

AI Without Access: Shipping Intelligence When You Can't See User Data

UXDX USA 2026 – May 13 2026

Frédéric Rivain

CTO Dashlane

AI Capabilities



Machine Learning



Data Processing



Generative AI

Locked Data



Personal Information



Financial Data



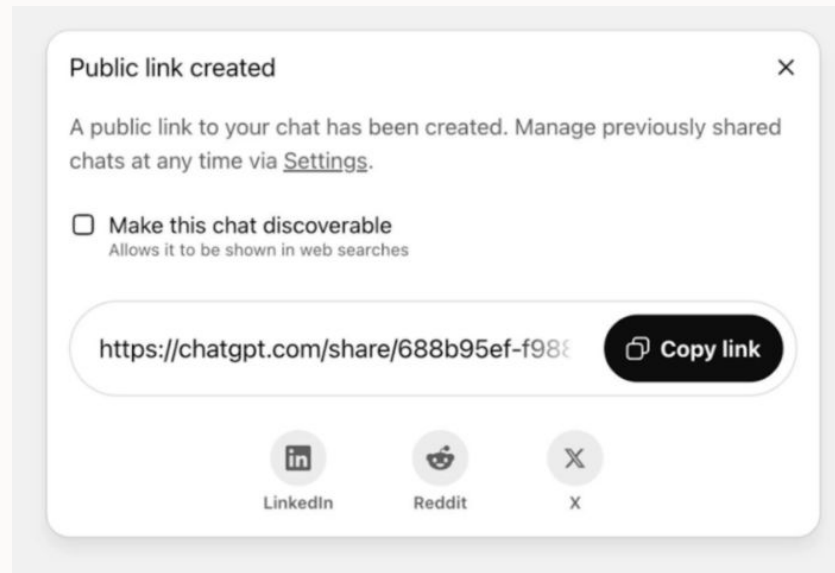
Encrypted Content





Nearly 100,000 ChatGPT conversations searchable on Google

- Summer 2025
- Thousands of private AI conversations indexed by search engines
- Users didn't realize their data could become public





Three forces you can't ignore


1. Customers → “Don’t look at my data”
2. Regulators → “Prove you don’t”
3. Enterprises → “Put it in the contract”

68% of consumers globally are (...) somewhat or very concerned about their privacy online, and 57% of consumers agree that AI poses a significant threat to their privacy.

Forbes, March 18 2025



What actually breaks



Encrypting data is not enough



*End-to-end encrypted...
except when it isn't.*



Building private AI by design, in practice



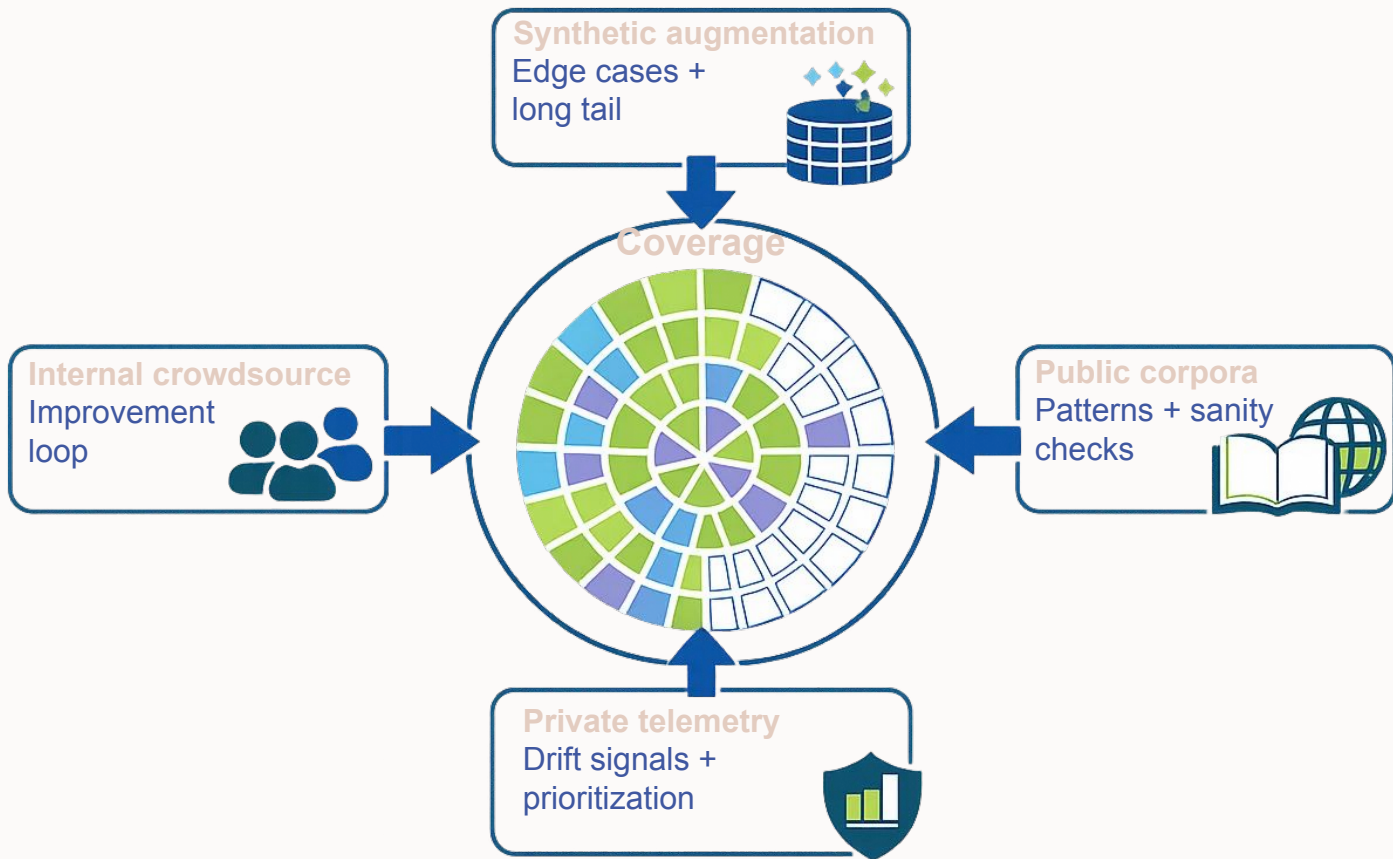
Use Case 1: Protecting users from phishing without seeing their browsing activity

- Data collection
- Model training
- Deployment

The screenshot shows a notification box with a light blue background. At the top left is a logo consisting of three vertical bars of varying heights. To its right, the text "Phishing attempt detected" is displayed. Below this is a red-bordered box containing a red circular icon with a white 'X' and the text "phishingwebsite.com" and "Suspicious URL". Underneath is a white-bordered box with a grey circular icon containing an 'i' and the text "This website was flagged by Dashlane's AI-powered phishing protection, and it may be trying to steal your data." At the bottom of the notification, there is a line of text: "Questions? Reach out to ithelpdek@verovenn.com". At the very bottom, there are two buttons: a grey button labeled "Not a scam" and a teal button labeled "Leave website".



Training a model without seeing the data





Frugal modeling

Start smaller than you think



Less data needed



Faster iteration loops




Clearer failure analysis



Lower cost + energy



Easier governance




On-device deploy



Lesson 1

- Data coverage is critical
- Frugal Model may be enough
- Run on device for maximum privacy

Use Case 2: AI Copilot on data we can never see



AI Advisor
CONFIDENTIAL AI ENGINE

Highest priority risk
4 employees in Sales were targeted by phishing on a domain mimicking your CRM login. 2 also have compromised credentials and haven't remediated in 14 days.
[Nudge all](#) [Investigate incident](#)

Hi! I've flagged your most critical risks above. Ask me anything about your credentials, phishing, or dark web exposure.

Show phishing incidents this month →

Which employees need urgent attention? →

Generate a SOC 2 compliance-ready report →

Which users need attention most?

3 employees are highest priority:

- **J. Mitchell** – 2 phishing hits, 12 compromised credentials
- **M. Santos** – no remediation in 14 days
- **T. Kim** – new dark web exposure today

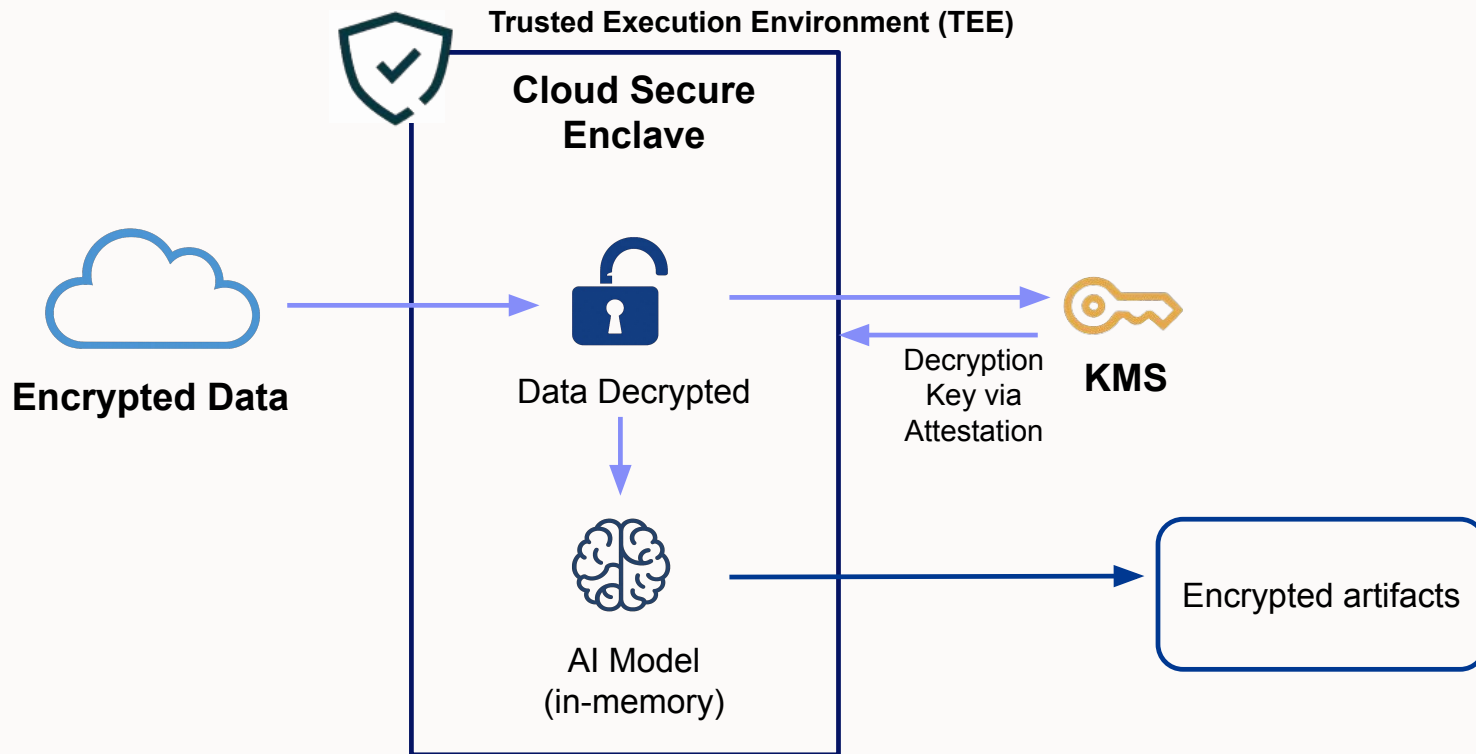
[View all at-risk employees →](#)

Ask about your security data... →

Processing happens inside a hardware-isolated boundary. Dashilane never sees your data in the clear.



What about Confidential Computing?





Deployment: Right Compute, Right place



ON DEVICE



Fast, Offline



Max privacy



Model size limited, feature engineering needed



CONFIDENTIAL COMPUTING



Bigger models, More context



Confidential boundary



Deployment only, training is limited



Lesson 2


- Feature Engineering may be enough
- End to end encryption: in transit, at rest, and in use leveraging Confidential Computing



A playbook for privacy-preserving AI




Privacy-by-design is a balancing act





Privacy constraints as an input, not a review gate

Privacy Stance
Zero-Knowledge Architecture
Privacy policy
Acceptable Use Policy
AI Policy





Where to start

Next Week

- Privacy One-Pager
- Cross-functional accountability

Next 3 months

- Launch acceptance criteria
- Training pipeline
- Deployment strategy
- First prototype

Next 6 months

- Pilot in production
- Repeat playbook



Takeaways

1. Start with the privacy boundary before the model choice.
2. Optimize training data collection for coverage, not just scale.
3. Use smaller models
4. Choose deployment based on the performance need
5. Ship incrementally. Don't wait for the "perfect" private AI architecture.

