

AI Cost Audit

Meridian Financial Group (Illustrative Sample) - FY26 Q2 - 48-72 hour AI Cost Audit

Prepared by TrustedAIGov(R) AUDIT AI (Port 8004) Generated 2026-06-17T09:55:09+00:00

AUDIT AI is advisory. Savings are estimates, not realised savings. Enforcement is Governance's.

1. Executive Summary

We reviewed USD 197000.0000 of AI spend and identified USD 51600.0000 (26.2%) in estimated cost-saving opportunities across 2 finding(s).

Total AI spend reviewed:	USD 197000.0000 (7080 records)
Estimated recoverable:	USD 51600.0000 (~26.2%, est.)
Optimization findings:	2

Executive Action Summary

What management should do. Across 8 AI asset(s): KEEP 4 SCALE 1 MONITOR 1 REVIEW 1 RETIRE 1. USD 78000.0000 of spend sits in assets flagged for review or retirement (~USD 51600.0000 estimated recoverable).

- [RETIRE] claims_summarizer - Candidate for Retirement (USD 34000.0000 spend, ~USD 34000.0000 recoverable)
 - Abandoned-agent pattern: sustained spend with negligible output (~USD 34000.0000 estimated recoverable).
- [REVIEW] fraud_copilot - Needs Owner Review (USD 44000.0000 spend, ~USD 17600.0000 recoverable)
 - Material spend (USD 44000.0000) with an open finding (~USD 17600.0000 estimated recoverable).
- [SCALE] policy_chatbot - Healthy - Expansion Candidate (USD 58000.0000 spend, ~USD 0.0000 recoverable)
 - Efficient and well-governed at material spend (USD 58000.0000); no findings.
 - Aligned to KPI 'Average Handling Time'. Expand only if the business case supports it (no ROI is asserted here).
- [MONITOR] marketing_gen - Watch (USD 14000.0000 spend, ~USD 0.0000 recoverable)
 - Accountability incomplete - assign owner / sponsor / KPI.
- [KEEP] hr_screener - Healthy (USD 7000.0000 spend, ~USD 0.0000 recoverable)
 - No findings; spend appears aligned to its business process.
 - No AI Economics Profile on file - add one for full accountability.
- [KEEP] uw_assistant - Healthy (USD 26000.0000 spend, ~USD 0.0000 recoverable)
 - No findings; spend appears aligned to its business process.
- [KEEP] dev_copilot - Healthy (USD 9000.0000 spend, ~USD 0.0000 recoverable)
 - No findings; spend appears aligned to its business process.
 - Accountability could be strengthened (add sponsor / KPI).
- [KEEP] recon_bot - Healthy (USD 5000.0000 spend, ~USD 0.0000 recoverable)
 - No findings; spend appears aligned to its business process.
 - Accountability could be strengthened (add sponsor / KPI).

Dispositions are advisory. AUDIT AI advises; Governance enforces; FinOps is the financial source of truth. No ROI is asserted.

2. Total Spend

USD 197000.0000 across 7080 usage records for FY26 Q2.

3. Provider Breakdown

Provider	Spend (USD)
openai	171000.0000
anthropic	26000.0000

4. Model Breakdown

Model	Spend (USD)
gpt-4o-mini-2024-07-18	127000.0000

gpt-4o-2024-08-06	44000.0000
claude-3-5-sonnet-20241022	26000.0000

5. Cost Centre Breakdown

Cost Centre	Spend (USD)
CC-CS-1100	58000.0000
CC-FR-4400	44000.0000
CC-CL-3200	34000.0000
CC-UW-2100	26000.0000
CC-MK-7200	14000.0000
CC-EN-5100	9000.0000
CC-HR-8100	7000.0000
CC-FN-6100	5000.0000

Spend by Business Process

Business Process	Spend (USD)
Customer Service	58000.0000
Fraud Detection	44000.0000
Claims Processing	34000.0000
Underwriting	26000.0000
Marketing Content	14000.0000
Software Engineering	9000.0000
Unmapped	7000.0000
Finance Operations	5000.0000

Accountability - Showback by Owner

Budget Owner	Spend (USD)
VP, Customer Service	58000.0000
Director, Financial Crime	44000.0000
Claims Operations Manager	34000.0000
Underwriting Lead	26000.0000
Unallocated	21000.0000
Engineering Manager	9000.0000
Financial Controller	5000.0000

AI Portfolio (ranked by spend)

- policy_chatbot (Customer Service / VP, Customer Service) USD 58000.0000
- fraud_copilot (Fraud Detection / Director, Financial Crime) USD 44000.0000 [FLAGGED]
- claims_summarizer (Claims Processing / Claims Operations Manager) USD 34000.0000 [FLAGGED]
- uw_assistant (Underwriting / Underwriting Lead) USD 26000.0000
- marketing_gen (Marketing Content / Unallocated) USD 14000.0000
- dev_copilot (Software Engineering / Engineering Manager) USD 9000.0000
- hr_screener (Unmapped / Unallocated) USD 7000.0000
- recon_bot (Finance Operations / Financial Controller) USD 5000.0000

Portfolio Governance Overview

Scores 0-100: Spend (relative size), Risk (worst finding + budget state), Waste (recoverable share), Accountability (profile completeness; higher is better).

Asset - status - action	Sp/Ri/Wa/Ac
claims_summarizer - Candidate for Retirement - RETIRE	59/100/100/50
fraud_copilot - Needs Owner Review - REVIEW	76/100/40/67
policy_chatbot - Healthy - Expansion Candidate - SCALE	100/0/0/100
marketing_gen - Watch - MONITOR	24/0/0/17
hr_screener - Healthy - KEEP	12/0/0/n/a
uw_assistant - Healthy - KEEP	45/0/0/83
dev_copilot - Healthy - KEEP	16/0/0/50
recon_bot - Healthy - KEEP	9/0/0/67

Assets Requiring Review

fraud_copilot (Fraud Detection / Director, Financial Crime) USD 44000.0000 spend, ~USD 17600.0000 recoverable
 - Material spend (USD 44000.0000) with an open finding (~USD 17600.0000 estimated recoverable).

Assets Recommended for Retirement

claims_summarizer (Claims Processing / Claims Operations Manager) USD 34000.0000 spend, ~USD 34000.0000 recoverable
 - Abandoned-agent pattern: sustained spend with negligible output (~USD 34000.0000 estimated recoverable).

Assets Recommended for Expansion

policy_chatbot (Customer Service / VP, Customer Service) USD 58000.0000 spend, ~USD 0.0000 recoverable
 - Efficient and well-governed at material spend (USD 58000.0000); no findings.
 - Aligned to KPI 'Average Handling Time'. Expand only if the business case supports it (no ROI is asserted here).

Decision Accountability

Every REVIEW / RETIRE / SCALE recommendation with a named owner, sponsor, and advisory review-by date. 3 decision(s):
 1 ready, 2 blocked.

[RETIRE] claims_summarizer - due 01-Jul-2026 (Missing Sponsor)

owner Claims Operations Manager sponsor UNASSIGNED status: Blocked - Awaiting Accountability

[REVIEW] fraud_copilot - due 17-Jul-2026 (Missing Sponsor)

owner Director, Financial Crime sponsor UNASSIGNED status: Blocked - Awaiting Accountability

[SCALE] policy_chatbot - due 01-Aug-2026 (Ready for Governance Review)

owner VP, Customer Service sponsor Chief Operating Officer status: Pending Governance Review

Decisions Requiring Executive Review

- policy_chatbot (SCALE, Customer Service) - owner VP, Customer Service, sponsor Chief Operating Officer, review by 01-Aug-2026.

Missing Accountability Information

Asset - action	Missing
claims_summarizer - RETIRE	Executive Sponsor, Business Owner
fraud_copilot - REVIEW	Executive Sponsor

Accountability views are ADVISORY. Owners and sponsors are read from the AI Economics Profile; due dates are suggested SLA targets, not commitments. AUDIT AI does not approve, schedule, or enforce - the Governance plane is the System of Control. No workflow or approval engine; nothing is persisted.

6. Optimization Findings

[CRITICAL] Inactive Resource Leakage (ABANDONED_AGENT) - claims_summarizer (~USD 34000.0000)

Agent 'claims_summarizer' spent USD 34000.0000 producing ~0 output (0.030% out/in). Review or terminate.

[CRITICAL] Oversized Prompt / Context Bloat (CONTEXT_WINDOW_OVERCONSUMPTION) - fraud_copilot (~USD 17600.0000)

Avg prompt 2,200,000 tokens on 'gpt-4o-2024-08-06' (in:out 55:1). Cache/compress static context; ~40% input cost recoverable.

7. Estimated Recoverable Spend

Total estimated recoverable: USD 51600.0000 (~26.2% of reviewed spend).

Finding / subject	Est. (USD)
Inactive Resource Leakage (ABANDONED_AGENT) - claims_summarizer	34000.0000
Oversized Prompt / Context Bloat (CONTEXT_WINDOW_OVERCONSUMPTION) - fraud_copilot	17600.0000

Estimates only. Realised savings are confirmed in the FinOps ledger after adoption.

8. Budget Risk Summary

Budget thresholds were not included in this ledger view. Budget monitoring is provided by AI FinOps (8003).

9. Data Quality Findings

Unallocated spend: USD 0.0000

Records missing attribution: 0

10. Areas Requiring Review

- CRITICAL - ABANDONED_AGENT - claims_summarizer (~USD 34000.0000)

- CRITICAL - CONTEXT_WINDOW_OVERCONSUMPTION - fraud_copilot (~USD 17600.0000)

11. Methodology

- Costs were read from the AI FinOps System of Record (Port 8003) through a read-only contract; AUDIT AI holds no ledger of its own and never modifies financial truth.

- Findings are produced by pattern-based detection over token and cost usage: abandoned agents (spend with negligible output), context-window over-consumption (oversized prompts), and vendor/model-tier misalignment (premium model on trivial work).

- Portfolio dispositions (KEEP/SCALE/MONITOR/REVIEW/RETIRE) and decision accountability (owner, sponsor, due date) are derived only from the FinOps ledger, the AI Economics Profile, and these findings. No ROI, benchmarking, transaction (OD-6) or outcome (OD-7) data is used; due dates are advisory SLA targets, not commitments.

- Savings figures are modelled ESTIMATES of recoverable spend, not realised savings. Realised savings are confirmed in the FinOps ledger after a change is adopted.

- This assessment is advisory: it recommends, it does not act. Enforcement (alerting, blocking, termination), approval, and scheduling are performed only by the Governance plane upon human decision; AUDIT AI runs no workflow or approval engine.

- Cost-per-successful-task and output-quality scoring are intentionally excluded pending a ratified quality-data source (OD-2).