SOUTHERN BANK SINGAPORE (FICTIONAL)

AI risk management evidence report

Against the MAS Guidelines on AI Risk Management

Institution	Southern Bank Singapore (fictional), bank, Singapore
Reporting period	1 Jul – 11 Sep 2026 (Jul–Sep 2026 to date)
Figures as of	11 Sep 2026, 09:41 SGT
Framework	MAS Guidelines on AI Risk Management, issued 7 Oct 2026 · due 7 Oct 2027 and 7 Oct 2028 · Notice 658 · Notice FSM-N05 · Fair Dealing Guidelines
Prepared with	Winsen Rho

READ THIS FIRST

- **This is a sample.** Southern Bank Singapore, its people, vendors, agents and every figure in this report are fictional. No real institution is described.
- **The Guidelines are guidance, not law.** MAS issued them on 7 October 2026. It expects oversight, identification, inventory and materiality by 7 October 2027, and lifecycle controls and capabilities by 7 October 2028, with high-risk uses sooner. Notices 658 and FSM-N05 and the PDPA bind; FEAT is voluntary.
- **A readiness view.** The evidence is Southern's board pack of 11 Sep 2026. It pre-dates the Guidelines and is read against the expectations that start in 2027 and 2028. Nothing here is a view on whether Southern meets the Guidelines.
- **No certification.** This report maps the bank's evidence to the Guidelines. It does not certify, validate or attest compliance, and neither does Winsen Rho. The mapping is the bank's own and the bank remains accountable.

CONTENTS

About this report

A record of what Southern Bank Singapore can show, from its own signed evidence, against each expectation in the MAS Guidelines on AI Risk Management. Where the evidence is incomplete, or where only the bank can answer, the report says so and names who owns the fix.

1	Summary for the Board Risk Committee	Estate, position, what to know, actions
2	Scope, basis and method	What applies and when, how evidence is drawn, status rules
3	What MAS expects	The expectations, the two dates, and how each is evidenced
4	Expectation by expectation	21 blocks across sections 3 to 6, plus notices and customers
5	Open gaps and remediation plan	Every gap and partial, with owner, age and the Guidelines paragraph
6	Annex	Vendor file, customers, evidence index, findings, glossary, sign-off

How to read it

Each expectation block gets one status. It is the weakest among the controls and requirements the bank maps to it; where Rho holds no record that answers an expectation, it is For the bank. One open item is enough to mark a block partial or gap, so read the status with the evidence beneath it.

Every figure carries its period and as-of date, and resolves to a record ID in the evidence index (Annex). Ages are days since the item was opened, counted to 11 Sep 2026. Paragraph numbers are the final Guidelines'.

STATUS	MEANING
● Evidence in place	Every mapped control has current evidence.
● Partial	Evidence exists but is incomplete for part of the estate, a use case or the period.
● Gap	A mapped control has no evidence. It has an owner and an age.
● For the bank	MAS expects it; the evidence or the judgement is the bank's own.

> Language used throughout: "MAS expects" and "should" for the Guidelines; "requires" only for binding notices and law; "association" or "coincides with", never a claim that one thing produced another. Rho supplies evidence. The bank reads it, decides, and remains accountable.

SECTION 1

Summary for the Board Risk Committee

Southern Bank Singapore runs 187 AI agents, 41 of them high-risk. MAS issued its final AI risk management Guidelines on 7 October 2026. Read against them, 1 of the 21 expectation blocks have evidence in place, 10 are partial, 5 have at least one gap and 5 rest with the bank. Every gap has a named owner.

ESTATE AT A GLANCE · AS OF 11 SEP 2026

187	41	9	14
AI agents in the register	**rated high-risk**	**vendor arrangements**	**open findings**
A 62 · B 41 · C1 71 · C2 13 +11 since 30 Jun (176)	183 of 187 observed by Rho 4 registered, never connected	2 material · 5 non-material 1 licensed · 1 under review	4 high · 6 medium · 4 low all owned · oldest 102 days

Classes: A built by Southern · B vendor technology run by Southern · C1 vendor-operated AI service · C2 AI embedded in a SaaS product. Source: register export inventory_2026-09-11.csv; pk_board_q3.

Position against the MAS Guidelines

21 expectation blocks across sections 3 to 6, read against the final Guidelines



1	10	5	5
Evidence in place	Partial	Gap	For the bank

5 blocks rest with the bank: Rho holds no record that answers them (section 3 and the gap list in section 4). Section 2 of the Guidelines, the basic route, does not apply: Southern's estate is material.

Control obligations evidenced MAS Guidelines, Notice 658, PDPA	**62 of 72 (86%)**
Control gaps · partials	**6 · 4**
Supervisor questions answerable internal question set v3, 118 questions	**97 of 118 (82%)**
Open findings with a named owner	**14 of 14**

OVERSIGHT · S3	AIRG 3.4 Partial	AIRG 3.5 Partial	AIRG 3.6 Bank
IDENTIFY · S4	AIRG 4.2 Partial	AIRG 4.5 Partial	AIRG 4.10 Partial
LIFECYCLE · S5	AIRG 5.3 Gap	AIRG 5.4 Bank	AIRG 5.5 Gap
	AIRG 5.7 Partial	AIRG 5.9 Partial	AIRG 5.10 Gap
	AIRG 5.14 In place	AIRG 5.16 Bank	AIRG 5.19 Partial
	AIRG 5.23 Partial	AIRG 5.24 Gap	AIRG 5.25 Gap
	AIRG 5.26 Bank		
CAPABILITIES · S6	AIRG 6.1 Partial	AIRG 6.3 Bank	

Section 1 (scope) and section 2 (proportionality) set who and how much; they are read in section 2 of this report. Dates: sections 3 and 4 by 7 Oct 2027; sections 5 and 6 by 7 Oct 2028; high-risk uses sooner.

Three things the Committee should know

a **Vendor X claims 94% task success for the agent. Measured on Southern's own customers, it is 71%.**

61,131 of 86,100 calls, 21 Jan – 31 Mar 2026, outcomes matured. Matched against Southern's own collectors on 8,412 accounts, the agent is cheaper per good outcome (S$5.09 against S$8.14) but keeps fewer promises (61% against 68%) and draws more complaints (14.2 against 9.1 per 10,000 contacts).

f_002 · Grace Tan · 41 d · to be raised at the Vendor X renewal, 31 Jan 2027

b **Vendor X changed the model on 14 Aug without the notice its contract requires, and conversion has since fallen below the floor.**

Version 4.1.8 became 4.2.1 at 01:30; the nightly probe flagged it at 03:00. Heartland promise-to-pay conversion fell from 0.51 to 0.44, below the pre-registered floor of 0.45. The fall coincides with the version change: on the same 2,000 cases the version is associated with 6 of the 7 points and customer mix with 1 (inv_0418). Recommend-only was proposed on 5 Sep and has waited 6 days for the owner.

f_001 · Siti Aminah · 28 d · f_005 · Grace Tan · 6 d

c **6 blocks fall due in 2027. None has full evidence yet, and none has a gap.**

Oversight and the inventory are the first dates MAS set. Southern evidences board reporting and escalation, and its estate is largely observed (183 of 187 agents), but 12 AI features inside SaaS products have no named owner, 14 purpose records are incomplete and the senior manager accountable for AI oversight is not yet designated. The harder expectations, third-party AI and change, fall due in 2028 but MAS expects high-risk uses to apply them sooner.

f_004 · Kenneth Ong · 40 d · f_013 · Mei Ling Wong · 51 d · AIRG 3.5, 4.2, 4.5

Headline actions

#	ACTION	OWNER	MAS AIRG
1	Decide the recommend-only proposal for the Collections Voice Agent (proposed 5 Sep).	Grace Tan	AIRG 5.23, 3.5(e)
2	Designate the senior manager accountable for AI oversight, and record the Board risk appetite for AI.	Board	AIRG 3.4, 2.5(a)
3	Name owners for the 12 unowned AI features inside SaaS products, and complete the 14 missing purpose records.	Kenneth Ong · Mei Ling Wong	AIRG 4.2, 4.5
4	Connect the 4 silent agents, and drill the kill switch on the Mortgage Pre-qualification Assistant.	Kenneth Ong · Wei Jie Lim	AIRG 5.3, 4.5
5	Re-validate the Collections Voice Agent on version 4.2.1 and complete the 2 overdue independent validations, due 30 Sep.	Mei Ling Wong	AIRG 5.19, 5.24
6	Obtain Vendor X's §8.4 reply and sub-contractor attestation; carry f_001, f_002 and f_009 into the 31 Jan 2027 renewal.	Siti Aminah	AIRG 5.10, 5.11, 5.25
7	Record customer AI disclosure for the CRM Next-Best-Action feature and Transaction Fraud Screening.	Grace Tan	AIRG 5.5
8	Record the reportability assessment for inc_0045 under FSM-N05.	Mei Ling Wong	FSM-N05

SECTION 2

Scope, basis and method

The Guidelines apply to every financial institution and every form of AI, including generative AI, agents and pilots. They are guidance MAS expects institutions to follow. This report reads Southern's evidence against them, section by section, ahead of the 7 October 2027 and 7 October 2028 dates.

2.1 Who and what the Guidelines cover

QUESTION	WHAT THE GUIDELINES SAY	PARAS
Who	Financial institutions as defined in the Financial Services and Markets Act 2022 s.2, including Singapore branches and subsidiaries of foreign groups. Locally incorporated institutions under MAS consolidated supervision, or that own critical information infrastructure, are covered on a group basis.	1.2, 1.5
What AI	AI that learns from data or premises: machine learning, generative AI, AI agents and hybrid systems, pilots included. Outputs from fixed programmed logic or rules alone are outside.	1.3
How much	Scaled to size, AI use and risk. An institution whose AI is unlikely to materially harm itself, its customers or others if it performs badly or is unavailable may follow the lighter basic-governance route. The test is for the institution, not tool by tool.	2.1–2.5
When	Sections 2 to 4 (oversight, identification, inventory, materiality, the basic route) by 7 October 2027. Sections 5 and 6 (lifecycle controls, capabilities) by 7 October 2028. For high-risk uses MAS expects appropriate lifecycle controls as soon as possible.	1.8
Force	Guidance. MAS "expects" and institutions "should". It sits beside binding law and supervisory material and does not replace them.	1.1–1.2

Southern is a material AI user: 41 of 187 agents are high-risk and the Collections Voice Agent is a high-risk, third-party use case. The basic route is therefore not available and the full approach in sections 3 to 6 applies. The sample is read as at 11 Sep 2026; MAS issued the Guidelines on 7 Oct 2026, so this is a readiness view.

2.2 What the report is read against

INSTRUMENT	STATUS	WHAT IT ASKS	IN THIS REPORT
MAS Guidelines on AI Risk Management	Guidance · issued 7 Oct 2026 · phased 2027 and 2028	Oversight and accountability; AI identification, inventory and materiality; lifecycle controls scaled to risk; capabilities and capacity.	Sections 3 and 4; the expectation blocks
MAS Notice 658 and Guidelines on Outsourcing (Banks)	Binding notice, effective 11 Dec 2024	A register of outsourced relevant services; the service provider notifies the bank of sub-contracting; adverse developments.	Section 4 and Annex B
Notice FSM-N05 and Circular MAS/TCRS/2025/08	Binding notice; circular from 1 Feb 2026	A relevant technology incident notified within 1 hour of discovery; para 8 report within 14 days.	Section 4; only if the bank assesses an event as reportable
MAS FEAT principles and Veritas	Principles, 2018; voluntary	Fairness, ethics, accountability and transparency in AI and data analytics. The Guidelines preserve FEAT and allow reference to it when defining fair outcomes.	Fairness (AIRG 5.7)
Fair Dealing Guidelines; **ABS Code of Consumer Banking Practice**	Guidelines (revised May 2024); industry code (Oct 2025)	MI on complaints; clear information; debt recovery conduct (s.15); complaints process (s.16).	Section 4, customers

INSTRUMENT	STATUS	WHAT IT ASKS	IN THIS REPORT
PDPA 2012; Banking Act s.47	Binding	Protection (s.24), retention limitation (s.25), transfer limitation (s.26); banking secrecy.	Section 4, customers

Context only, not binding: the MAS Information Paper on AI model risk management (Dec 2024), the IMDA Model AI Governance Framework for Agentic AI, MindForge and the proposed Third-Party Risk Management Guidelines (consultation P004-2026, not issued). The ~1% human sample, the 08:00–21:00 calling window and the ±8-point tolerance are Southern policy, not MAS thresholds.

2.3 How evidence is drawn

1. **Recorded where the agent acts.** The Rho Bridge runs inside Southern's network and records each decision: inputs, tools used, the authority in force and the action taken. Identifiers are tokenised in the bank before anything leaves.
2. **Outcomes from the bank's own systems.** Results come from Southern's loan management system, collections dispositions, CRM, complaints register, bureau snapshots and card disputes, not from vendor reports.
3. **Every figure is a counted rate.** Each carries numerator, denominator, cohort, period and as-of date.
4. **Signed and reproducible.** Evidence packs are signed by the named second-line officer, object-locked for 7 years from period end, and can be regenerated from stored inputs to show the same figures.
5. **Nothing is overwritten.** When a figure matures or is restated, both views stay signed and retrievable.

WHAT RHO DOES AND DOES NOT DO

Rho does	Inventory and discovery; decision records; outcome tracking; vendor scorecards; human-oversight records; alerts; recording control instructions for the bank's gateway; signed evidence packs.
Rho does not	Validate models, audit, certify or file with a regulator. Validation stays with the bank and parties independent of it. The bank's gateway enforces any pause or restriction; Rho records the instruction.
The bank does	Maps controls to the MAS Guidelines, the notices and the Fair Dealing Guidelines, decides on every finding, signs every pack and remains accountable for every agent, including vendor agents.

2.4 How a status is set

● Evidence in place	Every control Southern maps to the expectation has current evidence in Rho for this estate or use case. Accepted exceptions are stated in the text.
● Partial	Evidence exists, but at least one mapped control is incomplete for part of the estate, a customer segment, the third-party arrangement or the period. The shortfall has an owner.
● Gap	At least one mapped control has no evidence for part of the estate or the vendor. The gap has a named owner and an age.
● For the bank	An expectation, or an assessment, that only the bank can evidence or make. Either Rho holds no record that answers it (general data governance, system security, development documentation, secure retirement, resourcing), or the judgement is the bank's (materiality, whether an event is a relevant incident under FSM-N05). Rho shows what it holds and never pre-fills the bank's answer.

Source of mapping: Southern's AI risk policy, control framework (C-xx) and requirement register (R-xx), maintained by Southern compliance. A block's status is the weakest among the items mapped to it. Where Rho holds no record that answers an expectation, the status is For the bank and the block says what is missing.

SECTION 3

What MAS expects

The Guidelines expect institutions to know where they use AI, justify trust in each use in proportion to its risk, and stay accountable when a vendor supplies the AI. Seven expectations carry that, in two phases.

3.1 The dates

DATE	WHAT MAS EXPECTS BY THEN	SOUTHERN TODAY, AS OF 11 SEP 2026
7 Oct 2027	Oversight and accountability (section 3); AI identification, the AI inventory and materiality assessment (section 4); the basic-governance route for institutions that qualify (section 2).	6 blocks. In place 0 · Partial 5 · Gap 0 · For the bank 1.
7 Oct 2028	Lifecycle controls scaled to materiality (section 5): contingency, data, transparency, fairness, human oversight, third-party AI, testing, validation, monitoring, revalidation, change; and capabilities and capacity (section 6).	15 blocks. In place 1 · Partial 5 · Gap 5 · For the bank 4.
As soon as possible	For high-risk use cases, appropriate lifecycle controls without waiting for the 2028 date (Response 13.5 to 13.6). Southern has 41 high-risk agents; the Collections Voice Agent is one.	MAS gives no numeric period. The report treats high-risk uses as due now.

3.2 The seven expectations

EXPECTATION	IN SHORT	HOW SOUTHERN EVIDENCES IT	SEE
1 · Accountable oversight	The Board and senior management own AI risk, set appetite, and escalate.	The Board Risk Committee receives the AI estate pack each quarter (pk_board_q3, signed 11 Sep: agents 176 to 187, open findings 9 to 14, readiness 79% to 82%). The red breach was escalated by the ladder on 5 Sep. 14 of 14 open findings have a named owner. Gap: no senior manager is yet designated for AI oversight.	3.4–3.6
2 · Know your AI	Identify every AI use, keep an accurate inventory, rate its materiality.	187 agents are registered; 183 are observed; discovery found 196 against 187 declared. Every record carries purpose, scope, model, vendor, dependencies, lifecycle status, risk class, validation status and owner. Gaps: 12 AI features inside SaaS products without an owner, 4 agents that never connected, 14 incomplete purpose records.	4.2–4.13
3 · Controls in proportion to risk	Lifecycle controls follow materiality; high-risk uses get the full set.	Risk class A, B, C1 or C2 is recorded for every agent (62, 41, 71, 13) and approved by the second line. The Collections Voice Agent is class C1, high-risk and material: it carries pre-registered thresholds, a signed authority record, an independent validation and drilled stops.	5.1–5.2
4 · Trust justified for each use	Test before deployment; explain; check fairness; keep humans able to intervene.	Expectations and thresholds were signed by the owner, the vendor and the second line before the first production call. Humans review a stratified sample (470 of 39,140 calls, 1.2%); overrides were right on 80.71%. Fairness is tracked by cohort against a ±8-point tolerance: heartland cohorts sit outside it since 16 Aug. Gap: no AI disclosure record for 2 customer touchpoints.	5.5–5.9, 5.14

EXPECTATION	IN SHORT	HOW SOUTHERN EVIDENCES IT	SEE
5 · Reliance within appetite	Test vendor AI on your own data; manage change notices and concentration; plan for withdrawal.	Vendor X claims 94% task success; on Southern's own customers it is 71%. frontier-a sits under 18 of 41 high-risk agents with no shared fallback. Vendor X changed its model on 14 Aug without the notice its contract requires (f_001). Gaps: three AI clauses missing from the contract; a speech sub-contractor not disclosed (f_009).	5.10–5.11, 5.25
6 · Oversight that adapts	Monitor deployed AI against thresholds; handle incidents; revalidate high-risk AI; control change.	Nightly probes caught the 14 Aug change 1.5 h after it happened. Tiered warnings fired at amber 21 Aug and red 5 Sep. Median time to detect an incident is 0.4 h (31 h before Rho). Gaps: version 4.2.1 went live without re-approval or re-validation; 2 independent validations overdue.	5.19, 5.23–5.25
7 · Capability	People who build, deploy and use AI are competent; resources are adequate.	Competency is recorded for every reviewer of a high-risk agent: 3 of 4 current on the Collections Voice Agent; 4 reviewers across high-risk agents lack it. Resourcing and infrastructure adequacy are the bank's, not Rho's.	6.1–6.3

3.3 Where the Guidelines go beyond what Rho evidences

MAS also expects general data governance (5.4), system and technology security and security testing (5.16, 5.22), development documentation that allows replication (5.17), reviewer reporting to an approval body and its follow-through (5.21), logging of access to AI assets (5.23(d)), secure retirement (5.26), resourcing and infrastructure (6.1, 6.3), and local demonstration to MAS (3.6, Response 8.8). Rho evidences decision-to-source lineage, the Bridge's tokenisation and egress inspection, and approval trails, but holds no record that answers these in full. They are shown as For the bank, with the reason, and listed once in section 4. Southern's own systems and policies carry them.

SECTION 4

Expectation by expectation

Each block states what MAS expects, the status, the evidence Southern holds in Rho and any open item with its owner and age. Paragraph numbers are the final Guidelines'. Section 5 expectations are scaled to materiality; the evidence below is for the Collections Voice Agent and the estate.

Oversight

SECTION 3 · DUE 7 OCT 2027 · ACCOUNTABILITY AND ESCALATION

AIRG 3.4 Board oversight of AI risk

paras 3.4(a)-(e) · 3.2-3.3 · 2.5(a)

Partial

EXPECTS

The Board approves and regularly reviews the AI risk approach and strategy, addresses AI in the risk appetite in qualitative and quantitative terms where suitable, sets Board and senior-management roles, and holds enough AI understanding to challenge.

EVIDENCE

- The Board Risk Committee receives AI estate packs each quarter: pk_board_q2 (signed 14 Jul) and pk_board_q3 (signed 11 Sep, for the 18 Sep meeting).
- pk_board_q3: 3 of 178 pre-registered expectations in breach; open findings 9 to 14; readiness 79% to 82%; agents 176 to 187.

OPEN

Open — No senior manager is designated as accountable for AI oversight; the signature line is left blank. — Board · n/r

The Board-approved AI risk appetite and its strategy documents are bank documents, not Rho records. This report evidences how oversight operates, not the policy text.

AIRG 3.5 Senior management: frameworks, escalation, Board updates

paras 3.5(a)-(g) · 4.3

Partial

EXPECTS

Senior management implements and reviews the AI risk framework and lifecycle control strategy, assigns responsibilities with clear accountability and segregation of duties, runs an internal escalation process for material AI risks and exceptions, updates the Board in good time, and resources the work.

EVIDENCE

- Segregation of duties is recorded: the second line (Mei Ling Wong) approves pre-registrations, policy, promotions and risk class; the platform team (Wei Jie Lim) promotes; the business owner (Grace Tan) is accountable. 175 of 187 agents have a named owner.
- Escalation (3.5(e)): the ladder moved the Collections Voice Agent to a recommend-only proposal on 5 Sep after the red breach.
- Board updates (3.5(f)): the quarterly pack, signed by the named second-line officer.

OPEN

Open — The recommend-only proposal has waited 6 days for the owner; no decision recorded (f_005). — Grace Tan · 6 d

For the bank — Resourcing and capacity of the AI control functions (3.5(g)) are not evidenced in Rho. — Bank

AIRG 3.6 Group framework and local accountability

para 3.6 · 1.5

For the bank

EXPECTS

Where the institution relies on a group AI framework, local management ensures it meets the Guidelines and addresses Singapore requirements, has timely access to relevant reports, and can show its local oversight to MAS.

NOTE

Southern's group framework fit assessment, local management reporting lines and MAS demonstration are bank records. Rho holds none of them, so the status is For the bank.

AIRG 4.2 Identification of AI use paras 4.2 · 4.3 · 4.4

Partial

EXPECTS

Consistent identification of AI across business and functional areas, with clear definitions and criteria, including AI embedded in or delivered by material third-party services; a control function with independent oversight; residual risk from unidentified AI within appetite.

EVIDENCE

- Discovery from the Bridge and the vendor register: 187 declared, 196 observed; 106 agents on a direct feed, 64 on an adapter, 13 inferred, 4 never connected.
- Embedded AI inside SaaS products is found from vendor traffic, not only from declarations: 13 features identified.
- Undeclared AI is surfaced: svc-acct-8841 made 4,102 calls before it was registered (f_014).

OPEN

Partial	12 of the 13 embedded SaaS AI features have no named owner (f_004).	Kenneth Ong · 40 d
Gap	Undeclared AI: svc-acct-8841, 4,102 calls (f_014).	Wei Jie Lim · 39 d

AIRG 4.5 AI inventory and its attributes paras 4.5–4.9 · fn 24–25

Partial

EXPECTS

An accurate inventory of AI use cases, systems and models, kept current for new, updated and retired items; attributes such as purpose, approved scope, model type, data, dependencies, lifecycle status, materiality rating, validation status, owners and documentation; linked to other inventories; maintained under a control function with independent oversight.

EVIDENCE

- The register records each attribute for 187 agents; observed-only dependencies are listed beside the declared ones; every record states how it was observed.
- The second line approves the register entry and the risk class; changes create a new version, so retired and superseded entries stay retrievable.

OPEN

Gap	Purpose record incomplete on 14 agents (C-03, f_013).	Mei Ling Wong · 51 d
Gap	4 registered agents have never connected; the oldest has been silent for 102 days (f_003).	Kenneth Ong · 102 d

AIRG 4.10 Risk materiality: impact, complexity, reliance paras 4.10–4.13

Partial

EXPECTS

One methodology to assess the materiality of every AI use case, covering inherent and residual risk and at least impact, complexity and reliance (impact on the institution and customers, data sensitivity, visibility into third-party AI, autonomy and human involvement); residual risk within appetite before deployment; a control function that approves or arbitrates each rating.

EVIDENCE

- A risk class is recorded for every agent (A 62, B 41, C1 71, C2 13); 41 of 187 are high-risk. The second line approves it.
- The Collections Voice Agent: high-risk and material (class C1); third-party (Vendor X); autonomy recorded in its authority record; human involvement measured (section 5).

OPEN

Partial	Rho records the class and its owner; the bank holds the methodology document and the residual-risk acceptance. 14 purpose records are incomplete, which weakens the impact assessment for those agents.	Mei Ling Wong · 51 d

Materiality is the bank's judgement. Rho never pre-fills it.

AIRG 5.3 Contingency plans and the kill switch

para 5.3 · fn 29 · Response 10.70

Gap

EXPECTS

Use-case-specific contingency plans with operationally feasible fallback for AI failure or unexpected behaviour in high-risk uses, reviewed and tested; clear, regularly tested activation protocols where the AI has kill switches.

EVIDENCE

- Collections Voice Agent: kill-switch drills stopped it in 9.8 s, 4.2 s and 3.1 s (14 Jul); the fallback was drilled on 14 Jul.
- The business continuity drill fell back in 6 m 40 s.

OPEN

Gap	The kill switch on the Mortgage Pre-qualification Assistant was configured on 1 Jun and has never been drilled; the agent is not connected (C-15).	Wei Jie Lim · 102 d

AIRG 5.4 Data governance

paras 5.4(a)-(g) · Response 10.4-10.8

For the bank

EXPECTS

Data used by AI is fit for purpose and representative, quality-checked, classified, secured across its lifecycle, handled in line with privacy and consent requirements, and documented from source to use.

EVIDENCE

- Rho evidences decision-to-source lineage (every record states what it read and how it was observed), retention by record type (decision records 7 years, sampled transcripts 3 years) and tokenisation in the Bridge. Call audio stays inside the bank; all Rho Cloud storage is in the bank's region.

NOTE

Fitness, representativeness, quality, classification and secure lifecycle of training and operating data are the bank's data governance. Rho holds no record that answers them, so the status is For the bank.

AIRG 5.5 Transparency and explainability

paras 5.5-5.6 · fn 32-34 · Response 10.12-10.14

Gap

EXPECTS

Transparency and explainability proportionate to context and materiality: audience-appropriate explanations, customer information that AI is in use, the consequences of its decisions and redress.

EVIDENCE

- AI was disclosed within 10 s on 468 of 470 sampled calls; nm_0071 (38 calls late) was corrected in the call.
- Any decision can be reconstructed with the authority and policy in force at the time.

OPEN

Gap	No disclosure record for the CRM Next-Best-Action feature and Transaction Fraud Screening (C-23, R-62).	Grace Tan · 12 d
Partial	13 of 142 complaints cannot be linked to a decision (C-24).	Grace Tan · 24 d

AIRG 5.7 Fairness

paras 5.7-5.8 · fn 35-36 · Response 10.19-10.21

Partial

EXPECTS

Define fair outcomes in context; assess whether fairness principles are relevant to each use case; apply proportionate controls against harmful bias and discrimination; document results and mitigation.

EVIDENCE

- Cohorts against Southern's ±8-point tolerance (C-12), 16 Aug – 10 Sep against 15 Jul – 13 Aug: heartland Mandarin 0.49 to 0.40 (–15 against the central reference, outside); heartland other 0.52 to 0.46 (–9, outside).
- Counterfactual flips on 4.2.1 against 4.1.8: language 9.8% against 0.8%; name only 4% against 1.1%.

OPEN

Open	The disparity sits within f_005; inv_0418 is attached to f_001. Both are open.	Grace Tan · 6 d

The definition of fair outcomes and the tolerance are Southern's, referenced to FEAT; MAS sets no threshold.

AIRG 5.9 Human oversight

paras 5.9(a)-(d) · 5.23(c)-(e) · fn 37

Partial

EXPECTS

Oversight controls proportionate to materiality: roles and escalation, competence and authority to intervene, outputs routed to a human on set conditions, allowance for automation bias; records of human oversight reviewed for effectiveness.

EVIDENCE

- Reviewers were right on 2,013 of 2,494 overrides (80.71%) against the agent's 70.69% (Jan–Mar, matured). 470 of 39,140 calls were human-calibrated in Aug (1.2%, 8 strata): Southern policy.
- Override records are kept per reviewer, with outcome and competency status.

OPEN

Partial	Competency current for 3 of 4 reviewers of the agent; R. Ng expired 40 days. Across high-risk agents 4 reviewers lack current competency.	Mei Ling Wong · 40 d
Open	R. Quek left on 2 Sep with reviewer access still active 9 days.	No owner recorded · 9 d

AIRG 5.10 Third-party AI: testing, accountability, change, concentration

paras 5.10-5.11(a)-(h) · Response 9.5-9.21

Gap

EXPECTS

Test third-party AI in the institution's own context, using its own data where practicable, and compensate for gaps in disclosure; keep primary accountability; weigh transparency, supply chain, concentration, change notices, contingency, legal agreements, staff capability and complexity by materiality; where residual risk cannot be kept within appetite, limit, suspend or replace the provider.

EVIDENCE

- Own-data testing: 12,000 Southern calls scored where Vendor X withheld training-data provenance; outcomes measured on Southern's customers (71% against the claimed 94%). Residual risk accepted by Mei Ling Wong on 2 Jul (C-19).
- Concentration: frontier-a sits under 18 of the 41 high-risk agents, shared by Vendor X and Vendor W with no shared fallback.
- Change notices: Vendor Y and Vendor W notified all 7 of their changes this quarter; Vendor X notified 2.

OPEN

Gap	The 14 Aug model change came without notice (MSA §8.4) (f_001).	Siti Aminah · 28 d
Gap	Legal agreements (5.11(f)): sub-contractor disclosure, audit of sub-contractors and AI-specific clauses missing from MSA-2025-118.	Siti Aminah · at renewal
Gap	Speech sub-contractor observed in Bridge traffic, not notified; attestation requested 23 Aug, due 6 Sep, overdue (f_009).	Siti Aminah · 19 d

NOTE

A third-party AI service is not automatically a material outsourcing; whether the speech provider falls within the Notice 658 sub-contracting clause is the bank's assessment.

AIRG 5.14 Testing before deployment, with agreed thresholds

paras 5.12-5.15 · 5.18

Evidence in place

EXPECTS

Choose models and features deliberately; evaluate and test each system before deployment in its use context against clear, measurable thresholds agreed by business owners, developers and reviewers; mitigate known limitations with guardrails.

EVIDENCE

- Expectations and thresholds pre-registered before production for 178 of 187 agents (C-06). Collections Voice Agent: signed 14 Jan 09:00 by the owner (Grace Tan), the vendor (Ryan Chua, Vendor X) and the second line (Mei Ling Wong); sha256 e1c9…7b2f. That is the three-party agreement 5.14(b) describes.
- Promotion on 21 Jan followed a dry run against the Dec 2025 pilot (0 would have failed), second-line approval and a configured fallback.
- Pilots carry a success criterion and an expiry (the Hardship Triage pilot expires 30 Sep; f_011).

NOTE

9 of 187 agents are not yet pre-registered; the bank has accepted this as a minor exception.

AIRG 5.16 Security, development records, readiness

paras 5.16-5.17 · 5.22 · Response 10.36-10.44

For the bank

EXPECTS

Secure deployment environments, access and segregation for AI components, plugin and API governance, development documentation that allows the implementation and evaluation to be understood and replicated, and technology and security readiness review with adversarial testing.

EVIDENCE

- Rho's Bridge tokenises identifiers and inspects egress inside the bank, and Rho Cloud holds clean records only. Evidence packs are signed and regenerate to the same figures.

NOTE

System security, security testing, development documentation and readiness reviews are the bank's engineering controls, read with the MAS Technology Risk Management Guidelines. Rho holds no record that answers them, so the status is For the bank.

AIRG 5.19 Independent review and validation

paras 5.18-5.21 · Response 10.57-10.60

Partial

EXPECTS

A documented pre-deployment review by parties outside development; for high-risk uses, formal independent validation by competent staff independent of development and deployment, with effective challenge, covering design soundness, data, implementation, results, explainability and fairness; reviewers report to the approval body, which acts on their findings.

EVIDENCE

- Independent validation: 39 of 41 high-risk agents validated and current (pk_val_credit_q2 is one example).
- Collections Voice Agent: compensatory testing on 12,000 Southern calls, accepted by the second line on 2 Jul.

OPEN

Partial	2 independent validations of high-risk agents overdue, due 30 Sep (C-22, f_008).	Mei Ling Wong · 22 d
For the bank	Reviewer reporting to the approval body and its follow-through (5.21) are bank records.	Bank

---

AIRG 5.23 Ongoing monitoring, issues and incidents

paras 5.23(a)-(e) · Response 10.66-10.69

Partial

EXPECTS

Monitor all deployed AI, third-party AI included, with frequency and intensity proportionate to materiality: defined metrics, thresholds and early-warning tiers, drift checks, workflow monitoring; report, track, escalate and resolve issues and incidents; override or deactivation for high-risk uses; an accountable person; user feedback channels; trained monitors.

EVIDENCE

- Tiered warnings: amber 21 Aug (0.49), red 5 Sep (0.44 against a 0.45 floor). The nightly probe fell from 39/40 to 34/40 on 14 Aug, 1.5 h after the unnotified model change.
- Actions and tools monitored: access breadth 3.4 times baseline flagged (f_006); 12 episodes on policy v11 after v12 (f_012).
- Every incident and near miss is recorded with detection, containment and the internal report; median time to detect 0.4 h in Jul–Sep 2026, against 31 h in Oct–Dec 2025, before Rho.

OPEN

Partial	Hardship criterion below the judge threshold (κ 0.58) since 4 Sep; human scoring on all hardship cases (126 of 200 revalidation cases done).	Mei Ling Wong · 7 d

Baseline records of all monitoring activity and logging of access to AI assets (5.23(d)) extend beyond what Rho records.

AIRG 5.24 Revalidation of high-risk AI

para 5.24 · Response 10.68

Gap

EXPECTS

Regular revalidation of high-risk use cases by competent, objective parties independent of development and deployment, with frequency and rigour informed by materiality.

EVIDENCE
- Before the change: a dry run before production (0 would have failed). After: a 400-case parallel run of 4.2.1 against the in-house team.

OPEN
Gap — Version 4.2.1 went live without re-approval or independent re-validation; no re-validation record exists in the ledger. — Siti Aminah · 28 d

---

AIRG 5.25 Change management, including vendor-driven updates

para 5.25(a)-(c) · Response 10.73-10.74 — Gap

EXPECTS
Assess the significance of each change and review and re-approve it accordingly; for vendor-driven AI updates, compensate when a significant high-risk change proceeds without prior review; control and track internal changes.

EVIDENCE
- Change gate for vendor and internal changes: Vendor Z 2.3 to 2.4 passed 6 probes (8 Sep); credit policy v11 to v12 approved 9 Sep; Vendor Y 3.1 to 3.2 held in the gate.
- Draft policy v13 tested before approval: it would have failed 437 of 33,760 episodes (15 Aug – 10 Sep).

OPEN
Gap — Vendor X changed the model on 14 Aug without notice and without re-approval; compensating testing began after the fact (f_001, f_005). — Siti Aminah · 28 d

General internal change governance, prevention of unauthorised internal changes and rollback sit with the bank.

AIRG 5.26 Secure retirement

para 5.26 — For the bank

EXPECTS
Secure retirement of AI: dependencies handled, data retained or destroyed as required, stakeholders told.

NOTE
Rho keeps retired register entries and their evidence for the retention period. The retirement procedure, data destruction and stakeholder notices are the bank's, so the status is For the bank.

Capabilities

SECTION 6 · DUE 7 OCT 2028 · PEOPLE AND INFRASTRUCTURE

AIRG 6.1 Competence and conduct

paras 6.1-6.2 · 3.5(g) · Response 11.9 — Partial

EXPECTS
Ensure competence and proper conduct of developers, deployers and users through recruitment, training and support; review personnel capability and capacity and keep programmes current.

EVIDENCE
- Competency records for every reviewer of a high-risk agent. On the Collections Voice Agent, 3 of 4 reviewers hold current competency; the reviewer whose competency has lapsed was right on 411 of 612 overrides (67.2%), the lowest of the four.

OPEN
Partial — 4 reviewers of high-risk agents without current competency: R. Ng (expired 40 d), P. Neo (30 d), M. Ho (17 d), T. Sim (never trained). — Mei Ling Wong · 40 d

Board and senior-management training records and staff conduct reviews are held outside Rho and are not assessed here.

AIRG 6.3 Infrastructure and resources

para 6.3 · Response 11.11 — For the bank

EXPECTS
Adequate hardware, software, networks and data pipelines to support AI performance, scalability and resilience, and technology-neutral adequacy reviews.

NOTE
Infrastructure adequacy is the bank's. Rho records availability and fallback drills (AIRG 5.3) but not capacity or scalability, so the status is For the bank.

Binding notices

MAS NOTICES · BINDING · READ BESIDE THE GUIDELINES

N658 Outsourced relevant services and sub-contracting

Notice 658 · Guidelines on Outsourcing (Banks) — Gap

REQUIRES: Keep the register of outsourced relevant services; the service provider notifies the bank in writing before, or within a reasonable time of, engaging a sub-contractor; adverse developments are notified to MAS.

EVIDENCE
- Vendor X is on the register (9 arrangements: 2 material, 5 non-material, 1 licensed, 1 under review); its entry is extracted in Annex B.

OPEN

Gap	Speech sub-contractor observed in Bridge traffic, not notified; attestation overdue (f_009).	Siti Aminah · 19 d

NOTE: Whether the speech provider falls within the sub-contracting clause, and whether any event is an adverse development to notify, is Southern's assessment.

FSM-N05 Technology incident reporting

Notice FSM-N05 paras 7, 8 · MAS/TCRS/2025/08 — For the bank

REQUIRES: Notify MAS within 1 hour of discovering a relevant incident; the para 8 report within 14 days; initial report within 24 h and final within 14 days via MAS-Tx.

EVIDENCE
- inc_0042 and inc_0043 were handled on Southern's internal targets (inc_0043: internal outsourcing incident target 13 Sep).

NOTE: No event is recorded as assessed reportable. The 13 Aug 20:10 notice and the 20 Aug report are internal. Rho never pre-fills the assessment.

FD 1.5.1 Fair dealing: MI and complaints

Fair Dealing Guidelines 1.5.1 · Outcomes 4, 5 · ABS s.16 — Partial

EXPECTS: Management information on complaints and misconduct; clear information; complaints handled independently, effectively and promptly.

EVIDENCE
- The agent: 118 complaints in 95,240 contacts (12.4 per 10,000), 1 Jul – 11 Sep. Across the estate, 129 of 142 AI complaints linked to a decision; 33 led to a change.

OPEN

Partial	13 of 142 complaints cannot be linked to a decision.	Grace Tan · 24 d
Diverging	3 public claims diverge from measured results, including "every AI decision affecting a customer is reviewed by a human" against 1.2% measured.	No owner recorded · 29–41 d

ABS 15 Debt recovery conduct

ABS Code s.15(b)(iv), (d)(i)-(ii), (e)(ii)-(iii) — Partial

EXPECTS: Calls at suitable times and places; repayment options that account for the customer's situation; the caller and purpose identified; no recovery from referees, family or friends who are not guarantors.

EVIDENCE
- Contact outside the 08:00–21:00 window (Southern policy) is a pre-registered unacceptable failure: 41 accounts in Jan–Mar; 96 contacted after a complaint was on record.
- Hardship-flagged promise kept 41% against 63% in-house; draft policy v13 would have failed 331 episodes on hardship in a dry run. contact_third_party is prohibited; 0 attempts.

NOTE: The ABS Code sets no calling hours; 08:00 to 21:00 is Southern policy.

PDPA

Personal data

PDPA s.24, 25, 26 · Banking Act s.47

Partial

REQUIRES

Reasonable protection, retention only as long as needed, and comparable protection for transfers out of Singapore.

EVIDENCE

- Audio stays in the bank; records tokenised at the Rho Bridge; all Rho Cloud storage in the bank's region, cross-region 0.
- Retention: decision records 7 years, sampled transcripts 3 years. Disposal records are kept for each data request (dsr_0230: 159 fields disposed of across 41 episodes).

OPEN

Partial

14 purpose records incomplete (f_013).

Mei Ling Wong · 51 d

SECTION 5

Open gaps and remediation plan

Every gap and partial in this report, with owner and age in days to 11 Sep 2026 (n/r: none recorded) and the Guidelines paragraph it answers. All 14 findings are in Annex C.

REF	ITEM	MAS AIRG	OWNER	AGE	NEXT STEP
GAPS					
f_014	Undeclared AI: svc-acct-8841, 4,102 calls	AIRG 4.2	Wei Jie Lim	39 d	Register or retire; inv_0397 merged into discovery
C-15	Kill switch on the Mortgage Pre-qualification Assistant never drilled	AIRG 5.3	Wei Jie Lim	102 d	Drill once the agent is connected (f_003, Kenneth Ong)
C-23	No disclosure record for CRM Next-Best-Action and Transaction Fraud Screening (R-62)	AIRG 5.5	Grace Tan	12 d	Record each customer touchpoint and its disclosure
f_001	Version 4.2.1 live without re-approval or independent re-validation	AIRG 5.24	Mei Ling Wong	28 d	Re-validate; record the conclusion
f_001	Model change not notified (MSA §8.4); 4.2.1 not gated or revalidated	AIRG 5.11(d), 5.25	Siti Aminah	28 d	Obtain the reply; Mei Ling Wong decides validation or limits
f_009	Speech sub-contractor observed, not disclosed	Notice 658 · AIRG 5.11	Siti Aminah	19 d	Obtain the attestation; add the clause at renewal
C-16	Three AI clauses missing from MSA-2025-118 v1	AIRG 5.11(f)	Siti Aminah	n/r	Negotiate at the 31 Jan 2027 renewal
PARTIALS AND OPEN ITEMS					
f_005	Heartland conversion red at 0.44; recommend-only proposed	AIRG 5.23(b)	Grace Tan	6 d	Owner decides the proposal
f_006	Access breadth 3.4 times baseline	AIRG 5.23	Wei Jie Lim	11 d	Review scope against the authority record
f_012	12 episodes on policy v11 after v12	AIRG 5.23	Wei Jie Lim	8 d	Close the policy-currency gap in the gateway
C-10	Hardship criterion below judge threshold (κ 0.58)	AIRG 5.23(a)	Mei Ling Wong	7 d	Human scoring until 200 revalidation cases (126 done)
S33	Reviewer competency lapsed (R. Ng); departed reviewer access	AIRG 5.9, 6.1	Mei Ling Wong	40 d	Retrain or reassign; remove access
f_002	Vendor claim 94% against 71% measured	AIRG 5.10	Grace Tan	41 d	Raise at the 31 Jan 2027 renewal
S49	Claim of human review of every AI decision (1.2%)	FD Outcome 4	No owner recorded	29 d	Board to assign an owner
C-03	Purpose record incomplete on 14 agents (f_013)	AIRG 4.5	Mei Ling Wong	51 d	Complete the purpose fields; owner signs
C-04	12 embedded SaaS AI features without a named owner (f_004)	AIRG 4.2	Kenneth Ong	40 d	Assign owners from the discovery evidence; owners attest quarterly
C-22	2 independent validations of high-risk agents overdue (f_008)	AIRG 5.19	Mei Ling Wong	22 d	Complete by 30 Sep
C-24	13 of 142 complaints cannot be linked to a decision	AIRG 5.5	Grace Tan	24 d	Link each complaint to its decision record
f_003	4 registered agents never connected; oldest 102 days	AIRG 4.5	Kenneth Ong	102 d	Connect, or retire from the register

REF	ITEM	MAS AIRG	OWNER	AGE	NEXT STEP
FOR THE BANK · OUTSIDE WHAT RHO EVIDENCES					
AIRG 3.4	Designate the senior manager accountable for AI oversight; record the Board risk appetite for AI	AIRG 3.4, 2.5(a)	Board	n/r	Board to name the role; add to the sign-off
AIRG 5.4	General AI data governance: fitness, quality, representativeness, classification, lifecycle	AIRG 5.4	Bank data office	n/r	Bank policy; Rho evidences decision-to-source lineage and retention only
AIRG 5.16	System security, security testing, development documentation, readiness review	AIRG 5.16, 5.17, 5.22	Wei Jie Lim	n/r	Bank engineering controls, read with the TRM Guidelines
AIRG 5.26	Secure retirement procedure and resourcing and infrastructure adequacy	AIRG 5.26, 6.1, 6.3	Wei Jie Lim	n/r	Bank procedure; Rho keeps the retired record
FSM-N05	Reportability assessment for inc_0045	FSM-N05 para 7	Mei Ling Wong	n/r	The bank's assessment; Rho never pre-fills it

SECTION 6

Annex

A Third-party AI arrangement · Notice 658 register extract · Vendor X

Arrangement	Vendor X · MSA-2025-118 v1, effective 1 Jan 2025 · material · renewal 31 Jan 2027 (142 days)
Scope	15 agents, 9 of them high-risk; flagship Collections Voice Agent
Clauses	8 of 11 of Southern's AI clauses present. Missing: sub-contractor disclosure, right to audit sub-contractors, AI-specific liability.
Change notice	MSA-2025-118 v1 §8.4: notify before any material model or version change and obtain approval before production deployment (PRESENT, and breached on 14 Aug)
Change discipline	2 changes notified (12 May, 24 Jun), both through the gate; 1 unnotified (14 Aug, f_001). Letter ltr_vx_0815 sent 15 Aug; reply due 29 Aug; none received.
Sub-contractors	Third-party speech provider seen in Rho Bridge traffic, not declared (f_009); attestation requested 23 Aug, due 6 Sep, not received. mcp-crm-bridge v0.9 also observed, not declared.
Claim against measurement	Task success: 94% claimed (Vendor X's own benchmark, not Southern's population) against 71% measured (61,131 of 86,100, agt_coll_vx episodes 21 Jan – 31 Mar 2026, outcomes matured); f_002.
Compensatory testing	Could not obtain training data provenance and evaluation methodology. Tested instead: 12,000 Southern calls scored against rubric_collections_v7; acceptable on central areas, below floor on heartland. Residual risk (training population unknown; heartland underperformance unexplained) accepted by Mei Ling Wong on 2 Jul 2026.
Incidents	2 attributable (inc_0042 (0.4 h), inc_0043 (1.5 h)); mean time to detection 0.95 h against 31 h before Rho (Oct–Dec 2025); the vendor told Southern first in 0 of 2.
Concentration	Base model frontier-a is shared with Vendor W, with no shared fallback; 18 of the 41 high-risk agents run on it.
Renewal	Renewal brief brief_vx_renewal (draft, Siti Aminah) carries f_001, f_002 and f_009.
Register context	9 arrangements; the CRM SaaS provider is under review, disclosure request unanswered 37 days.
Adverse development	Notified to MAS: the bank's assessment (yes or no, date, basis). Not pre-filled.
Register submission	Kept under Notice 658 and submitted as MAS requires; semi-annual submission is the TPRM consultation's proposal.

B Customer outcomes and fair dealing

Calling window	08:00–21:00 is Southern's policy window (the agent's calling day is recorded separately); the ABS Code sets no hours.
What went wrong	Accounts where something went wrong are counted in accounts (Jan–Mar); complaints are counted per contact (Jul–Sep). The two units are not added.
Disposal records	Where a customer asks for data to be deleted, Rho keeps a disposal record; the PDPA sets no right to erasure.

ACCOUNTS WHERE SOMETHING WENT WRONG · 21 JAN – 31 MAR 2026

WHAT WENT WRONG	ACCOUNTS	OUTSTANDING BALANCE	RECOVERABLE
Promise recorded, no payment followed	142	S$108,000	partly · 38

WHAT WENT WRONG	ACCOUNTS	OUTSTANDING BALANCE	RECOVERABLE
Customer contacted after complaint on record	96	S$68,850	n/a · conduct
Fee waived without required approval (before the 1 Jul gate)	61	S$45,900	yes · 61
Contacted outside 08:00–21:00	41	S$25,650	n/a · conduct
All accounts	**340**	**S$248,400**	100% traced to a decision

21 Jan – 31 Mar 2026 · outcomes matured as of 11 Sep. Units are accounts, not complaints. The re-run reproduced 340 of 340 (signature and figures verified 10 Sep 17:40 by Mei Ling Wong).

COMPLAINTS · 1 JUL – 11 SEP 2026

AI-RELATED COMPLAINTS, 1 JUL – 11 SEP 2026	COUNT
All AI-related complaints	142
About the Collections Voice Agent	118
Linked to the decision they concern	129 (90.85%)
Cannot be linked to a decision	13
WHAT HAPPENED TO THEM	
Led to a change	33
Assessed, no change	70
Open	39
BY NATURE · ALL AI AGENTS	
Contact after complaint	51
Tone or pressure	35
Wrong amount	24
Not told it was AI	18
Other	14

The agent drew 118 complaints in 95,240 contacts (12.4 per 10,000). Examples: cmp_2211 (called after a complaint, open 2 d, Grace Tan); cmp_2190 (said it was a bank officer, not a machine; resolved by Southern's own script change on 8 Sep); cmp_2105 (settlement offered then withdrawn, inc_0042, resolved).

PUBLIC CLAIMS ABOUT THE BANK'S AI

CLAIM	MEASURED	STATUS	AGE
"Our AI resolves 90% of customer queries without human help"	74%	Diverging	34 d
"AI-assisted collections recover 30% more"	-10%	Diverging	41 d
"Every AI decision affecting a customer is reviewed by a human"	1.2%	Diverging	29 d
"Hardship triage live across all branches"		Outdated	12 d
"All of our AI models are built in-house"	62 of 187 class A	Withdrawn	

11 public claims in the register: 6 supported, 3 diverging, 1 outdated, 1 withdrawn. The claim "All of our AI models are built in-house" was withdrawn by an annual-report erratum on 2 Jul 2026.

C Evidence index

8 evidence packs were signed or filed from 1 Jul to 11 Sep 2026 (7 signed, 1 filed). Each is signed by Mei Ling Wong, Head of AI & Model Risk (second line), stored in ap-southeast-1 and object-locked for 7 years from the end of its period.

PACK	CONTENTS	PERIOD · AS OF	SIGNING AND REPRODUCTION
pk_board_q3	AI estate, Board Risk Committee	1 Jul – 11 Sep 2026	Signed 11 Sep 09:39, for the 18 Sep meeting · locked to 30 Sep 2033
pk_2026q1_coll_r2	Collections agent estate, supervisory evidence, restated (19 agents; 86,100 episodes)	21 Jan – 31 Mar 2026 · as of 11 Sep 08:00	Signed 11 Sep 09:00 · sha256 4a9c…e17b · reproduced 09:24 in 41 s, figures matched · locked to 31 Mar 2033 · held for inspection
pk_rec23_h1	Internal inventory report, half-yearly (176 agents)	1 Jan – 30 Jun 2026 · as of 1 Jul	Signed
pk_board_q2	AI estate, Board Risk Committee	1 Apr – 30 Jun 2026 · as of 14 Jul	Signed
pk_val_credit_q2	Retail Underwriting Assistant, validation evidence	1 Apr – 30 Jun 2026 · as of 6 Aug	Signed
pk_inc_0042	Internal incident report: settlement offered outside authority	Occurred 13 Aug · as of 20 Aug	Filed 20 Aug by the named officer
pk_rca_0043	Post-incident review: unnotified model change	As of 28 Aug	Signed 28 Aug
pk_vy_renewal	Vendor Y renewal brief	As of 2 Sep	Signed
RETAINED AND RELATED RECORDS · NOT COUNTED AS PACKS			
pk_2026q1_coll	Collections supervisory evidence, as issued (January promise kept 63%)	21 Jan – 31 Mar 2026 · as of 5 Apr	Signed 7 Apr · superseded by r2, retained and valid as issued · locked to 31 Mar 2033
inventory_2026-09-11.csv	Register export, 187 agents	As of 11 Sep 09:41	Signed by Mei Ling Wong · an export, not a pack
ltr_vx_0815	Letter to Vendor X citing MSA §8.4	15 Aug · reply due 29 Aug	From Siti Aminah · no reply received
brief_vx_renewal	Vendor X renewal brief, carries f_001, f_002, f_009	Renewal 31 Jan 2027	Draft · Siti Aminah
pre-registration	Collections Voice Agent expectations, 3 signatories	Signed 14 Jan 09:00	sha256 e1c9…7b2f · before the first production episode

Other record IDs cited: incidents inc_0042 to inc_0045 and near miss nm_0071; investigations inv_0418, inv_0402, inv_0397; red-team findings rt_041 and as_112; complaints cmp_2105 and cmp_2203; data request dsr_0230; authority record auth_collections_v7. Evidence reads, 1 Jul – 11 Sep: 386 (owners 355, Rho agents 28, non-owners 3), each logged.

D Open findings across the estate · 14, all owned

ID	SEVERITY	FINDING	OWNER	AGE
f_003	High	Registered agent has never emitted (Mortgage Pre-qualification Assistant)	Kenneth Ong	102 d
f_013	Low	Purpose record incomplete on 14 agents	Mei Ling Wong	51 d

ID	SEVERITY	FINDING	OWNER	AGE
f_002	High	Vendor claim diverges: 94% claimed, 71% measured	Grace Tan	41 d
f_004	High	Unowned agent in production (CRM Next-Best-Action)	Kenneth Ong	40 d
f_014	Low	Shadow inference endpoint: 4,102 calls	Wei Jie Lim	39 d
f_007	Medium	Unscoreable rate 21% above 15% ceiling (KYC Document Verification)	Ahmad Yusof	34 d
f_001	High	Unnotified model version change (Vendor X, 4.1.8 to 4.2.1)	Siti Aminah	28 d
f_008	Medium	Independent validation overdue (2 high-risk agents)	Mei Ling Wong	22 d
f_009	Medium	Undisclosed sub-contractor (speech provider)	Siti Aminah	19 d
f_010	Medium	Rubber-stamp rate 31% against 15% expected (Retail Underwriting Assistant)	Kenneth Ong	15 d
f_006	Medium	Scope creep: access breadth 3.4 times baseline	Wei Jie Lim	11 d
f_012	Low	Policy currency: 12 episodes on v11 after v12	Wei Jie Lim	8 d
f_005	Medium	Pre-registered threshold breached: red	Grace Tan	6 d
f_011	Low	Pilot expiring in 19 days (Hardship Triage)	Grace Tan	3 d

E Glossary and definitions

Agent	Any AI system that makes or shapes a decision, whether built by Southern or supplied by a vendor.
Class A, B, C1, C2	A: built by Southern. B: vendor technology run by Southern. C1: vendor-operated AI service. C2: AI embedded in a SaaS product.
Episode · decision record	One unit of agent work, such as a call, with the inputs, tools, authority and action recorded.
Pre-registration	Expectations, floors and early-warning levels written and signed before an agent's first production episode.
Authority record	What an agent may do alone, what needs approval and what is prohibited, in force at a given date.
Recommend-only	A mode in which the agent proposes and a person decides. A threshold breach can propose it; a person approves.
Evidence pack	A signed, object-locked set of records and figures for a period, which can be regenerated from stored inputs.
Rho Bridge · Rho Cloud	The Bridge runs inside the bank and records agent activity. Rho Cloud holds evidence, in ap-southeast-1 for Southern.
Judge j7.3 · κ	The model that scores episodes against a rubric. κ (kappa) measures its agreement with human raters.
Promise to pay · DPD	A customer's recorded commitment to pay; "kept" if paid within 45 days. DPD: days past due.
AIRG	The MAS Guidelines on Artificial Intelligence Risk Management for Financial Institutions, issued 7 October 2026. Guidance: MAS expects and institutions should.
Inherent and residual risk	Inherent risk is the risk of a use case before controls; residual risk is what remains after them. High-risk lifecycle controls are triggered by inherent risk (paras 4.10–4.11).
Basic governance route	A lighter route for institutions whose AI is unlikely to materially harm them, their customers or others if it performs badly or is unavailable (para 2.3). Southern does not qualify.
Relevant incident	A technology incident that meets the FSM-N05 criteria, as the bank assesses it.

F Sign-off

Mei Ling Wong
Head of AI & Model Risk (second line) · prepares and signs

Grace Tan
Head of Collections · first-line attestation (agent owner)

Siti Aminah
Head of Third-Party Risk · third-party annex

Senior manager for AI oversight (role)
To be designated by the bank · no name in the ledger

Contributing: Wei Jie Lim, Head of AI Platform. For: Board Risk Committee, 18 Sep 2026 (outcome to be minuted; secretariat Daniel Koh). Copy: Internal Audit (third line), read access logged. The bank holds the obligation; Rho holds the evidence.

Sources: MAS Guidelines on Artificial Intelligence Risk Management for Financial Institutions (final, 7 October 2026) and MAS's response to feedback; MAS Notice 658 and Guidelines on Outsourcing (Banks); Notice FSM-N05 and Circular MAS/TCRS/2025/08; MAS FEAT principles (2018); MAS Guidelines on Fair Dealing; ABS Code of Consumer Banking Practice (Oct 2025) s.15, s.16; PDPA 2012 s.24 to s.26; Banking Act s.47. Paraphrased; see the Guidelines for the full text. This sample report is illustrative; Southern Bank Singapore and all figures are fictional.