Winsen Rho

AI REGULATION GUIDE · SINGAPORE

What the Monetary Authority of Singapore expects of financial institutions using AI

A practitioner's guide to the Guidelines on Artificial Intelligence Risk Management issued on 7 October 2026, covering scope, the phased implementation dates, supervisory expectations, and the evidence a financial institution should be prepared to produce.

STATUS	TAKES EFFECT	AS OF
Final guidelines, issued 7 Oct 2026	7 Oct 2027 and 7 Oct 2028	20 October 2026

HOW TO READ THIS

Status, scope and conventions.

This guide explains the Monetary Authority of Singapore's Guidelines on Artificial Intelligence Risk Management as issued on 7 October 2026. Every date, threshold and paragraph number comes from MAS's own text, cited as it appears.

7 Oct 2026
FINAL GUIDELINES ISSUED

7 Oct 2027
SECTIONS 3 AND 4 APPLY

7 Oct 2028
SECTIONS 5 AND 6 APPLY

3
MINIMUM RISK DIMENSIONS

65
NUMBERED PARAGRAPHS

Status labels used throughout

Every instrument carries one label. The shape tells you the status even in black and white.

■ BINDING	Acts, subsidiary legislation and MAS Notices. Written as *requires* or *must*.
◆ GUIDANCE	MAS Guidelines. Supervisory expectations, not law. Written as *expects* or *should*.
○ VOLUNTARY	Principles, information papers and national frameworks MAS points to.
⬚ DRAFT · CONSULTATION	Proposals not yet issued in final form. Never presented as final.

> **The Guidelines are guidance.** MAS classifies guidelines as best-practice standards. Contravention is not a criminal offence and does not attract civil penalties, but the degree of observance may affect MAS's overall risk assessment of the institution.

Paragraph references such as para 4.12 point to the Guidelines. References to the response paper read Response 13.5.

Contents

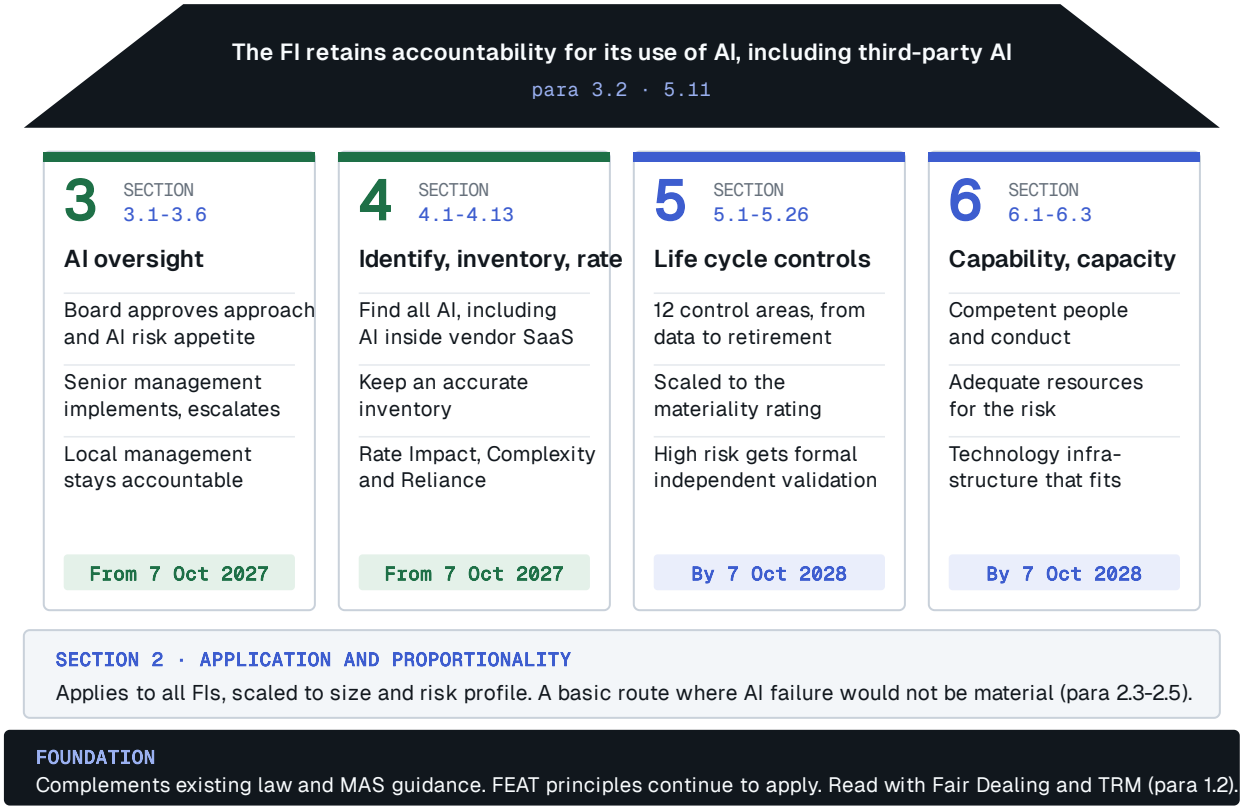
3	The stance on one page
4	Who and what is covered
5	Timeline
6	What is binding, and what is not
7	The expectations, section by section
16	What a supervisor will look for
17	Where institutions often fall short
18	How to get ready
20	Summary reference
21	What Winsen Rho can do for you
25	Sources and disclaimer

Intended readers. The person in a bank, insurer or other financial institution asked to make the Guidelines real, typically in model risk, technology risk or compliance. No prior reading of the MAS text is assumed.

THE STANCE ON ONE PAGE

Identify all AI, assess its materiality, and apply proportionate controls.

MAS sets outcomes rather than prescribing techniques. Each financial institution is expected to identify its use of AI, assess the risk materiality of every use case, and calibrate its controls accordingly, under the oversight of the board and senior management.



Structure of the Guidelines, Sections 1 to 6. Source MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 1.2, 1.8, 2.1-2.5, 3.2, 5.11. As of 20 Oct 2026.

Who is addressed

The board or a delegated committee, senior management, a designated control function for identification, inventory and materiality, the first, second and third lines, and local senior management of FIs in global groups.

What MAS does not prescribe

It does not mandate specific testing techniques or minimum controls for generative or agentic AI Response 12.7, nor set a universal threshold for high risk Response 7.19.

What MAS does not restrict

MAS does not intend to restrict any AI technology or use case, provided the risks are adequately addressed Response 12.3. Agents are in scope today para 1.5.

WHO AND WHAT IS COVERED

All financial institutions, all forms of AI, assessed at the level of the use case.

The Guidelines apply to all financial institutions as defined in section 2 of the Financial Services and Markets Act 2022, in a proportionate manner. Locally incorporated FIs under consolidated supervision, or owning critical information infrastructure, apply them on a group basis.

COUNTS AS AI

- Machine learning, supervised to reinforcement
- Deep learning
- Natural language processing
- Computer vision
- Generative AI, including LLMs
- AI agents, including multi-agent systems

NOT ORDINARILY AI

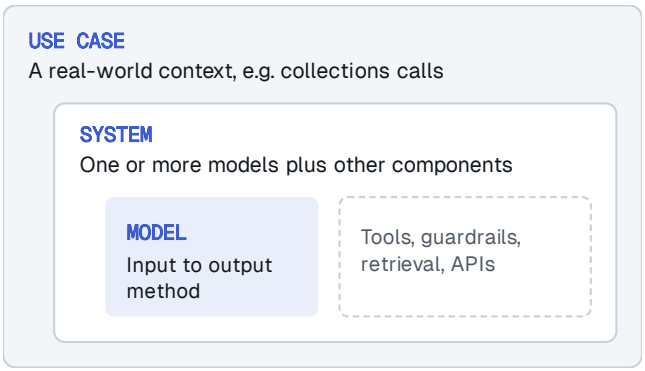
- Rule-based systems
- Explicit formulae specified by hand
- Robotic process automation without learning
- Calculators and fixed-logic tools

The test is whether outputs are derived through learned premises rather than predefined logic. The designated control function is the final arbiter of what counts as AI (para 4.3).

THIRD-PARTY AI INCLUDES

Procured models	Subscribed services	APIs you rely on	AI inside SaaS features

Definition of AI and third-party AI. Source MAS Guidelines on AI Risk Management, 7 Oct 2026, para 1.3, footnotes 6 and 10, para 4.3. As of 20 Oct 2026.



Model, system and use case as defined in para 1.4. Materiality is assessed per use case (para 4.10).

Why the use case is the unit of assessment

The same model can be harmless drafting internal notes and material when it decides a credit limit. MAS asks for the risk materiality assessment at the level of the use case, taking account of dependencies between systems and models (para 4.10, footnote 26).

Groups and branches

An FI in a global group may use the group framework if it meets the Guidelines (para 1.5). Local senior management remains accountable for Singapore-specific considerations and should be able to show MAS how it discharges oversight (para 3.6).

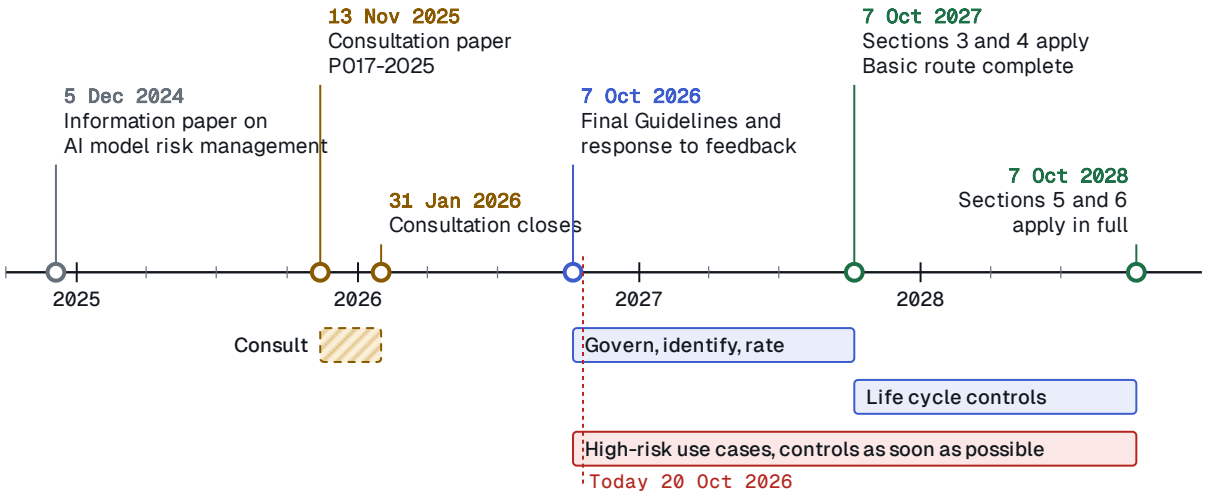
Agents

AI agents, which may be granted autonomy and tool access, can execute unauthorised or erroneous actions at scale (para 1.11). MAS intends to consult separately on further agentic guidance Response 12.9.

TIMELINE

Governance within twelve months, life cycle controls within twenty-four.

The Guidelines take effect on 7 October 2027. MAS allows a phased start, with governance and inventory first and life cycle controls a year later. One expectation applies sooner.



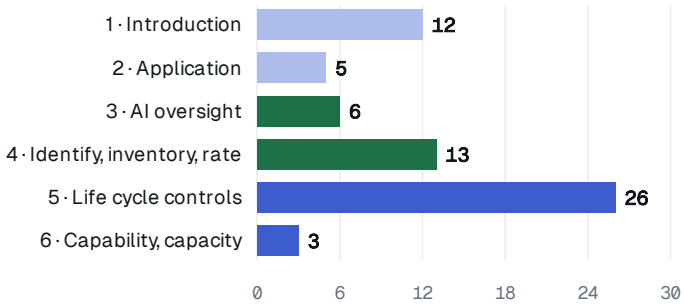
Dates from MAS Guidelines on AI Risk Management, 7 Oct 2026, para 1.8, and MAS Response to Feedback P017-2025, 7 Oct 2026, paras 1.1-1.2 and 13.5-13.6. Information paper of 5 Dec 2024 shown for context. Scale is linear. As of 20 Oct 2026.

Earlier expectation. MAS says an FI should apply appropriate life cycle controls for its high-risk use cases as soon as possible, and not wait for the end of the 24-month transition Response 13.5.

The basic route under para 2.3 has one date only. FIs that qualify should complete implementation within 12 months, by 7 October 2027 Response 13.6.

MAS originally proposed a single 12-month transition. Respondents most often asked for 18 to 24 months Response 13.1-13.2.

Distribution of paragraphs



Count of numbered paragraphs, Sections 1 to 6. Source MAS Guidelines on AI Risk Management, 7 Oct 2026. Green applies from 7 Oct 2027, blue by 7 Oct 2028. As of 20 Oct 2026.

WHAT IS BINDING, AND WHAT IS NOT

The AI Guidelines are supervisory guidance, set on a foundation of binding law.

The AI Guidelines sit on top of rules that already bind an FI. Where third-party AI is involved, MAS says its outsourcing and third-party requirements apply as well (footnote 39). A supervisor will read them together.

BINDING LAW Acts of Parliament. Must be met.	**Financial Services and Markets Act 2022** **Banking Act and sector Acts** **Personal Data Protection Act 2012**	defines FI statute PDPC
BINDING NOTICES MAS Notices. Must be met.	**Notice on Technology Risk Management** **Notice 658 on outsourcing (banks)**	FSM-N05 Notice 658
SUPERVISORY GUIDANCE Expected, not law. Observance weighs in risk assessment.	**Guidelines on AI Risk Management** **Technology Risk Management Guidelines** **Guidelines on Outsourcing** **Guidelines on Fair Dealing**	7 Oct 2026 2021 banks conduct
GOOD PRACTICE AND PRINCIPLES Voluntary. Referenced as useful sources.	**FEAT principles** **Information paper on AI model risk management** **IMDA Model AI Governance Framework for Agentic AI** **IMDA starter kit for testing LLM applications**	2018 · continue to apply Dec 2024 reference reference

Instruments named in MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 1.2, 5.16, footnotes 2, 3, 11, 15, 39, 43, 47. Notice dates per MAS notices pages, see Sources. Legal effect of guidelines per MAS's classification of instruments. As of 20 Oct 2026.

Language conventions

The Guidelines use *should* and *expects* throughout. This guide keeps that language. Where we say *must*, it refers to a notice or an Act.

Interaction with outsourcing rules

A vendor's AI can fall under the AI Guidelines and under outsourcing rules at once. Some services are exempt from outsourcing rules, so not every AI vendor is an outsourced relevant service.

Continuity with FEAT

FEAT, issued in 2018, continues to guide AI use (para 1.2). The Guidelines point to it when an FI defines what it considers fair outcomes (footnote 35).

The board sets direction and appetite. Senior management implements.

MAS does not require a new committee. An FI may use existing risk functions or a new central one, but it should be able to show a consolidated, coordinated view of AI risk across functions (para 3.3, Response 5.3).

BOARD OR DELEGATED COMMITTEE para 3.4	SENIOR MANAGEMENT para 3.5	LOCAL SENIOR MANAGEMENT para 3.6
Approve and review the governance approach and strategic direction Put AI risks in the risk appetite, qualitative and quantitative Set roles for board and senior management Hold enough AI understanding to challenge Keep it all under regular review	Implement frameworks across the FI Review their effectiveness Control risk over the whole life cycle Assign roles, including control functions for inventory and rating Run an escalation process for AI incidents and threshold breaches Update the board in a timely way Resource it, people and training	Stay accountable when a group framework is used See Singapore reports and data in time to act Have clear escalation paths for Singapore issues Show MAS how oversight is discharged

Responsibilities by role. Source MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 3.4-3.6. Applies from 7 Oct 2027. As of 20 Oct 2026.

What a quantitative AI risk appetite can look like

MAS gives examples of measures (footnote 18). Thresholds on the number or financial impact of AI-related incidents. The number of material AI use cases that depend on a single provider. The number of material AI use cases in breach of their performance thresholds.

Each of these is a count an FI can only report if it knows its estate, its incidents and its vendors.

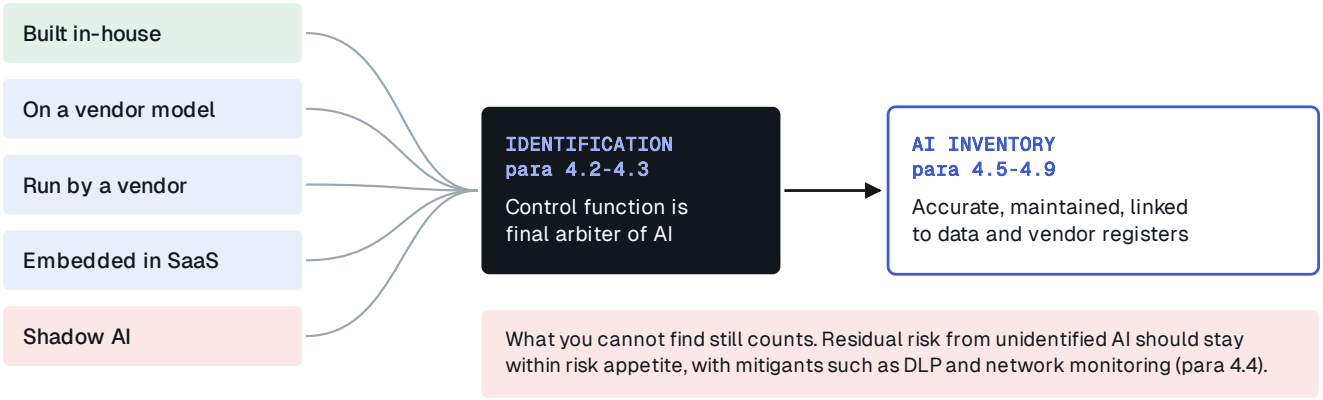
Three lines of defence

MAS suggests the three lines of defence may be used (footnote 19). First line, developers and users. Second line, independent challenge and validation. Third line, internal audit assurance over the framework.

AI incident means a failure, malfunction or poor performance of the underlying AI system or model (footnote 20).

Identification is the prerequisite for every other control.

MAS calls identification a critical prerequisite for everything else (para 4.2). Where AI sits inside a third-party service, identification should at least cover services from material third-party providers.



Identification and inventory flow. Source MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 4.2-4.9, footnotes 21-23. Shadow AI is named in para 4.4. As of 20 Oct 2026.

An inventory record, the way MAS describes it

PURPOSE AND DESCRIPTION	Collections outreach, early arrears
APPROVED SCOPE OF USE	Singapore, retail, English and Mandarin
MODEL TYPE	Vendor LLM agent with tools
DATA USED	Account status, contact history
DEPENDENCIES	Vendor X model v4.2, telephony API
LIFE CYCLE STATUS	Live
RISK MATERIALITY RATING	High, approved by control function
MODEL REVIEW STATUS	Re-validation due Q2
KEY ROLES	Owner, deputy, developer
LINKS TO DOCUMENTATION	Validation report, test results
EXTRA ATTRIBUTES FOR AI AGENTS · FOOTNOTE 25	
AGENT IDENTIFIER	agt-col-014
TOOLS AND SYSTEMS IT CAN REACH	CRM write, payment link
COMPONENTS	Planner, retriever, voice
GUARDRAILS	No settlement offers without approval

Attributes from para 4.7 and footnote 25. Example values are fictional. MAS lists attributes an FI may include, not a fixed schema.

Roles

A designated control function sets the identification and inventory policies and keeps independent oversight. Business units may do the actual identification and updates (paras 4.3, 4.9).

Maintenance

Set a frequency for updating the inventory as use cases are added, changed or decommissioned (para 4.5). Review the inventory design itself as AI changes, for agents in particular (para 4.8).

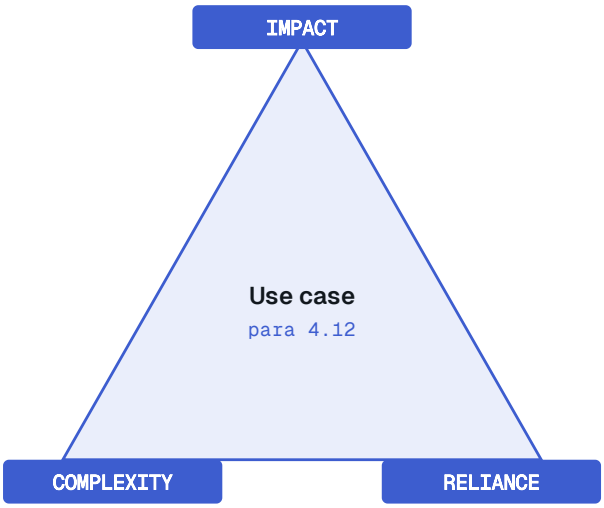
Linkage to other registers

Use consistent identifiers so the AI inventory cross-references data inventories and third-party or outsourcing registers (para 4.6, footnote 23).

> The inventory record anchors the framework. Ratings, controls, reviews and incidents are all referenced to it.

Impact, Complexity and Reliance, assessed before and after controls.

Each AI use case should be rated with one consistent methodology, before and after controls. The residual rating should sit within risk appetite before the use case goes live (para 4.10-4.11).



IMPACT

Consequences of failure for the FI, financial to reputational, and for customers, fairness to protection. Sensitivity of the data processed.

COMPLEXITY

Nature of the technology, novelty of the use, data, explainability. For vendor AI, how much you can see into the model and its data.

RELIANCE

How far the decision depends on AI output, the autonomy granted, and the human oversight in place, including reliance on a vendor.

Dimensions as defined in MAS Guidelines on AI Risk Management, 7 Oct 2026, para 4.12. MAS says these are the minimum, not an exhaustive list. As of 20 Oct 2026.



Assessment flow. Source paras 4.11 and 5.2. As of 20 Oct 2026.

Weighting is left to the institution

MAS will not prescribe how to combine the three dimensions into one rating, nor define a universal threshold for high Response 7.19. A simple model that can materially affect clients can be rated higher than a complex one that is not customer-facing.

Changes from the consultation

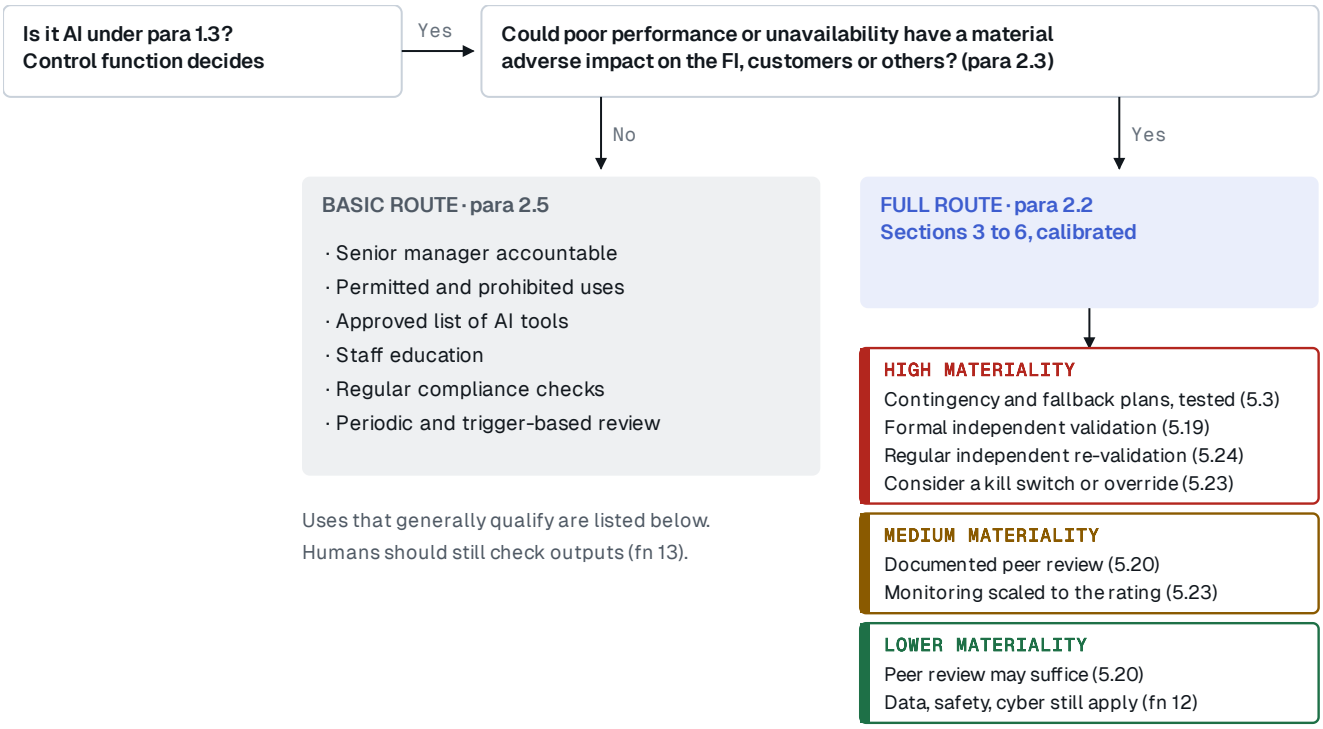
MAS removed availability of alternatives from Reliance. Reliance now centres on dependence on AI outputs, autonomy, human oversight and vendor dependence Response 7.13. Fallbacks still matter for high risk (para 5.3).

A designated control function

A designated control function sets the methodology and approves or arbitrates each rating. Business units may do the assessment, but the control function keeps it consistent and documented (para 4.13).

A basic governance route for AI of limited materiality.

If AI failing would not materially harm the FI or its customers, an FI may apply basic AI governance policies instead of the full set. Everything else is scaled to the materiality rating.



EXAMPLES THAT GENERALLY QUALIFY FOR THE BASIC ROUTE · PARA 2.4

- Drafting or proofreading emails to customers
- Summarising documents or notes for internal use
- First review of documents for internal use
- Formulas, charts or visuals for internal use
- Image generation for marketing material
- Chatbots that help staff find internal policies

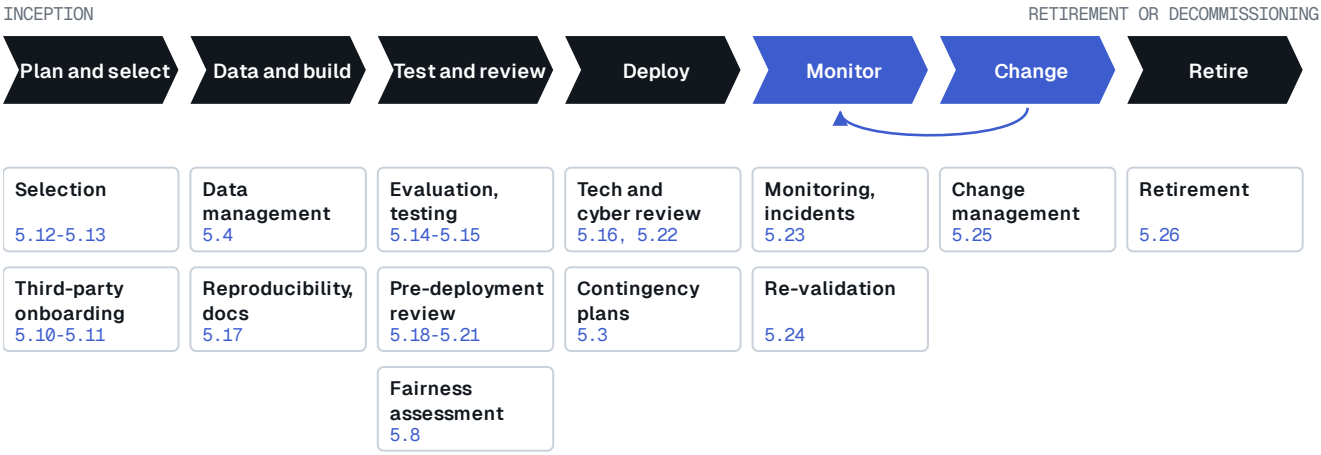
Decision path and controls ladder. Source MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 1.3, 2.2-2.5, 5.3, 5.19-5.20, 5.23-5.24, footnotes 12-13. Labels high, medium and lower are illustrative, MAS sets no fixed tiers. As of 20 Oct 2026.

Eligibility is reviewed, not fixed. The basic route is not a one-off classification. Para 2.5(f) expects periodic and trigger-based review of whether AI use still meets the criteria, and the basic route has to be complete by 7 October 2027.

Relevance as well as rating. Under the full route, some controls may matter less for some uses (para 2.2(b), footnote 12). Explainability and fairness weigh most in credit scoring, underwriting and advice. Data, safety and cybersecurity controls stay relevant even for copilots.

Controls across the life cycle, from inception to retirement.

MAS adopts the ISO/IEC 22989 view of the life cycle, from inception to retirement (footnote 27). Section 5 places twelve headed control areas on it, each applied in proportion to the use case's rating (para 5.2).



ACROSS THE WHOLE LIFE CYCLE	Human oversight, designed in and reviewed	5.9
ACROSS THE WHOLE LIFE CYCLE	Transparency and explainability to the audience	5.5-5.6
ACROSS THE WHOLE LIFE CYCLE	Fairness, defined by the FI	5.7
ACROSS THE WHOLE LIFE CYCLE	Partial deployments and pilots, with limits	fn 28

Placement of Section 5 controls on life cycle stages is Winsen Labs' reading. Paragraph numbers from MAS Guidelines on AI Risk Management, 7 Oct 2026. Expected by 7 Oct 2028, sooner for high-risk use cases. As of 20 Oct 2026.

Limited vendor disclosure

Identify the risks the gap creates and mitigate, for example by limiting use. Residual risk should still sit within appetite (para 5.2).

Pilots and partial deployments

Partial deployments may need adjusted controls, governed by clear policies with time and user limits, success criteria and close monitoring (footnote 28).

Records as a control

Keep records from development to incident management, and review human oversight decisions, including near misses (para 5.9(d)).

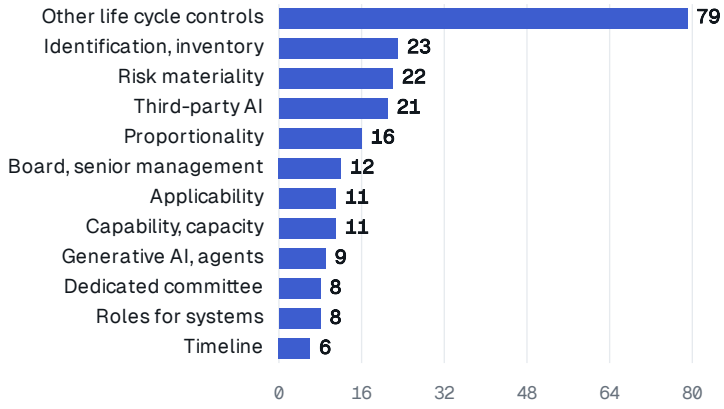
WHERE THE WEIGHT SITS

Expectations concentrate at planning and monitoring, and third-party AI spans every stage.

Reading the Guidelines theme by theme against the life cycle shows two clusters. Governance work concentrates at planning, and operating work at monitoring. Third-party AI is the only theme heavy at almost every stage.



Shade shows the strength of a theme's expectations at each stage, from none to strong (dots 0 to 3). This is Winsen Labs' editorial reading of paragraphs cited at right, not a MAS classification. Totals sum the dots per stage. As of 20 Oct 2026.



Count of numbered paragraphs per topic in MAS Response to Feedback P017-2025, 7 Oct 2026, excluding the preface. Total 226. A measure of where respondents pressed, not of importance.

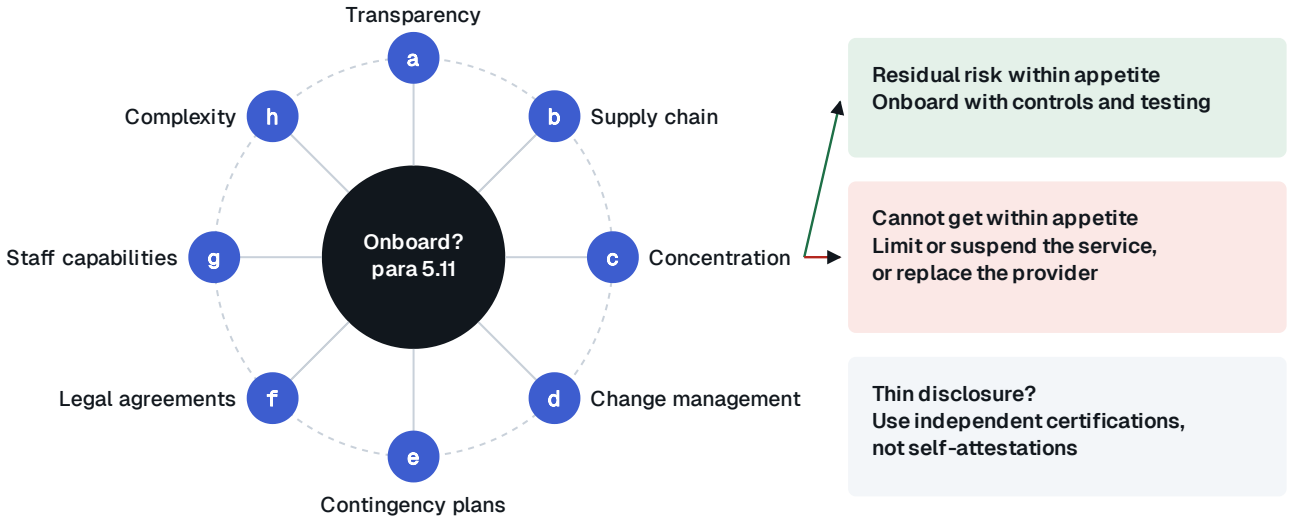
Reading the response paper

MAS answered 135 respondents, 40 of whom asked for confidentiality Response Annex A. Life cycle controls drew the most discussion, followed by identification, materiality and third-party AI.

These are the areas where industry found the expectations hardest to read, and where supervisors are likely to probe first.

Accountability for third-party AI remains with the institution.

The decision to onboard third-party AI is the FI's own, and the FI keeps primary accountability for its effects on customers and on its ability to comply (para 5.11).



Considerations (a) to (h) from MAS Guidelines on AI Risk Management, 7 Oct 2026, para 5.11. Outcome paths from para 5.11 and 5.11(a). As of 20 Oct 2026.

Testing on the institution's own data

Third-party AI should be reviewed in the context of the FI's own use cases, including with the FI's own data, with compensatory testing where the provider discloses too little. Keep the record of how it was tested and judged suitable (para 5.10).

Visibility of changes

Contracts should give risk-appropriate visibility over the introduction of AI and later updates (para 5.10). Consider clauses on performance, data protection, audit rights, notification of AI changes, and seeking the FI's agreement before adding AI (para 5.11(f)).

Concentration risk

Over-reliance on key providers, direct or indirect, is a named consideration (para 5.11(c)). MAS's own example of a risk appetite measure is the number of material use cases depending on a single provider (footnote 18).

Outsourcing requirements also apply. MAS outsourcing and third-party expectations apply to third-party AI as well (footnote 39). If an AI vendor is also an outsourced relevant service, Notice 658 duties apply in full.

Human oversight that can be evidenced, and monitoring that extends to agents.

Human oversight should be designed in, staffed by people with authority to intervene, and reviewed against what actually happened. Monitoring should cover every deployed AI, third-party included (paras 5.9, 5.23).

Human oversight, four parts · para 5.9

Roles	Who oversees, how it escalates, who decides.
Capabilities	Competent people with the authority and ability to intervene.
Design	Systems built to route outputs to human review when preset conditions are met.
Records and review	Logs kept, and oversight decisions and interventions reviewed, near misses included.

Post-deployment monitoring · para 5.23-5.24

Monitoring is critical because AI can go stale as data and concepts drift. Its frequency and intensity should scale with the rating, and it covers third-party AI used in the FI.

Consider automation bias and decision fatigue as AI scales. In the loop or over the loop, both count as human oversight (para 5.9, footnote 37).

Name an accountable person for monitoring and incidents, and train both monitors and users to report unintended behaviour (para 5.23(c), (e)).



Monitoring loop from MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 5.23(a)-(e) and 5.24. As of 20 Oct 2026.

AI agents

Monitor information flows and decision paths across workflows, such as reasoning, actions taken and tools used (para 5.23(a)). MAS says this helps pinpoint the source of a failure across chains of agents Response 12.8.

Records to retain

Monitoring results, issues, incidents and remediation. Strictly controlled and logged access to models, data, pipelines and configuration. For generative AI and agents, consider prompts, responses, model versions and reasoning (para 5.23(d)).

Kill switches and overrides

For high-materiality AI, consider a kill switch or override to deactivate rapidly when risk tolerances are exceeded (para 5.23(b)). Where one exists, test its activation protocol regularly (para 5.3).

Defined thresholds before deployment, and re-approval on significant change.

Evaluation should set clear, measurable thresholds agreed by business owners, developers and reviewers. Reviews before deployment scale in scope and independence with the rating (paras 5.14, 5.18-5.19).

LOWER MATERIALITY · para 5.20

Documented review, such as peer review by qualified people not involved in development or deployment, on key aspects of the process.

HIGH MATERIALITY · para 5.19

Formal independent validation before deployment, with effective challenge, covering design soundness, data, implementation, testing, explainability and fairness, and assumptions and limits. Then regular re-validation by independent parties (5.24).

Review intensity. Source MAS Guidelines on AI Risk Management, 7 Oct 2026, paras 5.18-5.21, 5.24. As of 20 Oct 2026.

Evaluation and testing · 5.14-5.15

- Measures aligned to objectives, with agreed thresholds
- Out-of-sample and out-of-time tests, sensitivity, stability, sub-populations, stress and adversarial cases
- Avoid overfitting, prefer simpler models unless uplift is clear
- For generative AI and agents, test key failure modes and guardrails

Documentation · 5.17

- Detailed enough for a reviewer or auditor to understand and potentially replicate
- Data, selection rationale, training code versions and settings
- Measures, thresholds, results, explainability and fairness work
- Assumptions, limitations, mitigants

Change · 5.25

- A framework to judge the significance of a change
- Significant changes trigger review and re-approval
- Where change can happen without prior review in high-materiality use, apply compensating controls such as enhanced monitoring
- Stricter controls for automatic updates

Review findings. Reviewers report limitations, remediation and conditions for use to the approval body, which should make sure they are actioned (para 5.21).

Security reviews before deployment. Technology and cyber reviews cover secure design, least privilege, deployment checklists, penetration testing and red teaming against manipulation, evasion and poisoning (para 5.22).

WHAT A SUPERVISOR WILL LOOK FOR

Likely lines of supervisory enquiry, and the evidence that addresses them.

MAS has not published examination procedures for these Guidelines. The questions below follow from the text itself, grouped by theme, with the evidence that would address each.

OVERSIGHT

Show me the board's approval of your AI governance approach and risk appetite.
EVIDENCE Board minutes, appetite statement with AI measures (3.4)

How do AI incidents and threshold breaches reach senior management and the board?
EVIDENCE Escalation procedure, incident log, board updates (3.5(e)-(f))

If you rely on a group framework, how does Singapore management see local AI risk?
EVIDENCE Local reports and escalation paths (3.6)

IDENTIFICATION AND INVENTORY

How did you find AI embedded in your material vendors' services?
EVIDENCE Identification method, vendor responses, gaps and mitigants (4.2, 4.4)

Who decides whether something is AI, and where is that decision recorded?
EVIDENCE Designated control function, documented rulings (4.3)

Walk me through one inventory record and its links to your outsourcing register.
EVIDENCE Inventory entry, linked identifiers (4.6-4.7)

MATERIALITY

How are Impact, Complexity and Reliance scored and combined?
EVIDENCE Methodology, weighting rationale (4.10, 4.12)

Show me a residual rating that was approved before deployment.
EVIDENCE Inherent and residual ratings, approval record (4.11, 4.13)

LIFE CYCLE

Which high-risk use cases had formal independent validation before go-live?
EVIDENCE Validation reports, independence evidence (5.19)

What thresholds trigger an alert, and what happened the last time one fired?
EVIDENCE Monitoring metrics, tiered thresholds, incident trail (5.23)

Show me reviews of human overrides, including near misses.
EVIDENCE Oversight logs and periodic reviews (5.9(d))

THIRD-PARTY AI

How did you test this vendor's AI on your own data?
EVIDENCE Test plan, results, suitability record (5.10)

How would you know if the vendor changed the model?
EVIDENCE Contract clauses, change notices, monitoring (5.11(d))

What is your concentration on any single AI provider?
EVIDENCE Dependency count by provider (5.11(c), fn 18)

Questions drafted by Winsen Labs from MAS Guidelines on AI Risk Management, 7 Oct 2026. Paragraph references show the source of each expectation. Not a MAS checklist. As of 20 Oct 2026.

WHERE INSTITUTIONS OFTEN FALL SHORT

Visibility, rather than control design, is where institutions most often fall short.

The response paper records the areas respondents found most challenging. Read with the Guidelines, it indicates recurring weaknesses. These describe patterns that typically appear, not findings about any institution.

■■■	**Embedded AI in vendor services stays undisclosed** Vendors do not say where AI sits in a SaaS product, so identification stops at what was declared.	4.2, 4.4, Response 13.2
■■□	**Shadow AI outside approved tools** Staff use public tools or features no one switched on deliberately.	4.4, fn 22
■■■	**Retrofitting AI already in production** Ratings and controls are applied to new use cases first, while deployed AI waits.	Response 13.2, 13.5
■■■	**Thin vendor disclosure on models and data** Complexity cannot be rated and testing is left to the provider's claims.	4.12(b), 5.10, 5.11(a)
■■□	**Talent for AI risk management** Second-line reviewers lack depth to challenge complex or generative AI.	6.1, Response 13.2
■■□	**Group frameworks that miss local needs** Local management lacks the data to show MAS how it oversees.	3.6, Response 13.2
■■□	**Oversight that cannot be evidenced** Reviews happen but decisions, overrides and near misses are not logged.	5.9(d)
■■□	**Monitoring that averages away agents** Model metrics are watched while agent actions and tool use are not.	5.23(a), Response 12.8

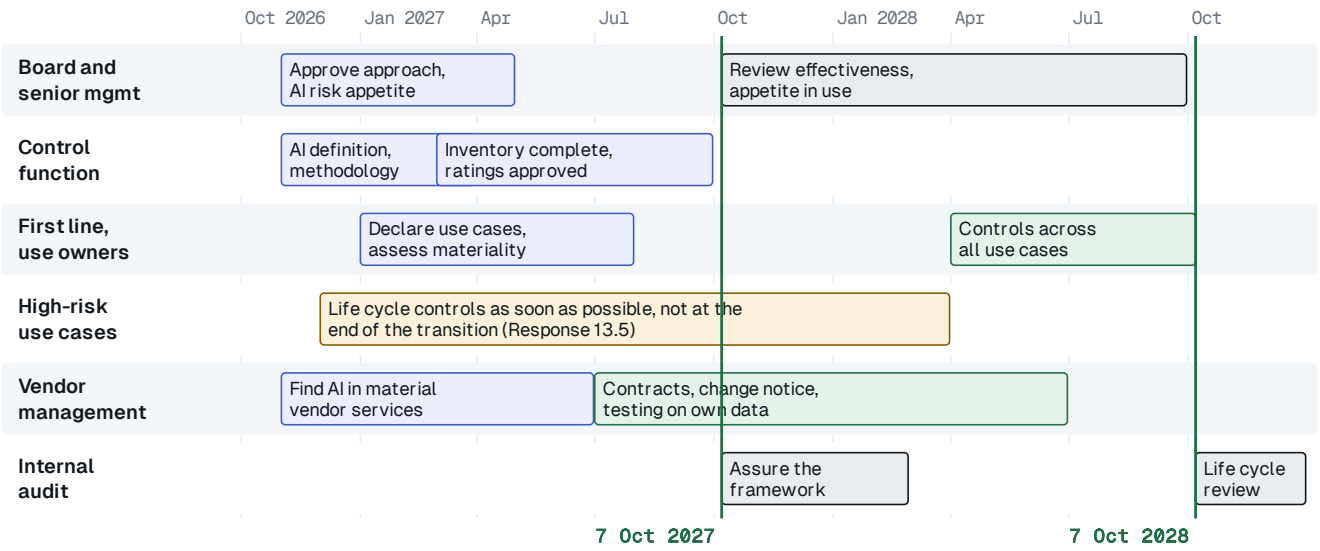
Gaps drawn from challenges respondents raised in MAS Response to Feedback P017-2025, 7 Oct 2026 (para 13.2 and others cited) and from constraints MAS itself names in the Guidelines. Bars show Winsen Labs' judgement of how often each typically appears, from occasional to common. Not survey data. As of 20 Oct 2026.

> **A common thread.** Most of these gaps are, at root, gaps in visibility. Where an institution cannot see a use case, a vendor's changes or the actions of its reviewers, documented controls have nothing to operate on.

Framing gaps internally. An identified gap is a work item with an owner and a target date, not a breach. The Guidelines set expectations, and supervisors are likely to look for a credible path as well as the end state.

A phased path to both implementation dates.

Ownership and discovery come first, as every later step depends on them. High-risk use cases should be brought forward, as MAS expects. The plan below is illustrative and should be adapted to the institution's estate.



Illustrative sequencing by Winsen Labs. Only the two green markers and the high-risk expectation come from MAS (para 1.8, Response 13.5). As of 20 Oct 2026.

Seven steps

1. **Assign ownership** 3.4, 3.5, 4.3
   Board or committee, accountable senior manager, designated control function.
2. **Define and identify AI** 1.3, 4.2
   Agree the definition, then sweep in-house, vendor and SaaS-embedded AI.
3. **Classify the estate** 2.3-2.5
   Separate basic-route uses from the rest, and write down why.
4. **Build the inventory** 4.5-4.7
   One record per use case, linked to data and vendor registers.
5. **Assess every use case** 4.10-4.13
   Impact, Complexity, Reliance, inherent and residual, approved.
6. **Prioritise high-risk use cases** 5.3, 5.19, 5.23
   Validation, contingency, monitoring and kill switch decisions.
7. **Evidence effectiveness** 5.9(d), 5.23(d)
   Logs, oversight reviews, incident trail, audit assurance.

The first ninety days

- Board paper proposing governance approach and owners
- Draft AI definition, tested on ten real cases
- Request to material vendors on AI in their services
- Inventory template with para 4.7 attributes
- Materiality method drafted and piloted on five use cases
- List of likely high-risk use cases with interim controls

Engage internal audit early on the evidence it will expect in late 2027, as this determines what should be recorded now.

SELF-ASSESSMENT

Twenty-four checks against the Guidelines.

Mark only what can be evidenced today, not what is planned. Each unmarked item should become a work item with an owner and a date. Items under Sections 3 and 4 apply from 7 October 2027.

SECTIONS 2 TO 4 · FROM 7 OCT 2027

☐	Board or committee has approved the AI governance approach	3.4(a)
☐	AI risks appear in the risk appetite, with measures	3.4(b)
☐	Escalation path for AI incidents and threshold breaches	3.5(e)
☐	Local management can evidence oversight to MAS	3.6
☐	Written AI definition and identification process	4.2
☐	Control function named as arbiter of what is AI	4.3
☐	Material vendors asked about embedded AI	4.2
☐	Mitigants for shadow AI and unidentified use	4.4
☐	Inventory with para 4.7 attributes, and update frequency	4.5-4.7
☐	Inventory linked to data and outsourcing registers	4.6
☐	Materiality method covering Impact, Complexity, Reliance	4.12
☐	Inherent and residual ratings, approved before go-live	4.11, 4.13
☐	Basic-route uses listed with review triggers	2.3-2.5

SECTIONS 5 AND 6 · BY 7 OCT 2028

☐	Contingency plans for high-risk use cases, tested	5.3
☐	Human oversight roles, design and records	5.9
☐	Third-party AI tested on own data, documented	5.10
☐	Contracts give visibility of AI changes	5.10, 5.11(f)
☐	Independent validation for high-risk use cases	5.19
☐	Monitoring with tiered thresholds and drift checks	5.23(a)
☐	Logs for generative AI and agents considered	5.23(d)
☐	Change framework with re-approval triggers	5.25
☐	Retirement controls, including data removal	5.26
☐	Competence and resources reviewed	6.1-6.2
☐	Infrastructure fit for each use case	6.3

Checks drafted by Winsen Labs from MAS Guidelines on AI Risk Management, 7 Oct 2026. As of 20 Oct 2026.

GLOSSARY

AI
Machine-based systems or models deriving outputs through learned premises (1.3).

Use case
A specific real-world context a model or system is applied to (1.4).

Third-party AI
AI developed, owned, operated or provided by an external party that the FI relies on (fn 10).

Risk materiality
The rating of an AI use case on Impact, Complexity and Reliance (4.12).

Residual risk
Materiality after controls. Should sit within appetite before deployment (4.11).

AI incident
Failure, malfunction or poor performance of AI (fn 20).

Kill switch
A mechanism to deactivate AI quickly when it exceeds tolerances (fn 29).

Shadow AI
AI used without approval or outside identified channels (1.10(f), 4.4).

SUMMARY REFERENCE

The MAS Guidelines on AI Risk Management in summary.

◆ GUIDANCE · FINAL | ISSUED 7 OCT 2026 | AS OF 20 OCT 2026

DATES

7 Oct 2027

Sections 3 and 4 apply, oversight, identification, inventory, materiality. Basic route complete.

7 Oct 2028

Sections 5 and 6 apply, life cycle controls and capability.

High-risk use cases as soon as possible.

WHO IT COVERS

All FIs under FSMA 2022, proportionately. Group basis for locally incorporated FIs under consolidated supervision or owning CII. All AI, generative AI and agents included. Rule-based tools and RPA ordinarily excluded.

LEGAL EFFECT

Guidance. Not a criminal offence or civil penalty, but observance weighs in MAS's risk assessment.

RATE EVERY USE CASE

Impact on FI and customers, data sensitivity.
Complexity of technology, data, explainability, vendor visibility.
Reliance on AI output, autonomy, human oversight.

Inherent and residual. Residual within appetite before deployment. Control function approves.

BASIC ROUTE · PARA 2.3-2.5

If AI failing would not be material. Accountable senior manager, permitted and prohibited uses, approved tools list, staff education, compliance checks, periodic review of eligibility.

HIGH MATERIALITY

Contingency plans tested. Formal independent validation before go-live. Regular independent re-validation. Consider a kill switch.

INVENTORY ATTRIBUTES · PARA 4.7

Purpose, approved scope, model type, data, dependencies, life cycle status, risk rating, review status, owners, documentation links. For agents add identifiers, tools and systems reachable, components, guardrails.

THIRD-PARTY AI · PARA 5.10-5.11

You stay accountable. Test on your own data. Get change visibility in contracts. Watch concentration. If not within appetite, limit, suspend or replace.

TWELVE CONTROL AREAS IN SECTION 5

Data 5.4	**Transparency** 5.5	**Fairness** 5.7	**Human oversight** 5.9	**Third-party** 5.10	**Selection** 5.12
Testing 5.14	**Tech and cyber** 5.16	**Reproducibility** 5.17	**Pre-deploy review** 5.18	**Monitoring** 5.23	**Change** 5.25

FIVE LIKELY SUPERVISORY QUESTIONS

1. Show the board's approval and AI risk appetite.
2. How did you find AI inside vendor services?
3. Walk me through one inventory record and rating.
4. Which high-risk use cases were independently validated?
5. How would you know if a vendor changed its model?

Summary of MAS Guidelines on AI Risk Management, 7 Oct 2026 and MAS Response to Feedback P017-2025, 7 Oct 2026. Not legal advice. As of 20 Oct 2026.

What Winsen Rho can do for you

Effortless accountability for every AI action.

You know what MAS expects. Rho records every AI agent you run, every decision it makes, and every outcome it leads to, so the evidence is already there when your board, your auditors or your supervisor ask.

01

Know your AI

Every agent, owned and accounted for, including the ones nobody declared.

02

Prove it

Every decision reopenable years later, signed and reproducible.

03

Answer anyone

Board, audit, supervisor, from one record.

> Evidence for your teams, not a certificate. Rho supplies the evidence. Your institution decides, and remains accountable.

What Winsen Rho can do for you

What Rho does, in six capabilities.

Each capability produces records that can be placed before a reviewer. The references indicate where MAS looks for evidence of this kind.

Inventory

A living register of every AI agent, built in-house, on a vendor's model, run by a vendor, or embedded in a SaaS product. Discovery finds agents nobody declared. Every entry has a named owner, a risk tier and a lifecycle state.

para 4.2-4.9

Decisions

What each agent saw, decided and did, under whose authority, and what happened next. Customer identity is tokenised inside your network. Any decision can be reconstructed as it was.

para 5.9(d), 5.23(d)

Cost and value

Each agent measured on outcomes from your own systems, compared fairly with your own team, with outcomes held back until they mature.

para 3.4(b), 5.14(a)

Vendors

Every vendor's agents measured on your customers, set against what the vendor claims and what the contract says. Unnotified model changes and concentration are surfaced.

para 5.10-5.11

Human oversight

How each agent is overseen, who reviewed and overrode what, and whether that review improved outcomes.

para 5.9

Drift

Nightly probes and thresholds that catch a silent model change, with tiered alerts and an owner on every finding.

para 5.23-5.25

References are to MAS Guidelines on AI Risk Management, 7 Oct 2026. A reference shows where the evidence is relevant, not that the expectation is met. Rho does not validate models, certify compliance or file reports.

THE SAME RECORD, FOR THE PRIVACY AND OUTSOURCING RULES AROUND IT

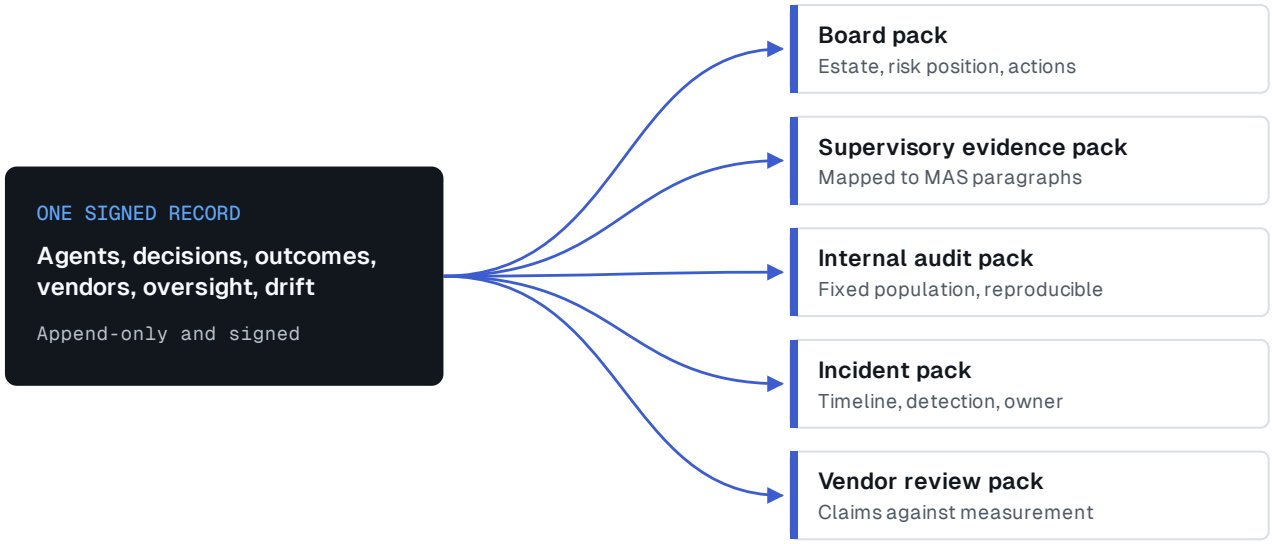
BINDING NOTICE **MAS Notice 658 and Guidelines on Outsourcing (banks)**	An inventory of vendor agents to feed your outsourcing register, vendor scorecards built only from your own outcomes, reproducible packs for auditors and MAS
BINDING NOTICE **MAS Notice FSM-N05 and Technology Risk Management Guidelines**	An inventory of agents and the systems they can reach, alerts when an agent uses permissions it should not have, signed incident packs for root-cause work
BINDING LAW **Personal Data Protection Act 2012**	Decision records showing which personal data each agent used, alerts on permission breaches involving personal data, packs for breach assessments and access requests

Evidence descriptions as published on winsenrho.com, reviewed 24 Sep 2026. Some services are exempt from outsourcing rules, so not every AI vendor is an outsourced relevant service. The duties, notifications and filings stay with your institution.

What Winsen Rho can do for you

From the record to the people who ask.

Each evidence pack is generated from the same signed record, so the board, internal audit and the supervisor work from identical figures.



Pack names as produced by Winsen Rho today. Figures in any pack resolve to record IDs in the signed record.

Signed, retained, reproducible

Packs are signed by the named second-line officer, retained for the retention period your institution sets, and can be regenerated later to show the same figures.

Sample reports

Two sample reports, prepared for the fictional Southern Bank, show the format and depth of a pack. The RBI FREE-AI sample evidence report. The MAS AIRG sample AI risk report.

BUILT FOR A BANK'S SECURITY REVIEW

Data stays in your region.
Only allowlisted traffic, to endpoints in your region.

Customer identity never leaves your bank.
Tokenised by the Bridge inside your network.

The keys stay with you.
Customer tokens are keyed in your HSM, so Rho Cloud cannot re-identify a person.

Nobody can quietly edit the record.
Append-only and signed. Every read is logged.

A pilot, the demo and the sample reports are available. Security documentation is shared under NDA as part of your third-party review.

What Winsen Rho can do for you

MAS's expectations, and where Rho fits.

THE EXPECTATION	The evidence Rho holds	WHAT STAYS WITH YOUR INSTITUTION
Identify all AI, including AI in vendor services para 4.2-4.4	Discovery of agents nobody declared, classed as in-house, on a vendor model, vendor-operated or embedded in SaaS	The AI definition, the control function's rulings, controls for shadow AI outside agent traffic
Keep an accurate AI inventory para 4.5-4.9	A living register with named owner and deputy, risk tier, life cycle state, model and version, permitted actions, record completeness	Inventory policy and update frequency, links to data and outsourcing registers
Rate materiality on Impact, Complexity, Reliance para 4.10-4.13	Autonomy, authority and oversight mode per agent, vendor visibility, the recorded tier and its approvals	The methodology, the ratings, and their approval by the control function
Oversight, appetite and escalation para 3.4-3.5	Board pack from the signed record, incidents with an owner, counts usable as appetite measures such as single-provider dependence	The appetite, roles, policies and every decision
Human oversight, recorded and reviewed para 5.9	Who reviewed, approved or overrode, whether overrides improved outcomes, reviewer training currency	The oversight design, authority to intervene, staffing
Third-party AI tested on your data, changes and concentration para 5.10-5.11	Vendor agents measured on your customers against claims and contract, unnotified model changes, concentration by vendor and model	Due diligence, contracts, the onboarding decision, compensatory testing design
Monitoring, thresholds, drift and agent paths para 5.23	Nightly probes, tiered thresholds, drift alerts, decision paths across agents, prompts, responses and versions	Choice of metrics and thresholds, remediation
Kill switch and contingency for high risk para 5.3, 5.23(b)	Pause or restrict an agent through your gateway, with the decision logged	Contingency plans, fallbacks, activation testing
Change management and re-validation triggers para 5.24-5.25	Model and version change record, alerts that can trigger re-validation, enhanced monitoring as compensating evidence	The change framework, re-approval, and validation itself

Rho does not cover these. They stay with you.

- Data management and governance (5.4)
- Customer transparency and explainability methods (5.5-5.6)
- Fairness definitions and assessments (5.7-5.8)
- Model selection (5.12-5.13)
- Pre-deployment testing and independent validation (5.14-5.21)
- System security and cyber reviews (5.16, 5.22)
- Development documentation (5.17)
- Secure retirement and data removal (5.26)
- Competence, resourcing, infrastructure (6.1-6.3)

NEXT STEP

See your own estate through MAS's lens.

Book a 30-minute walkthrough.

hello@winsenrho.com
winsenrho.com
demo.winsenrho.com

Mapping maintained by Winsen Labs against MAS Guidelines on AI Risk Management, 7 Oct 2026, as of 20 October 2026. It is a reading of where evidence is relevant. Your institution decides which controls answer which expectation.

SOURCES AND DISCLAIMER

Primary sources relied on in this guide.

Primary sources only. Each was read on 20 October 2026. Secondary commentary was used for orientation and is not relied on for any fact in this guide.

STATUS	SOURCE	NOTE
◆ GUIDANCE	**Guidelines on Artificial Intelligence Risk Management, 7 October 2026** https://www.mas.gov.sg/-/media/mas-media-library/regulation/guidelines/bd/guidelines-on-artificial-intelligence-risk-management-for-financial-institutions/guidelines-on-artificial-intelligence-risk-management.pdf	Read in full Read 20 Oct 2026
◆ GUIDANCE	**Response to Feedback Received on Guidelines on AI Risk Management, P017-2025, 7 October 2026** https://www.mas.gov.sg/-/media/mas-media-library/publications/consultations/bd/2025/response-to-feedback-received-on-guidelines-on-artificial-intelligence-risk-management.pdf	Read in full Read 20 Oct 2026
◆ GUIDANCE	**MAS media release, 7 October 2026** https://www.mas.gov.sg/news/media-releases/2026/mas-sets-out-supervisory-expectations-on-responsible-ai-adoption-by-financial-institutions	Via search copy, site blocks automated reads Read 20 Oct 2026
⬚ CONSULTATION	**Consultation Paper P017-2025, 13 November 2025** https://www.mas.gov.sg/-/media/mas-media-library/publications/consultations/bd/2025/final_consultation_paper_on_guidelines_on_ai_risk_management_forrelease.pdf	Superseded by final text Read 20 Oct 2026
○ VOLUNTARY	**Information Paper on AI Model Risk Management, 5 December 2024** https://www.mas.gov.sg/publications/monographs-or-information-paper/2024/artificial-intelligence-model-risk-management	Context only Read 20 Oct 2026
○ VOLUNTARY	**FEAT Principles, 2018** https://www.mas.gov.sg/publications/monographs-or-information-paper/2018/feat	Cited in para 1.2 Read 20 Oct 2026
■ BINDING	**Notice FSM-N05 on Technology Risk Management** https://www.mas.gov.sg/regulation/notices/notice-fsm-n05	Re-check, date from Winsen library Read 20 Oct 2026
■ BINDING	**Notice 658 on Management of Outsourced Relevant Services (banks)** https://www.mas.gov.sg/regulation/notices/notice-658	Re-check, date from Winsen library Read 20 Oct 2026
■ BINDING	**Personal Data Protection Act 2012** https://sso.agc.gov.sg/Act/PDPA2012	Statute Read 20 Oct 2026
◆ GUIDANCE	**MAS classification of regulatory instruments, legal effect of guidelines** https://www.mas.gov.sg/regulation	Re-check, read via reproduction Read 20 Oct 2026

Disclaimer. This guide is general information prepared by Winsen Labs, India, as of 20 October 2026. It is not legal, regulatory or compliance advice, and does not create any obligation or certify compliance. Regulatory texts change, so check the source before relying on any point. Questions drafted for supervisors, readiness plans, heat-grid weights and gap frequencies are Winsen Labs' judgement and are labelled as such. Examples are fictional and describe no real institution. Winsen Rho supplies evidence. Your institution decides, and remains accountable. Evidence for your teams, not a certificate.

Winsen Rho is built by Winsen Labs. winsenrho.com · hello@winsenrho.com