



AI GOVERNANCE · SYSTEMIC RISK · INSTITUTIONAL RESILIENCE

Public Comment on the CPMI-IOSCO Discussion Paper *FMI's reliance on third-party service providers: challenges and risks*

Submitted to the Secretariats of the Committee on Payments and Market Infrastructures and of IOSCO.

Questions 1, 2 and 5, with Questions 3 and 4 addressed together.

Respondent André Afonso Silva, M.Sc. · Rio de Janeiro, Brazil

Capacity Personal capacity. Not on behalf of any FMI or service provider

Submitted to cpmi@bis.org · cpmi-iosco@iosco.org

Consultation closes 1 December 2026 · **Prepared** 24 September 2026

RESPONDENT

Respondent

I am André A. Silva, founder and CEO of Essência A.I., an AI governance consultancy in Rio de Janeiro. I also serve as Vice-President of the Council on Innovation, Digital Transformation and Deep Techs of the Federation of Foreign Trade Chambers (FCCE), an affiliation stated for transparency only. I do not speak for the FCCE, any FMI or any third-party service provider.

Since October 2025 I have published *Essência AI Insights*, a weekly series on AI governance, systemic risk and institutional infrastructure. The series is listed in the AI Dialogue Partnerships Hub of the United Nations Global Dialogue on AI Governance.

I answer Questions 1, 2 and 5 and address Questions 3 and 4 together. I leave Questions 6 to 11 to FMIs and service providers, to whom they are specifically directed. The analytical distinctions below come from my published work. Where I transpose one to FMI third-party risk, I mark the step as mine. Empirical and regulatory claims are cited separately to the primary sources listed at the end.

RECOMMENDATIONS

Summary of recommendations

First, provider counts should not be read as failure-path counts. Separately contracted services can converge on one critical underlying component. Four provider attestations do not establish four independent failure paths.

Second, independence should be treated as time-dependent rather than settled at procurement. Architectures, subprocessors and configurations change between formal review cycles. A dated reconciliation can detect the loss of separation, although it cannot prove independence.

Third, contractual entitlement and exercised capability should be recorded as distinct evidentiary states. A right can be legally valid without ever having been exercised. For critical services, the contract should specify how the capability will be exercised and what evidence that exercise leaves behind.

QUESTION 1

Are the challenges accurate and comprehensive?

The six challenges are, in my reading, accurately identified. I do not propose a seventh. My concern is narrower and it is about the evidentiary reach of the mechanisms available to address them.

The paper already names the object at issue. Section 2.2.1 includes, among the forms of concentration risk, cases where multiple third-party service providers depend on the same key nth party ^[1]. The point is therefore not that shared dependencies were overlooked. The question is how the independence of separately contracted providers is established and how it is re-established over time.

The paper also records real limits on the evidence available for answering that question. Approximately one third of surveyed FMIs reported visibility only into their direct providers. Others reported visibility into fourth parties, with only a small portion reporting visibility beyond that point. Service providers, questioned separately, reported that visibility degrades quickly along their own supply chains [1].

These findings matter because contractual plurality and dependency plurality are different facts. Four contracts can still converge on one critical component. Four provider attestations do not, by themselves, establish four independent failure paths, because each attestation establishes, at most, the provider's own disclosed branch. Branch-level mapping establishes depth along a branch. It does not, by itself, establish cross-provider failure-path independence.

A further distinction belongs to the evidence rather than to the dependency, and I mark it as mine. Dependency information supplied by the inspected provider may be entirely accurate. Its origin still bears on the inference drawn from it and it remains subject to checks on integrity, completeness, configuration and access control. More to the point, no provider can establish another provider's independence from its own chain alone.

The cyber resilience toolkit supplies genuine mapping instruments for this work, including inventories, dependency mapping and software bills of materials [3]. Those instruments improve visibility into dependencies represented within their scope. An SBOM focused on software components need not surface a foundation-model dependency unless that dependency is represented in the inventory.

Artificial intelligence makes the layer problem easy to see. BIS Papers No 154 separates the AI supply chain into five layers: hardware, cloud infrastructure, training data, foundation models and applications [6]. Foundation models can support downstream applications that reach the buyer through separate commercial relationships. Two separately procured AI services can therefore depend on a common model layer or cloud layer. This is a structural possibility and not a claim about how often it occurs among FMIs. I have developed the distinction between the procured layer and the layer that carries the dependency in my published work [5].

FSI Occasional Paper No 28 approaches the same layer from the cyber side. Its authors identify common cloud, software and frontier AI providers as potential concentration points [7]. They also identify sovereign access risk. A single provider's disruption or policy decision can affect firms across jurisdictions. Those are the authors' findings. The paper states that they do not necessarily reflect the views of the BIS or of the Basel-based standard-setting bodies.

The same logic is not specific to AI. Any critical nth party can create convergence across contractual branches and AI simply makes the distance between product layer and substrate unusually visible. My answer to Question 1 is therefore qualified rather than negative. The six challenges are comprehensive at the category level. I did not identify in the paper a procedure expressly designed to test whether separately contracted providers still represent separate critical dependency paths.

QUESTION 2

What possible solutions could address the challenges?

I address three of the six challenges here: concentration, supply-chain opacity and ecosystem interconnection. I do not propose a separate solution for each challenge the paper identifies.

A dependency assessment is always dated. Provider architectures, subprocessors and technical configurations can change after it is completed. An observation of separation at one moment therefore cannot establish permanent independence. Execution and configuration move continuously while verification remains episodic and the interval between the two is where exposure accumulates. I develop this cadence asymmetry in a working paper, which is conceptual and architectural and which offers no quantitative validation or sector-specific case study [4].

I propose a risk-based dependency reconciliation exercise. Its object is the critical dependency path, not corporate or legal independence. The exercise asks one question: do services treated as operational alternatives share a material underlying dependency?

The starting point is the set of services that matter to resilience. For those, the FMI identifies the layers the service actually requires, which may include cloud infrastructure, model providers, critical software and material nth parties. Scope stays proportionate to criticality.

The exercise then reconciles the evidence available across providers. SBOM data supports software dependencies represented in the inventory. Cloud, model and contractual dependencies may require additional evidence, including targeted written requests. Where the available evidence cannot resolve a material dependency, the result should be recorded as unresolved rather than as absent.

The output is a dated dependency record. It should identify confirmed common dependencies, material unresolved dependencies, the evidence supporting each conclusion and who controlled the production of that evidence. Reassessment should be both periodic and event-driven, triggered by material provider, architecture or subprocessor changes.

Two limits belong in the record itself. Failing to identify a shared dependency is not proof that none exists. The record should say what the defined exercise found using the evidence available. The exercise is designed to detect the loss of separation. It does not guarantee detection.

Nothing here requires unsafe testing. The exercise does not involve creating, in a live environment, the dependency being examined. Reconciled inventories, targeted evidence requests and controlled technical validation can support the exercise.

The toolkit already contains adjacent machinery, since it recommends testing with ecosystem entities and reassessment after significant changes [3]. The addition proposed here gives that work a specific object: convergence across providers on a single critical dependency.

QUESTIONS 3 AND 4

Would further engagement or policy be helpful?

Yes, and I would favour focused work on evidence structure rather than an additional general policy layer.

The paper itself explains why. It notes that individual FMIs cannot obtain a system-wide perspective. It adds that supervisory information may be limited to the FMI sector and to a single jurisdiction [1]. Annex 1 records FMIs asking for greater industry-level collaboration on interconnectedness. They also noted that concentration should be assessed at ecosystem level, not only at FMI level [1]. Two different

units of analysis are observing the same infrastructure. The FMI asks whether its own operations survive a provider failure. The authority asks whether many FMIs depend on the same underlying component.

The conceptual work already exists, so the useful step is not another taxonomy. BIS Papers No 154 offers a five-layer decomposition of the AI supply chain [6]. The FSB has published monitoring work covering third-party dependencies [8]. It proposes indicators for criticality, concentration and substitutability. It includes a case study on AI-related third-party dependencies and service provider concentration. The gap I identify is a structure that lets dependency evidence travel between the firm-level and the system-level view without the two being confused.

Three implementation questions follow and they are where further CPMI-IOSCO engagement would be most useful. What minimum fields describe a critical dependency at FMI level? Which identifiers allow one underlying provider to be recognised across different commercial products? Which material changes should trigger an updated dependency record?

The first question belongs to the firm-level view, the second enables aggregation and the third connects both across time. I would not begin with a global repository, since jurisdictions face different confidentiality, legal and supervisory constraints. Comparable fields and identifiers are a narrower first step and they support aggregation wherever existing mandates already permit it.

QUESTION 5

Are there contractual aspects in addition to those addressed?

Yes. I suggest distinguishing contractual entitlement from exercised operational capability and recording the two as separate evidentiary states.

The paper already addresses audit rights, information-sharing provisions, access to testing and service-level expectations and it records practical limits on exercising them [1]. Annex 1 notes that the right to audit can be difficult to enforce, and that certain FMIs receive pre-set contracts with limited negotiating room [1]. My proposed addition concerns what evidence supports counting such a provision as a resilience capability.

A contractual right can be legally valid without ever having been exercised, and nothing here questions that validity. The clause establishes the entitlement. A dated exercise provides evidence about operational exercisability. The two answer different questions and should not be recorded as one.

The paper gives the distinction practical weight. It cites the Level 3 cyber review. That review found most FMIs had not included critical third-party service providers in testing their response, resumption and recovery plans [1]. That finding does not show those arrangements were incapable of working. It shows why documentary availability and exercised capability need to remain separate categories.

Exercises should be proportionate rather than literal. An audit right need not require a full audit each cycle, since a targeted information request can test access mechanics and response times. An exit plan need not require an actual exit, and the paper records that full exit may be unrealistic under stress. A technical portability test or a tabletop exercise can produce relevant evidence instead. A joint testing provision can be evidenced by a coordinated test, recorded with date, scope, expected result and observed result.

Evidentiary status should also be allowed to expire. After a material configuration change, a prior exercise remains historical evidence, and it should not silently become evidence for the new arrangement.

One published example shows what such a record looks like, and I offer it as an illustration of form rather than as evidence about FMIs. On 4 August 2026 the UK AI Security Institute published an incident report covering a cyber evaluation conducted in late July [9]. The report sets out the configuration used, including that internet access was deliberately enabled and that no sandbox escape occurred. It also records the detection route, the actions taken and the controls not yet implemented. It also records that the unsanctioned attempts were unsuccessful and that, to the institute's knowledge, no real-world harm resulted. The adequacy of that arrangement can be discussed from outside because the record exists and was published.

The disposition I propose is contractual and narrow. For critical services, the contract would specify five things. How each capability is to be exercised. The risk-based cadence of that exercise. What evidence is retained. Who may access it. And what follows when an exercise shows the capability does not work as assumed. The last element carries the most weight, because an exercise with no consequence attached produces a record and not a remedy.

The underlying proposition is about inference rather than about contracts. Where no record of exercise exists, the contractual text alone gives an external assessor no basis for a distinction. A capability that can be used and one that cannot look the same. That limit applies to the supervisor, to the counterparty and to the FMI's own board. Recording the two states separately removes it at low cost.

SCOPE

On Questions 6 to 11

I do not answer Questions 6 to 11. They are directed specifically at FMIs and at third-party service providers, and I am neither. I leave them to respondents with direct operational experience.

METHOD

Method note

This comment was developed with AI assistance, under my established disclosure practice. Source selection, analytical positions, constraints, judgement and final responsibility are mine. Every factual claim was checked against the primary source cited. Where I transpose a concept from my prior work, I identify that step as my own.

SOURCES

References

- [1] Committee on Payments and Market Infrastructures and Board of IOSCO, *FMI's reliance on third-party service providers: challenges and risks*, discussion paper, 8 September 2026.
- [2] Committee on Payments and Market Infrastructures and Board of IOSCO, *CPMI-IOSCO consultation on third-party risk at FMIs*, cover note, 8 September 2026.
- [3] Committee on Payments and Market Infrastructures and Board of IOSCO, *Cyber resilience toolkit: practical considerations for FMIs*, consultative report, 8 September 2026.
- [4] Silva, A. A., *Temporal Sovereignty and Systemic Augmented Intelligence: Rebuilding Institutional Authority in the Era of Continuous Execution and Quantum Decryption*, SSRN, 2026, DOI 10.2139/ssrn.6266298.
- [5] Silva, A. A., "The Layer That Does Not Appear on the Invoice", *Medium*, 12 August 2026.
- [6] Gambacorta, L. and Shreeti, V., *The AI supply chain*, BIS Papers No 154, Bank for International Settlements, 18 March 2025.
- [7] Crisanto, J. C., Currat, A. and Yong, J., *When machines attack: frontier AI cyber threats and policy responses in the financial sector*, FSI Occasional Paper No 28, Financial Stability Institute, Bank for International Settlements, 9 September 2026.
- [8] Financial Stability Board, *Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector*, 10 October 2025.
- [9] UK AI Security Institute, *Security Incident INC-2026-07-28-01*, 4 August 2026.

AUTHOR

Author profiles

LinkedIn [linkedin.com/in/andreglobal](https://www.linkedin.com/in/andreglobal)
Medium medium.com/@andreglobal
Substack andreglobal.substack.com

