

Proposed Additions to Exhibit B

AI Risk Evaluation Supplement, Version 5.0

Exhibit B: AIS Program (Checklist)

Submitted by the Center for AI Oversight, September 26, 2026

Companion to the Center's comment, "Exhibit B Is the Oversight Exhibit"

Summary

Five recommendations, expressed as one amendment to item 1 and seven new entries, plus one optional item. Each is answerable from a document the company either holds or does not. None asks an examiner to evaluate a model, a data source or a testing method. None prescribes governance content.

Proposed ref	Item	Completes
Item 1, amended	What the program governs, its granularity, who approved that as appropriate, the standards elected and the depth	What the program covers, and whether anyone decided it was the right depth for this company
2b, new	Committee and officer accountability, with delegation	Who answers, at two levels
2c, new	Competency at the accountable level	Whether the recipient can interrogate the report
2d, new	AI Systems declared out of scope, and who accepted them	The boundary of the program
3q, new	Ratified AI risk tolerance	Whether the tolerance is a decision
3r, new	Authority to operate, dated and expiring	Who authorized automation
3s, new	The escalation line	What happens when a system runs outside tolerance
3t, new	The disclosure position	What the company will report, and when
3u, new, optional	Independent assurance over the program	Whether the commitment was tested

On reference letters. These are drafted against the version 5.0 Checklist. The 2 series contains items 2 and 2a, so the accountability items are proposed as 2b, 2c and 2d. The 3 series runs continuously from 3a to 3p, with 3o and 3p new at this exposure, so the proposed 3 series items continue at 3q. The letters are a convenience rather than a recommendation, and the Center has no view on placement. Note that the Narrative option uses its own numbering, in which 2a to 2e already exist and address a different subject.

On the two options. Version 5.0 offers Exhibit B as a Narrative or a Checklist. These items are drafted for the Checklist. Part 3 notes where the Narrative already reaches nearby ground and what it leaves open, so the Working Group can decide whether to mirror the additions there.

Items in the 3 series complete the existing stem, *"Reference the processes and procedures of the Company AIS Program that addresses the following: How the Insurance Company..."*

Part 1. Clean text for insertion

Amendment to item 1

Existing (v5.0, Checklist item 1): "Has the company adopted a written AIS Program? If yes, when was it adopted and what is the frequency of review for updating?"

As amended:

Has the company adopted a written AIS Program? If yes, when was it adopted and what is the frequency of review for updating? State what the AIS Program governs and the granularity at which it operates, identify the individual or body that approved that scope and granularity as appropriate to the company's size, complexity and use of AI Systems, and state the date the scope and granularity are next reviewed for continued appropriateness. Identify the external frameworks or standards the company elected to apply to the AIS Program, the depth of application elected, and the individual or body that approved the election, with the date of approval.

Note on placement. As amended this item carries two distinct questions, one about the Program's own scope and granularity and one about the external instruments elected. The Working Group may prefer to split it into item 1 and a new item 1a. The Center has no view on which.

New items, 2 series

Ref	AI Systems Use Questions for Company	Company Response
2b	Identify the committee or governing body that holds oversight of the AIS Program by charter. Identify by name and role the officer accountable to it for the AIS Program, together with any delegation of that accountability and the date the delegation was recorded.	
2c	Describe how the company establishes that the accountable officer and the committee hold the competency required to interrogate what the AIS Program reports to them.	
2d	Identify any AI Systems the company has declared outside the scope of the AIS Program, the named individual who accepted each exclusion, and the date each exclusion is scheduled for review.	

New items, 3 series

Completing the existing stem, *"Reference the processes and procedures of the Company AIS Program that addresses the following: How the Insurance Company..."*

Ref	AI Systems Use Questions for Company	Company Response
3q	Ratifies AI System risk tolerance at the accountable level, identifying the body that ratified it, the date of ratification, and the date of next scheduled review	
3r	Grants, bounds and periodically re-ratifies authority to operate for AI Systems recorded as automating a decision, recording the authorizing individual or body, the date of authorization, its scope, and its expiration or re-ratification date	
3s	Escalates an AI System operating outside stated tolerance, identifying (i) the standing cadence on which the AIS Program reports to the accountable committee independent of events; (ii) whether that reporting carries adverse information, incidents, findings, regulatory inquiries and deviations from tolerance and not only progress; (iii) the escalation triggers and materiality thresholds fixed in advance; (iv) the party with authority to continue, constrain or withdraw the System; (v) whether at least one reporting pathway to the governing body does not depend on management's discretion; and (vi) whether the response to an escalated matter is recorded	
3t	Fixes the company's disclosure position for AI System matters, identifying the criteria for disclosure to the department, to consumers and to the public, and any clock that statute or regulation attaches to it	

Optional item

Ref	AI Systems Use Questions for Company	Company Response
3u	Obtains independent assurance over the AIS Program, identifying the provider, the scope of the review and the date of the most recent review. Where assurance has been obtained, the object of review is whether the company applied the frameworks and depth of application it recorded in response to item 1.	

A blank response to the optional item is information rather than a deficiency. Independent assurance over an AI systems program is not yet common practice, and the value of the question in the near term is that it establishes a baseline.

Part 2. Basis for each item

Not proposed for inclusion. Recorded for the Working Group's consideration.

The amendment to item 1. The existing item establishes that a program exists and when it was adopted. It does not establish what the program governs, or what the company undertook to do.

The object clause is four words in the exhibit and it separates two different documents. Version 5.0 defines an AIS Program as "the controls and processes that an insurer adopts and implements to ensure the

responsible use of AI Systems," which "may be comprised of general guidelines, governance framework, risk management and internal controls, and third-party oversight." That is a definition of what a company does, and Exhibits A, C and D already reach that ground thoroughly. A program whose stated object is instead the company's capacity to make informed decisions about those systems is an oversight program, which is what this exhibit is positioned to test. Without the clause the examiner cannot tell which one was filed, and a company that has built only the first will answer the rest of the exhibit in terms of its model inventory.

"Depth of application elected" is deliberate: a company that adopts a framework in part, deferring specified elements, has made a legitimate decision that should be recorded rather than obscured. The approver is asked for because adoption with no identified approver is an administrative act, not a governance decision.

2b, committee and officer. Four existing questions come close. Checklist item 2 asks whether the Board of Directors or management was "involved in the adoption," and item 2a asks "what is the role of the Board of Directors or management in the AIS Program." Narrative item 1a asks "what role maintains the AIS Program." Narrative item 1e asks how "responsibility for governance within the organization is assigned."

A company can answer all four accurately and leave the examiner unable to name anyone. Maintaining a program is an administrative function, assigning responsibility describes an organizational chart, involvement in adoption is a historical fact, and a role is a description rather than a name. None of the four resolves to a party who answers.

Insurance regulation already has the naming mechanism in two places, quoted in the accompanying comment: the CGAD signature under Model Act #305 Section 3.B, and the ORSA Summary Report signature under Part C of the Guidance Manual's Introduction. This item asks the AIS Program equivalent of a question the sector already asks on adjacent subjects. It asks for a role as well as a name, because the role survives personnel changes and can be tested against the company's own delegation records at the next examination.

2c, competency. The ORSA attestation confirms that a copy of the report was "provided to the insurer's board of directors or the appropriate committee." Provision is receipt. What makes receipt into oversight is the capacity to interrogate what was received. The item does not ask the accountable officer to validate a model. It asks whether the company can show the officer is equipped to question a report about one.

2d, exclusions. The item most likely to produce a finding, and the one with no existing analogue in either option. Every AI systems program has a boundary. Systems fall outside it because they are low risk, because they are legacy, because they were acquired with a book of business, or because no one has decided. The first three are decisions. The fourth is not. A company answering Exhibits A and C in good faith may simply not report systems it considers out of scope, leaving the examiner unaware that a scoping judgment was made at all. A company that cannot name an acceptor has not made a decision about the system. It has omitted it.

3q, ratification. Existing item 3k asks how the company "Quantifies AI System risk levels," and new item 3o asks how it "Determined 'materiality' and what threshold was used for assessing AI Systems risk." Both ask for the assessment and how it was derived. Assessment is a technical act performed by a modeling or risk function, and it may be qualitative as readily as quantitative. Ratification is a governance act performed by a body that can be held to it, and a tolerance assessed but never ratified is a working assumption the company is not committed to.

This item is also what makes Exhibit C's provision that "AI model risk criteria are set by the insurance company" supportable. The Center supports that provision, because allowing the company to set its own

criteria is what lets the Supplement scale from a small domestic carrier to a global group. But criteria set by the company are a decision by the company, and this is where that decision is recorded.

3r, authority to operate. Existing item 3e asks how the company "Evaluates whether AI Systems are suitable for their intended use and should continue to be used as designed," and Narrative item 1g asks how it "assess[es] autonomy, reversibility, and reporting impact risk." Both are the right questions asked in the passive register of evaluation rather than the active register of authorization. Assessing the risk of autonomy is not the same as authorizing it.

Exhibit C item 10 already records AI model type as automate, augment or support, and the Supplement defines Automation as "an AI System that does not involve human intervention." Removing the human is a decision, not a design outcome. The four attributes requested, authorizing party, date, scope and expiration, are the minimum that makes an authorization testable. An authority that does not expire is not an authority.

3s and 3t, escalation and disclosure. Narrative item 1b asks the company to discuss "the governance structure, Board reporting and frequency," which establishes that reports move and how often. Checklist items 3m and 3n are consumer-facing. None of them establishes what the report carries, what triggers one outside the cadence, who can act on it, or what was done.

Escalation is not incident response. Incident response addresses events already recognized as incidents. Escalation addresses the system behaving as designed and outside the boundary the company set for it, which is the harder case and the one that produces supervisory surprise.

Each element of 3s is testable from the company's own records. A standing cadence establishes that reporting does not wait on something going wrong. Adverse information establishes that the report is decision-useful rather than a progress update. Triggers fixed in advance establish that the threshold was not set after the event. A destination with authority to withdraw establishes that escalation resolves to an action. A pathway independent of management's discretion establishes that the governing body can learn what management would prefer it not learn. And a recorded response establishes what was done, because a red flag received without one is, on the record, a red flag ignored.

Disclosure is proposed as a separate item at 3t rather than a seventh element of 3s, because it is answered by a different document and often by a different function. A company that has not fixed its criteria in advance will fix them under pressure, and the department is the party most affected by that.

3u, independent assurance. Narrative item 1d asks how "the effectiveness of the AIS Program and individual models are assessed and modified," and Narrative item 3a asks whether validation and testing are "performed by someone that is independent from development." The Working Group has therefore already applied the independence principle, but to model validation rather than to the program.

The object of assurance stated in the item is deliberate. The useful question is not whether a program conforms to an external standard in the abstract. It is whether the company applied what it said it would apply, at the depth it recorded in item 1. That is testable, it is within the competence of an internal audit function or an external provider, and it does not require the assurance provider to evaluate a model. ISO/IEC 42006:2025, published July 7, 2025, gives this market a published accreditation standard, which allows the Working Group to reference the availability of independent assurance without mandating its use.

Part 3. Three observations on the existing exhibit

No new items required.

Items 2 and 2a both read "Board of Directors or management." The disjunctive is understandable in an instrument accommodating a range of company sizes and structures, and item 2a already asks the right question in asking for the role. It also permits a response that is accurate and uninformative, because a role can be described without anyone being named. If the Working Group prefers not to add a separate accountability item, the narrower alternative is to amend 2a to ask which body holds the role, under what instrument it holds it, and who answers for the program by name.

The 3 series stem is "How the Insurance Company..." Every item in the series is therefore a process question, which is appropriate for most of the series. Items 3q through 3t are decision questions rather than process questions. The Working Group may prefer to group them under a separate stem, for example *"Identify the following decisions and the party that made them,"* rather than extending the existing series.

The Narrative option asks companies to discuss; the Checklist asks them to point at a document. That difference is the Working Group's own design and the Center does not propose changing it. It is worth noting only that a discussion is not a record. Where the Working Group elects to mirror these additions in the Narrative, the Center would suggest preserving the identifying verbs, so that the company identifies a committee, an officer, a body, a date and a party rather than discussing how accountability is approached.

Proportionality

For a small or resource-constrained insurer, 2b and 3s carry most of the value, with 3r added wherever a system automates a decision. That minimum viable form is why an oversight instrument scales from a small domestic carrier to a global group without a separate small-entity version.