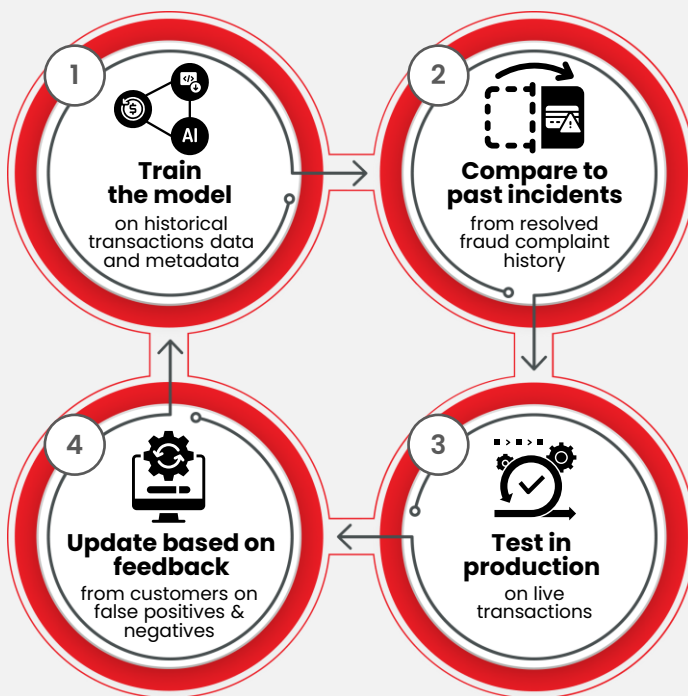


Industry Example: Finance

One of the largest financial services networks in the world built and deployed an AI that detects and prevents fraud in real time. The solution involves regular rather than agentic AI because of the latter's latency, since fraudulent transactions happen within seconds so speed is a key consideration.

The system ingests historical transaction data (authorizations, approvals, declines, etc.) coupled with relevant metadata (such as time stamps, merchant information, location, whether PIN or card was used, etc.) and runs those against a history of customer complaints to identify fraud patterns. The training is then applied to detect fraud on live transactions in close feedback from account holders, who flag false positives and negatives (Fig 1).



The combination of robust historical information and feedback from customers made the AI detection system more accurate over time; once corrected, it didn't make the same mistake twice.

As the system grew more effective, the number of false positives and negatives kept going down along with actual fraud incidence. Customer satisfaction improved, service agents received fewer complaints, and profitability increased.

Fig. 1: Feedback loop that improves model performance over time.

Our work for a major fintech shows the importance of proper data curation in the success of AI systems.

We started by delivering a secure, ETL pipeline that ingested the client's historical records into a data lake. With that foundation in place, the next stage was building an extensible AI agent factory, with shared orchestration and composable components that enable adding new assistants without rebuilding the system.

As a result, purpose-built agents support the client by handling 15 to 50k cases a day, depending on the context.

[Read the full case study](#) ➔