Truveil

AI ACCOUNTABILITY AUDIT REPORT

MarketMind Equity Intelligence Agent

Audit of MarketMind Equity Intelligence Agent

REPORT ID	SR-155
GENERATED	Wed, 02 Sep 2026 21:35:03 GMT
SUBJECT AGENT	MarketMind Equity Intelligence Agent
AGENT CATEGORY	2 (Declared at Registration)
RISK CLASSIFICATION	MEDIUM
FRAMEWORKS APPLIED	Truveil methodology, DPDP Rules 2025

Report Contents

SECTION I

FOR CXO

Auditor Opinion

C

Accountability Grade · **62.8/100**

MEDIUM RISK

This Category 2 investment suggestion agent scores 62.8/100 (Grade C, MEDIUM risk). Transparency is strong; accountability and data trust are partial; reversibility is the weakest dimension. Critical gaps include absent AI disclosure (informed by EU AI Act Article 50), missing groundedness checks (DPDP Rules 2025 Rule 6), and no appeal or correction workflows (DPDP Rules 2025 Rule 14). Immediate remediation is required before production scaling.

TRANSPARENCY	ACCOUNTABILITY	DATA TRUST	REVERSIBILITY
91.7	64.3	55.6	46.7

ISSUED BY

This report has been generated by Truveil, an AI agent accountability platform. The accountability scoring is rule-based and deterministic, calibrated against primary regulatory text. The narrative interpretation in this report is produced by Anthropic Claude under structured constraints maintained by Truveil. No portion of the scoring is delegated to a generative model.

SECTION II

FOR AUDITOR

Scope and Methodology

What this report covers

This report covers a single execution of the agent identified on the cover page. The audit assesses the logged actions, decisions, validations, and human checkpoints recorded during that execution against four accountability dimensions: *Transparency*, *Accountability*, *Data Trust*, and *Reversibility*. Each dimension corresponds to a category of regulatory obligation recurring across the frameworks Truveil applies.

How Truveil scores

Truveil's scoring engine is rule-based and deterministic. The same evidence, scored against the same framework, produces the same result every time. Each accountability signal is anchored to specific provisions in primary regulatory text. Caps and floors apply where a binding obligation is unmet, these are not weighted heuristics but explicit thresholds defined in the underlying regulation.

Narrative interpretation in this report is produced by a language model under structured constraints. The model summarises findings already determined by the engine. It does not score, classify, or rank the agent. Citations in this report are restricted to provisions present in the engine's citation index for this run.

Regulatory anchors for this audit

Truveil methodology, DPDP Rules 2025

What this report does not certify

This report is an evidence-based assessment of one execution of one agent. It does not constitute a regulatory filing, a legal opinion, or a certification of conformity. It is not a substitute for accredited third-party assurance such as ISO/IEC 42001 certification. Truveil's role is to surface accountability gaps that would otherwise remain undocumented, against the regulatory provisions identified in the engine's citation index.

SECTION III

FOR AUDITOR

Agent System Description

Agent and run metadata

AGENT	MarketMind Equity Intelligence Agent
CATEGORY	2 (Declared at Registration)
CATEGORY BASIS	Category 2 declared at agent registration.
RUN ID	20260610_AM
EXECUTION TIMESTAMP	2026-06-09T22:24:34.332Z
TOTAL LOGGED EVENTS	16

Step-by-step execution

> Reconstructed from 16 logged events via the SDK; SDK instrumentation observes events the agent code explicitly logs.

01 Market Research Dispatch

Agent completed market research for AM slot, identifying 20–24 notable stocks across three cycles.

OK

02 Investment Suggestion Generation

Agent generated BUY suggestions for multiple securities with allocation percentages and time horizons.

WATCH

03 Brief Persistence and Human Review Queue

Briefs persisted with status pending, explicitly awaiting human review before execution.

OK

04 User Bookmark Activity

User bookmarked ICICIBANK suggestions across two events; no approval or override action recorded.

WATCH

05 Approval and Override Gate INFERRED

No formal approval event or override capability was logged at any workflow stage.

RISK

SECTION IV

FOR AUDITOR

Findings

Findings index

#	SEVERITY	FINDING	FRAMEWORK
1	HIGH	Missing Reversibility Controls: Override, Appeal, and Correction Workflows	Truveil methodology, informed by EU AI Act Article 14(4)(d) and NIST AI RMF as reference standards
2	HIGH	Absent AI Disclosure to End User	Truveil methodology, informed by EU AI Act Article 50 and NIST AI RMF as reference standards
3	HIGH	Groundedness Check Absent	DPDP Rules 2025 Rule 6, effective 13 May 2027
4	MEDIUM	Model Owner Unnamed and No Approval Event	Truveil methodology, informed by EU AI Act Article 26 and NIST AI RMF as reference standards
5	MEDIUM	Input Validation Not Logged	DPDP Rules 2025 Rule 6, effective 13 May 2027

Findings in full

Finding 1 — Missing Reversibility Controls: Override, Appeal, and Correction Workflows

HIGH

Override controls, appeal channel, and correction workflow are all absent — three required signals for Category 2 agents with no logged remediation path.

FRAMEWORK	Truveil methodology, informed by EU AI Act Article 14(4)(d) and NIST AI RMF as reference standards
EVIDENCE	absent; required for Category 2 agents
AFFECTED STEP	REGISTRATION

Finding 2 Absent AI Disclosure to End User HIGH

No AI disclosure event was logged. Users receiving investment suggestions were not notified of AI involvement, required for Category 2 agents.

FRAMEWORK	Truveil methodology, informed by EU AI Act Article 50 and NIST AI RMF as reference standards
EVIDENCE	absent; required for Category 2 agents
AFFECTED STEP	Step 02: Investment Suggestion Generation

Finding 3 Groundedness Check Absent HIGH

No groundedness signal was logged for any suggestion. This is a required data trust control for Category 2 agents under DPDP Rules 2025 Rule 6.

FRAMEWORK	DPDP Rules 2025 Rule 6, effective 13 May 2027
EVIDENCE	absent; required for Category 2 agents
AFFECTED STEP	Step 02: Investment Suggestion Generation

Finding 4 Model Owner Unnamed and No Approval Event MEDIUM

No model owner or deployer is identified in any log. No approval decision was recorded despite briefs being queued for human review.

FRAMEWORK	Truveil methodology, informed by EU AI Act Article 26 and NIST AI RMF as reference standards
EVIDENCE	no owner or deployer mention in any log detail
AFFECTED STEP	REGISTRATION

Finding 5 Input Validation Not Logged MEDIUM

No validation or verify event was recorded for market research inputs feeding into suggestion generation, weakening data trust.

FRAMEWORK	DPDP Rules 2025 Rule 6, effective 13 May 2027
EVIDENCE	no validation or verify event
AFFECTED STEP	Step 01: Market Research Dispatch

SECTION V

FOR CXO

Recommended Actions

Actions are ordered by priority. *Now* actions remediate findings that breach binding obligations under the frameworks applied. *Soon* actions address material gaps that fall short of best practice without breaching a binding rule. *Later* actions improve accountability posture for forward-looking compliance.

01 **Implement Override, Appeal, and Correction Workflow Logging**

Register and log override controls, an appeal channel, and a correction workflow for all Category 2 suggestion cycles before production release.

NOW

02 **Add User-Facing AI Disclosure Event**

Emit a structured AI disclosure log event and surface a disclosure notice to users before any investment suggestion is presented.

NOW

03 **Integrate Groundedness Verification into Suggestion Pipeline**

Log a groundedness check result for every suggestion generated, anchored to cited source data, aligned to DPDP Rules 2025 Rule 6.

NOW

04 **Register Model Owner and Log Approval Decisions**

Add model owner and deployer to registration metadata; log structured approval events when human reviewers act on pending briefs.

SOON

05 **Add Input Validation and Versioned Output Logging**

Emit validation events after each research cycle and attach version identifiers to all persisted suggestion briefs.

SOON

SECTION VI

FOR ENGINEER

Technical Detail

This section is intended for engineering and platform teams responsible for the agent. Each finding from Section IV is restated with its developer-view structured fields and suggested remediation steps. Affected steps refer to the trace numbering in Section III.

Finding 1 · Missing Reversibility Controls: Override, Appeal, and Correction Workflows

FAILURE TYPE	absent_reversibility_controls
AFFECTED STEP	REGISTRATION
CONFIDENCE	HIGH

SUGGESTED FIXES

1. Implement and log an override control event for every suggestion cycle.
2. Register an appeal channel endpoint and log its availability at agent startup.

Finding 2 · Absent AI Disclosure to End User

FAILURE TYPE	missing_disclosure_event
AFFECTED STEP	Step 02: Investment Suggestion Generation
CONFIDENCE	HIGH

SUGGESTED FIXES

1. Emit an AI disclosure log event before each suggestion is surfaced to the user.
2. Add a user-facing disclosure banner tied to session initiation and log its trigger.

Finding 3 · Groundedness Check Absent

FAILURE TYPE	absent_groundedness_check
AFFECTED STEP	Step 02: Investment Suggestion Generation
CONFIDENCE	HIGH

SUGGESTED FIXES

1. Log a groundedness verification event for each generated suggestion.
2. Validate output claims against sourced data before persisting the brief.

Finding 4 · Model Owner Unnamed and No Approval Event

FAILURE TYPE	missing_owner_and_approval_event
AFFECTED STEP	REGISTRATION
CONFIDENCE	HIGH

SUGGESTED FIXES

1. Add model_owner and deployer fields to agent registration metadata.
2. Log a structured approval event when a human reviewer acts on a pending brief.

Finding 5 · Input Validation Not Logged

FAILURE TYPE	absent_input_validation_event
AFFECTED STEP	Step 01: Market Research Dispatch
CONFIDENCE	MEDIUM

SUGGESTED FIXES

1. Emit a validation event after each MARKET_RESEARCH_COMPLETED with schema check results.
2. Gate suggestion generation on a passed validation flag in the research output.

> *Address all NOW-priority actions before production scaling; reversibility and disclosure gaps pose the greatest regulatory exposure.*