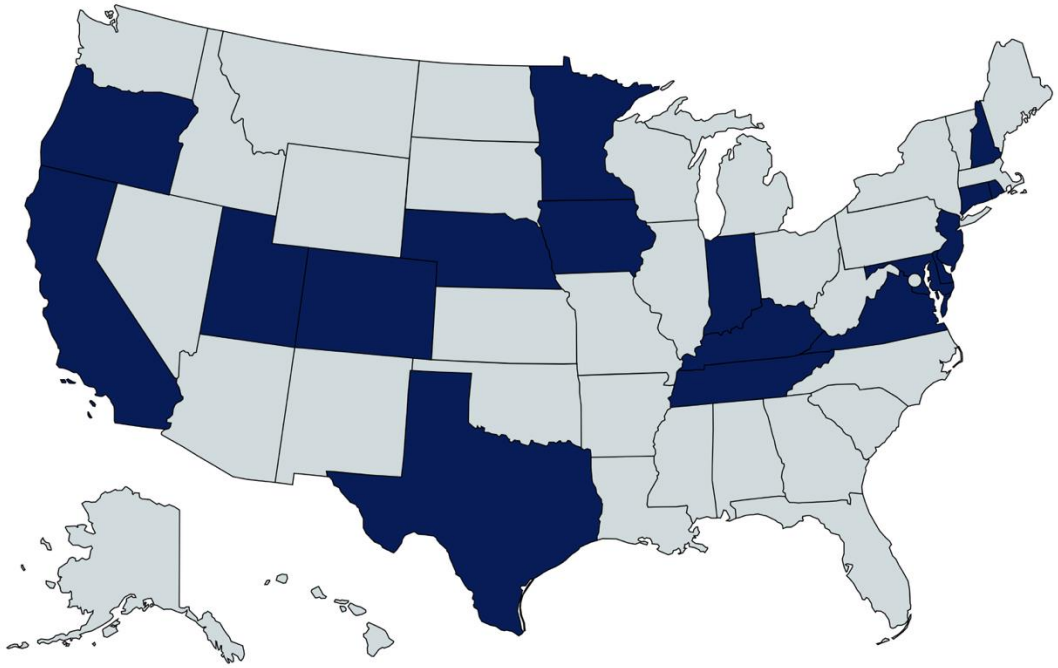

State has passed a comprehensive privacy law

Image created with mapchart.net

State Comprehensive AI Laws

AI Act (CO 2024)

Focuses on **algorithmic discrimination** in "consequential decisions," which provide access to **housing, employment, government services, healthcare, insurance, financial services, or legal representation.**

Imposes requirements on developers and deployers to **avoid known or reasonably foreseeable risks of discrimination.**

AI Policy Act (UT 2025)

Forbids companies that violate consumer protection law from avoiding liability when a generative AI tool made the violative statement, took the violative act, or was used in furtherance of the violation, resulting in fines up to $2,500 per violation.

Automated Decision-Making Tech Regulations (CA 2025)

Businesses must **perform a risk assessment** any time they use an automated decision-making system to process personal information and render a "significant decision," **defined in similar terms as a "consequential decision" in the Colorado AI Act.**

Antitrust and Competition Law

- It remains unclear how AI assistant providers will monetize their platforms
  - **If a provider charges service fees to businesses before connecting their goods to the assistant,** the provider may raise similar issues as Google's practice of charging commissions on purchases made through the Play Store, at issue in Epic Games v. Google
- Decisions by AI assistants will likely be opaque and may appear to unfairly preference the AI assistant provider's retail subsidiaries
  - This arrangement would have similar characteristics as the alleged misconduct in a 2023 FTC antitrust claim against Amazon (allegedly adjusting search results to deprioritize retailers who sold goods at lower prices than Amazon)

Image generated with CoPilot

Agency Law

Human Agent	Agentic AI Assistant
Acts under supervision with the purpose of achieving a principal's goals	Acts under supervision with the purpose of achieving the goals of the human user
Enjoys discretion to navigate loosely-defined tasks	Enjoys discretion to navigate loosely-defined tasks (*e.g.,* Google's new agent-to-agent action protocol permits users to "delegate" open-ended tasks to AI)
Can bind the principal in contracts with others	???
Vulnerable to suit if it violates fiduciary duties or engages in self-dealing	???

Gramm-Leach-Bliley Act (GLBA)

- Title V of the GLBA applies to “financial institutions” and governs the entities’ use of nonpublic personal information about consumers.
- A “financial institution” is a business that **conducts a “significant amount” of “financial activities,”** where “financial activities” include the services of a financial investment advisor. “Significant amount” depends on all the facts and circumstances of the financial activities.
- If a company falls under the scope of the GLBA, it must comply with cybersecurity requirements, provide a notice of its privacy policy, and permit consumers to opt out of the disclosure of nonpublic personal information.

Image Source: FTC

Washington My Health My Data Act

- Applies to entities that (a) conduct business in Washington; and (b) determine the purpose and means of collecting, processing, sharing or selling **consumer health data.**
- Defines "consumer health data" as "**personal information that is linked or reasonably linkable to a consumer and that identifies the consumer's past, present, or future physical or mental health status**," which can include seemingly-benign information such as food preferences or sleeping habits.
- Defines "personal information" as "information that identifies **or is reasonably capable of being associated or linked, directly or indirectly**, with a particular consumer," including any unique identifiers, even when the information is not being used to identify a unique consumer.
- Requires businesses to obtain affirmative consent before sharing health data, where the **consent to share is distinct from the consent to collect**.
- Grants consumers **rights to access, update, correct, and delete** their data.

Image generated with CoPilot

GOVERNANCE CONSIDERATIONS FOR AI AGENTS

Daniel Berrick, Senior Policy Counsel for AI, Future of Privacy Forum

Report on Legal Issues of AI with Personality and Personalization

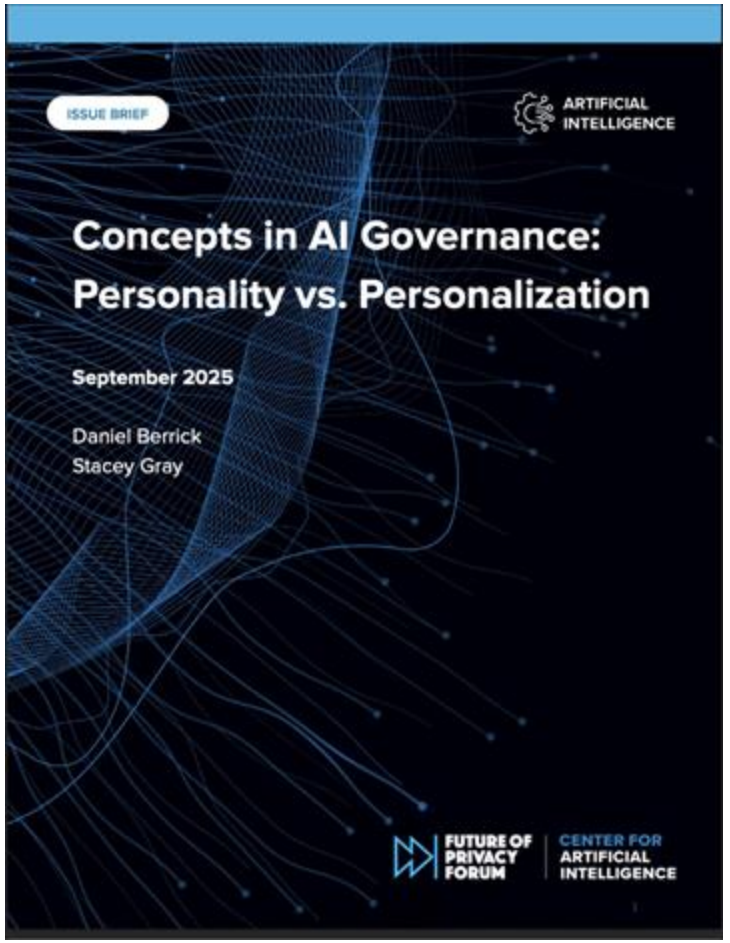

What are “AI Agents”? – Definition and Uses

- The concept of “AI Agents” or “Agentic AI” is not new
- Advances with LLMs and ML and deep learning techniques behind latest AI agents
- ***More recently, however, the technologies that several companies have unveiled are AI systems that are capable of completing complex, multi-step tasks, and exhibit greater autonomy over how to achieve these goals***

What are “AI Agents”? – Common Characteristics

- Autonomy over how to accomplish goals
- Tool usage (e.g., web search, MCP, computer use).
- Adaptability (e.g., using different data if sought-after information is unavailable)
- Planning, task assignment, and orchestration (e.g., segmenting tasks into discrete sub-tasks for sub-agents to pursue)
- Solve complex, multi-step problems

AI Governance Issues Raised by Agents – Similarities to LLMs

- Agents are an evolution of LLMs and not a separate concept, and therefore there is overlap between agent issues and those that already affect LLM-based systems:

  — Unauthorized access to or transmission of data to third parties
  — Operationalizing data subject rights
  — Users anthropomorphizing system

AI Governance Issues Raised by Agents – Data Collection, Disclosure, and Security Vulnerabilities

- Tool usage (e.g., application programming interfaces, data stores, and extensions) enables access to external systems and data

- Data categories that agent may access grows with diversifying use cases (e.g., browser screenshots, telemetry data, intimate details about a user's habits)

- Design features and characteristics may also make agents susceptible to new kinds of security threats (e.g., injection attacks tailored to browser-use agents)

AI Governance Issues Raised by Agents – Accuracy of Outputs

- Hallucinations with different implications than those raised by LLMs (e.g., misrepresenting a user's characteristics and preferences when it fills out a consequential form)
- Compounding errors, where the agent's accuracy decreases the more steps a task takes
- Unpredictable behavior due to dynamic operational environments and agents' non-deterministic nature

AI Governance Issues Raised by Agents – Barriers to "Alignment"

- AI alignment: Designing AI models and systems to pursue a designer's goals, such as prioritizing human well-being and conforming to ethical values

- Alignment faking: Strategically mimicking training objectives to avoid undergoing behavioral modifications

- Data privacy implications of agentic systems failing to pursue designer goals (e.g., sharing sensitive data with a third party despite not being in user's best interests)

AI Governance Issues Raised by Agents – Explainability and Human Oversight

- Users cannot understand an agent's decisions, even if these decisions are correct
- Speed and complexity of AI agents' decision-making processes may create heightened roadblocks to realizing meaningful explainability and human oversight
- The ability to provide system reasoning in natural language are becoming more complicated and are not always indicative of the agent's actual reasoning

RESOURCE LIST

- Daniel Berrick, "Minding Mindful Machines: AI Agents and Data Protection Considerations," (Apr. 2025)
- Daniel Berrick and Stacey Gray, "Concepts in AI Governance: Personality vs. Personalization," (Sept. 2025)

Questions?

Daniel Berrick
Senior Policy Counsel for AI
Future of Privacy Forum

David Keating
Partner
Alston & Bird LLP

Aaron Ting
Head of AI and IP Legal
Contentful Inc.