Future-Proofing AI/ML Compliance Through Strong Data Privacy Foundations

October 28, San Diego

About Me

Vasudha Hegde
Sr. Privacy Program Manager
DoorDash

- Background in Telecommunications Engineering – 3 years as an Audio Broadcast Engineer
- Masters in Information Systems from UMD, College Park
- 5 years in Privacy Consulting at PwC
- 5 years in Privacy Program Management with a focus in Governance programs

Why This Matters Now: Innovation Vs. Compliance

Innovation

- GenAI Deployment, LLMs and Chatbots
- ML Automation
- Predictive Analytics
- AI-driven personalization

Compliance

- EU AI Act compliance and regulatory monitoring
- Privacy by Design
- Risk and Impact Assessments
- Transparency and Accountability

Polling Question #1

When you hear about new AI regulations in the US (like the Texas Responsible AI Governance Act), what's your first reaction?

a. "We're on top of it!"

b. "It's important but complicated."

c. "It feels overwhelming, but I'm trying to keep up."

d. "It doesn't matter to us; we do no AI development."

The Surge in AI/ML Oversight in the Last Year

broad scope

Q3 2024

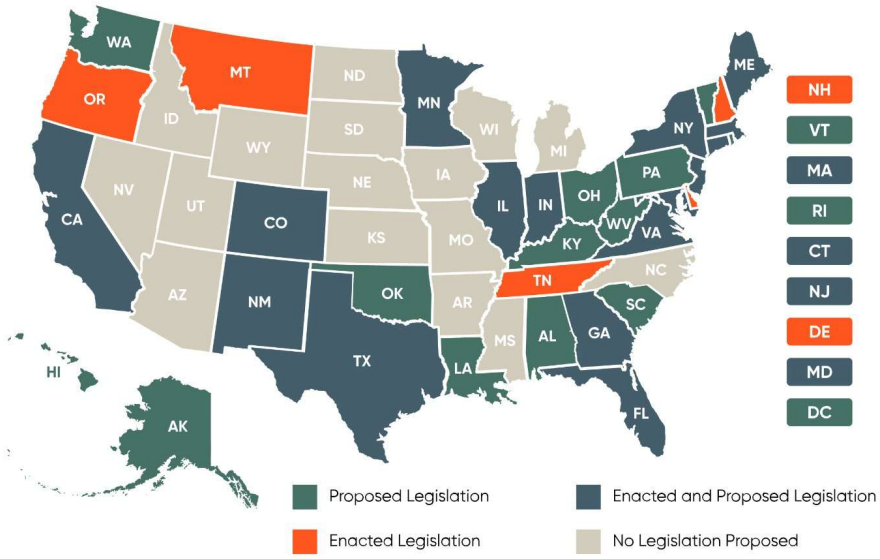

Q3 2025

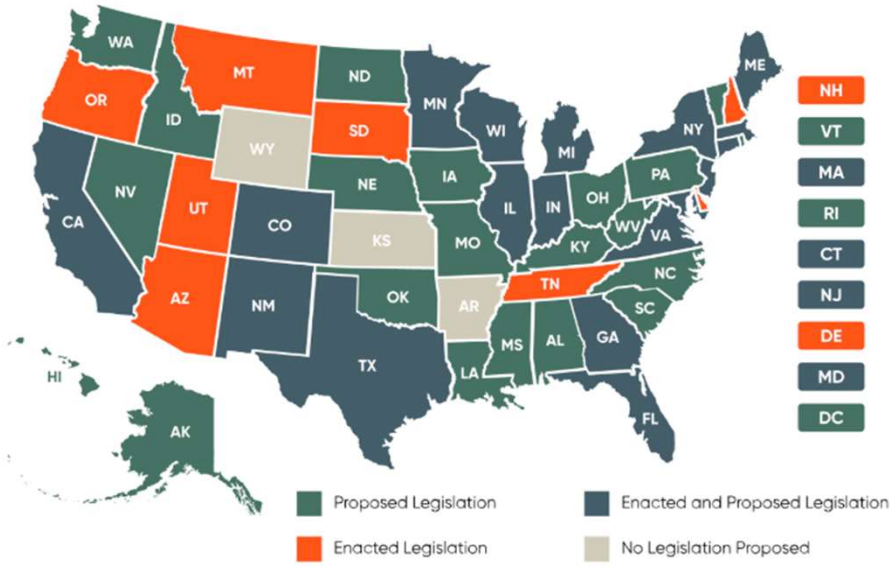

Source: bclplaw.com

Where AI/ML Laws Embed Privacy Principles

EU AI Act

Data quality, transparency, risk management

Colorado AI Act

Right to opt-out of automated decisions

Canada's AIDA

Impact assessments focused on data and discrimination risks

Common Theme: Privacy-by-Design for AI Systems

The Core Privacy Risks in AI/ML

- Unintended Personal Data Use
- Lack of consent and transparency in training data
- Data biases and fairness risks
- Data minimization vs. model performance improvement

Polling Question #2

Who is responsible for AI/ML governance in your organizations?

a. The Privacy team

b. The Legal/Security team

c. A dedicated/newly formed AI/ML Governance team

d. None of the above – Engineers/Product teams have to ensure responsible tech on their own

e. I'm not sure!

Why Privacy Practices Future-Proof AI/ML Regulatory Compliance

Robust Data Governance	Easier to comply with regulations and laws
Consent, Transparency and Fairness Controls	Built-in safeguards for AI/ML models and use cases
Documentation and Risk Assessments	Easier to meet auditability requirements
Trust with Users and Regulators	Competitive advantage, more user adoption and higher brand reputation

What ”Good” Looks Like: Key Pillars of a Privacy-Centered AI Governance Program

01 DATA GOVERNANCE

Data inventory and data labeling/classification hygiene

02 LEGAL BASIS

Purpose limitation and consent management

03 TRANSPARENCY

Bias audits, model explainability and documentation

04 RISK MANAGEMENT

Privacy and AI/ML risk assessments, baked into development lifecycle

05 ACCOUNTABILITY

Ownership, responsible development and oversight

06 MAINTENANCE

Periodic reviews, auditability and continuous performance monitoring

Polling Question #3

Which area feels like the biggest challenge for future AI compliance at your organization?

a. Collecting and using the right data

b. Explaining how the AI system makes decisions

c. Managing risks like bias or unfair outcomes

d. Meeting new transparency and reporting requirements

e. I'm not sure yet – it's all moving fast!

Common Pitfalls to Avoid

- ✗ Building first, retrofitting privacy and “responsible” controls later
- ✗ Siloed privacy teams (vs. embedded in AI/ML dev lifecycle)
- ✗ Assuming anonymized data means no risk
- ✗ Treating privacy as a “checkbox” rather than as a design principle

Quick Wins to Get Starts on AI/ML Compliance

- Build a central AI/ML system inventory with appropriate metadata
- Integrate risk or privacy impact assessments into model development lifecycle
- Develop guidelines for consent, transparency, and opt-out early to ensure responsible design and development
- Stand up cross-functional AI governance committees including privacy, product, legal and security teams

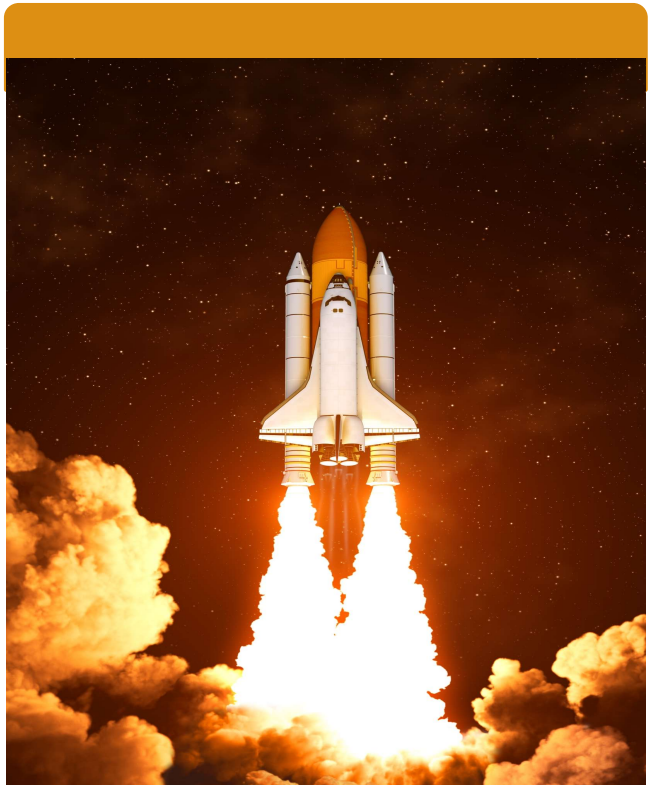

Case Study – Building With Vs. Without Privacy by Design

Case: A global company deploys an AI chatbot using customer data.

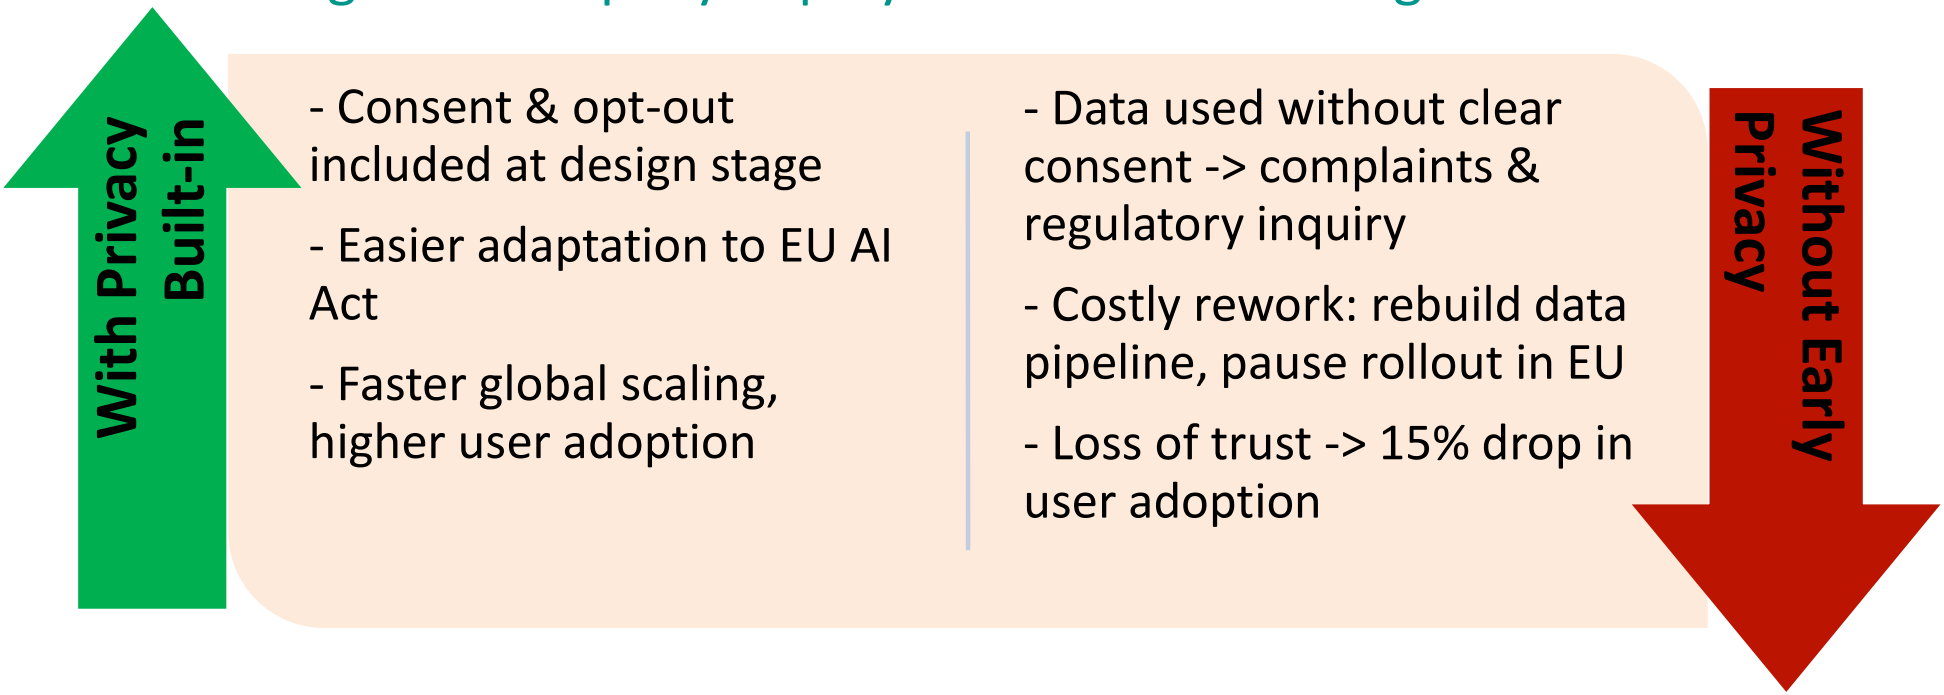

Exercise – Would you approve this AI tool?

Scenario: Your HR team wants to launch and use an AI-powered hiring tool

Features of the tool:

- Trains on 5 years of resumes
- Optimizes for “culture fit”
- Automated rejections/interview scheduling
- No bias audit results/testing

1. What are the top privacy risks here?
2. How would you adjust the development or procurement process to embed privacy early?
3. Who in your org needs to be involved?

Final Takeaways

01 Future compliance depends on today's privacy hygiene

02 Privacy is a lever for trust, responsible AI, and global scalability

03 Organizations that embed privacy now will move faster and safer later

Questions?

Thank you for attending!

Twitter: @InfoGovWorld
LinkedIn: @InfoGov World Magazine

www.InfoGovWorldConference.com
AIWorldconference.ai

Appendix

Operationalizing Privacy in AI Development

Step 1: Map & Inventory - Know your AI/ML systems, data flows, and purposes.

Step 2: Define Privacy Triggers - Clear rules for when privacy reviews are needed.

Step 3: Integrate Reviews Early - Embed privacy into intake, design, and procurement processes.

Step 4: Cross-Functional Governance - Privacy, Legal, Security, Product in one committee.

Step 5: Document & Communicate - Risk assessments, model explainability docs, transparent user notices.

Step 6: Monitor & Evolve - Ongoing audits, lessons learned, adapt to new regulations.

Common Points of Friction in Operationalization (1/2)

Late involvement of Privacy teams and processes; Lack of early engagement during Product deployment	• Establish clear **triggers** for when a **privacy review is needed** • Integrate privacy review **into the intake or procurement process** with automated flags • Establish other channels where privacy questions can be asked and answered (office hours, email alias, slack channel etc.)
Duplicate or Fragmented Review Processes leading to confusion amongst Product/Business teams	• Build a **centralized intake form** that routes to the right teams • Align review teams on **shared language and review criteria** to reduce confusion and eliminate redundancy • Create an **info page/KB with instructions** to product teams on the end-to-end review process and requirements
Privacy compliance is often seen as a blocker for product deployment, not an enabler	• Use a **tiered review approach** (light-touch for test environments, full review for production) • Be transparent about **SLA timelines** and parallel processing • Provide early privacy guidance to inform product teams on **"how to ship responsibly"** to speed up reviews and avoid re-work

Common Points of Friction in Operationalization (2/2)

Poor Visibility into Systems or Data Use

- Maintain a **central system and data inventory** with sufficient details about data use, retention, integrations and other information gathered during previous reviews
- Enforce **standardized documentation of systems**, with examples and templates

Incomplete, and/or vaguely or incorrectly answered questionnaires

- **Provide context-specific guidance** (examples or tooltips) to clarify response expectations
- **Enable collaborative submission** between product, engineering, and data teams to complete the form together, especially for technical questions