Winsen Rho

AI REGULATION GUIDE · INDIA

What the Reserve Bank of India expects of regulated entities using AI

A practitioner's guide to the FREE-AI framework, the draft Guidance on Regulatory Principles for Model Risk Management, and the binding rules that already apply, covering status, scope, supervisory expectations, and the evidence a regulated entity should be prepared to produce.

STATUS	WATCH FOR	AS OF
No binding AI direction yet	Final model risk guidance	20 October 2026

HOW TO READ THIS

Status, scope and conventions.

India has no binding, AI-specific direction for regulated entities as of 20 October 2026. What exists is a committee report, a draft guidance, and binding rules on outsourcing and personal data that already apply to AI. This guide keeps those three apart on every page.

0
BINDING AI-SPECIFIC DIRECTIONS

26
FREE-AI RECOMMENDATIONS

13
ADDRESSED TO RES

64
PARAGRAPHS IN THE DRAFT

13 May 2027
DPDP MAIN DUTIES APPLY

Status labels used throughout

Every statement is tied to its instrument. The shape tells you the status even in black and white.

■ BINDING	Acts, rules and RBI Directions. Written as *requires* or *must*.
⬚ DRAFT	The draft Model Risk Management Guidance and the draft Data Governance Guidance. Written as *the draft proposes*.
○ RECOMMENDATIONS	The FREE-AI committee report. Written as *the committee recommends*.

FREE-AI is not a direction. It is a report of recommendations by a committee set up by the RBI. It creates no obligation on its own, even where its text says *must*. This guide reports that wording as a recommendation.

References read `Rec 23` or `FREE-AI 4.4.69` for the report, and `Draft 21` for paragraphs of the draft guidance.

Contents

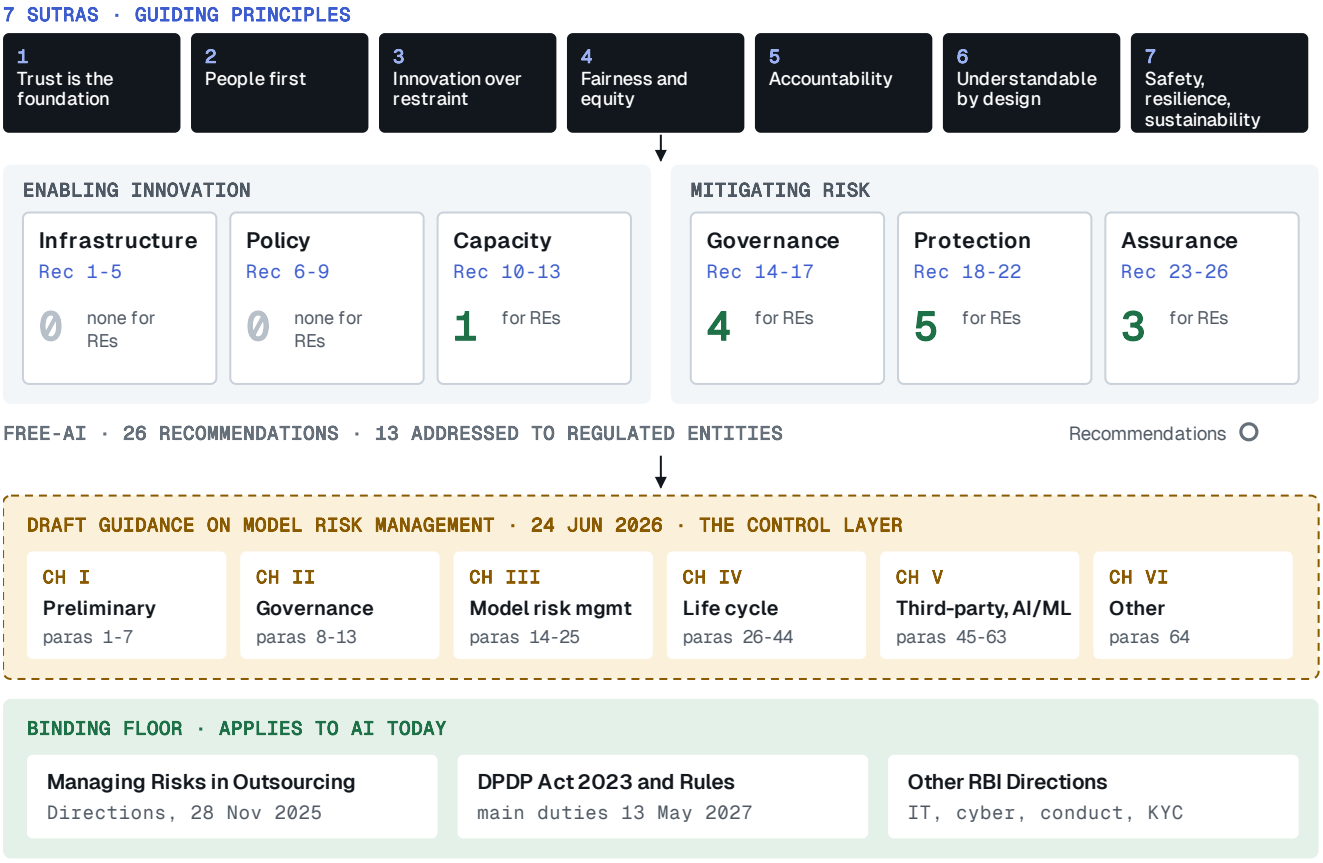
3	The stance on one page
4	Who and what is covered
5	Timeline
6	What is binding, and what is not
7	The expectations, theme by theme
17	What a supervisor will look for
18	Where institutions often fall short
19	How to get ready
21	Summary reference
22	What Winsen Rho can do for you
26	Sources and disclaimer

Intended readers. The person in a bank, NBFC or other regulated entity asked to work out what the RBI expects on AI, typically in model risk, technology risk or compliance. No prior reading is assumed.

THE STANCE ON ONE PAGE

Committee principles, draft controls, and binding law beneath both.

The RBI's position is built in layers. The FREE-AI report sets direction through seven principles and 26 recommendations. The draft model risk guidance proposes the controls. Binding outsourcing and data-protection rules already apply.



Structure of RBI FREE-AI Committee Report, 13 Aug 2025 (Sutras 4.3, pillars 4.4, Recommendations 1-26) and RBI draft Guidance on Model Risk Management, 24 Jun 2026 (chapters I-VI). Counts of RE-addressed recommendations from the report's own tags. As of 20 Oct 2026.

Accountability

Sutra 5 says entities deploying AI remain fully accountable whatever the level of automation. The draft proposes the same for every model, bought or built Draft 8, 45.

An enabling stance

The committee's third sutra asks that responsible innovation be prioritised over cautionary restraint, other things being equal. The tone is enabling, the controls proportionate.

Further requirements

The draft notes that further requirements for AI models, if any, may be issued later, citing Utkarsh 2029 Draft 2.

WHO AND WHAT IS COVERED

Broad coverage of entities, and of models whether or not they use AI.

FREE-AI addresses regulated entities, regulators, government and industry in turn. The draft guidance applies to eleven types of RE and to all models they use, built internally, bought, or both (Draft 4, 6).

DRAFT MRM GUIDANCE APPLIES TO 11 TYPES OF RE · DRAFT 4

Commercial and foreign banks	Small finance banks	Payments banks
Local area banks	Regional rural banks	Urban co-operative banks
Rural co-operative banks	NBFCs, base to top layer	All-India financial institutions
Asset reconstruction companies	Credit information companies	

Applicability list from RBI draft Guidance on Model Risk Management, 24 Jun 2026, para 4. Draft, not final. As of 20 Oct 2026.

What counts as a model · Draft 7(3)

INPUT
Data and assumptions, empirical or judgement-based

→

PROCESSING
Statistical, financial or other techniques, AI and ML included

→

OUTPUT
Results used for business decisions

Includes algorithms, rule engines, interfaces and tools that materially affect decisions, whether or not the RE calls them models. A spreadsheet that sets lending rates counts (Draft 7(3), illustration).

Definition paraphrased from RBI draft Guidance on Model Risk Management, 24 Jun 2026, para 7(3). As of 20 Oct 2026.

A broad definition

The draft covers rule-based and spreadsheet tools as well as AI. For AI and ML models, including foundation and frontier models, it proposes added controls on top (Draft 49). Inventories built only for AI will miss part of the estate it describes.

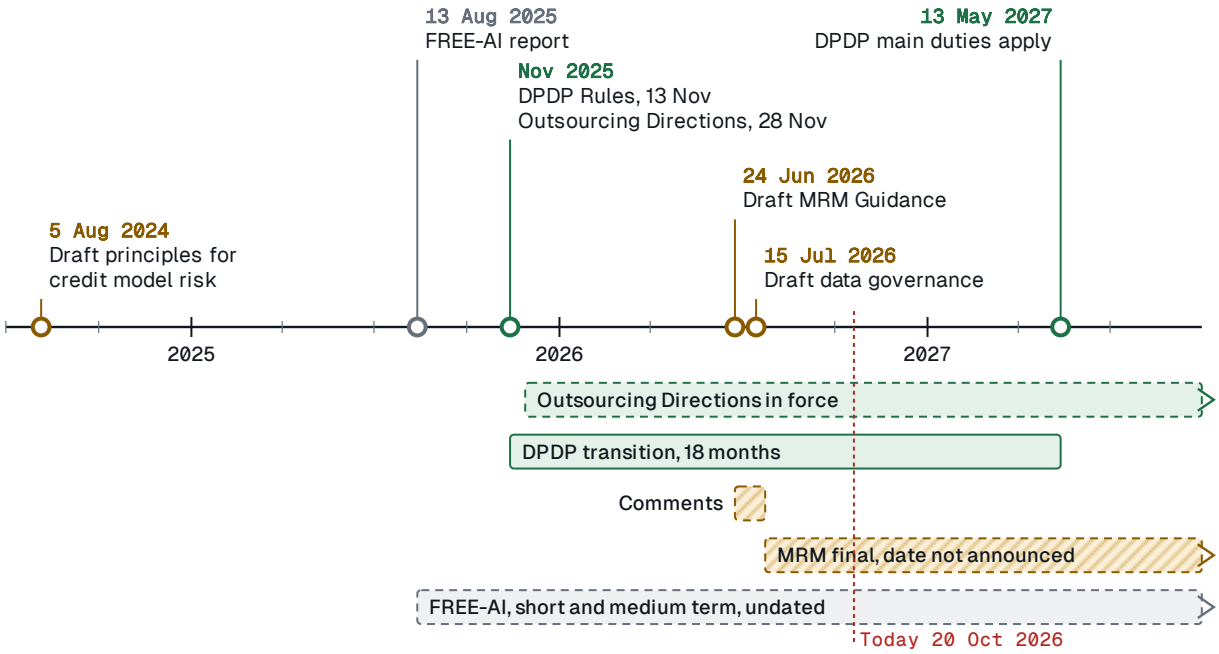
Providers and FinTechs

Providers are not addressed directly by the draft, but an RE that acquires, uses or relies on a third-party model is accountable for its outcomes at any stage of the life cycle (Draft 45). FREE-AI proposes incident reporting for FinTechs too Rec 22.

TIMELINE

Binding rules in force, and guidance yet to be finalised.

Two binding rules already reach AI. The model risk guidance is the next event, and no date has been set for it. FREE-AI tags its recommendations short term or medium term, without dates.

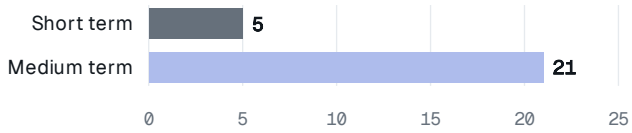


Dates from the RBI press release of 24 Jun 2026 (which also dates the 5 Aug 2024 draft), the RBI drafts listing (15 Jul 2026), RBI FREE-AI Committee Report, 13 Aug 2025, RE-wise Managing Risks in Outsourcing Directions of 28 Nov 2025, and DPDP Rules G.S.R. 846(E) of 13 Nov 2025. Dashed bars have no end date. Scale is linear. As of 20 Oct 2026.

No implementation clock has started for AI controls. Comments on the draft closed on 24 July 2026. As of 20 October 2026 we found no final text on rbi.org.in. Implementation dates will be set in the final guidance.

When final, the guidance would supersede Chapter 3 on credit risk models of the 2002 Guidance Note on Credit Risk Management Draft 64.

Recommendations by term



Count of the 26 recommendations by the term tag the report attaches to each. Source RBI FREE-AI Committee Report, 13 Aug 2025, section 4.4. Short-term items are Recs 1, 2, 9, 23 and 25.

The AI-specific texts are not yet binding. The surrounding rules are.

A supervisor today can test an RE against its outsourcing and data duties, and will read the AI texts as statements of what good looks like. The draft also says that where it conflicts with a Direction, the Direction prevails (Draft 5).

BINDING LAW AND DIRECTIONS Must be met today. Already reach AI used by or for an RE.	**RBI Managing Risks in Outsourcing Directions, 2025**	RE-wise, 28 Nov 2025
	Digital Personal Data Protection Act 2023	MeitY
	DPDP Rules 2025	13 Nov 2025, phased
	Other RBI Directions on IT, cyber, conduct	as applicable
DRAFT GUIDANCE Proposals, not final. Signal what RBI will look for.	**Guidance on Regulatory Principles for Model Risk Management**	24 Jun 2026
	Guidance on Regulatory Expectations for Data Governance	15 Jul 2026
COMMITTEE RECOMMENDATIONS Not a direction. Sets direction of travel.	**FREE-AI Committee Report, 7 Sutras, 26 Recs**	13 Aug 2025
	Annexure IV, suggested AI changes to Directions	for reference

Instruments as listed on rbi.org.in and meity.gov.in, read 20 Oct 2026. The draft guidance has a comments-closed status, with no final text found. FREE-AI Annexure IV suggests AI-specific changes to Master Directions, which RBI has not adopted as such.

Language conventions

FREE-AI often says *must*. It is still a recommendation. This guide uses *recommends* for the report, *proposes* for drafts, and *requires* only for Directions and the DPDP Act.

Scope of the Outsourcing Directions

The Outsourcing Directions cover material outsourcing of IT services. Licensed or off-the-shelf software sits outside, so not every AI vendor is an outsourced service. Accountability for the outcome stays either way.

Personal data

AI that processes personal data is subject to the DPDP Act. The Rules bring notice, security safeguards and breach intimation duties into force on 13 May 2027.

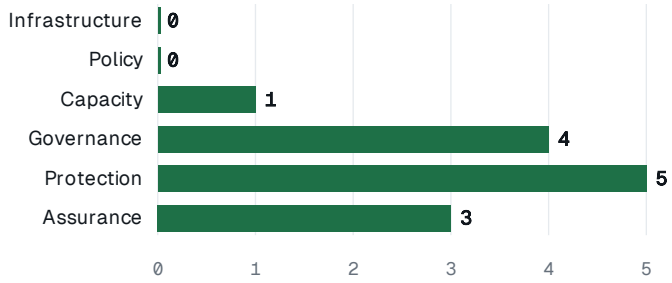
Thirteen recommendations are addressed to regulated entities, twelve of them on risk mitigation.

The report's 26 recommendations are tagged with who should act and when. Read only the green ones as being about your institution. The rest ask regulators, government or industry to act.

INFRASTRUCTURE	POLICY	CAPACITY	GOVERNANCE	PROTECTION	ASSURANCE
Rec 1 S Data infrastructure	Rec 6 M Enabling policies	**Rec 10 M Capacity in REs**	**Rec 14 M Board AI policy**	**Rec 18 M Consumer protection**	**Rec 23 S AI inventory**
Rec 2 S Innovation sandbox	Rec 7 M Affirmative action	Rec 11 M Regulator capacity	**Rec 15 M Data life cycle**	**Rec 19 M Cybersecurity**	**Rec 24 M AI audit**
Rec 3 M Incentives, funding	Rec 8 M Liability framework	Rec 12 M Sharing practices	**Rec 16 M AI system governance**	**Rec 20 M Red teaming**	**Rec 25 S Disclosures**
Rec 4 M Indigenous models	Rec 9 S Institutional framework	Rec 13 M Reward innovation	**Rec 17 M Product approval**	**Rec 21 M AI continuity plans**	Rec 26 M Compliance toolkit
Rec 5 M AI with DPI				**Rec 22 M Incident reporting**	

■ Addressed to REs, alone or jointly □ Addressed to regulators, government or industry S short term · M

Addressee and term tags as printed after each recommendation in RBI FREE-AI Committee Report, 13 Aug 2025, section 4.4. Short labels are Winsen Labs' abbreviations. As of 20 Oct 2026.



Count of recommendations whose tag includes REs, per pillar. Source as above.

Shared recommendations

Recs 22, 23, 24 and 25 are shared between REs and regulators or supervisors. This guide covers the RE part. Rec 10, capacity building within REs, is the only RE item on the enabling side.

The sample evidence report by Winsen Rho uses the same 13.

GOVERNANCE

Board-approved frameworks for model risk and for AI.

The draft proposes a Board-approved Model Risk Management Framework covering all models, AI and ML included, with three lines of defence (Draft 9-15). FREE-AI separately recommends a board-approved AI policy (Rec 14).

BOARD Draft 11	RISK MANAGEMENT COMMITTEE Draft 12	SENIOR MANAGEMENT Draft 13
Oversee the model risk framework Approve and review it, delegate to the RMCB as needed Approve model risk appetite, forward-looking, stress-tested Approve policies, tiering included	Review validation reports of high-risk models, approve deployment Review tiering reports, at least annually Oversee exception, third-party and AI models, review breaches	Resource and operationalise the framework Implement risk-based tiering Maintain the inventory and documentation Review policies, report to RMCB

Roles from RBI draft Guidance on Model Risk Management, 24 Jun 2026, paras 11-13. Draft, not final. As of 20 Oct 2026.

DRAFT · WHAT THE FRAMEWORK COVERS · DRAFT 10

Model taxonomy, governance structure, scope of use, tiering method, inventory and documentation standards, and policies for every life cycle stage from selection to decommissioning, with exceptions and their mitigants.

If a model's assessed risk exceeds appetite, the draft proposes timely action and a report to the RMCB (Draft 14).

RECOMMENDATION · WHAT THE AI POLICY COVERS · REC 14

Governance structure, accountability, risk appetite, operational safeguards, auditability, consumer protection, AI disclosures, a model life cycle framework and a liability framework. Industry bodies should give smaller entities a template.

Rec 10 adds AI competence for the board and C-suite, and training for staff who use AI.

Three lines, named. Model owners are the first line, an independent model risk management and validation function the second, and internal audit the third (Draft 15).

INVENTORY

Both texts treat the inventory as the foundation of oversight.

FREE-AI and the draft both make the inventory the base of oversight. Read together, the draft is stricter on scope and retention, and FREE-AI adds grievances and target groups.

	FREE-AI · REC 23, 4.4.68-4.4.69	DRAFT · PARAS 17, 21-24
SCOPE	All AI systems in use, models and use cases	All models, active, inactive, under development and decommissioned
KEY FIELDS	Model type and purpose, use cases, target groups, dependencies, risk category, grievances	Owners, developers, validators, approvers, risk tier, intended use, upstream and downstream dependencies, findings
UPDATE CADENCE	At least half-yearly	Up to date, tier reviewed at least annually
GATEKEEPING	Ready for supervisory inspection and audit	No model used, relied on or deployed unless in the inventory
RETENTION	Not specified	Decommissioned models kept for at least ten years
SECTOR VIEW	Regulators to build a sector-wide AI repository	Not addressed

Side-by-side comparison. Sources RBI FREE-AI Committee Report, 13 Aug 2025, Rec 23 and paras 4.4.68-4.4.71, and RBI draft Guidance on Model Risk Management, 24 Jun 2026, paras 17, 21-24. As of 20 Oct 2026.

A MODEL'S LIFE IN THE INVENTORY · DRAFT 21, 23



Ten years runs from decommissioning, or from when the model stops serving as a backup or benchmark, whichever is later, or longer if law requires. Documentation should be kept for the same period (Draft 23-24).

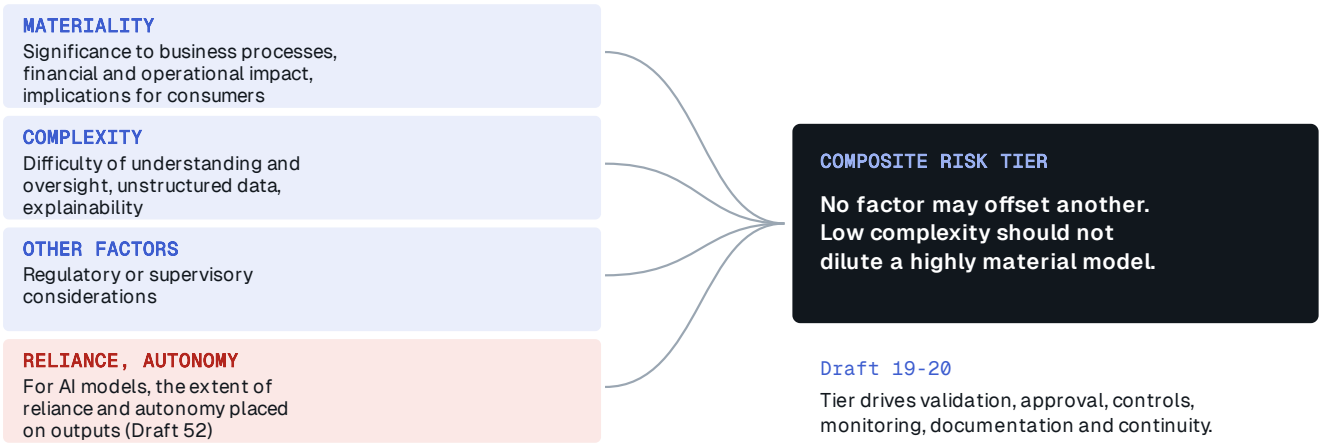
Grievances. FREE-AI asks for the volume and nature of grievances on each AI system, how they were resolved, and whether the system changed as a result (FREE-AI 4.4.68).

AI agents. Under the draft's definition, an agent that drives decisions is a model in the inventory, with its upstream and downstream dependencies recorded (Draft 7(3), 22).

RISK TIERING

Composite risk tiering for every model, without offsetting.

The draft proposes a risk-based tiering of all models in the inventory, reviewed at least annually. For AI, reliance and autonomy are added to the factors (Draft 17-20, 52). FREE-AI asks for each AI system to be categorised high, medium or low under board-approved policy.



Factors from RBI draft Guidance on Model Risk Management, 24 Jun 2026, paras 18-20 and 52. Draft, not final. As of 20 Oct 2026.

What the tier determines

- How deep, how often and how a model is validated
- Who approves, with high or equivalent tier models approved by the RMCB
- Controls, monitoring scope and reporting
- Inventory detail, documentation and continuity planning

Draft 18

Third-party models

Whatever their tier, third-party models would get independent validation by the RE and enhanced RMCB oversight (Draft 46). A low tier does not lighten that.

Where risk exceeds appetite

Enhanced controls, restrictions on use, remediation or decommissioning, with a report to the RMCB (Draft 14). For AI where risk cannot be managed, deploy only where it can (Draft 50).

PROPORTIONALITY

Testing, oversight and approval scale with risk.

Both texts scale effort with risk. FREE-AI does it with concrete frequencies for red teaming and audit. The draft does it through the model tier, which sets approval and validation intensity.

	LOW RISK	MEDIUM RISK	HIGH RISK
○ Red teaming · Rec 20	At least before deployment	At least half-yearly, plus triggers	At least half-yearly, plus triggers
○ Human oversight · Rec 16	Autonomy may be considered for routine tasks	Human oversight critical	Human oversight critical
○ AI audit · Rec 24	Internal audit, minimal	Internal audit, proportionate	Detailed, plus independent third-party audit
⬚ Approval · Draft 18	Delegated approval possible	Delegated approval possible	RMCB approves deployment
⬚ Explainability · Draft 54(1)	Thresholds set for all AI	Thresholds set for all AI	Higher thresholds, or compensating controls

Rows marked with a ring are FREE-AI recommendations (paras 4.4.49, 4.4.60, 4.4.74). Rows with a dashed square are draft proposals (Draft 18, 54(1)). Low, medium and high follow FREE-AI's labels. The draft speaks of high or equivalent tiers. As of 20 Oct 2026.

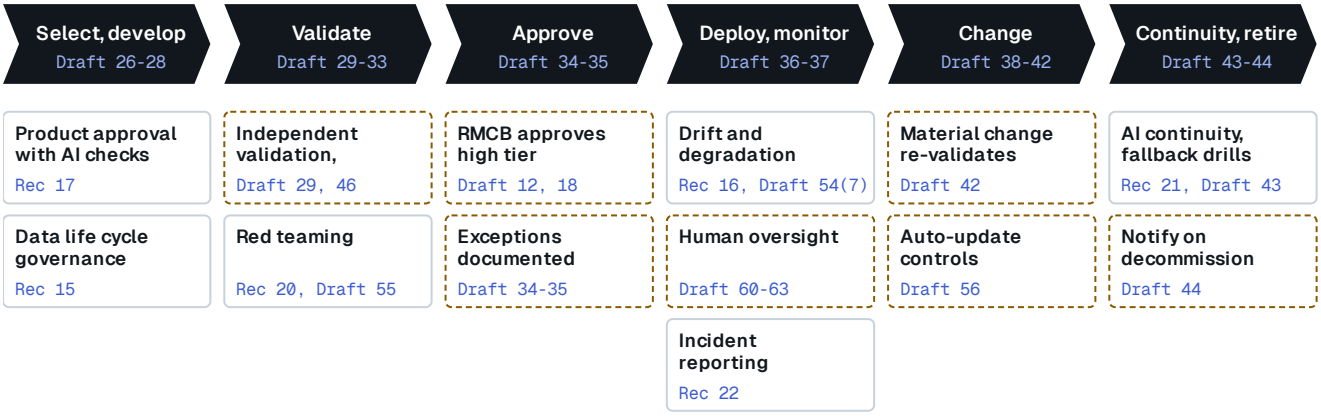
Limited scope for lighter treatment. FREE-AI allows that autonomous AI may be considered for routine or low-risk tasks, and lighter audits for low-risk uses such as document summarisation (4.4.49, 4.4.74). Neither text offers a general exemption like a basic-policies route.

Scale of the entity. The draft is meant to be applied in a manner commensurate with the RE's nature, scale and complexity, and the materiality and risk of its models (Draft 3).

LIFE CYCLE

The draft defines the stages. FREE-AI informs the controls at each.

Chapter IV of the draft sets out the model life cycle from selection to decommissioning. FREE-AI's governance and protection recommendations fall naturally onto it.



Stages from RBI draft Guidance on Model Risk Management, 24 Jun 2026, Chapter IV. Placement of FREE-AI recommendations is Winsen Labs' reading. Dashed amber boxes are draft proposals, solid boxes are recommendations. As of 20 Oct 2026.

Validation

Before and after deployment, after modification, on triggers and periodically. Reports to the RMCB or delegate within three months of completion (Draft 30, 33).

Change management

Every change preceded by a documented impact assessment, with a log of changes, versions and approvals. Crossing the material-change threshold restarts validation and approval (Draft 40-42).

Silent degradation

FREE-AI warns that AI can keep running while outputs degrade. It gives periodic human checks on a sample of decisions, such as 1%, as an example (4.4.62).

WHERE THE WEIGHT SITS

Expectations concentrate at validation and monitoring, and third-party models span every stage.

Read across both texts, the heaviest expectations fall at validation and in operation. Third-party models are the only theme strong at nearly every stage, because the draft applies the whole framework to them and more.



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 the paragraphs cited at right, across FREE-AI (recommendations) and the draft (proposals). Not an RBI classification. As of 20 Oct 2026.

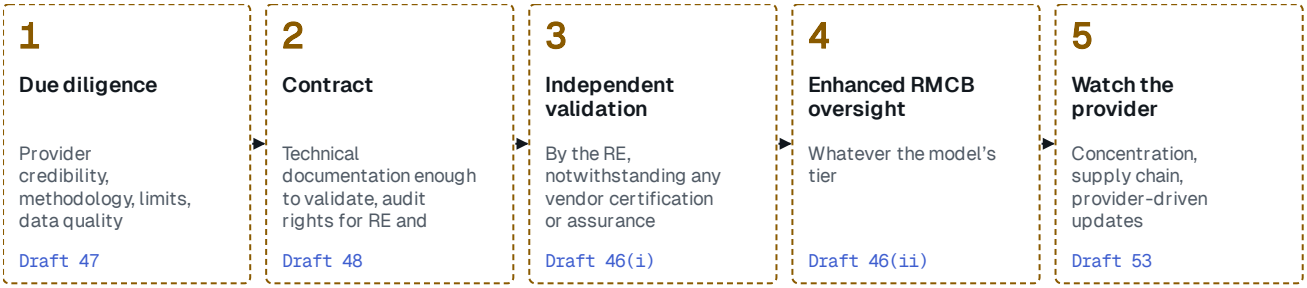
Paragraph counts. The draft has 64 paragraphs. Chapter IV on the life cycle has 19, and Chapter V on third-party and AI models has 19. Together they are 38 of 64, close to six in ten.

Starting point. Inventory and tiering are prerequisites for every other row. The draft says no model should be used unless it is in the inventory (Draft 21).

THIRD-PARTY MODELS

Vendor assurance does not replace the entity's own validation.

An RE that acquires, uses or relies on a third-party model is accountable for its outcomes (Draft 45). The draft proposes that every part of the framework applies to vendor models, plus independent validation by the RE and extra RMCB oversight (Draft 46).



ALREADY BINDING · MANAGING RISKS IN OUTSOURCING DIRECTIONS, 2025

For material outsourcing of IT services, the board stays accountable, arrangements are risk-assessed and approved, and vendors are monitored. Where an AI service is such an outsourcing, these duties apply today.

Controls from RBI draft Guidance on Model Risk Management, 24 Jun 2026, paras 45-48, 51, 53. Outsourcing summary from RE-wise Managing Risks in Outsourcing Directions of 28 Nov 2025. As of 20 Oct 2026.

Limited vendor disclosure

Where a provider does not disclose enough about an AI model, identify the resulting risks and mitigate them, for example by limiting use (Draft 51).

Provider-driven updates

The draft names changes in model behaviour or capability from provider-driven updates as a risk to consider for material third-party AI (Draft 53).

AI in the supply chain

FREE-AI Annexure IV suggests outsourcing rules require disclosure of AI use by vendors and their subcontractors. This is a suggestion, not adopted text.

AI-SPECIFIC CONTROLS AND HUMAN OVERSIGHT

Explainability, stability and the ability to intervene.

Section B of Chapter V of the draft adds controls for AI and ML models on top of the general framework. FREE-AI makes a parallel point for autonomous AI, which it recommends only after clear safeguards, with humans defining what AI may do alone (4.4.49).

Explainability thresholds Draft 54(1)
Set for all AI, higher for material decisions. Where full explainability is not achievable, compensating controls.

Hallucination boundaries Draft 54(2)
System-level controls for generative models that drive customer interaction or decisions.

Bias and fairness Draft 54(3)
Identify discriminatory outputs, assess fairness, recalibrate or redesign.

Generalisation Draft 54(4)-(5)
No overfitting, out-of-sample tests, no reliance on spurious correlations.

Output variability Draft 54(6)
Similar inputs should not give unexplained different outputs. Manage with confidence scores.

Data risks and drift Draft 54(7)
Quality, representativeness, IP, and ongoing data and concept drift.

Red teaming Draft 55
Structured challenge, especially for customer-facing or generative models.

Automatic updates Draft 56
Defined scope, strict justification, stronger data checks and monitoring.

Customer-facing safeguards Draft 59
Prompt-injection controls, AI disclosure, and a switch to a human on request.

Controls from RBI draft Guidance on Model Risk Management, 24 Jun 2026, paras 54-59. Draft, not final. As of 20 Oct 2026.

Human oversight · Draft 60-63

Human-in-command
In the loop, on the loop, or other mechanisms
Draft 60(i)

Override and kill switch
Suspend or deactivate the model
Draft 60(ii)

Periodic human review
Of outputs and model-driven decisions
Draft 60(iii)

Review the oversight
Decisions, overrides, incidents, near misses
Draft 63

Oversight should account for automation bias, over-reliance and decision fatigue, and the people doing it should be able to challenge, override or escalate (Draft 61-62).

> **Agents in FREE-AI.** The committee recommends that REs define clearly which tasks AI may perform autonomously and when human oversight is required, and remain liable for the actions and outcomes of autonomous AI they deploy (4.4.49, Rec 16).

PROTECTION AND ASSURANCE

Consumer protection, resilience, audit and disclosure.

Most of FREE-AI's protection and assurance pillars are addressed to REs. They describe what the committee expects an RE to be able to show when something goes wrong, and what it says publicly about its AI.

Consumer protection

REC 18

A board-approved framework on transparency, fairness and recourse. Customers told when they deal with AI and able to reach a human (4.4.54).

AI continuity plans

REC 21

BCPs that cover AI-specific degradation as well as outages, with fallbacks and drills. A model should be able to declare itself unavailable (4.4.62).

Incident reporting

REC 22

Regulators to set up AI incident reporting with a tolerant, good-faith approach. Where customers are harmed, the RE compensates (4.4.64).

AI audit

REC 24

Risk-based audit of data, model and outputs. Third-party audits for high-risk or complex use. Framework reviewed at least every two years.

Disclosures

REC 25

AI-related disclosures in annual reports and on websites. Regulators to specify a disclosure framework.

Cybersecurity

REC 19

Identify AI security risks and strengthen defences. AI systems able to be terminated instantly where significant harm is at risk (4.4.58).

Summaries of RBI FREE-AI Committee Report, 13 Aug 2025, Recs 18, 19, 21, 22, 24, 25 and the paragraphs cited. All are recommendations, tagged medium term except Rec 25, short term. As of 20 Oct 2026.

Consistent with the draft. It proposes that no model harming consumers be used, that grievance redress cover consumer-facing models, and that customers can switch to a human on request (Draft 25, 59(iii)). DRAFT

Binding data protection duties. From 13 May 2027 the Rules require reasonable security safeguards and intimation of personal data breaches. AI processing personal data is inside that. BINDING

WHAT A SUPERVISOR WILL LOOK FOR

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

FREE-AI recommends that supervisors adopt standardised AI checklists and model risk templates (4.4.75). None has yet been published. The questions below follow from the texts, with the evidence that would address each.

GOVERNANCE

Show me the board-approved AI policy and model risk appetite.
EVIDENCE Policy, board minutes, appetite statement (Rec 14, Draft 11)

What does the RMCB see about AI and third-party models, and how often?
EVIDENCE RMCB papers, tiering reports (Draft 12)

INVENTORY AND TIERING

Is every AI system and model in the inventory, including vendor and retired ones?
EVIDENCE Inventory extract with retired models (Rec 23, Draft 21-23)

When was the inventory last updated, and what changed?
EVIDENCE Half-yearly update log (Rec 23)

How were reliance and autonomy reflected in this model's tier?
EVIDENCE Tiering methodology and rationale (Draft 19-20, 52)

VALIDATION AND CHANGE

Who validated this vendor model, independently of the vendor?
EVIDENCE Validation report, independence (Draft 29, 46)

What counts as a material change, and show me the last one.
EVIDENCE Threshold, change log, re-validation (Draft 40-42)

OPERATION

How do you detect drift or silent degradation?
EVIDENCE Monitoring metrics, drift alerts, sampled human checks (Rec 16, 21)

Show me an override and the review of overrides.
EVIDENCE Oversight logs and periodic review (Draft 60, 63)

How would you stop this model today?
EVIDENCE Kill-switch procedure and test record (Draft 60(ii))

VENDORS AND CUSTOMERS

What can you see of this vendor model, and what audit rights do you hold?
EVIDENCE Contract clauses, documentation (Draft 48)

How would you learn of a provider-driven update?
EVIDENCE Change notices, behaviour monitoring (Draft 53)

Do customers know they are dealing with AI, and can they reach a human?
EVIDENCE Disclosures, channel design, grievances (Rec 18, Draft 59)

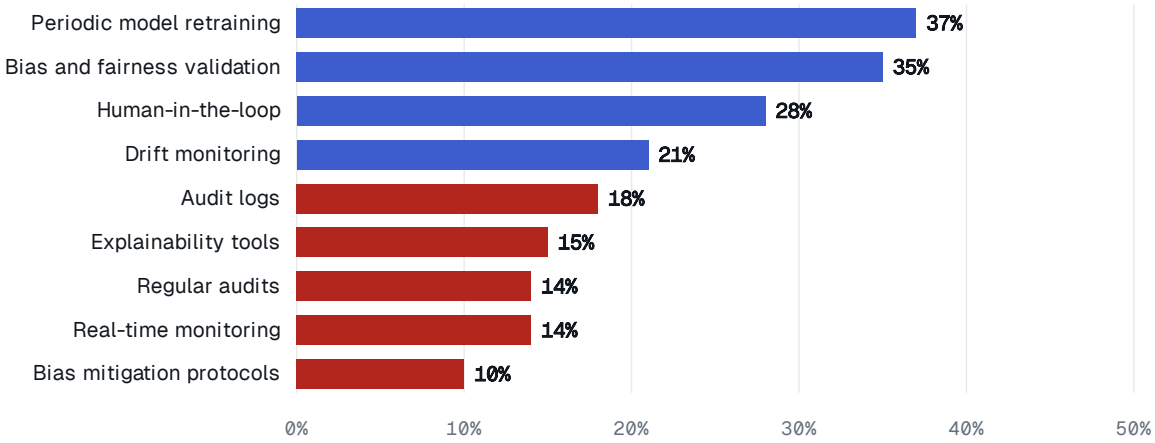
Which outsourced IT arrangements involve AI?
EVIDENCE Outsourcing register (Outsourcing Directions)

Questions drafted by Winsen Labs from FREE-AI, the draft guidance and the Outsourcing Directions. Not an RBI checklist. As of 20 Oct 2026.

WHERE INSTITUTIONS OFTEN FALL SHORT

Few entities using AI reported controls that evidence its behaviour in production.

In its 2025 survey, the RBI found that 127 of 612 supervised entities, or 20.8%, were using or developing AI. Among them, practices that evidence the behaviour of AI in operation were the least common.



Reported by the 127 entities using or developing AI, RBI Department of Supervision survey of Feb-May 2025, as published in RBI FREE-AI Committee Report, 13 Aug 2025, paras 3.3.2 and 3.3.13. Red marks practices below 20%. Sector-level figures, not about any institution. Survey date predates this guide by over a year.

1 in 3

BOARD-LEVEL AI FRAMEWORK

Only about a third of respondents, mainly large banks, reported some board-level AI oversight framework (3.3.11).

1 in 4

FORMAL AI INCIDENT PROCESS

About a quarter had formal processes for AI incidents or failures (3.3.11).

Few

DATA LINEAGE

Most had not put in place data lineage and traceability, and handled the AI data life cycle in a fragmented way (3.3.12).

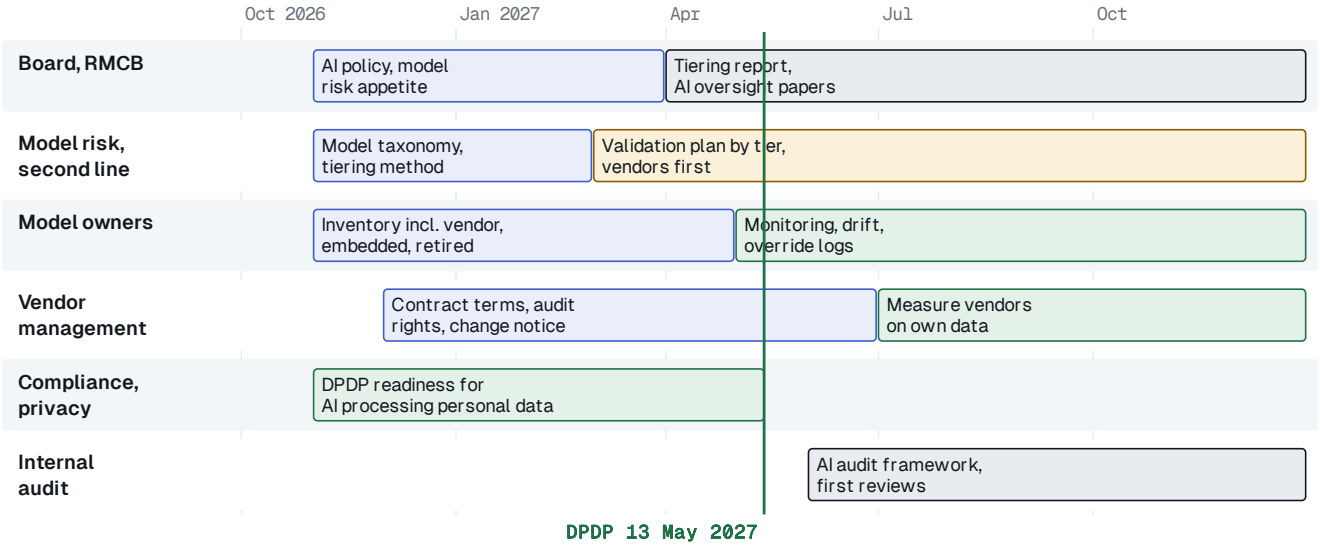
Typical implications. Inventories often stop at declared, in-house models. Vendor and embedded AI, decommissioned models and agents are often missing. Monitoring often ends at retraining, without drift checks or logs a reviewer can reopen.

Interpreting the data. These are self-reported sector figures from early 2025. Practice has likely moved since. Use them to ask where your own institution stands, not as a benchmark you can claim against.

HOW TO GET READY

Preparing against the draft ahead of the final text.

The draft is consistent with established model risk practice. Because inventory, tiering and vendor validation take several months, work begun against the draft reduces the risk of a compressed timeline once the final text is issued.



Illustrative sequencing by Winsen Labs. Only the DPDP date is fixed in law. The final model risk guidance may set its own dates. As of 20 Oct 2026.

Seven steps

1. **Establish the status of each text** pp. 6
   Treat Directions and DPDP as binding, the draft as the likely standard, and FREE-AI as direction of travel.
2. **Complete the inventory** Draft 21-23
   All models, AI or not, vendor and embedded included, retired ones kept.
3. **Tier with reliance and autonomy** Draft 19-20, 52
   Add reliance and autonomy for AI, and allow no offsetting.
4. **Validate third-party models independently** Draft 46
   Plan independent validation of third-party models first.
5. **Evidence human oversight** Draft 60-63
   Override, kill-switch and review records for AI decisions.
6. **Monitor in production** Draft 54, Rec 21
   Drift, variability and sampled human checks.
7. **Prepare for disclosure** Rec 22, 25
   Annual report and website disclosures, incident trail.

The first ninety days

- Board note on FREE-AI, the draft and binding rules, with owners
- Model taxonomy that includes rule engines and spreadsheets
- Inventory sweep including vendor, embedded and retired models
- Tiering method piloted on ten models, AI and non-AI
- Vendor letter on AI use, updates and audit rights
- List of AI touching personal data, mapped to DPDP duties

Track the RBI site for the final guidance and any AI-specific Directions. When they land, re-check this plan against the dates.

SELF-ASSESSMENT

Twenty-four checks against the texts.

Mark only what can be evidenced today. Each item cites its source, so that an unmarked item can be read as a gap against law, a draft, or a recommendation.

GOVERNANCE, INVENTORY, TIERING

☐	Board-approved AI policy covering Rec 14 areas	Rec 14
☐	Board-approved model risk framework for all models	Draft 9-10
☐	Model risk appetite approved by the board	Draft 11
☐	RMCB reviews tiering at least annually	Draft 12
☐	Three lines of defence defined for models	Draft 15
☐	Inventory of all models, AI and non-AI	Draft 21
☐	Vendor, embedded and retired models included	Draft 21, 23
☐	Inventory updated at least half-yearly	Rec 23
☐	Grievances recorded against AI systems	Rec 23
☐	Tiering includes reliance and autonomy	Draft 52
☐	No factor offsets another in tiering	Draft 20
☐	AI in product approval	Rec 17

LIFE CYCLE, OPERATION, CUSTOMERS

☐	Independent validation, vendor models included	Draft 29, 46
☐	Contracts give documentation and audit rights	Draft 48
☐	Provider-driven updates are detected	Draft 53
☐	Explainability thresholds set	Draft 54(1)
☐	Drift and output variability monitored	Draft 54(6)-(7)
☐	Red teaming in proportion to risk	Rec 20, Draft 55
☐	Override and kill switch, tested	Draft 60
☐	Oversight decisions reviewed	Draft 63
☐	AI-specific continuity plans and drills	Rec 21
☐	AI incidents recorded and escalated	Rec 22
☐	Customers told of AI, can reach a human	Rec 18, Draft 59
☐	AI processing personal data mapped to DPDP	DPDP

Checks drafted by Winsen Labs. Rec items are recommendations, Draft items are proposals, DPDP is binding. As of 20 Oct 2026.

GLOSSARY

RE
Regulated entity of the RBI, such as a bank, NBFC or AIFI.

FREE-AI
Framework for Responsible and Ethical Enablement of AI, a committee report.

MRMF
Model Risk Management Framework, board-approved under the draft.

RMCB
Risk Management Committee of the Board.

Model
Any system turning inputs into outputs used for decisions, AI or not (Draft 7(3)).

Tier
A model's composite risk rating, driving controls (Draft 17-20).

Kill switch
A mechanism to suspend or deactivate a model (Draft 60(ii)).

Red teaming
Adversarial testing to reveal hidden weaknesses (FREE-AI 4.4.59).

SUMMARY REFERENCE

The RBI's approach to AI in summary.

FREE-AI · RECOMMENDATIONS | MRM GUIDANCE · DRAFT | OUTSOURCING, DPDP · BINDING

STATUS

No binding AI-specific direction. FREE-AI, 13 Aug 2025, recommends. Draft MRM Guidance, 24 Jun 2026, proposes, comments closed 24 Jul 2026, final not yet issued. Outsourcing Directions, 28 Nov 2025, and DPDP require.

DATES THAT ARE FIXED

13 May 2027

DPDP main duties, security safeguards and breach intimation.

Outsourcing Directions in force since 28 Nov 2025. MRM final, date not announced.

SCOPE OF THE DRAFT

Eleven types of RE, from commercial banks to NBFCs and credit information companies. Every model, built or bought, AI or not, a spreadsheet pricing tool included.

INVENTORY

All models, active to decommissioned, retired ones kept ten years. Owners, validators, tier, use, dependencies, findings. Nothing used unless inventoried (Draft 21-23). FREE-AI adds grievances, target groups, half-yearly updates (Rec 23).

TIERING

Materiality, complexity, other factors, plus reliance and autonomy for AI. No offsetting. High tier needs RMCB approval.

VENDORS

You are accountable. Independent validation by you, whatever the vendor certifies. Enhanced RMCB oversight at any tier. Contracts with documentation, audit rights for you and RBI, exit. Watch provider-driven updates and concentration (Draft 45-48, 53).

OVERSIGHT

Human-in-command, override and kill switch, periodic review of outputs and of the oversight itself (Draft 60-63).

THE 13 FREE-AI RECOMMENDATIONS FOR RES

Rec 10 Capacity	Rec 14 Board AI policy	Rec 15 Data life cycle	Rec 16 System governance	Rec 17 Product approval	Rec 18 Consumers	Rec 19 Cybersecurity
Rec 20 Red teaming	Rec 21 AI continuity	Rec 22 Incidents	Rec 23 Inventory	Rec 24 Audit	Rec 25 Disclosures	

FIVE LIKELY SUPERVISORY QUESTIONS

1. Show the board-approved AI policy and model risk appetite.
2. Is every model in the inventory, vendor and retired included?
3. Who validated this vendor model, independently?
4. How would you stop this model today?
5. How do you detect drift and silent degradation?

Summary of RBI FREE-AI Committee Report, 13 Aug 2025, RBI draft Guidance on Model Risk Management, 24 Jun 2026, the Outsourcing Directions and the DPDP Rules. Not legal advice. As of 20 Oct 2026.

What Winsen Rho can do for you

Effortless accountability for every AI action.

You know what the RBI 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 Rho does, in six capabilities.

Each capability produces records that can be placed before a reviewer. The references indicate where the RBI 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.

Rec 23 · Draft 21-23

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.

Draft 41, 57, 60(iii)

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.

Draft 14, 37 · Rec 25

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.

Draft 46-48, 53

Human oversight

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

Draft 60-63 · Rec 16

Drift

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

Draft 37, 54(7) · Rec 21

References are to FREE-AI (recommendations) and the draft MRM Guidance (draft). 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

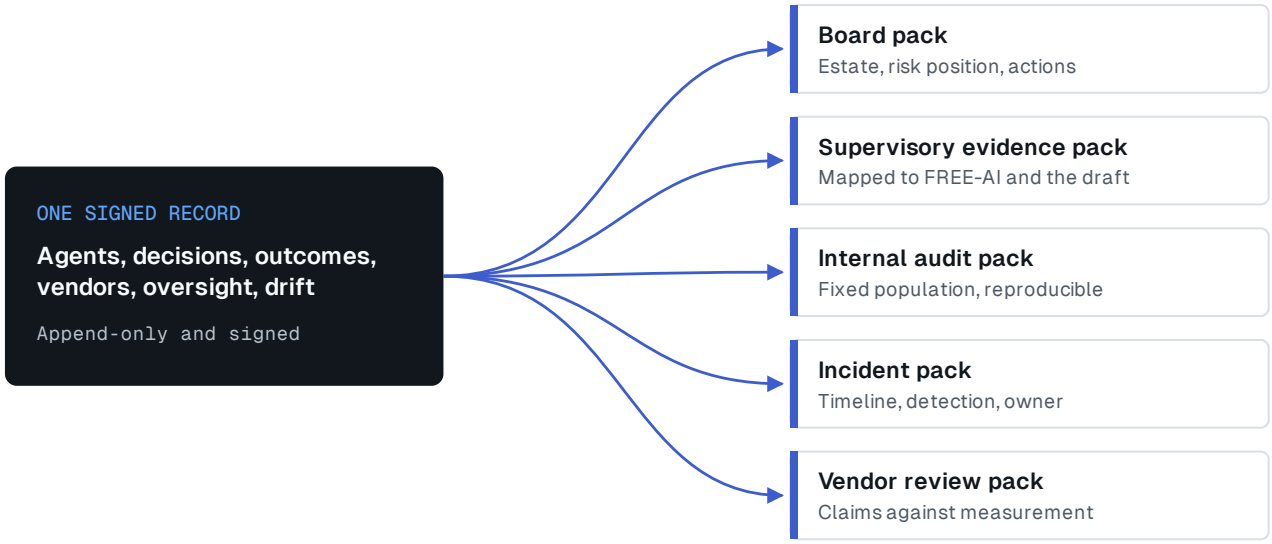
BINDING DIRECTION **RBI Managing Risks in Outsourcing Directions, 2025**	An inventory of third-party agents and the vendors behind them, each vendor scored on one measure from your own outcome data, alerts and incident packs to support your reporting clock
BINDING LAW **Digital Personal Data Protection Act 2023 and Rules 2025**	Decision records showing what personal data each agent saw and used, alerts when an agent goes beyond its permitted data or scope, packs for breach reports and audits

Evidence descriptions as published on winsenrho.com, reviewed 24 Sep 2026. Licensed and off-the-shelf software is outside the Outsourcing Directions, so not every AI vendor is covered. Reporting and notification decisions 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

The RBI's expectations, and where Rho fits.

THE EXPECTATION	The evidence Rho holds	WHAT STAYS WITH YOUR INSTITUTION
AI inventory, half-yearly, ready for supervisors Rec 23 · Draft 21-23	A living register of every agent, vendor and embedded ones included, with named owner, risk tier, life cycle state and retired agents kept	Inventory of non-agent models, taxonomy, update cadence, retention policy
Board policy, appetite and RMCB oversight Rec 14 · Draft 11-12	Board pack from the signed record, tier and approval history, incidents with an owner	The policy, the framework, appetite and every RMCB decision
Tiering with reliance and autonomy Draft 19-20, 52	Autonomy, authority and oversight mode per agent, the recorded tier and its approvals	The tiering method and the tiers
Human oversight, override, kill switch Draft 60-63 · Rec 16	Records of reviews, overrides and sampled checks, whether overrides helped, pause or restrict through your gateway with the decision logged	Oversight design, authority, kill-switch testing
Third-party models and provider updates Draft 45-48, 53	Vendor agents measured on your customers against claims and contract, unnotified model changes, concentration by vendor and model	Due diligence, contracts, and independent validation itself
Monitoring, drift and output variability Draft 37, 54(6)-(7) · Rec 16	Nightly probes, tiered thresholds, drift alerts, the same question getting different answers	Metrics, thresholds, remediation
Change records and material change Draft 41-42, 56	Model and version change record, alerts that can trigger re-validation	Change framework, materiality threshold, approvals
Incidents, continuity and audit Rec 21, 22, 24	Incident pack with time to detection, internal audit pack that reproduces exactly	Reporting decisions, BCP drills, audit plan and third-party audits
Disclosures Rec 25	Estate and outcome figures from the signed record	The disclosure itself

Rho does not cover these. They stay with you.

- Data life cycle governance (Rec 15, Draft 28)
- Cybersecurity and deployment security (Rec 19, Draft 58-59)
- Red teaming (Rec 20, Draft 55)
- Product approval (Rec 17)
- Model development and documentation (Draft 26-27, 57)
- Independent validation (Draft 29-33)
- Customer disclosures and grievance redress (Rec 18, Draft 59)
- Decommissioning (Draft 44)
- Capacity building (Rec 10)

NEXT STEP

See your own estate through the RBI's lens.

Book a 30-minute walkthrough.

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

Mapping maintained by Winsen Labs against FREE-AI (recommendations) and the draft MRM Guidance (draft), 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
RECOMMENDATIONS	**FREE-AI Committee Report, 13 August 2025** https://rbidocs.rbi.org.in/rdocs/PublicationReport/Pdfs/FREEAIR130820250A24FF2D4578453F824C72ED9F5D5851.PDF	Read in full Read 20 Oct 2026
DRAFT	**Draft Guidance on Regulatory Principles for Model Risk Management, 2026** https://www.rbi.org.in/Scripts/bs_viewcontent.aspx?Id=5089	Read in full Read 20 Oct 2026
DRAFT	**RBI press release on the draft, 24 June 2026** https://www.rbi.org.in/Scripts/BS_PressReleaseDisplay.aspx?prid=63006	Comments to 24 Jul 2026. Also dates the Aug 2024 credit draft Read 20 Oct 2026
DRAFT	**RBI Draft Notifications and Guidelines listing** https://www.rbi.org.in/Scripts/DraftNotificationsGuildelines.aspx	Data governance draft, 15 Jul 2026 Read 20 Oct 2026
BINDING	**RBI notifications and press releases, checked for a final MRM text** https://www.rbi.org.in/Scripts/NotificationUser.aspx	None found to 19 Oct 2026 Read 20 Oct 2026
BINDING	**RE-wise Managing Risks in Outsourcing Directions, 2025, 28 November 2025** https://www.rbi.org.in/Scripts/NotificationUser.aspx	Re-check RE-specific text Read 20 Oct 2026
BINDING	**Digital Personal Data Protection Rules, 2025, G.S.R. 846(E), 13 November 2025** https://www.meity.gov.in/static/uploads/2025/11/53450e6e5dc0bfa85ebd78686cadad39.pdf	Commencement dates 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. The draft guidance may change before it is final, and FREE-AI is a committee report, not a direction. Check the source before relying on any point. Supervisor questions, readiness plans, heat-grid weights and placements are Winsen Labs' judgement and are labelled as such. Survey figures are the RBI's, as published. 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