Conversations as critical infrastructure

Building ethical AI systems that unlock human knowledge

Sam Udotong | CTO & Co-founder, Fireflies.ai
sam@fireflies.ai | linkedin.com/in/sudotong

fireflies.ai + THE AI CONFERENCE

Fireflies: Your AI teammate for meetings

- Serving **20M+** people in over **500,000** orgs

- Users at >**75%** of the Fortune 500

- **3B+** meeting minutes processed

- Profitable since **2023** (raised **$19M** Seed & Series A)

- **120** employees across **20** countries, fully remote

- **70%** of team on R&D — no VP of Marketing or Sales

The paradox

- We spend 23 hours per week in meetings

- 70-80% of knowledge vanishes as soon as we leave them

The ethical dilemma: How do we capture knowledge — without compromising privacy, trust, or human agency?
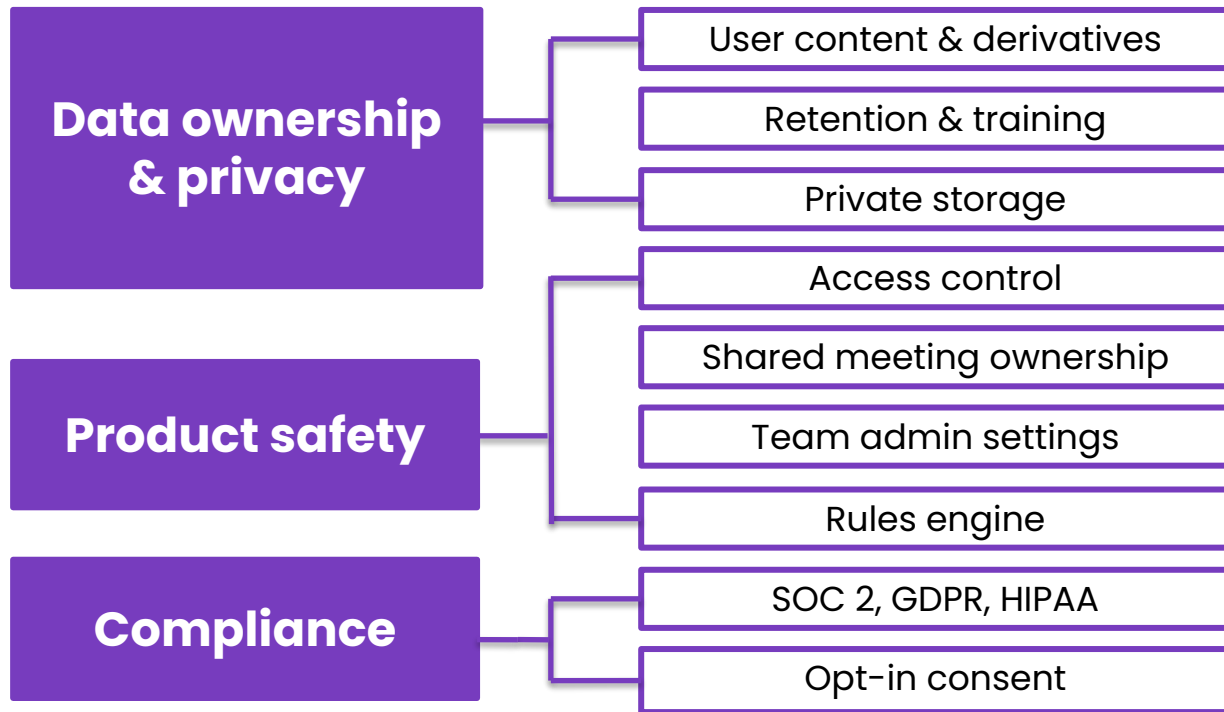
What we'll talk about today

- Practical frameworks for ethical AI development

- Real-world implementation lessons

- Actionable insights you can apply immediately

- How to build AI that creates value *and* maintains trust

Three pillars of trustworthy AI

Three pillars of trustworthy AI

Data ownership & privacy
- User content & derivatives
- Retention & training
- Private storage

Product safety
- Access control
- Shared meeting ownership
- Team admin settings
- Rules engine

Compliance
- SOC 2, GDPR, HIPAA
- Opt-in consent

Pillar 1: Data ownership & privacy

Principle: Users must have true control over their data — not just promises

Fireflies approach:

- Private storage — customers own their data

- Never training on customer data

- One meeting, one recording — no duplicates

Takeaway: Ask "Who can access this data?" before writing any code

THE AI CONFERENCE

Pillar 2: Product safety

Principle: Safety isn't about limiting functionality

— it's about thoughtful design

Fireflies approach:

- Team permission models reflecting real org structures
- Admin controls prevent accidental data exposure
- Rules engine for automation & extensibility

Takeaway: Design against "What's the worst that could happen?"

Pillar 3: Compliance as competitive advantage

Principle: Industry standards are the floor, not the ceiling

Fireflies approach:

- HIPAA compliant → unlocked healthcare market
- Consent as a first-class citizen
- Adaptable infrastructure that exceeds requirements

Takeaway: Design beyond the strictest customer today to win the rest tomorrow

Case study: From principles to product

Evolution of ethical AI at Fireflies

- 2019: Simple transcription — the "original" AI agent

- Challenge: How to expand value without compromising trust?

- Solution: Layer-by-layer approach to features

The layered approach

- **Core layer:** Secure transcription & notes

- **Layer 2:** Department-specific AI skills (post-meeting automation)

- **Layer 3:** Real-time intelligence (Talk to Fireflies)

- Each layer maintains same security principles

Critical decision points

- **When offered data to train models:** We said no, choosing user trust

- **When building for teams vs. individuals:** We chose shared access

- **When entering regulated industries:** Got compliant instead of avoiding complexity

- **When scaling:** We hired across 20 countries — a worldwide company builds a global product

Practical frameworks for your AI journey

THE AI CONFERENCE

Framework 1: Trust-first architecture

Before any feature: "How does this impact trust?"

Questions to ask:

- Who owns the data?
- Can users delete it?
- Would I want my most sensitive conversations processed this way?

Example: AI app store | user owns derivatives of meeting content

Framework 2: Value-ethics matrix

Example: Private storage implementation (high value, high risk)

Framework 3: Build vs. buy vs. partner

Evaluation criteria:

- **Build in-house:** When it's core to your competitive advantage and you need full control
- **Buy/integrate:** When proven solutions exist and speed to market matters more than customization
- **Partner strategically:** When you need specialized capabilities but want to maintain focus on your core product

Example: Partnership with Perplexity for in-meeting web search

Lessons from the trenches

THE AI CONFERENCE

What we learned from building Fireflies

Constraints breed innovation: Limited funding forced creative solutions

Trust compounds: Early security investments pay massive dividends

Talk to your users: Iterate quickly based on user feedback

Global teams create global products: Diversity isn't just ethical, it's practical

The future we're building

AI teammates that amplify human potential

Enhancing connection, not replacing it

Conversations are critical infrastructure — build responsibly

Build AI you'd trust with your most important conversations

Thank you!

Sam Udotong | CTO & Co-founder, Fireflies.ai
sam@fireflies.ai | linkedin.com/in/sudotong

THE AI
CONFERENCE