CONFIDENTIAL · FOR REGULATOR & BOARD ONLY

QUARTERLY GOVERNANCE PACK · Q1 2026

MAS AI Risk Management *Evidence Pack.*

A signed, examiner-ready report of **Apex Bank Singapore Pte. Ltd.**'s AI agent estate, governance posture, life-cycle controls, and remediation roadmap — produced for the Monetary Authority of Singapore in accordance with the AI Risk Management Guidelines (consultation paper, November 2025), the Veritas FEAT principles, and the 2024 Information Paper on AI Model Risk Management.

CUSTOMER	PERIOD	REGULATOR	GENERATED
Apex Bank SINGAPORE PTE. LTD.	Q1 *2026* 1 JAN – 31 MAR	*MAS* SINGAPORE · PRIMARY	23 May 2026 · 09:42 SGT

MAPPED TO THE FOLLOWING INSTRUMENTS

- **MAS AIRG (2025)** — Consultation Paper on Guidelines on AI Risk Management for Financial Institutions, 13 Nov 2025; 12-month transition
- **MAS Info Paper (2024)** — AI Model Risk Management, thematic-review good practices, 5 Dec 2024
- **MAS TRM (2021)** — Technology Risk Management Guidelines (relevant excerpts)
- **MAS FEAT (2018)** — Principles to Promote Fairness, Ethics, Accountability & Transparency in AI & Data Analytics
- **MindForge Toolkit (2026)** — Project MindForge AI Risk Management Toolkit, 20 Mar 2026
- **MAS Notice 626 / 656** — AML/CFT, where AI affects in-scope controls

DOCUMENT & EXECUTIVE SUMMARY

At a glance — *posture, evidence, sign-off.*

This pack consolidates the bank's AI agent inventory, materiality assessment, life-cycle control coverage, FEAT scoring, red-team findings, incidents, and remediation roadmap for the period — produced by AgentGuardian and counter-signed by accountable executives.

DOCUMENT	MAS AI Risk Management — Quarterly Governance Pack
CUSTOMER	Apex Bank Singapore Pte. Ltd. (MAS ID: 200X1234A · Locally-incorporated Bank)
REPORTING PERIOD	1 January 2026 — 31 March 2026 (Q1 2026)
PACK ID	APEX-MAS-2026-Q1-001
REGULATOR SCOPE	MAS (primary) · APRA (cross-border subsidiary in Sydney)
GENERATED	23 May 2026 09:42 SGT · AgentGuardian v2.4.1
OWNERS	CRO: Dr. Aishwarya Nair · CISO: Marcus Tan · Head of AI Governance: Lin Wei Ming
GLACIEN ATTESTATION	Glacien Pte. Ltd. · UEN 202548227C · Counter-signed
STATUS	● Examiner-ready · All sign-offs received

AGENTS GOVERNED	T1 CRITICAL	OPEN FINDINGS	POLICY COVERAGE	FEAT COMPOSITE
1,247	38	*7*	96%	92/100
Across 5 BUs, 3 cloud accounts	All HITL gated, all attested	0 unmitigated P1 · all on plan	By MAS AIRG life-cycle area	Weighted across F·E·A·T
▲ +98 vs Q4 2025	— no change	▼ -16 vs Q4 2025	▲ +11 pts QoQ	▲ +4 vs Q4 2025

EXECUTIVE SUMMARY — CRO

The bank's AI agent estate is now *under continuous governance.* No unmitigated P1 findings remain at quarter-end.

Q1 2026 was the first full quarter operating under AgentGuardian. We discovered and registered **1,247 AI agents** across Retail Banking, Corporate Banking, Wealth, Treasury, and Operations — including **62 shadow agents** which have since been assigned owners, risk-tiered, and brought under policy. Of **23 findings open at the start of the period**, 16 are now closed and 7 remain on plan with named owners and Q2 deadlines.

Against the proposed MAS AIRG Guidelines, the bank's posture across the three pillars (Elevated Oversight, Structured Risk Systems, Life-Cycle Controls) is summarised *green*; against FEAT, the composite of **92/100** reflects strengthening in Accountability (98) and Fairness (94), with Transparency (87) as the standing improvement priority. The bank is on track to meet the 12-month AIRG transition window with margin.

SECTION 01 · MAS AIRG PILLAR 1

Elevated *oversight & governance.*

Board and Senior Management oversight, AI Risk Committee composition, policy framework, and FEAT principles attestation — mapped to MAS AIRG §2 (Oversight) and the Dec 2024 Information Paper on AI MRM.

Board oversight checklist

All met

CONTROL	STATUS	EVIDENCE
AI strategy tabled to Board, Q1 2026	Approved	BMR-2026-02-14
Risk appetite for AI use approved	Approved	RAC-2026-01-22
Materiality framework ratified	Ratified	RAC-2026-01-22
AI Risk Committee charter in force	In force	CHRT-AIRC-v3
Quarterly AI risk readout to Board	Delivered	BMR-2026-03-28
Independent assurance arranged for FY26	Engaged	PROC-AUDIT-26

AI Risk Committee — Q1 attendance

100% quorum

MEMBER	ROLE	ATTEND.
Dr. A. Nair	CRO · Chair	3 / 3
M. Tan	CISO	3 / 3
L. W. Ming	Head of AI Governance	3 / 3
J. Pereira	Head of Compliance	3 / 3
R. Kumaran	Head of Internal Audit	3 / 3
S. Yong	Head of Data · CDO	2 / 3
(External) Glacien	Co-delivery · Advisory	3 / 3

Policy framework coverage

Twelve policies form the bank's AI governance framework. All are in force; quarterly review schedule maintained.

POLICY	SCOPE (PER MAS AIRG)	OWNER	LAST REVIEWED	NEXT REVIEW	STATUS
AI Use Policy	Permitted & prohibited use cases	CRO	2026-01-15	2026-07-15	In force
AI Data Management Standard	Data sourcing, lineage, retention	CDO	2026-02-22	2026-08-22	In force
AI Model Risk Management Std.	Validation, monitoring, change control	Head of AI Gov.	2026-03-04	2026-09-04	In force
AI Human Oversight Standard	HITL gating by materiality tier	Head of AI Gov.	2026-02-11	2026-08-11	In force
Agentic AI Controls Annex	Autonomy, tool use, multi-agent flows	CISO & AI Gov.	2026-03-19	2026-06-19	In force

SECTION 02 · MAS AIRG PILLAR 2

AI inventory & *materiality assessment.*

Continuous, cross-cloud discovery feeding a single inventory of record. Every agent assessed across the three MAS dimensions — Impact, Complexity, Reliance — and assigned a materiality tier (T1 critical to T4 low).

Risk distribution — 1,247 agents

Q1 2026

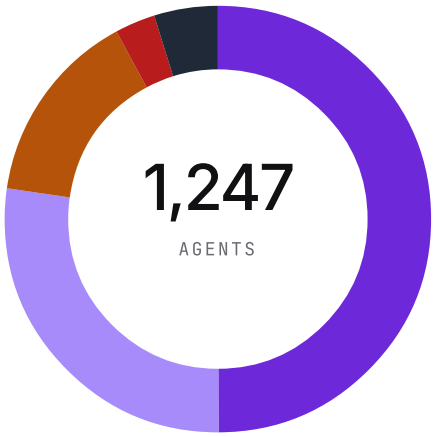

T1 Critical	38	3.0%
T2 High	184	14.8%
T3 Medium	622	49.9%
T4 Low	341	27.3%
Shadow (newly registered)	62	5.0%

3-axis materiality

MAS AIRG §3.2

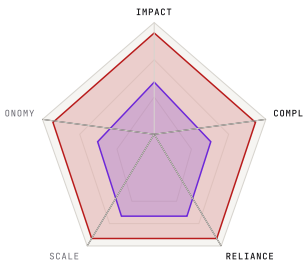

T1 envelope

T3 envelope

Inventory by business unit

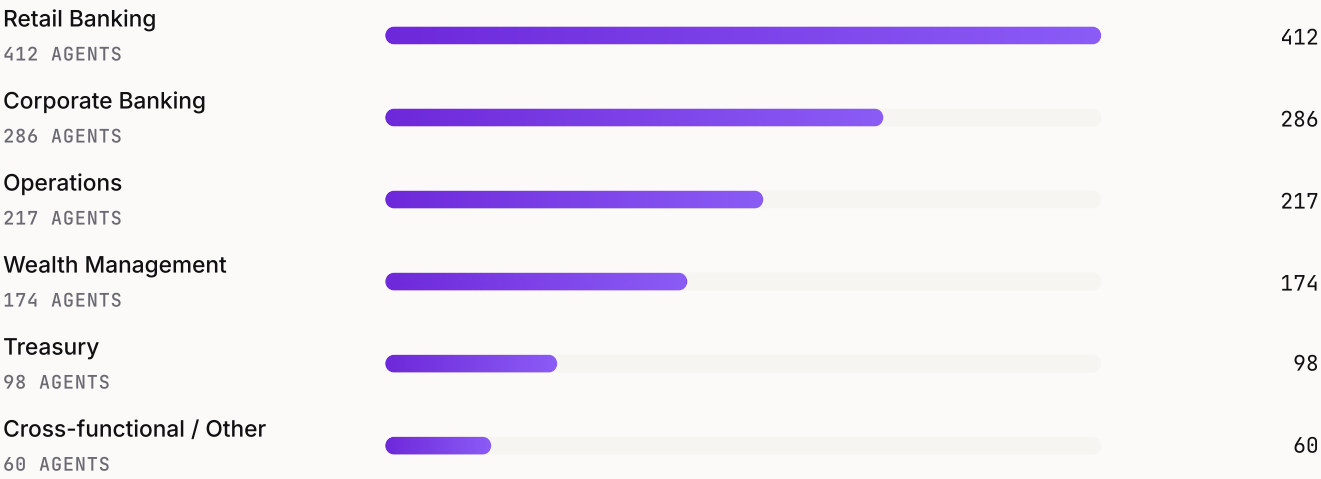

SECTION 03 · MAS AIRG PILLAR 3

Life-cycle controls — *five control families.*

Per MAS AIRG §4, applied proportionately to materiality. Coverage is the percentage of agents within scope that have the corresponding control evidenced and tested in-period.

Data Management AIRG §4.1	Provenance, quality, lineage, retention, sensitive-data handling, training-data attribution. Veritas Toolkit v2 fairness checks applied at ingest for T1/T2.	98% COVERAGE	Green
Fairness AIRG §4.2 · FEAT-F	Disparate-impact testing, demographic parity, equal-opportunity gates. Quarterly bias scan on credit, claims, wealth recommendations. Veritas fairness methodology.	94% COVERAGE	Green
Transparency & Explainability AIRG §4.3 · FEAT-T	Model cards, decision rationale capture, customer-facing explainability for T1/T2 use cases. Standard Chartered/HSBC/Truera transparency methodology applied.	87% COVERAGE	Amber
Human Oversight AIRG §4.4	HITL gating by materiality tier. T1 100% HITL, T2 87%, T3 42%, T4 0%. Reviewer training current. Override audit trail captured per decision.	96% WEIGHTED	Green
Third-Party AI Management AIRG §4.5	12 third-party AI vendors in inventory. 10 attested for FY26 (SOC 2 / ISO 42001 / FEAT alignment). 2 in re-review following capability change. No critical-risk vendor unattested.	10/12 ATTESTED	Amber

Life-cycle coverage by business unit

Heatmap of life-cycle control coverage by business unit. Cells show coverage band (green ≥ 90%, amber 70–89%, red < 70%) and the underlying control area.

BUSINESS UNIT	DATA MGMT	FAIRNESS	TRANSPARENCY	HUMAN OVERSIGHT	3RD-PARTY AI
Retail Banking 412 AGENTS	99%	96%	91%	98%	100%
Corporate Banking 286 AGENTS	98%	94%	82%	97%	90%
Operations 217 AGENTS	98%	95%	90%	97%	100%
Wealth Management 174 AGENTS	99%	90%	79%	100%	92%
Treasury 98 AGENTS	96%	93%	92%	94%	100%
Cross-functional 60 AGENTS	95%	87%	85%	93%	88%

SECTION 04 · FEAT & AGENTIC AI

FEAT alignment & *agentic AI controls.*

Scoring against the four Veritas FEAT principles (Fairness, Ethics, Accountability, Transparency) and the agentic-AI-specific controls from the MAS-led MindForge AI Risk Management Toolkit (March 2026).

VERITAS · FEAT

Fairness

F

94 / 100

Quarterly fairness scan on T1/T2 use cases. Demographic parity within tolerance for all credit and claims agents. Two amber items in Wealth Management on plan.

VERITAS · FEAT

Ethics

E

91 / 100

AI Use Policy in force. Prohibited-use register maintained. No customer-facing autonomous decisioning above risk-appetite threshold. Ethics committee escalations: 4 in Q1, all resolved.

VERITAS · FEAT

Accountability

A

98 / 100

Every agent has named owner, named approver, named control function. Decision audit trail captured for 100% of T1/T2 agent actions. RACI in force across all 5 BUs.

VERITAS · FEAT

Transparency

T

87 / 100

Model cards complete for T1 (100%), T2 (92%). Customer-facing explainability remediation in flight for two Wealth use cases. Standing Q2 priority.

Agentic AI-specific controls MAS MINDFORGE TOOLKIT · 20 MAR 2026

AGENTIC RISK	CONTROL IN PLACE	COVERAGE	STATUS
Tool / system access (autonomy risk)	Cedar policy enforced at AgentCore Gateway. Per-agent OAuth scoping. On-behalf-of identity.	100%	Green
Memory poisoning	AgentCore Memory provenance + rollback. Anti-poisoning detector on long-term memory writes.	98%	Green
Prompt injection	Bedrock Guardrails + AgentGuardian injection classifier. Continuous red-team battery.	100%	Green
Multi-agent collusion / cascade	A2A flow tracing. Cross-agent policy mediator. Cap on chained autonomous actions per session.	92%	Amber
Hallucination / unsafe output	Continuous LLM-as-judge evaluations. Output safety monitor with auto-quarantine on high-risk verdicts.	95%	Green
Sensitive-data exfiltration & agent identity	PII redaction at boundary. Egress monitor. NHI governance with per-agent signing certificate and provenance chain.	100%	Green

SECTION 05 · FINDINGS & RED-TEAM

Findings, red-team results, & *top-risk agents.*

Open findings as at 31 March 2026, with owner, severity, MAS area, and remediation deadline. All findings tracked to closure in AgentGuardian and reported into the quarterly Board AI risk readout.

#	AGENT	FINDING	SEV	TIER	MAS AREA	OWNER	DUE	STATUS
F-01	credit-decision-copilot	PII over-exposure in trace export	P1	T1	Data Mgmt §4.1	S. Yong	2026-04-30	Remediating
F-02	contract-review-agent	Tool hijack risk on unsanctioned MCP server	P2	T2	Agentic Annex	M. Tan	2026-05-31	Remediating
F-03	wealth-advisor-agent	Bias drift on age cohort > tolerance	P2	T2	Fairness §4.2	L. W. Ming	2026-05-15	Open · on plan
F-04	support-copilot-v3	Long-term memory anti-poison detector miss-rate	P2	T3	Agentic Annex	M. Tan	2026-06-15	Open · on plan
F-05	mortgage-calc-bot	Customer-facing explainability gap	P3	T2	Transparency §4.3	L. W. Ming	2026-06-30	Open · on plan
F-06	fraud-pattern-agent	Chained autonomous-action cap exceeded once	P3	T2	Agentic Annex	CISO	2026-04-22	Remediating
F-07	kyc-doc-extractor	Third-party AI vendor SOC 2 re-attestation pending	P3	T3	3rd-Party AI §4.5	Head of VRM	2026-05-10	Open · on plan

Red-team battery — Q1 2026 (top-38 T1/T2 agents) Pass

PROMPT INJECTION	MEMORY POISON	TOOL HIJACK	DATA LEAK	AUTONOMY CAP	BIAS / FAIRNESS
100%	98/96%	100%	100%	100/97%	100/94%
DETECT · BLOCK	DETECT · BLOCK	DETECT · BLOCK	DETECT · BLOCK	DETECT · BLOCK	DETECT · BLOCK

Frameworks: InjecAgent · MINJA · OWASP ASI · MITRE ATLAS v5.4. Top-risk agents: credit-decision-copilot (T1, 9.4), claim-triage-bot (T1, 9.1), wealth-advisor-agent (T2, 8.7), contract-review-agent (T2, 8.3), aml-screening-helper (T1, 8.1).

SECTION 06 · INCIDENTS & ROADMAP

Incidents this quarter, *remediation next quarter.*

Q1 2026 incident timeline (all contained, no customer impact, no regulator-reportable event) and the committed 90-day remediation roadmap closing all open findings on plan.

Q1 2026 incident timeline

2026-01-14	**INC-2026-014 · wealth-advisor-agent** — bias drift on age cohort detected by Veritas fairness scan. Auto-quarantined, model retrained, returned to production with monitoring tightened.	MTTR 4h · Closed
2026-02-09	**INC-2026-031 · claim-triage-bot** — prompt-injection attempt via inbound document. Blocked at Guardrail. No data exposure. Not customer-facing. Preventive, not reactive.	MTTR 0h · Closed
2026-03-22	**INC-2026-058 · corporate-loan-helper** — tool misuse: agent attempted to access an unsanctioned MCP server. Cedar policy denied. Owner notified, tool inventory tightened.	MTTR 2h · Closed

90-day remediation roadmap — Q2 2026

WORKSTREAM	APR 2026	MAY 2026	JUN 2026	OWNER
PII trace redaction F-01 · §4.1	■			CDO
Cedar tool allow-list tightening F-02, F-06 · AGENTIC	■	■		CISO
Wealth fairness re-train F-03 · §4.2	▫	■		Head AI Gov.
Memory anti-poison upgrade F-04 · AGENTIC	▫	■	■	CISO
Customer explainability v2 F-05 · §4.3		■	■	Head AI Gov.
3rd-party AI re-attestation F-07 · §4.5	■	■		VRM

SOLID — in execution · FAINT — preparatory / scoping

SECTION 07 · SIGN-OFF

Attestation & *sign-off.*

The undersigned accountable executives confirm that the information in this pack reflects, to the best of their knowledge, the state of the bank's AI agent estate as at 31 March 2026, and that the governance framework operates in alignment with MAS AIRG, FEAT, and the MAS Information Paper on AI MRM.

ACCOUNTABLE · 1ST LINE

Dr. Aishwarya Nair

Chief Risk Officer · Apex Bank Singapore Pte. Ltd.

✓ SIGNED · 23 MAY 2026

SIGNATURE ON FILE · AGENTGUARDIAN SIGCHAIN

ACCOUNTABLE · 2ND LINE

Marcus Tan

Chief Information Security Officer · Apex Bank Singapore Pte. Ltd.

✓ SIGNED · 23 MAY 2026

SIGNATURE ON FILE · AGENTGUARDIAN SIGCHAIN

INDEPENDENT · GLACIEN

Glacien Pte. Ltd.

AgentGuardian platform attestor · UEN 202548227C

✓ COUNTER-SIGNED · 23 MAY 2026

GLACIEN SIGCHAIN · PLATFORM-SIGNED

TAMPER-EVIDENT ATTESTATION

This pack is signed and tamper-evident. Any modification after the timestamp below invalidates the signature. AgentGuardian produces the underlying signature chain (sigchain) from the bank's AWS account, with customer-managed KMS keys; Glacien holds an independent counter-signing key for platform attestation.

PACK HASH (SHA-256)

`4f6a 7e2c 0b8d 9143 d51e 0bf8 7a91 2c4d 3f70 91ee aa12 5f04 e88d 4c0a 31b6 9c2f`

TIMESTAMP AUTHORITY

RFC 3161 · DigiCert TSA · 2026-05-23T09:42:18+08:00

VERIFICATION

verify.glacien.ai/pack/APEX-MAS-2026-Q1-001

ABOUT THIS DOCUMENT

AgentGuardian quarterly governance packs are produced for every regulated entity using AgentGuardian as their AI agent governance platform. Packs are jurisdiction-specific (this pack: MAS Singapore) and reflect the entity's AI agent estate, materiality, life-cycle control coverage, and remediation status. They are designed to satisfy regulator and board reporting expectations without requiring manual collation.

ABOUT GLACIEN

Glacien Pte. Ltd. is an AWS Select Partner with Agentic AI specialisation, Singapore-headquartered, supporting regulated enterprises across APAC. AgentGuardian is Glacien's AI Agent Governance and Evidence Platform, built natively on Amazon Bedrock AgentCore. info@glacien.ai · +65 8737 0905.