
Agentic AI in the Payments Division

Where it adds value, what it costs, and where to start. A twelve-month view for the executive team.

Prepared for the Payments Division executive team • July 2026

Three places agentic AI pays for itself in your Payments Division within 12 months.

We prototyped the flagship, so you can start where the value is, with the risk already worked down.

01

Intelligent reconciliation

Match payments across systems and investigate the breaks.

Low risk • runs in parallel

02

Fraud alert triage **PROTOTYPED**

Investigate every alert and hand analysts a decision-ready case.

Prototyped • where we start

03

AML compliance co-pilot

Draft suspicious-matter reports for an officer to review and lodge.

Highest stakes • last

THE OPPORTUNITY

The bottleneck isn't detection. It's the investigation behind every alert.

30k

Fraud alerts a month

~90%

Are false positives

~12_{min}

Of manual work, per alert

Your fraud engine already raises the alerts. Agentic AI earns its place by automating the human investigation behind each one, not by replacing the detection you already trust. Figures are illustrative of a large Australian bank.

A regulated bank climbs the capability ladder **deliberately**. You earn the right to more autonomy by proving governance at each rung.

The strategic principle. Why we don't ship the most autonomous system first

Intelligent reconciliation and exceptions

TODAY

- Ops staff match payments across core banking, scheme files and the ledger by hand.
- Breaks are worked in spreadsheets, and the same exception types recur every cycle.
- Month-end is a scramble, and the backlog is invisible until it is late.

WITH AN AGENT

- Matches across systems and investigates each break by pulling the underlying records.
- Drafts the correction with its reasoning attached, then queues it for an approver.
- Learns the recurring break types, so the queue shrinks rather than just moving.

VALUE LEVER

Cost-out, back office

RISK RUNG

Low

No customer impact, fully reversible

INDICATIVE ROI

30 to 40% less manual exception handling

Illustrative, to be calibrated in Discovery

Phase 1: one high-volume break type, on one rail, running in shadow mode.

Fraud alert triage

THE ONE WE BUILT

TODAY

- Analysts open every alert the engine raises, roughly 30,000 a month.
- Each one means pulling history, profile and device data from three systems.
- About 12 minutes per alert, and around 90% turn out to be false positives.

WITH AN AGENT

- Investigates each alert proactively, gathering the same context without being asked.
- Assembles a decision-ready case: signals, precedent, and a scored recommendation.
- Stops at a human gate. It recommends, and an analyst decides and is accountable.

VALUE LEVER

Productivity and loss avoidance

RISK RUNG

Medium

Recommends only, never acts

INDICATIVE ROI

~\$2.5M a year of analyst capacity

Illustrative, method on the ROI slide

Phase 1: a six-week shadow pilot on live alerts, measured against analyst decisions.

Real-time compliance and AML co-pilot

TODAY

- Officers review monitoring hits against AML typologies case by case.
- Suspicious-matter reports are written by hand, each one assembled from scratch.
- Backlogs push against AUSTRAC reporting deadlines rather than sitting clear of them.

WITH AN AGENT

- Monitors patterns against typologies and assembles the supporting evidence.
- Drafts the report with citations back to the underlying transactions.
- Routes it to an officer to review and lodge. It never lodges anything itself.

VALUE LEVER

Regulatory-risk reduction

RISK RUNG

High

AUSTRAC-facing, highest stakes

INDICATIVE ROI

Faster turnaround, fewer late lodgements


Illustrative, risk-avoidance not cost-out

Phase 1: drafting assistance only, with the officer reviewing and lodging every report.

Three rungs, twelve months


Normally you would start at the lowest-risk rung. The prototype changes that.

Triage's method is already proven, so you start where the value is bigger without taking more risk. Investment bands are illustrative.



The human gate and audit trail built at triage are reused by both rungs behind it, so the governance cost is paid once.

The triage assistant: four specialist agents, a human gate, and an audit trail



Legend: □ Deterministic (rules, no model) □ Model-driven (LLM) □ Human gate □ Tool □ Audit store

Prototype = orchestrator-workers + human gate. Production swaps stubs for MCP integrations, LangGraph/ADK orchestration, and an eval harness.

We didn't just recommend it. We built it.

Escalate

100 / 100 • high confidence

Likely account takeover

New device, foreign IP, new payee, contact just changed.

- **A human gate.** The agent recommends. Nothing is actioned until a person decides.

Clear

0 / 100 • high confidence

Clean false positive

Known device, home city, small amount above average.

- **A full audit trail.** Every step, tool call and reason logged for a regulator to read.

Monitor

35 / 100 • low confidence

Asks for human judgement

Crypto payee, and precedent genuinely conflicts.

- **Runs on a laptop.** A local model, one dependency, no API key and no data leaving the bank.

Consider it the **starting point for discovery**, not a demo.

Governance is the licence to operate, not the brake

The agent recommends. A person decides

- Under the Financial Accountability Regime, a named executive is personally accountable
- That accountability cannot be delegated to a model

Built for the Australian stack

- An operational-risk fallback under APRA CPS 230
- Where the model runs is a design choice (CPS 234, Privacy)

A regulator-readable audit trail

- Every step, tool call and reason is logged
- It turns a black box into something you can deploy

It lands on a live tailwind

- The 2025 Scams Prevention Framework raises the bar
- Banks are being told to reduce scam losses now

The economics of the fraud-triage rung

~\$2.5M

Analyst capacity recovered per year

~50% less handle time across ~30,000 alerts a month, redeployed to complex cases. Illustrative.

Months

Payback on a six-figure pilot

Pays back on recovered capacity alone, before counting any fraud-loss reduction. Illustrative.

Upside

Loss reduction and fewer disputes

Faster action on irreversible rails, plus fewer complaints reaching AFCA. Illustrative.

All figures are illustrative and would be calibrated in a Discovery. The method is on the appendix slide so you can challenge it.

This isn't theoretical. It is proven in Australian financial services

95%

lower cost per claim

HBF automated health-claims assessment with **agentic AI**, a “Missing Information Agent” that reasons over incomplete claims. Turnaround fell from **weeks to seconds** at **94%+ extraction accuracy, above human reviewers**, with zero disruption to core systems.

HBF • Mantel Group case study

Australia's first agentic AI home loan, shipped in **16 weeks**, with **audit trails and human-override** built into every workflow.

Lendi's “Guardian”: 10+ agents, “compliance by design”. The same human-in-the-loop and audit-trail model this proposal is built on, proven in a regulated Australian lender.

Lendi Group • Mantel Group case study

Start where the proof already exists. A **six-week engagement** that takes the triage prototype into your environment.

Your data, your analysts, shadow mode. Fixed price, and a decision gate at the end

THANK YOU · QUESTIONS WELCOME

The agent investigates and recommends. A human decides and is accountable.

The triage prototype is yours to keep.

It runs on a laptop against a local model. No API key, and no data leaves the bank

Nhial Majok · `agentic-ai-payments-division.vercel.app`