SINGAPORE REINSURERS' ASSOCIATION

SRA Reinsurers' AI Handbook

A reinsurer-specific companion to the MindForge AI Risk Management Toolkit

Practical guidance for assessing AI-related risks that may affect reinsurance relationships and decisions



SRA
SINGAPORE REINSURERS' ASSOCIATION

Updated October 2026

About this Handbook

WHY THIS MATTERS

Existing AI governance frameworks such as those issued by the Monetary Authority of Singapore (MAS) and the MindForge AI Risk Management Toolkit provide guidance on the governance, development and deployment of AI. This Handbook complements those frameworks and addresses a different question: **how can reinsurers assess AI-related risks arising through the business they reinsure and how might those considerations inform reinsurance decisions?**

It focuses on AI used by cedants and other parties where that use may affect risk selection, pricing, claims outcomes or portfolio performance. It complements, rather than replaces, existing AI governance frameworks and provides a reinsurance-specific perspective on considerations that may be relevant to reinsurers.

Who this Handbook is for

This Handbook is intended as a reference for reinsurance practitioners involved in functions such as underwriting, claims, risk management, legal, compliance, technology and AI governance. It may also serve as a reference for senior management overseeing AI-related risks within the reinsurance business.

What this Handbook covers

This Handbook sets out illustrative considerations for assessing AI-related risks in the reinsurance context and considering how those risks may influence reinsurance decisions. It covers key AI risk areas relevant to reinsurers, a structured lifecycle assessment and organisational enablers that support consistent and proportionate AI risk assessment. It is not intended to prescribe a particular approach, or to validate or provide assurance over a cedant's AI governance framework or its compliance with regulatory requirements or industry standards. Information on a cedant's AI governance may be considered, where relevant to understanding how AI could affect the reinsured risk.

How to use this Handbook

This Handbook may be read alongside existing AI governance frameworks. The relevance of individual considerations will depend on the circumstances, including the relevance of the AI use to the reinsured risk and the information reasonably available to the reinsurer. This Handbook will be reviewed once the final MAS Guidelines are issued and updated as appropriate.

The Handbook is structured as follows:

- **Section 1** frames the reinsurer's position and the principle of applying existing frameworks.
- **Section 2** sets out the key AI risk areas; data, model and decision reliability; third-party AI dependencies; underwriting, claims and concentration.
- **Section 3** illustrates a five-stage lifecycle for considering relevant AI risks within existing reinsurance assessments and decision-making processes.
- **Section 4** highlights enablers; skills, responsible AI use and assessment tools, market intelligence and portfolio monitoring.
- **Section 5** lists references.

This Handbook is published by the Singapore Reinsurers' Association (SRA) for informational purposes only. It provides non-binding, principles-based guidance on assessing AI-related risks in the reinsurance context and does not constitute legal, regulatory, technical or professional advice. Users remain responsible for complying with applicable laws, regulations, supervisory expectations and internal policies.

The guidance should be applied having regard to the nature and materiality of the relevant risks and should be read alongside applicable AI governance and risk management frameworks. No representation or warranty, express or implied, is made as to the accuracy, completeness, currency or suitability of the information contained in this Handbook, and neither the SRA nor any contributor accepts liability for any loss directly or indirectly arising from reliance on this Handbook.

This Handbook may be updated from time to time or withdrawn without further notice.

Table of Contents

About this Handbook 1

Who this Handbook is for 1

What this Handbook covers 1

How to use this Handbook 1

1 Context & principles 3

1.1 Reinsurers' role in the AI value chain 3

1.2 Applying existing AI governance frameworks in reinsurance 3

1.3 AI risks within existing risk management frameworks 4

2 Key reinsurer AI risk areas 4

2.1 Data, model and decision reliability 4

2.2 Third-party AI dependencies 6

2.3 Underwriting AI risks 7

2.4 Claims AI risks 8

2.5 Concentration risks 9

3 Reinsurer assessment lifecycle 11

3.1 Identify 12

3.2 Assess 12

3.3 Decide 13

3.4 Monitor 14

3.5 Respond 14

4 Enablers 14

4.1 Responsible AI use, skills and assessment tools 14

4.2 Portfolio monitoring and market intelligence 15

5 References 16

1 Context & Principles

1.1 Reinsurers' role in the AI value chain

WHY THIS MATTERS

Reinsurers may develop and use AI within their own organisations, for which existing AI governance frameworks such as the MindForge Toolkit remain directly relevant. They also face a distinct challenge when AI used by cedants or other participants in the insurance value chain affects the risks transferred under reinsurance arrangements.

The reinsurer's role is not to govern or validate the cedant's AI, but to consider, based on information reasonably available, whether and how relevant AI use may affect risk selection, pricing, claims outcomes or the performance of the reinsured portfolio, and how this may inform reinsurance decisions.

KEY CONSIDERATIONS

- Distinguish between AI used within the reinsurer's own organisation and AI that affects risks assumed through reinsurance.
- Focus the assessment on whether and how AI may change the nature, selection, management or emergence of the underlying reinsured risk, rather than assessing AI governance for its own sake.
- Translate relevant findings into underwriting, pricing, treaty, claims and portfolio decisions.
- Apply the assessment proportionately, considering the relevance of the AI use to the reinsured risk and the information reasonably available to the reinsurer.

1.2 Applying existing AI governance frameworks in reinsurance

WHY THIS MATTERS

Existing AI governance frameworks provide the common foundation for responsible AI governance across financial institutions. This Handbook does not introduce another governance framework. Existing principles and concepts may instead provide shared vocabulary and useful reference points for understanding AI-related risks relevant to reinsurance. Information arising from a cedant's existing governance processes may also inform the reinsurer's assessment, without the reinsurer validating or providing assurance over the cedant's AI governance framework.

KEY CONSIDERATIONS

- Use existing AI governance frameworks as a reference point and shared language for discussions with cedants, rather than as a compliance checklist for cedants.
- Focus any information requests on matters relevant to understanding the reinsured risk and the reinsurance decision, taking into account what information is reasonably available.
- Apply the depth of assessment proportionately to the relevance and potential significance of the AI use to the reinsured risk and reinsurance decision.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- What information is available to understand how the relevant AI use may affect the reinsured risk?
- What existing governance, control or assurance information may help inform that assessment?
- Are there material information gaps or uncertainties that should be reflected in the reinsurance assessment or decision?
- Is any additional information proportionate and necessary for the reinsurance decision?

1.3 AI risks within existing risk management frameworks

WHY THIS MATTERS

AI is not necessarily a separate category of risk. Rather, it may influence existing risk drivers such as data quality, model risk, underwriting selection, claims handling, operational resilience and third-party dependency. Relevant AI-risk considerations may therefore be incorporated into existing reinsurance underwriting, claims, portfolio management and risk management processes, rather than requiring a standalone AI assessment.

KEY CONSIDERATIONS

- Consider which existing risk drivers may be affected by the relevant AI use.
- Consider the potential impact on exposure, experience, claims and portfolio performance.
- Consider whether existing reinsurance risk-management levers could address findings or uncertainties that warrant further action.
- Consider whether identified AI-related risks or uncertainties warrant further assessment, monitoring or escalation.
- Consider potential accumulation across cedants, vendors, models and data sources.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- How could the relevant AI use affect the risk being reinsured?
- Have those implications been reflected in the relevant underwriting, pricing, claims or portfolio assessment?
- Do any information gaps or uncertainties require further consideration, monitoring or action?
- Could the AI use contribute to concentration or accumulation risk across the portfolio?

2 Key reinsurer AI risk areas

This section identifies five areas where AI use by cedants or other parties may affect the risks assumed by reinsurers. Each subsection considers why the risk matters, key considerations and questions to inform the assessment, and potential reinsurer actions. The relevance of these considerations will depend on the circumstances, including significance to the reinsured risk and the information reasonably available to the reinsurer. They are not intended to imply validation of the cedant's AI governance or controls.

2.1 Data, model and decision reliability

Scope: Data, models and AI-supported decisions used by cedants where these may affect the business being reinsured.

WHY THIS MATTERS

Where AI influences a cedant's risk selection, pricing or claims decisions in a way that may be relevant to the business being reinsured, weaknesses in the underlying data, model or outputs may affect the quality and performance of that business. The reinsurer should therefore consider limitations and uncertainties that may affect the risk being assumed.

KEY CONSIDERATIONS

- **Data quality and completeness:** Consider whether cedant data is accurate, complete and consistently structured, as gaps or inconsistencies can materially distort outputs from cedants' AI systems.
- **Privacy, confidentiality and permitted use:** Consider whether restrictions on collection, processing, retention or sharing of information create customer, regulatory or operational consequences relevant to the business being reinsured.
- **Performance, representativeness and sampling issues:** Assess whether cedants' AI-supported outcomes reflect underlying portfolios and whether they correspond to the intended underwriting or claims outcomes. Seek evidence of AI model performance and how this aligns with existing practices, while identifying any limitations, uncertainties, biases, variation in performance, and inconsistency over time.
- **Decision and model appropriateness:** Examine whether the cedant's approach to AI implementation and level of automation are appropriate for the task, considering the need for predictability, consistency and human judgement. AI involvement in key or sensitive decision-making warrants further scrutiny, including the use of decision thresholds, guardrails and human intervention.

PRACTICAL EXAMPLES

- **Ongoing calibration:** If a cedant's AI-supported claims-cost estimates deteriorate significantly as claims inflation changes, the cedant should establish whether inflation alone explains it, or whether model drift, changes in input data, or design limitations also contribute. The key issue is whether any residual uncertainty materially affects the risk being ceded.
- **Maintaining reliable information as AI use increases**: As cedants increasingly use AI to process data and support underwriting, pricing and claims decisions, reinsurers may receive information that has already been transformed by upstream AI supported processes. This heightens the importance of consistent, traceable and well structured data to understand the business ultimately ceded.

  The ACORD/SCOR case study (August 2025) illustrates one practical approach: applying common data standards (e.g. ACORD's Reinsurance & Large Commercial Data Standards) together with AI based extraction tools to submissions, bordereaux and claims advice significantly improved data quality while reducing manual effort and payment turnaround. Reliable, consistently structured information also helps identify changes in the underlying business and whether AI influenced outcomes at the cedant are affecting the risk transferred.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- Is the underlying data sufficiently complete, reliable and representative to support the relevant AI use?
- How does the cedant assess the accuracy, robustness and ongoing performance of AI-supported outputs?
- How are AI-supported decisions reviewed, challenged or overridden where necessary?
- Are there any known limitations or uncertainties that should influence the reinsurance decision?
- Are there any privacy, legal, confidentiality or contractual considerations affecting the information available for the reinsurance assessment?
- Are the AI approach and level of automation appropriate to the nature and significance of the decision it supports?

POTENTIAL REINSURER ACTIONS

- Seek targeted information on relevant data, model or output limitations where existing information is insufficient.
- Compare AI-influenced outcomes with relevant experience or benchmarks where appropriate.
- Reflect significant uncertainty through existing underwriting, pricing, terms, monitoring or engagement processes.

2.2 Third-party AI dependencies

Scope: Third-party AI relied upon by cedants or other parties where those dependencies may affect the business being reinsured. Common dependencies identified at individual cedant/treaty level may also inform the portfolio-level concentration assessment in Section 2.5. This section focuses primarily on the indirect route. The relevance of third-party AI dependencies will vary depending on the nature of the business and reinsurance arrangement.

WHY THIS MATTERS

Reinsurers may be exposed to third-party AI risk through two main routes:

- **Direct** (vendor tools used by the reinsurer): Most reinsurers buy an existing AI solution rather than build one, so their AI risk is largely a vendor-governance question. These risks are generally addressed through the reinsurer's existing AI governance and third-party risk management frameworks and are not the primary focus of this section.
- **Indirect** (AI models deployed by cedant): Reinsurers may assume risks from cedants who rely on third-party AI providers for underwriting, pricing, claims and other core insurance processes. While responsibility for managing those providers rests with the cedant, failures, disruptions or changes in those dependencies may still affect the underlying exposure, claims outcome or performance of the reinsured portfolio. The reinsurer's focus is therefore on understanding relevant third-party AI dependencies and how they may affect the risk being assumed. Where the same provider is relied upon across multiple cedants, the dependency may also contribute to correlated or concentration risk (see 2.5).

KEY CONSIDERATIONS

- Identify relevant third-party AI dependencies that support the cedant's underwriting, pricing or claims processes where these may affect the reinsured business. Consider what information is available on how those dependencies are managed and monitored.
- Consider whether third-party AI dependencies could introduce correlated risks across multiple cedants or portfolios.
- Consider whether any material limitations or uncertainties relating to a third-party AI dependency could affect the reinsurance assessment or decision.
- Consider whether contractual arrangements provide sufficient transparency, auditability and notification of material changes or incidents and are sufficiently tailored for AI.
- Monitor whether reliance on common AI providers may contribute to emerging concentration or operational resilience risks to the reinsurer.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- Which third-party AI providers support AI use relevant to the reinsured business?
- What information is available on how the cedant manages and monitors relevant third-party AI dependencies?
- What visibility is available over material model changes, incidents or service disruptions?
- Could reliance on a common AI provider create correlated exposure across multiple cedants or portfolios?
- Where relevant, what information or notification arrangements are available for material AI-related changes or incidents?

POTENTIAL REINSURER ACTIONS

- Seek targeted information on relevant third-party AI dependencies where needed to understand the reinsured risk.
- Consider appropriate information or notification arrangements for significant changes or incidents, where warranted.
- Reflect significant dependency-related uncertainty through existing underwriting, terms, monitoring or engagement processes.
- Consider common dependencies across cedants as part of the assessment in Section 2.5.

2.3 Underwriting AI risks

Scope: AI used by cedants in underwriting or pricing where it may affect the business being reinsured.

WHY THIS MATTERS

Cedants may use AI to support risk selection, underwriting and pricing. Where this influences the business being reinsured, it may affect the nature, quality or composition of the ceded portfolio and ultimately its performance. The reinsurer should therefore consider how relevant AI use may affect the risk being assumed and whether this should influence the reinsurance decision.

2.3.1 Decision ownership and human accountability

- **Decision ownership:** Consider whether an appropriate business owner or subject-matter expert retains accountability for AI-influenced underwriting decisions, including appropriate authority to review, challenge or override AI-supported outputs where necessary.
- **Capability:** Where relevant to the reinsured risk, consider whether those relying on AI-supported underwriting outputs understand their limitations, the extent to which they can reasonably rely on such outputs and exercise appropriate judgement.
- **Decision evidence:** Consider whether AI-influenced underwriting decisions can be sufficiently understood or appropriately evidenced by the cedant where relevant to the reinsurance assessment.

2.3.2 Pricing

- **Transparency:** Consider whether there is sufficient understanding of how AI influences pricing or risk selection where relevant to the reinsurance assessment, e.g. documented model assumptions, variables, and decision logic.
- **Scenario analysis:** Consider whether the cedant's scenario testing covers AI-specific risks such as adversarial attacks and model drift, and what these could mean for pricing.
- **Benchmarking:** Where appropriate, consider whether the cedant's AI-supported pricing is benchmarked against traditional methods and industry benchmarks to validate reasonableness. AI pricing engines can embed bias or feedback loops that concentrate risk in ways the treaty never anticipated.

2.3.3 Wording review

AI may be used to draft or review wordings. Separately, AI-related events or failures may create exposure under existing policy or treaty wordings even when AI is not expressly addressed (often called "silent AI"). Reinsurers should therefore consider whether relevant AI-related exposure is understood and appropriately reflected in the reinsurance assessment and wording, where applicable.

While AI-related events may overlap with cyber exposures in certain scenarios, silent AI should not be regarded as synonymous with silent cyber, as many AI-related liabilities can arise independently of any cyber-attack, system intrusion or information security failure.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- Where and to what extent does AI influence risk selection, underwriting or pricing on the business being reinsured?
- Who retains business accountability for relevant AI-influenced underwriting decisions?
- Are there known limitations, changes or uncertainties that could affect the characteristics or performance of the ceded business?
- Where appropriate, how does AI-influenced pricing or experience compare with relevant benchmarks or expectations?
- Could AI-related events or failures create exposure that is not clearly addressed in relevant wordings?

POTENTIAL REINSURER ACTIONS

- As part of the assessment, consider understanding underwriting authorities, override activity and model governance arrangements, including the management of significant model changes.
- Consider whether unclear accountability or explainability creates material uncertainty for the reinsured risk and whether this warrants adjustments to risk appetite, terms and conditions, or pricing.
- Benchmark AI-priced books against traditionally priced experience and challenge outliers.
- Reflect unexplained or unmitigated gaps in appetite, terms or wordings.

2.4 Claims AI risks

Scope: AI used by cedants in claims handling where it may affect claims outcomes on the business being reinsured.

WHY THIS MATTERS

Cedant claims-AI practices flow directly into ceded results. Weak validation or oversight could produce claims leakage, and shared triage and fraud tools can mis-settle claims in correlated ways across cedants.

Silent AI in claims wordings may also add exposure that was never priced. The reinsurer should therefore consider whether relevant AI use could materially affect claims outcomes on the business being reinsured and what information is available to assess that risk.

2.4.1 Claims assessment

- **Model reliability:** Consider whether known limitations, changes or systematic effects in AI-supported claims assessment could materially affect claims outcomes on the reinsured business.
- **Transparency and explainability:** Consider whether there is sufficient understanding of how AI influences relevant claims outcomes where this is necessary for the reinsurance assessment or claims review.
- **Unusual or complex claims:** Consider whether reliance on AI could produce different or less reliable outcomes for unusual, complex or changing claims patterns relevant to the reinsured business.
- **Emerging experience:** Consider whether claims experience indicates systematic changes in frequency, severity, settlement patterns or other outcomes associated with AI use.

2.4.2 Human review and oversight

- **Decision ownership:** Consider whether appropriate claims professionals retain responsibility and judgement for material AI-influenced claims decisions.
- **Human judgement:** Where relevant, consider whether claims professionals using AI-supported outputs can understand their limitations and review or challenge the outputs where necessary.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- Where and to what extent does AI influence claims triage, assessment or settlement on the business being reinsured?
- Who retains business responsibility and judgement for relevant AI-influenced claims decisions?
- Are there known limitations, changes or uncertainties that could affect claims outcomes?
- Does emerging claims experience indicate systematic changes in frequency, severity or settlement patterns associated with AI use?
- Could common AI-supported claims processes or providers contribute to correlated exposure across cedants?

POTENTIAL REINSURER ACTIONS

- Consider targeted review of a sample of large claims where AI materially influenced the outcome.
- Monitor frequency and severity drift on AI-triaged books and query systematic shifts.
- Seek targeted information on how AI materially influences claims handling and any known limitations or uncertainties relevant to the reinsured business.

2.5 Concentration risks

Scope: Portfolio-level risk arising from cedants' and vendors' shared AI infrastructure, a view the reinsurer may particularly be well-positioned to identify across its book.

WHY THIS MATTERS

By the very nature of its business, a reinsurer is exposed to multiple insurers, brokers, their underlying operations and portfolios across diverse geographies.

At the same time, the modern AI ecosystem is increasingly built on shared infrastructure. Where multiple cedants rely on common cloud infrastructure, models, data or vendors, a failure or disruption could create correlated effects across the reinsured portfolio. Reinsurers may therefore be particularly well positioned to identify and assess these common dependencies across their books.

Common use of the same provider, model or infrastructure does not, by itself, indicate concentration risk. The relevance of any shared dependency depends on whether it could give rise to correlated and material impacts across the reinsured portfolio. The same consideration informs which concentration indicators are relevant for portfolio monitoring in section 3.4.

Key risks and considerations: common AI vendors, correlated model behaviour and operational resilience. As Table 1 shows, the shared stack creates a correlation channel at each layer – a fault hits everyone above it at once.

Table 1

Shared layer	Illustrative scenario*	Potential failure mode	Correlates across
Cloud & compute	Several cedants rely on the same cloud providers (e.g. AWS, Microsoft Azure)	Outage; region failure	All cedants on that cloud
Foundation models	Multiple cedants use the same underlying model (e.g. OpenAI (GPT), Anthropic (Claude))	Defect; drift; recall	All apps on that model
Data & analytics	Several cedants rely on the same analytical models (e.g. Verisk, LexisNexis)	Data error or bias	All models fed by it
Application vendors	Multiple cedants rely on the same application (e.g. Akur8, Shift Technology)	Logic flaw; miscalibration	All carriers using it

*Provider references are illustrative of current market structure and imply no deficiency or expectation of failure in any named provider.

- **Correlated model behaviour:** A study by Gallagher Re warns that current AI model evaluation frameworks do not adequately assess concentration risk, meaning a failure of a single foundation model could affect thousands of policies Benchmark contamination may also create a false sense of model performance. Further, around 95% of Insurtech funding now flows to AI start-ups, which concentrates reliance on a small number of providers.
- **Operational resilience:** Reliance on common cloud providers, regions or other infrastructure may cause a single disruption to affect multiple cedants, creating potential accumulation risk for the reinsurer.
- **Silent AI accumulation:** Exposure neither modelled nor priced becomes "accidental capacity" for a correlated event. A few vendors (e.g. Verisk) are already issuing generative-AI exclusion endorsements.

AI-related events and incidents can also give rise to reinsurance claims, whether arising from affirmative AI coverage or through silent AI exposure embedded within existing policies.

PRACTICAL EXAMPLES

A foundation-model failure: A widely used model develops a flaw after an update. Cedant A's claims tool, B's underwriting assistant and C's pricing add-on all run on it. Within days they mishandle risk in correlated ways; weeks later, AI-liability and E&O policies start to generate notifications. A reinsurer that has mapped foundation-model exposure (see 3.1) can answer "who else is exposed?" in hours; one that has not will find the accumulation only as losses arrive.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- Which cloud providers, foundation models and application vendors sit behind the business you cede?
- What is your operational-resilience plan if a key AI provider or cloud region fails?
- Do your AI-exposed wordings address correlated AI failure – covered, excluded or sub-limited?

POTENTIAL REINSURER ACTIONS

- Map relevant common AI dependencies across cedants, where reasonably possible.
- Consider correlated AI failure scenarios as part of existing accumulation or portfolio analysis where relevant.
- Consider whether significant concentration findings should influence appetite, capacity, terms, wordings or monitoring.
- Incorporate relevant dependency information into existing underwriting or renewal processes where appropriate.
- Escalate significant AI concentration findings through existing accumulation or emerging-risk processes.

3 Reinsurer assessment lifecycle

WHY THIS MATTERS

Table 2 sets out a repeatable five-stage loop for managing AI risk under limited visibility. It is a loop, not a checklist – exposure changes between renewals, so the stages feed back into one another.

Identify and Assess (3.1–3.2) may form part of the existing submission and renewal process. Monitor (3.4) draws on information already flowing to the reinsurer, with the frequency and extent of reassessment depending on the relevance of the AI use. [JL4.1]Trigger-based reassessment may also be appropriate where significant changes or events occur.

Table 2

Stage	Core question	What you produce
3.1 Identify	Where is AI used and what exposure exists?	AI exposure map across the book
3.2 Assess	What to ask, and what evidence to seek?	Assessment tier (information confidence) per cedant/treaty
3.3 Decide	What action is appropriate?	Pricing / terms decision + rationale
3.4 Monitor	How are risks tracked over time?	Indicators + portfolio dashboard
3.5 Respond	How are incidents handled?	Escalation path + feedback to 3.1

3.1 Identify

- **Objective:** Build a defensible map of where AI sits in the ceded portfolio (akin to MindForge Toolkit's AI inventory), capturing AI presence at submission/renewal rather than inferring it, and flagging where vendors/models recur.
- **Practical questions:** Where is AI used in the business you cede? Built or bought, and from whom? Any fully automated decisions, and at what threshold?
- **Evidence to request:** A cedant AI/model inventory, named vendors and whether their framework maps to recognised AI governance frameworks (e.g. MindForge Toolkit/MAS/NIST)
- **Example:** A motor reinsurer maps its top 20 cedants and finds six share one pricing vendor – a concentration finding invisible treaty-by-treaty.

3.2 Assess

- **Objective:** Form a view, based on information reasonably available, of how the cedant's relevant AI use may affect the risk being reinsured and whether any uncertainty remains. Convert that view into an assessment tier that guides the depth of further enquiry and the choice of reinsurance levers (see 3.3).

  The assessment tiers in Table 3 below are intended to reflect the reinsurer's confidence in the information available, i.e. a decision-support tool for the reinsurer's own process, not a validation or rating of the cedant's AI governance. This mirrors the MindForge Toolkit's risk-materiality principle.

Table 3

Tier	What it means for reinsurers' assessment	Possible responses include
1: Well-evidenced	Clear understanding of how the AI use may affect the reinsured risk; no material unexplained uncertainty. Evidence may include observed performance and experience, the nature and materiality of the AI use, and relevant governance information (e.g. owner, validation, monitoring)	Standard terms; periodic review
2: Partially evidenced	Partial understanding; some material questions remain open. Available evidence is incomplete or mixed.	Engage; targeted conditions or info rights
3: Limited evidence / uncertain	No reliable view can yet be formed; material uncertainty remains (e.g. limited information after reasonable enquiry, reliance on opaque vendor models, unexplained changes in experience).	Consider whether residual uncertainty warrants further enquiry, monitoring or reflection in the reinsurance decision, depending on materiality and circumstances

Factors to inform the tier:

- **Practical questions:**
  - Who is accountable?
  - How are models validated and revalidated?
  - How do you test for bias?
  - How do you oversee AI vendors?
  - Any AI incident in 24 months?
- **Appropriateness:** Consider whether the AI approach and level of automation are appropriate to the decision being supported (see 2.1).
- **Decision ownership:** Consider whether there is clear business ownership of AI-influenced decisions; unclear ownership adds to the uncertainty the assessment reflects.

- **Request information:** AI governance policy, model inventory and validation summary, bias-testing results, vendor due-diligence records.
- **Red flags:** "We rely on the vendor" with no oversight; bias never tested.
- **Operating with incomplete information:** Assess using the information available at the time:
  - Cedants may not always be able to provide complete AI-related information; an information gap does not of itself indicate heightened underlying risk. The reinsurer should first consider whether the gap is material to understanding the reinsured risk. Where it is, reflect the resulting uncertainty proportionately in the assessment tier and, where relevant, the decision (see 3.3).
  - Absence of evidence after reasonable enquiry supports a more conservative tier rather than no tier. Where information is unavailable, reasonable proxies such as known vendor relationships, public disclosures, incident history or broker knowledge may inform the assessment.
  - Unanswered questions should be recorded so gaps are revisited at renewal. Residual uncertainty flows through to the decision (see 3.3) as pricing, terms, information rights or appetite.

3.3 Decide

- **Objective:** Connect the assessment tier to the levers available to the reinsurer, proportionate to materiality. Document the rationale supporting the reinsurance decision, including the basis on which any information gaps were treated (see 3.2).
- **Levers: Appetite ▪ Price ▪ Terms & conditions ▪ Information rights ▪ Engagement.**
  - Each lever draws on an existing reinsurance mechanism. Table 4 expands on each lever:

Table 4

Lever	Existing reinsurance mechanisms – AI application
Appetite	Line size, shares and capacity deployed; declinature where material automated decisioning cannot be assessed.
Price	Uncertainty loads on AI-priced or AI-assessed books, analogous to data-quality loadings.
Terms & conditions	Exclusions, sub-limits or event definitions addressing correlated AI failure; sunset/review clauses; claims-cooperation clauses for AI-assessed claims; deciding whether 'silent AI' is consciously covered, excluded or sub-limited.
Information rights	Warranties or conditions requiring notification of material AI incidents and changes of material vendor or model; periodic AI-governance attestations or updated questionnaire responses at renewal.
Engagement	Agreed improvement plans with reassessment milestones; joint reviews of AI-related experience; escalation to senior/portfolio level where concentration concerns arise (see 2.5).

- **Practical questions:** Is the exposure material? Is the cedant improving? Is AI failure covered, excluded or unaddressed in the wordings?
- **Example consideration:** For a Tier 3 cedant with material automated pricing: load for uncertainty, add information rights and an improvement plan, or restrict appetite and decide whether "silent AI" is covered, excluded or sub-limited.

3.4 Monitor

- **Objective:** Keep the assessment alive at two levels – the cedant (has its AI use, maturity or incidents changed?) and the portfolio (is concentration building?).
- **Indicators: Relevant where an underlying dependency could create correlated, material impacts on the reinsured portfolio (see 2.5).**
  - Vendor concentration (share of exposure behind each application vendor).
  - Foundation-model exposure (e.g. GPT/Claude/Gemini/Llama).
  - Cloud/region exposure (e.g. AWS/Azure/Google Cloud).
  - Assessment-tier trend.
  - AI incidents.
  - Loss-ratio drift on AI-priced books.
  - External signals (outages, recalls, regulatory action).
- **Suggested practices:** Set a reassessment cadence and triggers that force a return to 3.2 or 3.3 and feed concentration metrics into accumulation work (see 2.5).
  - **Periodic reassessment:** Proportionate to materiality and tier – e.g. at each renewal for Tier 3 cedants with material AI use; a light-touch confirmation of no material change for Tier 1.
  - **Trigger-based reassessment:** A defined event forces a return to 3.2 or 3.3 between cycles: a significant AI incident at the cedant or a shared vendor/model; a material change in the cedant's AI stack; AI adopted in a new material process; loss-ratio drift or assessment-tier deterioration; or external signals such as a major outage, model recall or regulatory action (see 2.5).
- Systemic triggers (shared vendor or model) additionally invoke the "who else is exposed?" response from 3.5.

3.5 Respond

- **Objective:** Have a pre-agreed way to respond to an AI incident and to feed lessons back into identification and assessment.
- **Idiosyncratic vs systemic:** For a single-cedant event, ask which treaties respond and what the loss is; for a systemic event, ask which cedants share the failed model/vendor and what the total accumulation is.
- **Practical questions / evidence:** Are escalation arrangements defined within the reinsurer's existing governance framework? Can the exposure map answer "who else is exposed?" quickly? Are AI losses notified under treaty terms? Stay forward-looking on emerging risks (agentic AI, model collapse, regulatory change).

4 Enablers

4.1 Responsible AI use, skills and assessment tools

WHY THIS MATTERS

Capabilities that may support reinsurance assessment include the ability to interpret relevant cedant AI information, understand potential treaty or wording implications and identify common dependencies or accumulation across cedants. Existing underwriting or assessment tools may also be adapted, where useful, to incorporate relevant AI considerations.

KEY CONSIDERATIONS

Skills and capability

- Identify the lines of business or cedants where detailed AI risk assessments are necessary for underwriting the reinsured risk.
- Build capability to interpret cedant AI disclosures and governance evidence (model inventories, validation summaries, bias-testing results).
- Develop understanding of the treaty-wording implications of AI failure (see 2.3.3) and of accumulation analytics across cedants (see 2.5).

Assessment tools

- Develop standard assessment questionnaires and due-diligence templates for treaty and facultative submissions, aligned to the assessment tier in 3.2.
- Maintain documentation supporting AI risk assessments.
- Periodically review and update assessment tools as AI practices evolve.

PRACTICAL EXAMPLE

Develop a standard AI assessment questionnaire to support consistent AI risk assessments across underwriting teams while allowing professional judgement for individual treaties.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- Do employees have the knowledge and tools needed to assess AI-related risks?
- Are AI assessments applied consistently across business functions?
- Are assessment tools reviewed and updated regularly?

4.2 Portfolio monitoring and market intelligence

WHY THIS MATTERS

A reinsurer's portfolio perspective may provide visibility into common AI dependencies and potential accumulation across multiple cedants. Existing portfolio and emerging-risk monitoring may therefore consider common providers, models or infrastructure where relevant.

KEY CONSIDERATIONS

Portfolio monitoring

- Monitor AI adoption across cedants and lines of business.
- Identify dependencies on common AI providers, models and technologies.
- Monitor emerging accumulation and concentration risks.
- Track the concentration indicators in 3.4 (vendor, foundation-model and cloud exposure) and feed them into accumulation management (see 2.5).

Continuous improvement

- Monitor regulatory developments, industry guidance and significant AI-related incidents.
- Update assessment approaches and share lessons learned across the organisation.
- Focus market intelligence on the (re)insurance AI vendor ecosystem – vendor incidents, model recalls and wording developments such as AI exclusions – relying on the MindForge Toolkit and institution-level frameworks for general regulatory watch.

PRACTICAL EXAMPLE

Following new guidance on agentic AI, a reinsurer updates its AI assessment questionnaire and underwriting due diligence to reflect emerging risks.

QUESTIONS TO INFORM REINSURER ASSESSMENT

- How is AI adoption monitored across the portfolio?
- Are common AI dependencies or concentration risks increasing?
- How are emerging AI developments reflected in your assessment approach?

5 References

- Mindforge AI Risk Management Toolkit – [JL7.1](Executive; Operationalisation, Jan 2026; Implementation Examples, Jan 2026); seven-dimension taxonomy (Monitoring & Stability; Cyber & Data Security); 3.2 Data Acquisition & Processing; 3.5 Monitoring
- MAS Guidelines on AI Risk Management (consultation) – supervisory expectations for FIs (to 31 Jan 2026); final Guidelines pending.
- MAS FEAT Principles (2018) and the Veritas Initiative – the foundation Project MindForge builds on.
- NIST AI Risk Management Framework 1.0 – Govern, Map, Measure, Manage.
- EU AI Act – high-risk classification of certain insurance pricing/underwriting.
- Gallagher Re – AI model evaluation & concentration risk; Swiss Re – AI boom reshapes risk.
- ACORD – data standards for AI in insurance; vendor references: Akur8, Cytora, Gradient AI, Shift Technology, Tractable, CCC, Verisk, LexisNexis.
- ACORD/SCOR case study – Reinsurance News (19 Aug 2025)



SINGAPORE REINSURERS' ASSOCIATION

SINGAPORE REINSURERS' ASSOCIATION
180 Cecil St, #16-04 Bangkok Bank Building,
Singapore 069546