




AI Playbook 2026

FSI RESILIENCE IN THE AI ERA



OVERVIEW

The FSI AI Playbook examines how far AI has moved into financial decision-making, and what is holding it back.

Click [here](#) to find out about how Elastic can support you.

THE CORE PREMISE

In the AI era, the defining challenge for financial institutions is no longer scale; it is **clarity at scale**. CIOs and CISOs are prioritising to modernise faster, defend more aggressively, and earn trust continuously.

An ET CIO x Elastic FSI AI Playbook is built on the premise that data is the common thread across resilience, risk, and growth. By unifying search, observability, and security on a single platform, leaders gain real-time insight into systems, threats, and customer experiences.

About the Survey

100+ senior leaders across India's financial services sector

RESPONDENT PROFILE

80+

organisations

across India's financial services sector

WHO

CIOs, CTOs, Heads of IT, CISOs and security heads and data, AI, risk, fraud and compliance leaders.

WHAT

Questions across seven themes: AI decision-making, the next divide in financial services, AI in financial decisions, trust at transaction speed, AI vs fraud, observability of money in motion, and data retrieval and context

HOW

Conducted by ET CIO in partnership with Elastic. Responses are anonymised and reported only as combined results.

Executive Summary

KEY INSIGHT

The biggest gaps sit beneath the model. Leaders trace AI shortfalls to data access, fraud outcomes and visibility far more often than to the model itself.

- **AI quality is decided in the data layer.** 78% of leaders trace poor AI outputs to how data is retrieved and delivered to the model, while only 22% point to the model itself.
- **Real-time decisions are held back by connected data.** 64% name siloed systems or slow data delivery as their biggest real-time AI constraint, sixteen times the share that cite model accuracy.
- **Customer trust is won or lost in fraud detection.** 66% see fraud-detection failures as the biggest threat, split almost evenly between fraud that gets through and fraud caught too late or too noisily.
- **Few AI decisions can be proven under audit.** Most fraud decisions can be explained with effort, but only 16% of leaders describe them as fully traceable and evidence-backed.
- **Visibility gaps slow down every investigation.** 70% lack end-to-end visibility into their AI agents, and 68% say finding the root cause across dependencies takes the longest when a transaction fails.

CHAPTER 01

The Data *Foundation*

How data access shapes AI output quality and real-time decisions.

SURVEY THEMES

The Next Divide in Financial Services

Data Retrieval & Context

78% Trace Poor AI Outputs to Data, Not the Model

KEY INSIGHT

AI quality begins with the data, not the model. Nearly four in five leaders trace output quality to how data is retrieved and delivered, not to the model itself.

78%

attribute low-quality AI outputs primarily to the data surrounding the model

WHAT DRIVES AI OUTPUT QUALITY

42%

Incorrect or irrelevant retrieval

+

36%

Relevant data exists but is inaccessible to AI when needed

THE MODEL ITSELF

22%

Hallucination, reasoning errors or bias

In Real Time, AI Can Only Decide With the Data It Sees

KEY INSIGHT

Decisions are only as complete as the data behind them. If the right data is fragmented, inaccessible or slow to arrive, even a powerful AI system can make a decision with an incomplete picture.

EXAMPLE: *A card payment at checkout*

01 **Transaction history**

02 **Identity details**

03 **Fraud signals**

04 **Customer records**

THE AI MUST DECIDE:

**Approve or
stop**


the payment, using all four sources at once

64% See the Data Layer as Their Biggest Real-Time AI Constraint

KEY INSIGHT

Real-time AI is constrained by the data layer. 64% cite siloed systems or slow data delivery as their biggest real-time constraint, against 4% who cite model accuracy.

BIGGEST CONSTRAINTS ON REAL TIME AI DECISIONS



CHAPTER 02

The Trust *Test*

Fraud detection, customer trust and how defensible AI decisions are under audit.

SURVEY THEMES

Trust at Transaction Speed

AI vs Fraud

66.4% See Fraud-Detection Failures as the Biggest Threat to Customer Trust

KEY INSIGHT

Fraud risk runs in both directions. Concern is split almost evenly: 33.7% cite fraud caught too late or with too many false positives, and 32.7% cite fraud that gets approved.

WHAT'S BREAKING CUSTOMER TRUST?

66.4%

identify failures in fraud detection as the biggest threat to customer trust

WHAT MAKES UP THE 66.4%

33.7%

Fraud flagged too late or with too many false positives

+

32.7%


Fraudulent transactions approved because they looked legitimate

Only 16% Can Fully Trace Their AI Fraud Decisions

KEY INSIGHT

Explainable is not the same as defensible. 64% say fraud decisions are explainable with effort, but only 16% describe them as fully traceable and evidence-backed.

HOW DEFENSIBLE ARE AI FRAUD DECISIONS UNDER AUDIT



• 16%

Fully traceable and evidence-backed

• 64%

Mostly explainable, with effort

• 21%

Hard to reconstruct or a regulatory risk

CHAPTER 03

The Visibility *Gap*

What leaders can and cannot see inside their AI systems, and what it takes to close the gap.

SURVEY THEME

Observability of Money in Motion

70% Can't Fully See What Their AI Agents Do


KEY INSIGHT

Visibility gaps slow down resolution. 70% lack end-to-end visibility into their AI agents, and 68% say identifying the root cause across dependencies takes the longest when a transaction fails.

WHAT MUST BE TRACEABLE IN EVERY AI DECISION

- 01 What data the AI accessed
- 02 Whether it was relevant and current
- 03 How it influenced the decision
- 04 What happened across the systems around it

VISIBILITY ACROSS AUTONOMOUS AI AGENTS



70%

Incomplete, fragmented or with telemetry gaps



30%

Fully end-to-end visibility

68%

say identifying the root cause across multiple dependencies takes the longest when a transaction or trade fails

CONCLUSION

The Way *Forward*

What it takes to protect trust as AI moves financial decisions.

Protecting Trust Needs More Than Better Models

It requires a stronger data and visibility foundation underneath them.

01

Connect the data

Bring together data across applications, transactions, security systems and third-party environments, and make it searchable, accessible and usable in real time.

02

Observe the AI

Search, analytics and observability give AI the context for better decisions, and give teams the visibility to understand what happened when decisions go wrong.

03

Prove the decisions

Make every AI decision traceable and evidence-backed, so it can be explained to customers and regulators. Today, only 15.7% can do that for fraud decisions.

Not simply to make AI smarter, but to make it trustworthy.

Five questions to ask your team

- | | | |
|---|---|---------|
| 1 | Can we name every data source our fraud and credit models draw on at decision time? | CONNECT |
| 2 | How fresh is that data when the AI reads it, and do we measure it? | CONNECT |
| 3 | Can we see what our AI agents did across systems, not just inside the model? | OBSERVE |
| 4 | When a transaction fails, how long does it take to find the root cause across dependencies? | OBSERVE |
| 5 | If an auditor picked one AI decision from last quarter, how long would it take to show exactly what data it used? | PROVE |

If the honest answer to any of these is "not sure", you are not alone. Only **16%** of your peers can fully trace their AI fraud decisions today.

Click [here](#) to find out more about how Elastic can support you.