iapp

Agentic AI:

Navigating the tension between privacy and the next generation of AI

Wednesday, 23 July
08:00–09:00 PDT
11:00–12:00 EDT
17:00–18:00 CEST

Welcome and introductions

Ojas Rege, CIPP/E, CIPM

SVP, Privacy & Data Governance
OneTrust

Agenda

- Introduction to agentic AI
  - Architecture, guardrails, regulatory implications
  - Myths to bust, blind spots to expose
- Putting it into practice
  - Governing traditional AI vs. agentic AI
  - Learnings, best practices, pitfalls to avoid

Blurring boundaries …

Each is fundamentally a data problem, with different but related goals and constraints

AI is an amplifier – it amplifies the risk and impact of existing privacy and data governance gaps in a company

Responsible use is the principle that sustains the long-term business value of emerging technologies

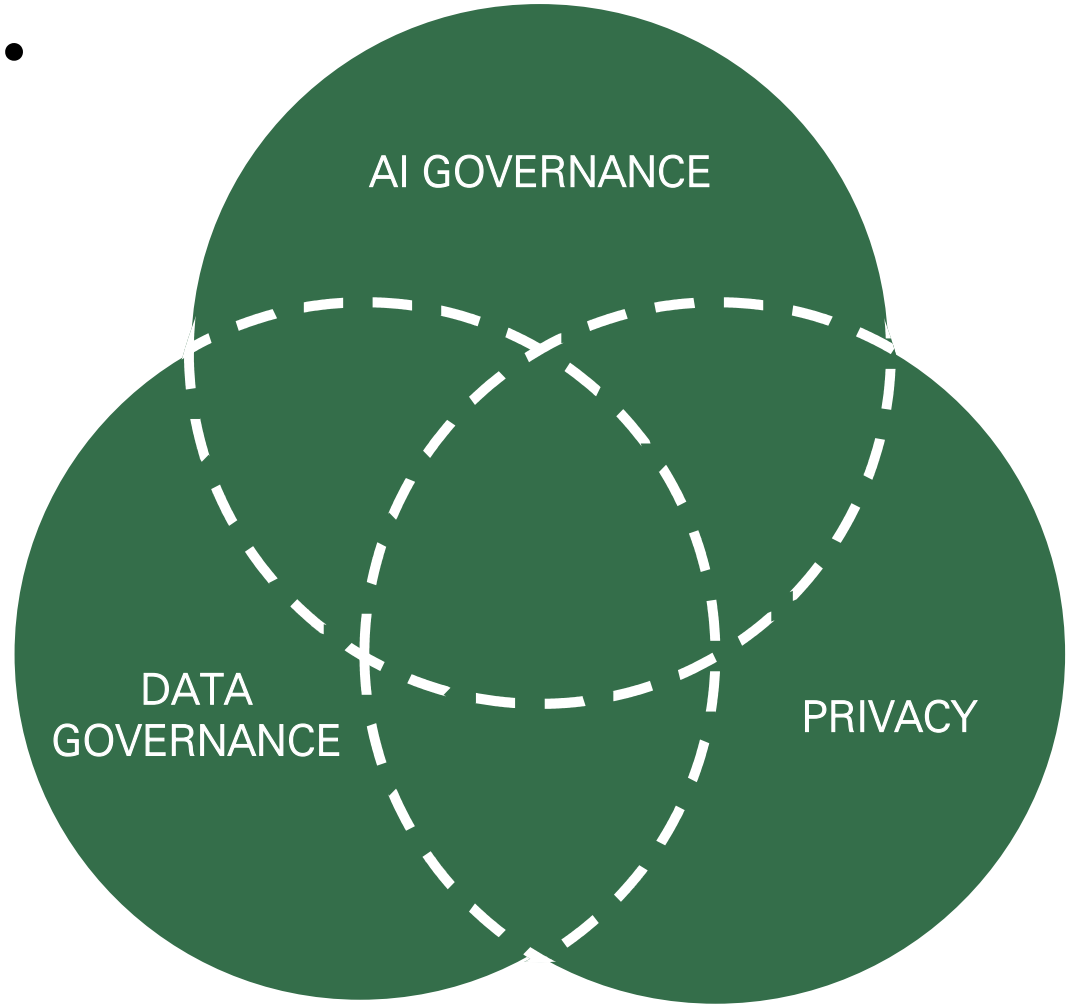

Risk teams

Security, Privacy, Governance, Risk, Ethics, Compliance

Under pressure to avoid trust-breaking events, unintended consequences, and enforcement actions

Shared goal

How do we enable the responsible use of data across all risk domains

…at the speed and volume demanded by the business?

Product & GTM teams

Marketing, Sales, Product, Data, BI, Emerging Tech / AI

Under pressure to accelerate AI, data, and other technology initiatives to stay ahead of competitors

Now, let's go down the AI rabbit hole!

Image: https://www.claudialamoreaux.com/rabbit-hole-this-way/

The world is going agentic …

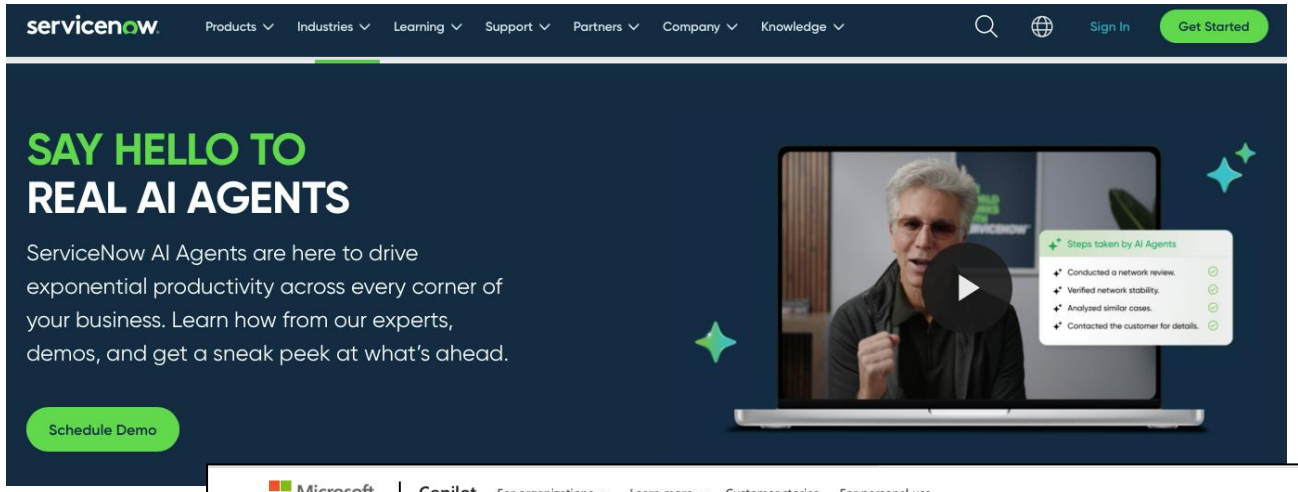

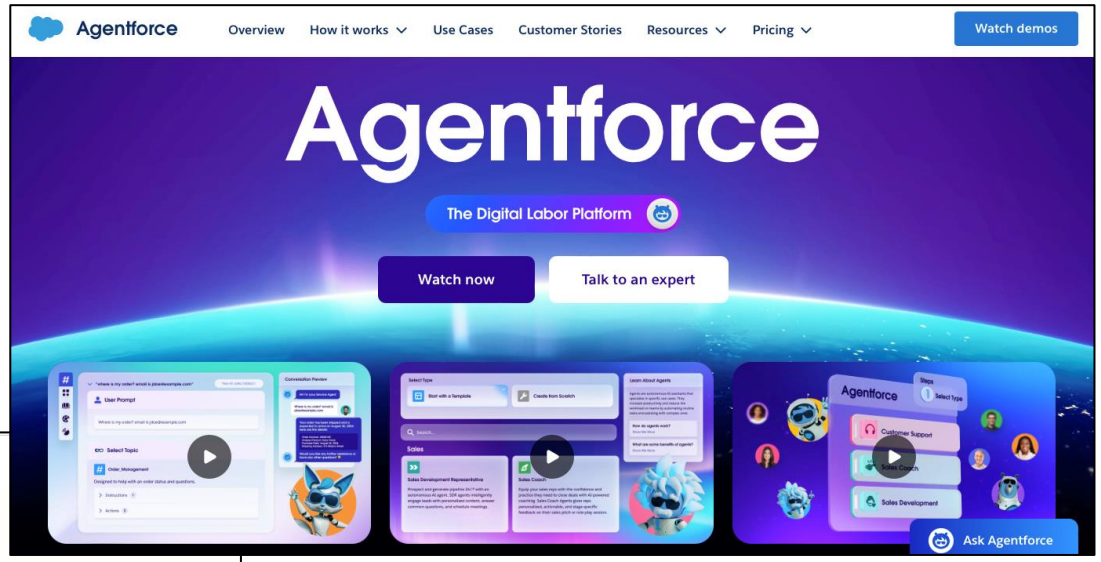

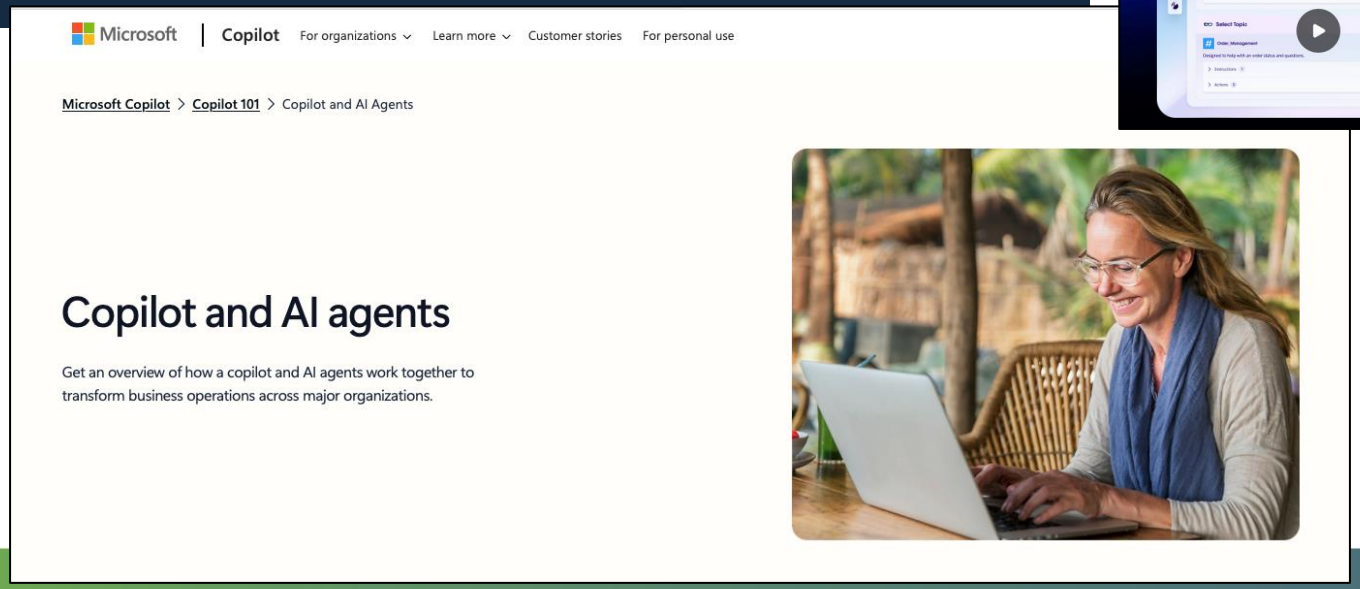

https://www.onetrust.com/news/onetrust-announces-its-first-data-privacy-agent/

March 24, 2025: AI Security Agents announcement (OneTrust / Microsoft)

Privacy at the Speed and Scale of AI: OneTrust Announces its First Data Privacy Agent

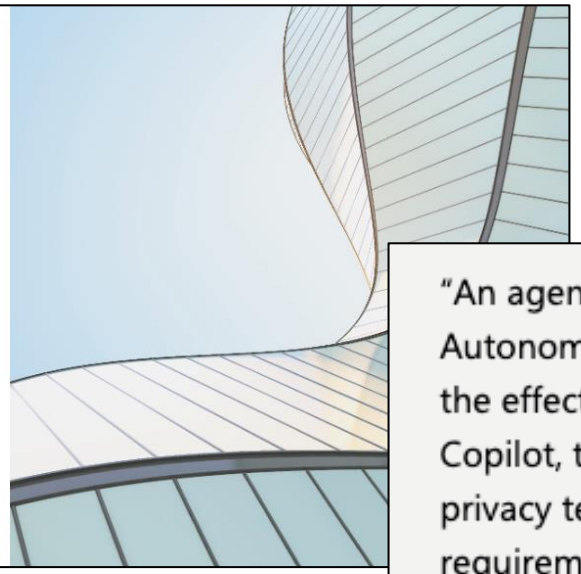

Built with Microsoft Security Copilot, OneTrust's new AI agent demonstrates the power of an agentic approach to privacy

March 24, 2025

> "An agentic approach to privacy will be game-changing for the industry. Autonomous AI agents will help our customers scale, augment, and increase the effectiveness of their privacy operations. Built using Microsoft Security Copilot, the OneTrust Privacy Breach Response Agent demonstrates how privacy teams can analyze and meet increasingly complex regulatory requirements in a fraction of the time required historically."
>
> —Blake Brannon, Chief Product and Strategy Officer, OneTrust

Microsoft | Microsoft Security Solutions Products More

Events AI and machine learning Microsoft Security Copilot

7 min read

Microsoft unveils Microsoft Security Copilot agents and new protections for AI

By Vasu Jakkal, Corporate Vice President, Microsoft Security

https://www.microsoft.com/en-us/security/blog/2025/03/24/microsoft-unveils-microsoft-security-copilot-agents-and-new-protections-for-ai/

What are agents?

Skills

+

APIs

+

Reasoning

Autonomous, adaptive, non-linear

An autonomous, adaptive, non-linear thinker

Managing agentic risk

CONTROLS

HUMAN IN THE MIDDLE

SAFEGUARD AGENTS

The stakes are increasing …

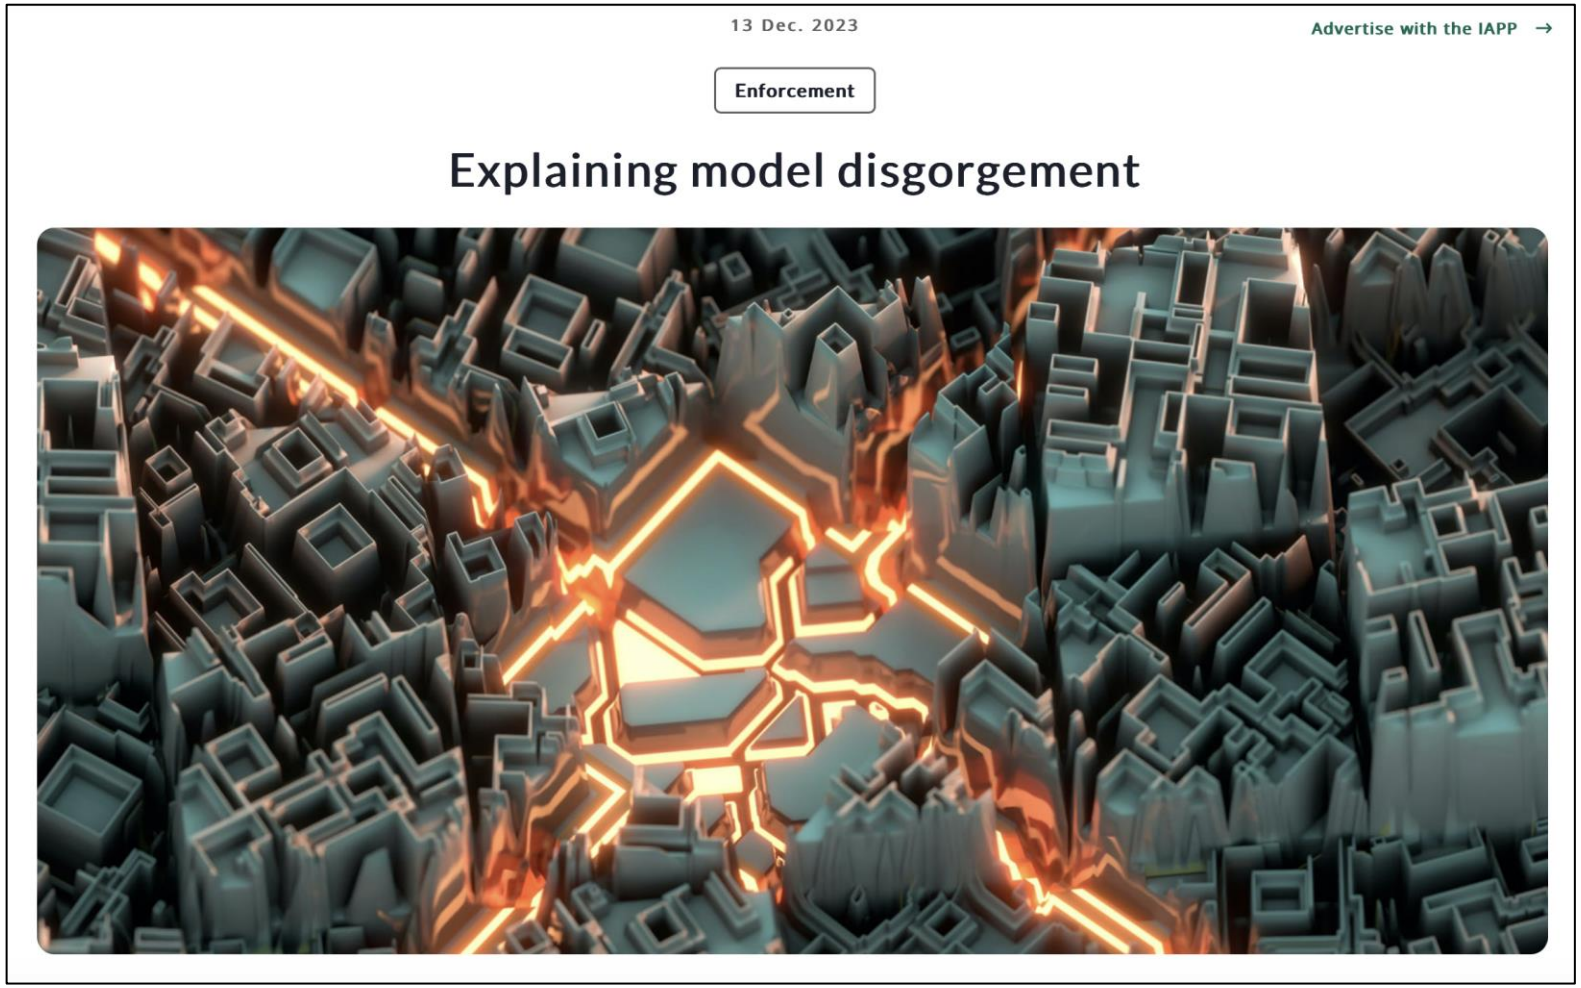

… and privacy is central

https://iapp.org/news/a/explaining-model-disgorgement

Tensions between AI and privacy

Purpose limitation

General-purpose AI models may use data in ways not originally agreed upon by consumers.

Right to be forgotten

AI is data hungry and machine learning models do not “forget” without major business impact.

Transparency

Transparency and explainability are challenging, especially with non-linear models.

Privacy foundational to AI data architectures

Key Risks In Real-Time Data Pipelines

Bias That Grows Over Time: AI models often use historical data. If that data is skewed, the bias can multiply as your system processes transactions in real time. For example, a credit-scoring model might penalize certain ZIP codes because the training data was unbalanced. When you're handling thousands of transactions a minute, a small bias can quickly become a major ethical and reputational problem.

Governance Gaps: Real-time data environments change quickly—sometimes so fast that governance rules struggle to keep up. Basic security measures like encryption and robust data catalogs can fall behind in the rush for real-time insights. If these protections aren't in place, sensitive information might end up exposed and you could lose customer trust or even run afoul of regulations.

Privacy and Compliance Roadblocks: Handling real-time data doesn't mean you get a free pass on privacy laws like the General Data Protection Regulation (GDPR). Managing consent, handling deletion requests and keeping proper records all get more complicated when data never stops moving. If your systems aren't built for compliance from the start, you'll have trouble meeting regulatory standards.

The "Black Box" Effect: Many AI models are hard to interpret, and real-time decisions can add another layer of complexity. If your team can't explain why a transaction was flagged as fraud or why a certain customer got a special offer, it's tough to fix mistakes or maintain transparency. A lack of explanation leads to skepticism, which can quickly undermine customer confidence.

Forbes

March 24, 2025

Designing Ethical, Real-Time Architecture

Privacy by Design: Start thinking about privacy at the beginning of every project. Use data encryption, limit access to sensitive fields and consider data masking for personally identifiable information (PII). Automating these processes reduces human error, which is critical in fast-moving environments.

https://www.forbes.com/councils/forbestechcouncil/2025/03/24/building-trust-in-motion-ethical-data-and-responsible-ai/

AI workloads need AI-ready data

Quality	Will my data deliver valid outcomes?
Security	Is my data secured against threats?
★ Fit for activation	Is my use of data permissible?

AI-driven innovation is the business case for privacy by design

Agenda

- Introduction to agentic AI
  - Architecture, guardrails, regulatory implications
  - Myths to bust, blind spots to expose

- Putting it into practice
  - Governing traditional AI vs. agentic AI
  - Learnings, best practices, pitfalls to avoid

Responsible use touches every stage of the data lifecycle

How do you delete or modify data when **retention** periods expire, consumers exercise their **rights** to withdraw consent, the purpose of collection is no longer valid, or data are redundant? What happens to AI models trained on that data?

Are you transparent about how you collect and use personal data? Do you have **consent** from the consumer, based on clear **purpose**, that complies with local regulation? Can you support web, mobile apps, and other connected devices?

Delivering the data and privacy promise: What **actually** happens

Making the data and privacy promise: What **should** happen

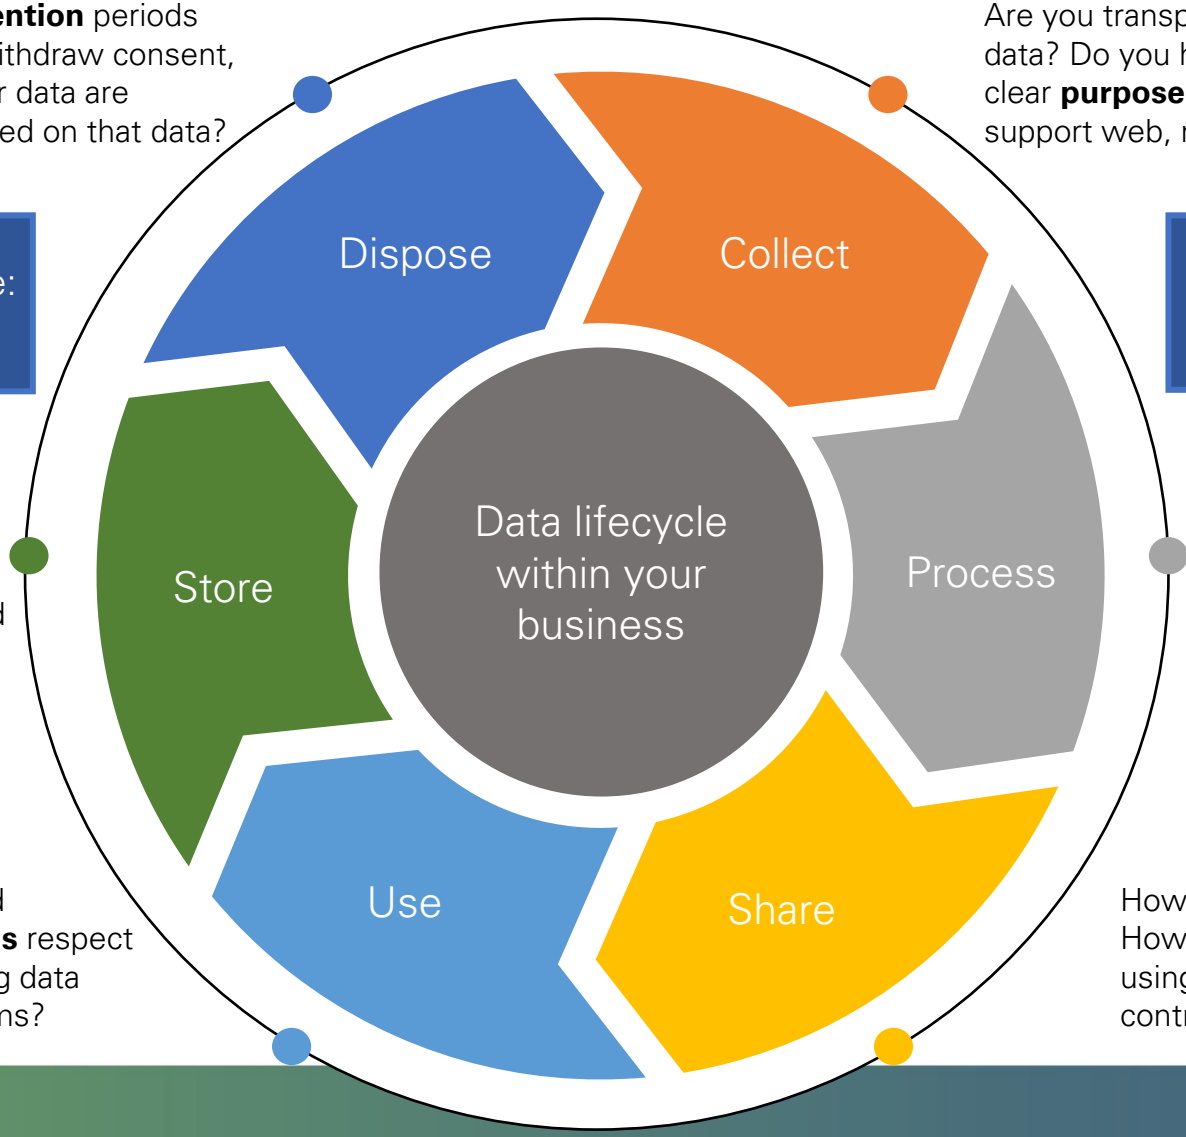

Can you locate personal data across your data estate? Are the right data **controls** in place? Do you know if personal data are being used to train AI models? Are you mitigating risk and resolving **drift** in data use?

How will all the functions, **activities**, and AI / behavioral models in your organization consume, process, and manage personal data? Is the purpose of each clear? Have you identified where **risk** is the highest?

How do you ensure personal data are used responsibly? Do your downstream **systems** respect consent and purpose **signals** when making data available for use by individuals or AI systems?

How are you sharing data with **third parties**? How are they managing that data? Are they using that data to train **AI models**? How are you controlling international data transfers?

Thinking through governance models

AI you buy

SHADOW applications

SCENARIO	GOVERNANCE
Employees use consumerised AI apps like ChatGPT for work, with no procurement or IT oversight. Employees are not aware of data risks associated with the tool.	Awareness and training Policies and employee attestation Third-party risk management Technical controls

PROCURED applications

SCENARIO	GOVERNANCE
Employees use apps procured by their company that utilise AI in their feature sets. Vendors may limited visibility into the AI tech used by the apps they have built.	Contract review • How built? • What tech? • Where is our data stored? • Is it siloed? • Is our data used to train? • Who benefits?

AI you build

THIRD-PARTY models

SCENARIO	GOVERNANCE
Internal data science teams use third-party foundational models for development. Oversight comes too late in the development lifecycle.	Shift left approach Data visibility and governance Conformity assessments Whitelisting Attestation

FIRST-PARTY models

SCENARIO	GOVERNANCE
Advanced data science teams build their own models. Oversight comes too late in the very expensive development lifecycle.	Shift left approach Data visibility and governance Conformity assessments Model risk mgmt Disclosure

Baseline governance requirements

Build an AI inventory

… to manage and monitor AI and its technical components, across the AI lifecycle.

Evaluate AI Risk

… consistently, and in line with global laws, standards, and organizational policies.

Monitor AI systems

… to foster collaboration and reduce administrative burden on technical resources.

Demonstrate transparency

… with key stakeholders to promote collaboration, trust, and compliance.

Dramatically expand AI and data literacy

Changing the discussion

Privacy impact assessments → Data enablement plans

What are your data ambitions?

What must I do to enable those?

Questions and answers

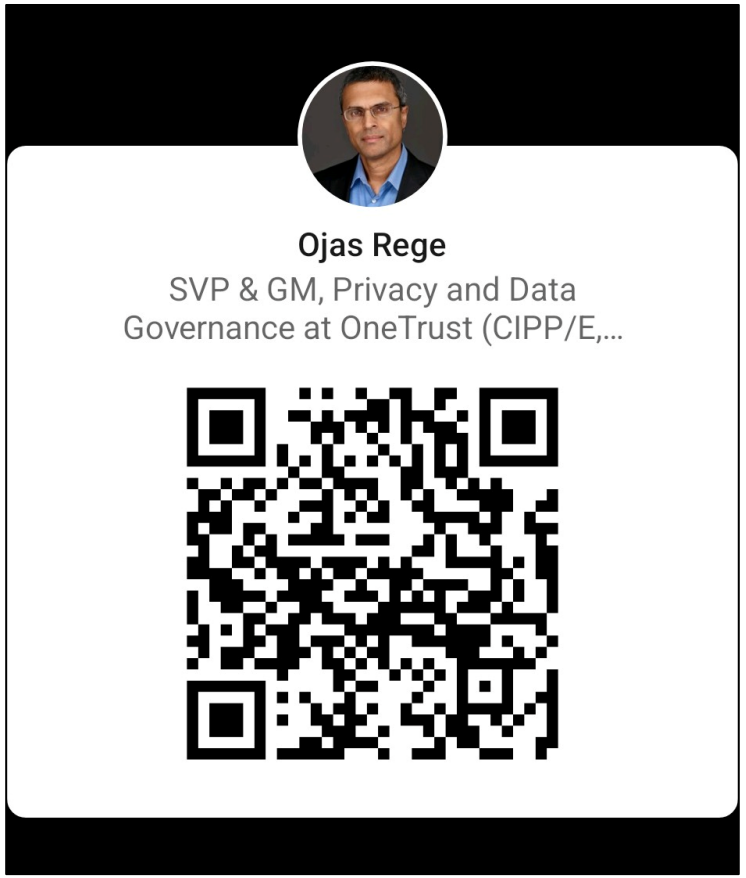

Thank you for your time!

Web Conference
Participant Feedback Survey

Please take this quick (2 minute) survey to let us know how satisfied you were with this program and to provide us with suggestions for future improvement.

Click here: https://iapp.questionpro.com/t/ACtQeZ6dfO

Thank you in advance!

For more information: www.iapp.org

Attention IAPP Certified Privacy Professionals:

This IAPP web conference may be applied toward the continuing privacy education (CPE) requirements of your CIPP/US, CIPP/E, CIPP/A, CIPP/C, CIPT or CIPM credential worth 1.0 credit hour. IAPP-certified professionals who are the named participant of the registration will automatically receive credit. If another certified professional has participated in the program but is not the named participant then the individual may submit for credit by submitting the continuing education application form here: submit for CPE credits.

Continuing Legal Education Credits:

The IAPP provides certificates of attendance to web conference attendees. Certificates must be self-submitted to the appropriate jurisdiction for continuing education credits. Please consult your specific governing body's rules and regulations to confirm if a web conference is an eligible format for attaining credits. Each IAPP web conference offers either 60 or 90 minutes of programming.

For questions on this or other
IAPP Web Conferences or recordings
or to obtain a copy of the slide presentation
please contact:

livewebconteam@iapp.org