King Yew Choo

London
kingyewchoo.co.uk

CPMI Secretariat and IOSCO Secretariat
By email to cpmi@bis.org and cpmi-iosco@iosco.org

12 November 2026

CPMI-IOSCO CONSULTATIVE REPORT, SEPTEMBER 2026 · CYBER RESILIENCE TOOLKIT

Authority and resubmission in FMI recovery and testing

Comment on sections 2.3, 3.3.3, 3.4 and 4.5.2, Box 4, Annex 2A and Tables 5 and 8

Respondent	King Yew Choo, London, United Kingdom
Capacity	Personal capacity
Interest	Founder, True Primary; declared at paragraph 12
Contact	Contact details are in the covering email, and at kingyewchoo.co.uk
Publication	No part is confidential. It may be published in full, with my name.
Structure	Numbered paragraphs 1 to 12, with a summary above and the reference below

WHAT I PROPOSE

1. **Log the authority behind an AI system's action.** Where an AI system can act in an FMI's ICT systems, section 3.3.3 should add two event types: the authority under which it acted, and the identifier of the transaction or instruction the action concerns (paragraphs 1 to 3).
2. **Test replay and resubmission together.** The toolkit names both as recovery steps. Annex 2A should carry a stylised scenario in which the FMI replays lost transactions while participants' automated systems resubmit the same instructions. A recovery plan tested against it passes when each instruction settles once (paragraphs 4 to 8).
3. **Enforce a red team agent's limits in the tooling that runs it.** For an AI agent, the permissions in Table 8 should hold while it runs: scope, rate limits, an end time and an immediate stop held by the control team, with every action logged (paragraphs 9 to 11).

Section 3.3.3: the authority behind an AI system's action

1 Section 3.3.3 states that "AI use in ICT systems can benefit from logging the following event types to provide visibility into systems: input, output, model version, confidence score, action taken and human overrides".[1] These event types record what the system did. A review after an incident also needs to establish whether the system was permitted to do it.

2 Where an AI system can take actions in the FMI's ICT systems, such as changing a configuration or isolating a system, the list should add two event types. The first is the authority under which the system acted: who granted it, its scope and limits, when it ends, and whether it had been revoked when the action occurred. The second, where the action concerns a transaction or instruction, is that item's identifier. Where the grant already sits in the FMI's access management, a reference to it is enough. Section 3.3.4 applies the same test to people: documented and tested pre-authorised emergency procedures let FMIs verify that stakeholders understand "who has the authority to do so, what actions are permitted".

3 With both fields, an FMI can separate an authorised action from an unauthorised one after the event, and can link every action taken on one instruction, which matters most in the recovery described in paragraphs 4 to 8. The authority record needs the same protection as the transaction logs. The second stylised scenario in Annex 2A has an insider whose changes are "hidden through log manipulation".

Recovery: replay and resubmission of the same instruction

4 The toolkit describes two routes by which a transaction can reach an FMI again after an integrity breach. The FMI may replay it. A fix-backward recovery in section 3.3.2 restores the FMI to its golden point and "may require thorough data reconciliation to solve data loss in relation to both internal and external systems". Box 4 lists "replay of (lost) transactions" among the reconciliation processes, and section 3.3.3 describes an event-driven architecture that "may include maintaining transaction replay capabilities". A participant may also resubmit it: section 3.4 lists the involvement of interconnected entities in resumption operations, "eg for coordinated actions, transaction resubmission etc". Table 5 adds pre-agreed backup channels "for settlement instructions if primary communication networks are unavailable".

5 Taken together, these routes can bring one instruction to the FMI twice. The risk exists today wherever participants' payment systems resubmit automatically on a timeout, or staff re-enter instructions through a backup channel. In a payment system where settlement is final and the rulebook provides no reversal, a duplicate is recovered by a return from the receiving participant, and the paying participant carries the liquidity cost until then. The same question arises for settlement instructions at a CSD or securities settlement system and for trade submissions to a CCP.

6 Annex 2A should add an eighth stylised example in its existing form, **Duplicate instructions on resumption**:

- Timing: a high-volume intraday window. All test payments are valid, funded and due to settle.
- Threat: malware corrupts the FMI's settlement database. The FMI halts processing and restores to a golden point from 40 minutes earlier, verified to predate the attacker's presence.
- Disruption: instructions received after the golden point are held in the FMI's separately stored transaction record, and some were never acknowledged. While processing is halted, participants' automated systems resubmit those instructions, some through the pre-agreed backup channel and some under new message references. On resumption, the FMI replays the instructions from its record and processes the resubmissions.
- Impact: instructions received both ways settle twice. The FMI's records and its participants' records diverge, and reconciliation has to identify each duplicate and arrange its return before the end-of-day cycle closes.

7 Section 2.3 asks FMIs to "Test existing strategies and response, resumption and recovery plans at regular intervals" against their scenarios. Tested against this one, a plan passes when each instruction settles once, whichever channel or message reference carried the resubmission. Three controls make that achievable. First, the FMI recognises a resubmission as the same instruction. Payment messaging already provides end-to-end references and possible-duplicate flags; the test shows whether they survive a replay and a backup channel, whose procedure should capture the original reference when an instruction is re-entered from secure email or phone. A duplicate check that relies on the message reference alone misses a resubmission sent under a new one. A resubmission that carries no original reference goes to the pending queue for confirmation with the participant, since matching on content alone risks rejecting a genuine repeat payment.

8 Second, a recognised resubmission settles nothing new: the FMI rejects it with a reason that points to the original, or reports the original's outcome. Third, an instruction whose outcome is unknown at recovery is held as pending, and is not re-executed, until reconciliation establishes its outcome. A time limit and a fallback agreed with participants in advance keep the pending queue from blocking completion of settlement by the end of the day. Participants share the control: during a declared halt, the FMI's incident communication tells them whether to suspend automatic resubmission and how to resubmit, so that resubmission stays coordinated as section 3.4 envisages.

Section 4.5.2 and Table 8: limits on a red team's AI agent

9 Section 4.5.2 states that red teams' use of AI tools such as AI agents, "if not subject to sufficient guardrails, could pose new risks to FMIs given their autonomous nature and decision-making capabilities". Table 8 already asks FMIs to "define red team actions requiring permission from the entity" and, in contracts with red team providers, to define out-of-scope systems and the "actions not allowed to be performed".

10 Those controls sit in the test plan and in the provider's contract. An AI agent can take many actions between two human reviews, and its access in the target is whatever it gains during the test. Its limits therefore need to be enforced in the tooling that runs the agent, while it runs: the systems and actions in scope, limits on volume and rate, an end time, and a stop that the control team can trigger at once. The provider contract in Table 8 is the place to require this, together with a log of every action the agent takes, the authority it ran under and any human override. The control team should test the stop before the active phase and receive the agent's action log from the tooling at closure, so that the check does not rest on the provider's report alone.

11 Set out this way, the guardrail can be checked before the test and after it. Before, the FMI confirms that the agent's scope falls within the "informed, exercise-specific written consent" secured from potentially affected participants under section 4.5.2. After, it confirms from the log that every action fell inside the scope and, where the control team used the stop, that it took effect at once.

Capacity and interest

12 I respond in a personal capacity. I founded True Primary, which builds infrastructure for the agent-initiated purchase of expert work, and declare its commercial interest in the safe adoption of agentic payments. True Primary is not an FMI and is not a participant in one.

King Yew Choo

London, United Kingdom

REFERENCES

Web addresses were accessed on 2 November 2026.

1. CPMI-IOSCO, Cyber resilience toolkit: practical considerations for FMIs, consultative report, September 2026, published 8 September 2026: section 2.3, page 19; Annex 2A, pages 20 and 21; section 3.3.2, page 28; Box 4, page 30; section 3.3.3, page 31; section 3.3.4, page 33; section 3.4 and Table 5, page 34; section 4.5.2 and Table 8, pages 55 and 56. https://www.bis.org/publications/cpmi-iosco-cyber-resilience-toolkit-practical-considerations-fmis.pdf