Private AI: Building Trust Through Verifiable Computation

Mingshen Sun, Mateus Guzzo

TikTok

PEPR 2026

Data and compute power demands in the AI era

AI has transformed how people learn, work and live.

The most powerful AI today runs on server-class hardware.

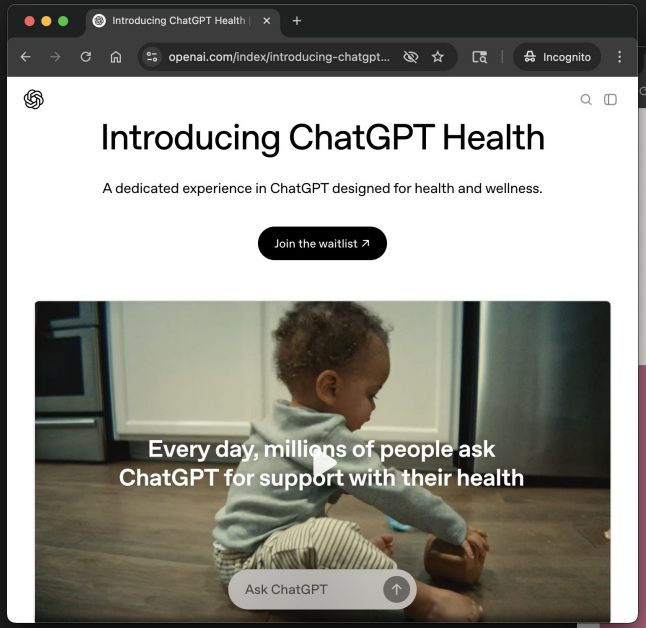

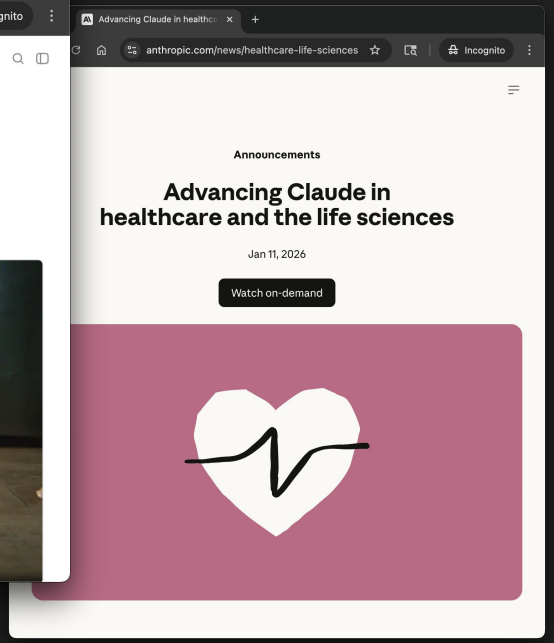

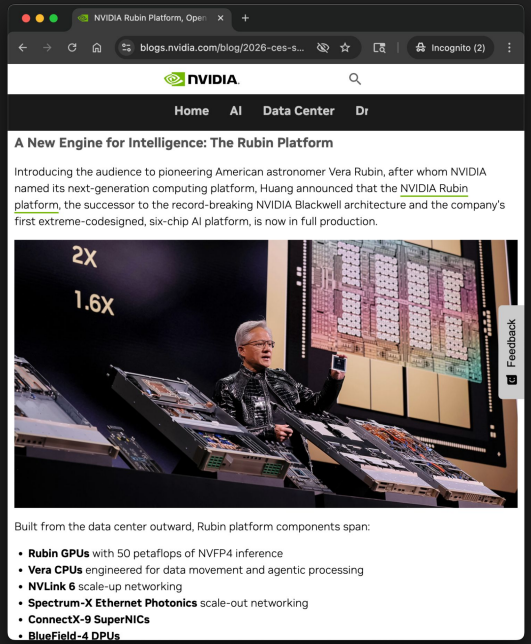

Sensitive user prompts and context sent to the service provider to be processed

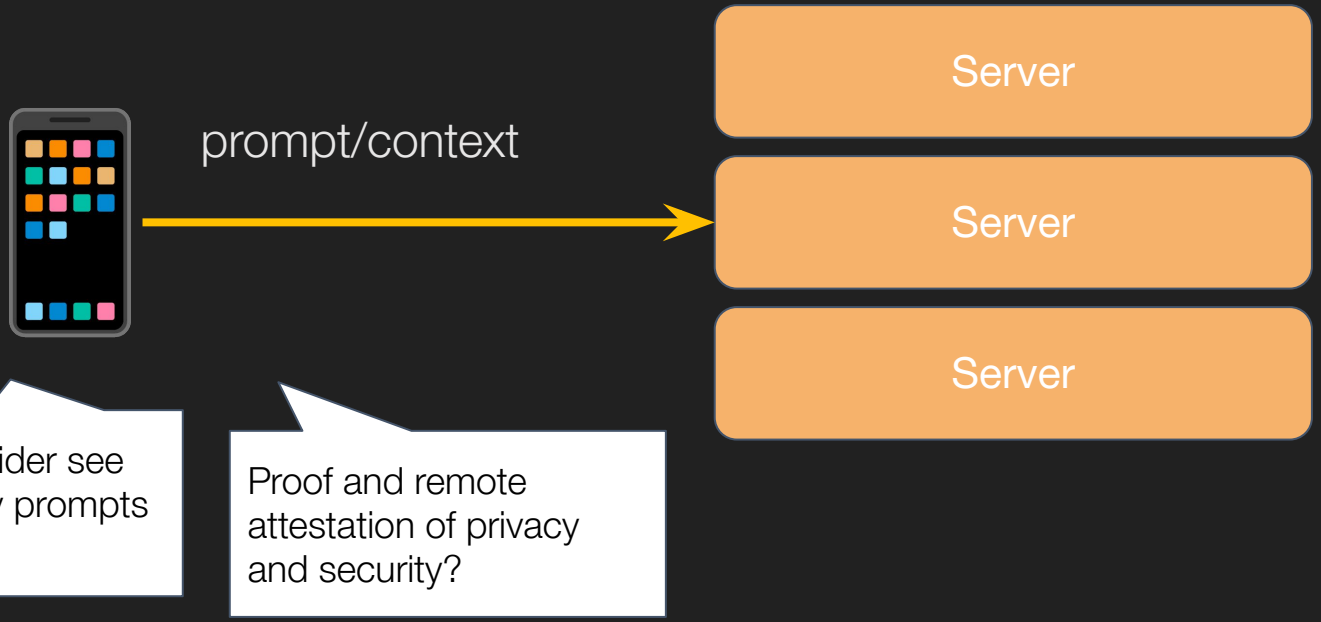

Sensitive user prompts and context *protected* by Trusted Execution Environments

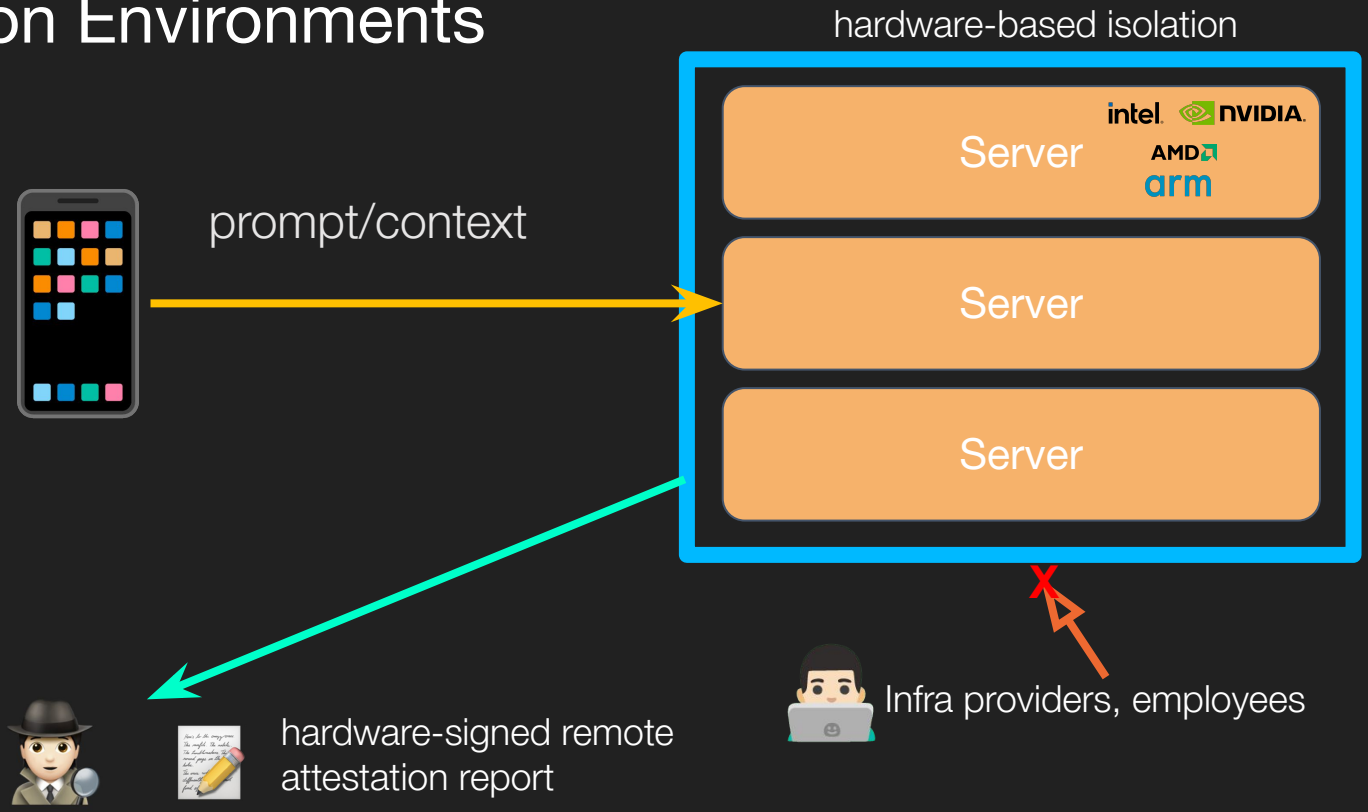

High-level goals of Private Verifiable Compute

Use powerful cloud AI on sensitive data, while making sure no one — not even the platform provider — can see who you are or what you sent. And let you verify that for yourself.

- Private — your data is protected the whole way through.
- Anonymous — your identity is separated from your request.
- Verifiable — you don’t have to take our word for it; you can check.

Design goals and privacy & security guarantees

Private processing

- Data Security
- *Privacy preserving*
- *Enforcement guarantees*
- *No privileged access*

Private storage

- *User-controlled encryption key*

Verifiable transparency

- *Code transparency and assurance*
- *Verifiable privacy*
- *Remote attestation*

How Private Verifiable Compute works?

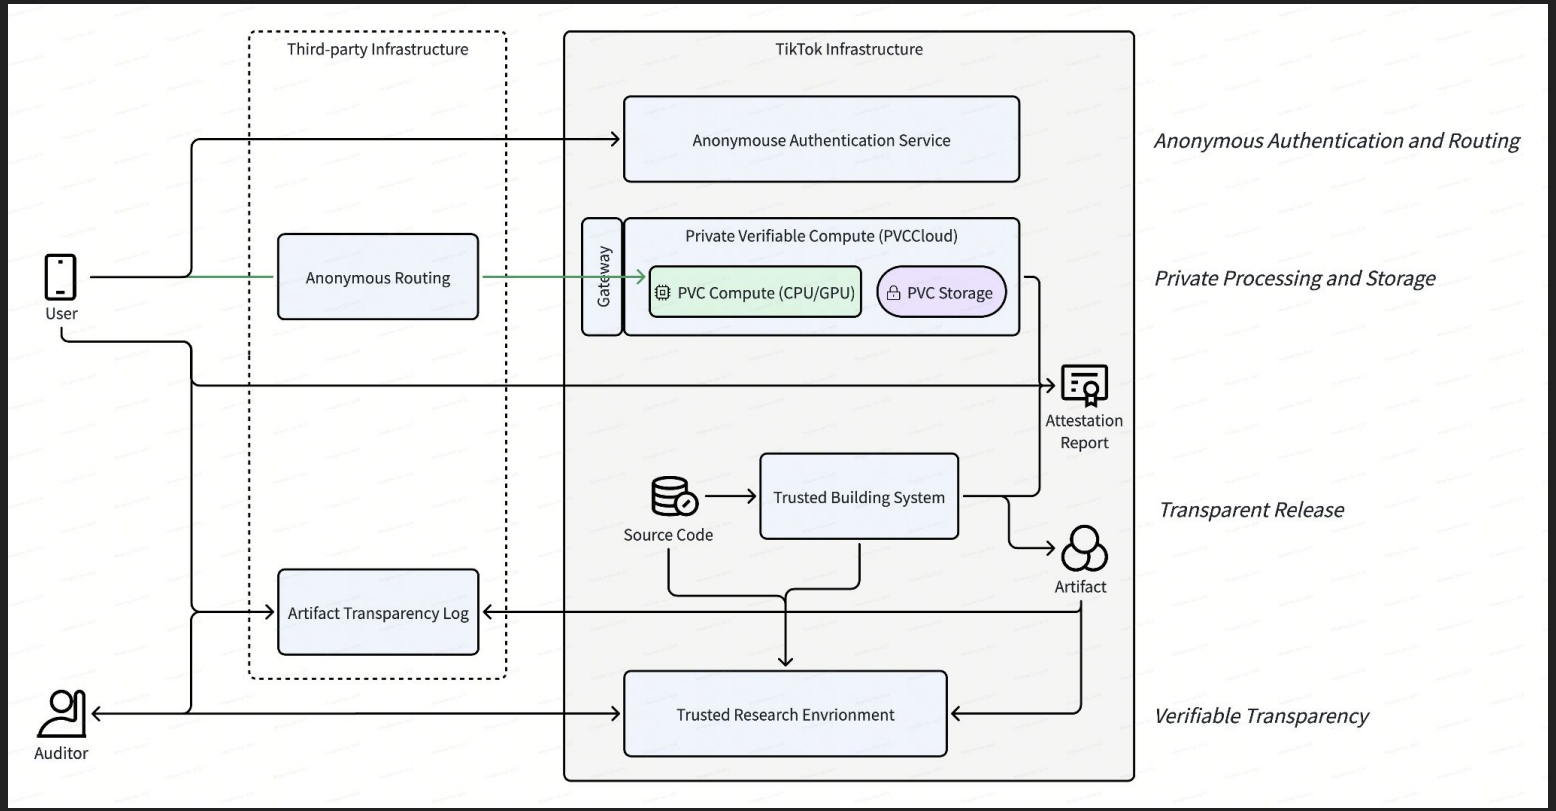

Value proposition for users and business

Build user trust

PVC leverages multiple innovations of Privacy Enhancing Technologies. With the nature of privacy, transparency and verifiability, it showcases the technical advancements of PET in private AI processing.

Mitigate risk

Due to the design of private processing and storage, sensitive data is not visible even to the service provider during processing, and additionally, this can be remotely attested by reviewers independently.

Unblock business innovation

New business scenarios with AI features relying on sensitive user data might face privacy and compliance challenges. Privacy and security guarantees ensured by PVC could potentially unblock valuable business innovations.

An example use case: private AI chatbot

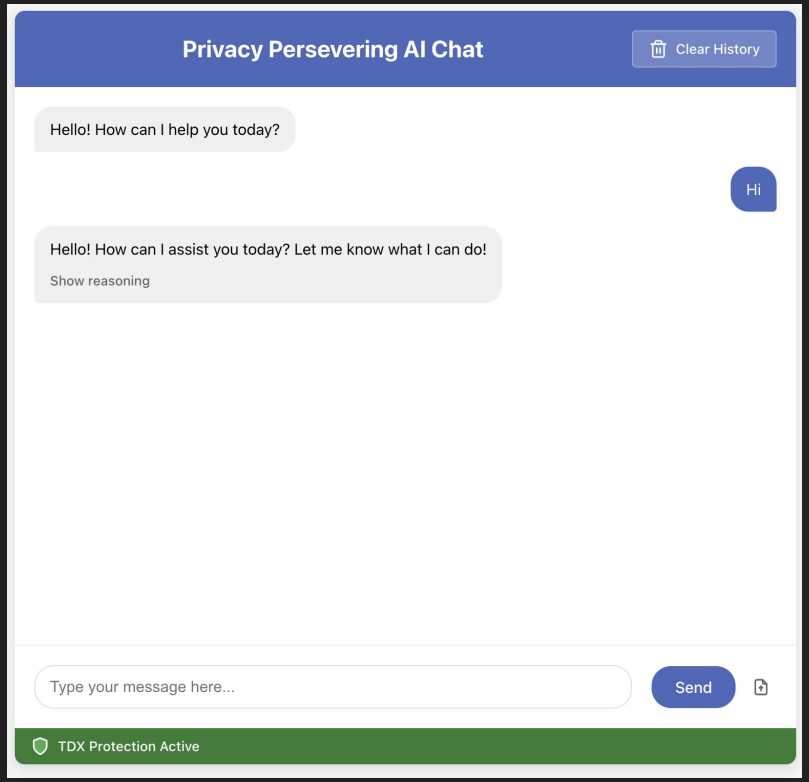

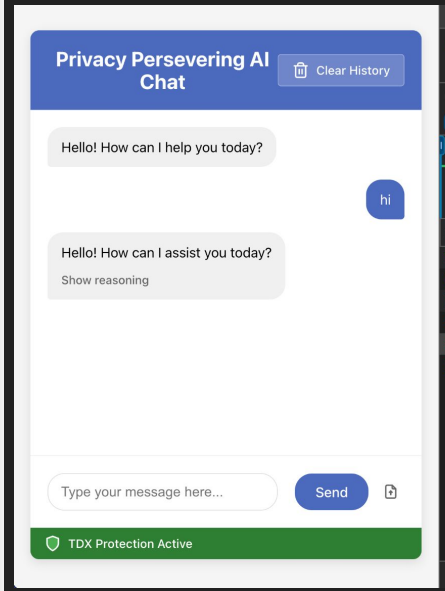

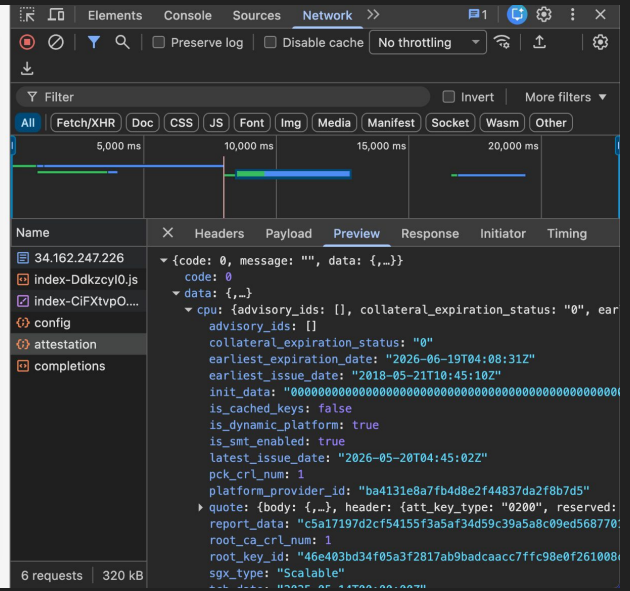

What does the remote attestation report look like?

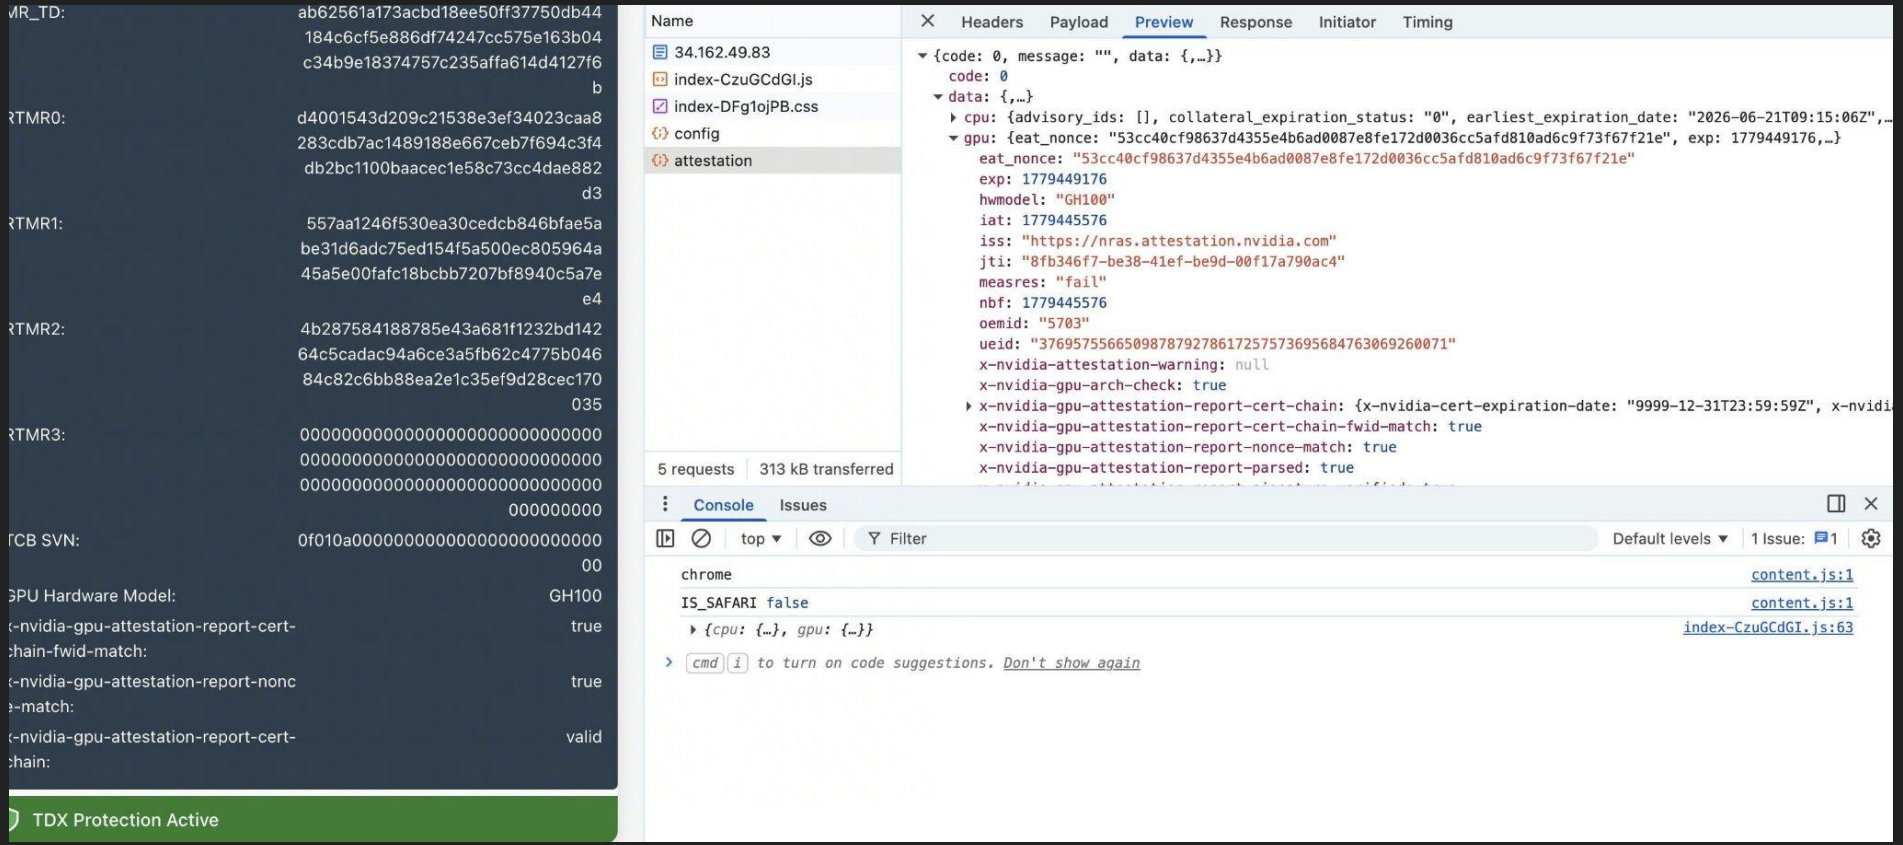

Open source for transparency and verifiability

Private Verifiable Compute is an open source project, published in Dec 2025.

- Cloud, on-premise, and hybrid deployment
- Open & transparent code release

https://github.com/tiktok-privacy-innovation/private-verifiable-compute

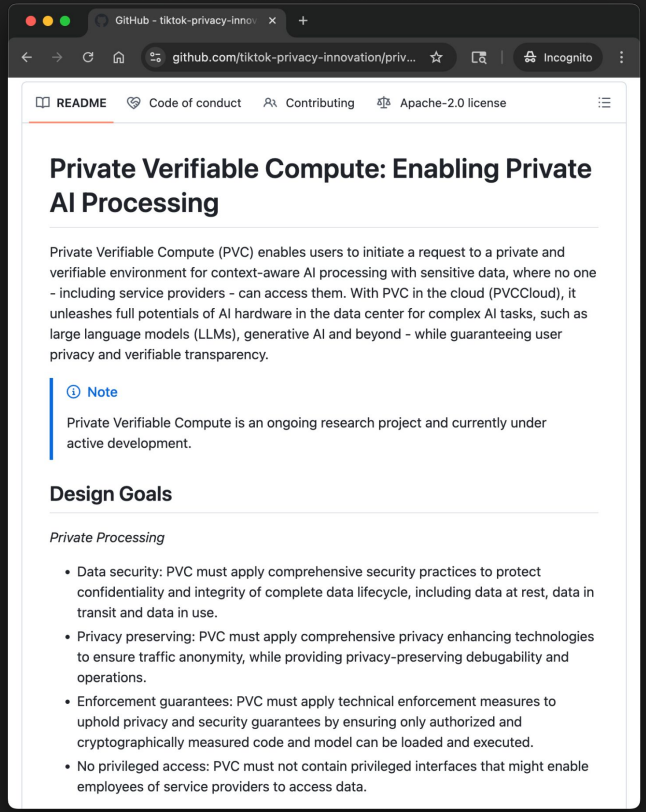

Thanks!