

All things AI

Charu Sharma & Michael Mernagh, Fenrock AI 2026

What to expect

- 1 Our story and BSA/Fraud use cases**
- 2 AI concepts**
- 3 Controls & guardrails - a practical checklist**
- 4 Q&A**

Introductions

Charu Sharma, Cofounder & CEO

- Built one of the largest healthcare API companies
- Named a top AI leader to follow by IBM and advises Fortune 500 execs on AI
- Forbes 30 under 30; Indian Taekwondo team; National Bravery Award

Michael Mernagh, Cofounder & CTO

- AI veteran from Google, Apple, Stanford
- Built one of the first AI agents in the world
- Invented and scaled Apple's first privacy-preserving AI

We care about controlled AI for banks

We are privacy freaks.

When we first dated...

When we started talking to banks...

Employees entering PII in AI platforms that are not controlled, governed and audited by the bank.

Data silos and point solutions.

CONTROLLED AI

GUARDRAILS

AUDITABILITY

Central AI layer; many use cases

Approved models

Bank data + connectors

Permissions + PII rules

Logging + audit trail

AI GOVERNANCE & CONTROL PLANE

BSA

Fraud

Lending

CSR

Finance

People use “AI” to mean many different things

Chatbot

You ask; it answers.
Pre-coded decision tree.

Copilot


Answers questions.

Agent

Performs tasks.

“What does our AML policy say about closing EDD cases when BO documents are missing or not updated in 12 months?”

“Help me perform EDD review for this customer.”



High-leverage AI use cases in BSA and fraud

Start where the output can be reviewed and defended

Alert investigation prep

QA on investigation notes

EDD / periodic reviews

Policy Q&A with citations

SAR narrative drafting

Exam / audit preparation

CTR review support

Fraud pattern synthesis

“Where can AI help me with information gathering, research and analysis, reporting, reconciliation so I can spend more time exercising judgment and being proactive?”

Run the enhanced due-diligence review for Whitcomb Scrap & Salvage LLC



How AI fails inside a bank

Hallucination- confident but unsupported answers

Data leakage- PII/SAR/internal info in wrong place

Bias/inconsistency- automates messy historical patterns

Silent change- model behavior changes after update

Each risk has a control. The question is whether the system enforces it.

Safe AI framework: four pillars

Boundaries

What can AI see, access, communicate with, change or execute?

Evidence

Can every material claim trace back to source data, policy or document?

Approval


When does a human review, decide or approve before action?

Audit

Can you reconstruct what happened six months later?

**If it cannot be explained, tested, approved and reconstructed,
it is not ready for high-consequence work.**

Specify your risk appetite workflow by workflow



Different workflows deserve different stopping points.

Alert notes, SAR drafting, customer communications, account restrictions and payment actions should not share the same generic AI approval threshold.

Five questions to ask any AI vendor

1. What customer/bank data does the AI see, and where does it go?

2. Can every material answer show source evidence?

3. What can the AI recommend versus actually do?

4. Can we reconstruct the model, data, output, action and approver later?

5. How do you test behavior when models, data, policies or threats change?

Ask us anything about AI

charu@fenrock.ai

mm@fenrock.ai



Charu Sharma

Cofounder & CEO at Fenrock AI (YC W26) | Ex
LinkedIn, SignalFire, Forbes 30U30

