

W O R K I N G P A P E R

International Financial Centers · Part V

Cyber Resilience and Financial Centers

Dr. Jochen Biedermann*, Marc Krisjanous**

October 2026

Keywords: international financial centers, cybersecurity, resilience,
artificial intelligence, digital assets

** The author is the Managing Director of the World Alliance of International Financial Centers.
This paper reflects the views of the author and not those of the World Alliance.*

*** The author is the Founder and Principal Consultant of Zanarc Limited.*

Table of Contents

1.	Introduction: a collective action problem	4
1.1.	Report Structure	4
2.	What has changed	5
2.1.	Artificial intelligence: the cost curve, not the attack	5
2.2.	Agentic AI: risk inside the transaction	6
2.3.	Digital assets and tokenization: from code to keys	10
2.4.	The cryptographic floor is moving	12
2.5.	What the evidence does not yet show	12
3.	The concern this paper digs into: third-party dependency	13
3.1.	Where incidents actually originate	13
3.2.	Why the regulatory perimeter cannot close the gap	14
3.3.	Concentration is now reappearing in the AI layer	14
3.4.	What firms are not doing, and what supervision will not ask them to do	15
4.	Why this is a financial center concern.....	16
4.1.	The confidence channel: real and to be argued carefully.....	16
4.2.	The operational channel: better measured, and more alarming	16
4.3.	The jurisdictional dimension	17
5.	Is promoting standards enough?.....	18
6.	The practical role: seven functions for a financial-center organization.....	20
6.1.	Map dependence on common providers and services.....	20
6.2.	Connect existing information-sharing arrangements.....	22
6.3.	Coordinate cross-firm exercises.....	24
6.4.	Help smaller firms improve.....	25
6.5.	Support communication during major incidents	25
6.6.	Build talent as center infrastructure.....	26
6.7.	Speak collectively, and carry lessons between centers	27
7.	A resilience maturity model	28
7.1.	Applying the model: assessment method	30
7.2.	Demonstrating progress through evidence	31
8.	Continuous review and improvement.....	33
9.	Conclusion	35
	Bibliography.....	36

Abstract

Cybersecurity has been a supervisory subject in finance for a decade, and the regulatory architecture is now dense. It binds financial institutions. It does not, and largely cannot, reach the ecosystem in which those firms operate: the vendors, technology providers, and small licensees through which most consequential incidents now arrive. That gap is where international financial centers have a role neither the regulator nor the individual firm can fill.

Five findings for financial centers to consider:

1. AI can reduce the cost and effort needed to carry out familiar attacks. The IMF, the BIS, the Bank of England, and the UK National Cyber Security Council converge on this: the cost, speed, and scale of intrusion have changed, not the existence of new attack classes. Combined with the concentration of financial institutions on a small number of model and cloud providers, the systemic claim that follows is *asymmetry plus concentration*, not novelty.

2. The point of failure has moved outward and downward. In the incidents that have most damaged market confidence in the last three years, the breached entity was rarely a systemically important bank. It was a wallet provider, a payments connectivity vendor, a statement printer, a post-trade platform, a security software vendor. Research at the Federal Reserve Bank of New York shows that banks with under USD 10 billion in assets can trigger significant system-wide payment impairment through third-party linkages.

3. Supervision is closing part of this gap, but only part. The EU and the UK have both created oversight regimes for critical ICT third parties, designating around 20 and 4 providers, respectively. These reach the largest providers. They do not reach the long tail, and the European Supervisory Authorities have no power to fine designated providers. The IMF's 2025 assessments of both the euro area and Switzerland document supervisory capacity that is thin relative to the surface being supervised.

4. Financial center organizations need to help members implement standards and obtain evidence that controls work. The EU Roundtable of Financial Centers, meeting in Frankfurt in October 2025, called for harmonized implementation of DORA and common metrics for cyber preparedness. The diagnosis is inconsistent implementation and the absence of measurement, not a shortage of rules. Threat-led testing results in the UK point the same way: current certification coexists with basic control failures.

5. A financial center organization can perform a definable set of practical functions. Seven are set out here: mapping dependence on common providers and services, connecting the information-sharing arrangements that already exist, coordinating cross-firm exercises, helping smaller firms improve, supporting communication during major incidents, building talent as center infrastructure, and speaking collectively while carrying lessons between financial centers. None of them makes a center a supervisor or an operational cyber-response body. Each has working precedent in at least one center today. None has been fully assembled anywhere.

The paper closes with a maturity model across nine dimensions and a cycle of thematic reviews. Both are self-assessment instruments, supported by documentary evidence rather than self-declaration. Neither is a league table. Financial centers differ in mandate, powers, and access to firms, so the model measures each center against its own previous position rather than its peers.

1. Introduction: a collective action problem

In 2020, Tim Maurer and Arthur Nelson argued for an international strategy to protect the financial system from cyber threats (Maurer & Nelson, 2021). Their central observation has aged well:

"Most of the financial services community does not lack resources or the ability to implement technical solutions. The main issue is a collective action problem."

Collective action problems have a characteristic signature. Every participant would be better off with a shared capability. No participant has sufficient private incentive to build it. Supervisors act on institutions one at a time and are constrained to their perimeter. Firms internalize their own losses but not the externalities they impose on the surrounding market. Shared capabilities can remain undeveloped because individual institutions lack sufficient incentive to establish them.

Financial centers sit precisely at that level. A center, whether a promotion agency, an industry association, a cluster body, or a public-private partnership, cannot supervise, and should not attempt to. But it naturally owns inherently collective capabilities: shared intelligence, market-wide exercise, a rehearsed public voice, a talent pipeline, a security baseline for the smallest firms, and a common position before policymakers.

Reputational damage from a major incident doesn't stay with the breached institution; it attaches to the jurisdiction. Every institution in a center bears part of the cost of an incident at any one of them, and at any of their suppliers. Where the loss is collective, the defense must be collective.

1.1. Report Structure

Section 2 sets out what has actually changed in the threat landscape, including AI, agentic AI, digital assets and the cryptographic transition, and is explicit about where the evidence base is thinner than the discourse. Section 3 examines dependence on common providers and services in depth, as the single resilience concern that most sharply defines the limits of firm-level supervision. Sections 4 and 5 establish why this is a financial center question and why promoting standards alone is insufficient. Section 6 sets out seven practical functions with existing precedent, Section 7 offers a maturity model for evidence-supported self-assessment, and Section 8 proposes a cycle of thematic reviews through which participating centers can compare findings and test whether conditions improve over time.

A note on scope. Resilience, security, and infrastructure quality already factor into financial center benchmarking through established instruments, and this paper deliberately does not propose additional index components. Its contribution is to go deeper into one concern rather than broader across many.

A note on role. Nothing proposed here asks a financial center organization to supervise firms, set binding requirements, examine controls, or act as an operational incident-response body. Regulators, central banks, and national cyber agencies are equipped for those functions and accountable for them. The proposed role is coordination, testing, learning, and support: the work that falls between firms, service providers, authorities, and jurisdictions, and that no single actor currently owns.

2. What has changed

2.1. Artificial intelligence: the cost curve, not the attack

The institutional consensus on AI and cyber risk is unusually consistent and runs counter to much of the commercial discourse.

In June 2026, the IMF concluded that "the main financial stability concern lies less in new types of cyberattacks than in the scale effects AI can unleash." The UK's National Cyber Security Center judged in May 2025 that AI-assisted vulnerability research and exploit development would be the most significant development, while fully automated end-to-end advanced attacks remained unlikely through 2027. In July 2026, the BIS estimated that the cost of mounting an attack fell from roughly USD 5,000–10,000 to potentially USD 50–100, and observed that critical vulnerability disclosures roughly doubled in daily volume from April 2026.

Practitioners read the direction the same way. In the World Economic Forum's *Global Cybersecurity Outlook 2026*, 94 percent of respondents named AI as the most significant driver of change in cybersecurity, and 87 percent identified AI-related vulnerabilities as the fastest-growing cyber risk of 2025.

The second half of the argument is concentration. The Bank of England's July 2026 *Financial Stability Report* observes that when multiple firms rely on the same providers, software components, or essential services, a vulnerability or a defensive shutdown at a common supplier can affect several institutions at once. The Cambridge Center for Alternative Finance's 2026 global survey provides the underlying numbers: 76 percent of industry respondents use one model provider, and the top three cloud providers serve over 80 percent.

The defensible systemic claim is therefore **asymmetry plus concentration**: attack costs are falling faster than defensive costs, while the attack surface is converging. Novelty is not the claim, and a paper that rests on novelty will be contradicted by the primary sources it cites.

Supervisors have already moved from analysis to instruction. On 7 July 2026, the ECB wrote to the chief executives of significant institutions (SSM-2026-0301), describing AI-enabled attack capability as "a long-term shift in the threat landscape rather than a temporary phenomenon," and requiring action plans to Joint Supervisory Teams by 31 October 2026. The expectations span accelerated vulnerability and patch management at scale, AI-enabled detection for internet-facing assets, verification that outsourced ICT providers are prepared for compressed timelines, defense-in-depth and zero-trust architecture, and strengthened incident response, crisis management and information-sharing arrangements. In Hong Kong, the HKMA and the SFC issued parallel circulars on 2 June 2026 addressing the same capability shift, with the HKMA signaling a pilot for a cyber resilience testing framework in late 2026. In the United Kingdom, the Bank of England, the FCA and HM Treasury issued a joint statement on frontier AI and cyber resilience on 15 May 2026, notably reinforcing existing operational resilience rules rather than creating new ones.

Two observations for a center audience. First, this is a live compliance cycle, not a forecast. Second, the ECB's last expectation, information-sharing arrangements, is not something an individual bank can build on its own. It is a collective capability, requested of firms individually.

2.2. Agentic AI: risk inside the transaction

Where generative AI changed the attacker's economics, agentic AI changes the institution's own control surface.

The IMF's April 2026 note on agentic payments states the structural mismatch precisely: payment systems assume explicit, transaction-level instructions, whereas agents rely on probabilistic reasoning and adaptive decision-making, so that "individual transactions may not correspond to explicit, transaction-level instructions." Three consequences follow. Identity frameworks presume human actors, and how to identify an agent (its identity, mandate, and authority) remains unresolved. Liability between principal and agent is undefined, and supervisors will struggle to distinguish unauthorized use from user negligence when an agent misdirects funds. Because agents are fed by models from a small number of providers, correlated agent behavior becomes a new channel for synchronized payment flows and intraday liquidity stress.

The IMF's proposed architecture is a useful anchor for center-level advocacy because it is simple and testable: a probabilistic intent and orchestration layer that performs no authorization; a deterministic control and authorization layer enforcing explicit user mandates; and a settlement layer that executes without reinterpretation. In short, keep probabilistic reasoning out of authorization and settlement.

The urgency reflects a readiness gap rather than a documented loss. The Cambridge survey finds that 52 percent of industry respondents actively adopt agentic AI, compared with 28 percent of regulators: a gap the report itself describes as a potential emergent supervisory risk. The Financial Stability Board's June 2026 consultation on responsible AI adoption identifies the specific agentic vulnerabilities: unauthorized autonomous actions, errors that surface only after deployment, data exposure through agent-enabled access, collusion among multiple agents, and the impracticality of real-time human monitoring at scale.

This is a window in which a center can be useful cheaply. Guidance on agent authorization and mandate design, developed with members before agentic payment volume arrives, is worth more than remediation afterward.

An assessment of AI-related cyber risk should distinguish three points of exposure: development, live operations, and third parties. Development covers AI-generated code, dependencies, tests, and deployment decisions. Operations covers agents and models that can influence customer, treasury, or administrative actions. Third-party exposure includes dependencies on external model providers, agent frameworks, retrieval sources, hosted tools, and memory services. Treating these as one general 'AI risk' category makes ownership and testing unclear.

For development, security-critical code and dependency changes should have traceable provenance, independent review, and tests not generated solely by the same model. Deployment controls should account for immutable or difficult-to-upgrade components, with a defined rollback or migration route. Treat a model's output as untrusted input until normal secure-development checks are complete.

For live operations, an agent should have a named identity, a narrowly defined mandate, short-lived credentials, and explicit limits on value, frequency, counterparties, and permitted tools. Render transaction intent in a trusted, human-readable form before signing. Require independent approval for unusual or irreversible actions, and keep monitoring, a kill switch, and a tested recovery route outside the agent's control.

Two input-path failures deserve separate treatment, because neither is caught by controls aimed at the agent's own permissions.

The first is adversarial content arriving through a normal business process. When a model reads documents supplied by a customer (e.g., an identity document, a utility bill, a company filing), that document becomes an input path the adversary controls. Testing of AI agents has shown that instructions concealed in the material an agent reads can change what the agent does. (Debenedetti, et al., 2024)

The control is to treat customer-supplied content as untrusted input, separate instruction from data in the prompt path, and limit what an agent may do based on a document alone.

The second is fabricated output that reads as competent. NIST uses the term confabulation for model output that is confidently stated and false. (NIST, 2024)




Figure 1: Adversarial input and fabricated output in AI-assisted decisions. The figure shows how instructions embedded in customer documents and false model output can influence approvals or actions, alongside proposed controls for handling untrusted content and verifying claims against underlying records.

A model may provide a plausible but invented explanation for an unusual payment, or summarize an ownership chain while omitting a beneficial owner. The risk is not the error itself, but that it arrives in the form an approver is least likely to challenge. Human review protects only when the approver can access the underlying records and the control expects them to.

Third-party controls should record the provider, model and version used; the data, tools and permissions exposed; the route by which model or framework changes reach production; and the fallback if the service is withdrawn or behaves unexpectedly. The institution should test material updates in its own environment before promotion. The same evidence should be available when a third-party product or agent framework embeds a model.

Change management deserves a sharper statement than “test material updates,” because the characteristic failure requires no action by the firm at all.

A vendor may update a tool during a contract, and the update can change how the tool behaves. A screening tool scores how closely a name resembles an entry on a sanctions list and raises an alert only above a set level; if an update raises that level, near-matches that previously produced an alert no longer do. An update can also change how a tool handles names transliterated from another script, or the proportion of cases routed to a staff member for review. (Krisjanous, 2026)

None of this requires any action by the firm, and none of it produces an incident. A firm that assessed and tested a tool once, at purchase, may be relying on a version it has never tested.




Figure 2: How vendor updates can leave a firm relying on untested controls. The firm tests version 1.0 before approval, but later updates change how the model matches names and refers cases for review. Nothing prompts a new test, so by version 1.3 controls may fail unnoticed. Illustrative example.

Where several institutions in a financial center use the same vendor, an upstream change can weaken the same control across them simultaneously. Such changes may occur without an attack and may not trigger an alert or be recognized as an incident.

This is the concentration argument of Section 2.1 applied to defenses rather than to attacks, and it is a question an aggregate view can answer where no individual firm can: which controls across the center rest on the same vendor's settings, and who would notice if those settings changed.

A financial center organization can add value through an aggregate view. An anonymized inventory can show whether firms depend on the same foundation model, agent framework, identity service, tool gateway, or hosted provider. The center can use that map to select common failure scenarios, including an unannounced model change, a corrupted retrieval source, provider loss, or simultaneous misuse of delegated authority. Results should inform member guidance and exercises without disclosing firm-level architectures.

The control objective should extend beyond least privilege to least agency. An institution should first determine which steps require probabilistic reasoning and which can remain deterministic. An agent acting for an authorized user should not automatically inherit that user's full authority. At execution, check each material action against the original instruction, the agent's current mandate, approved transaction parameters, and applicable policy limits. Delegated authority should never exceed the delegating agent's own mandate.

Human approval provides protection only when it operates as an independent control. The approver should receive a trusted record of the proposed action, the originating instruction, material data and tools used, policy exceptions, and any departure from the original request. The approval interface should sit outside the agent's control. Institutions should also monitor approval patterns for habitual acceptance, as repeated approval without meaningful review can convert a human checkpoint into a procedural formality.

Records should permit attribution and reconstruction of every material agent action. They should identify the agent, model, and configuration version; the human or system principal; the delegated mandate; the provenance of material inputs and retrieved data; tool calls and parameters; policy decisions; approvals; and the final outcome. These records should be protected from alteration by the agent. Investigators should be able to reconstruct the action sequence from preserved records without relying on the model's retrospective explanation. Re-running a probabilistic model may produce a different result and cannot establish what happened in the original execution.

A financial center organization could turn these principles into a common agent transaction control profile for member firms. The profile could define common control expectations for agent identity, delegated authority, action-level policy enforcement, approval evidence, logging, and emergency revocation. Joint exercises could test indirect prompt injection, poisoned tool descriptions, cross-agent privilege escalation, loss of a shared model provider, and revocation of agent credentials during an incident. Anonymized results would reveal whether the same weaknesses or dependencies exist across several institutions.

2.3. Digital assets and tokenization: from code to keys

Native cryptoassets, stablecoins, tokenized deposits and tokenized securities differ in their issuer, settlement asset, governance and recovery options. For tokenization arrangements, the FSB's 2024 analysis identifies the reference asset, participants, technology and links to legacy systems as sources of risk (FSB, 2024). A control failure may irreversibly move value. Financial centers with active digital-asset licensing or tokenization programs should examine privileged access, transaction controls, service-provider dependencies, and recovery. Elsewhere, shared infrastructure may require greater attention.

Controls must reflect the asset and operating model. Reserve-backed stablecoins require controls over reserves and redemption. Tokenized bank deposits remain liabilities of the issuing bank; their controls should address deposit records, transfer arrangements and settlement (Garratt & Shin, 2023). Tokenized securities and real-world assets require the token, underlying legal claim, and authoritative records to remain aligned. Cross-platform arrangements introduce bridges and messaging services, while native cryptoassets may have no issuer or administrator able to reverse a transfer.

In TRM Labs' H1 2026 dataset, smart-contract exploits accounted for 125 of 207 incidents (TRM Labs, 2026). Infrastructure and operational compromises accounted for about 15 percent of incidents but about 76 percent of reported losses. Code flaws were more frequent, but compromises affecting credentials, signing systems, and operational infrastructure caused most reported losses.

Bybit and KelpDAO show how transaction controls can fail even when asset-signing keys are not compromised. Sygnia's March 2025 analysis reports that the February 2025 Bybit attackers compromised a Safe developer workstation and cloud environment, injected malicious JavaScript into the frontend, and caused signers to approve data that differed from the transaction displayed (Sygnia, 2025). The approval process did not detect the difference between the displayed transaction and the transaction being authorized.

In the April 2026 KelpDAO incident, attackers used a developer's session keys to compromise two internal RPC nodes and disrupted external RPC services. The legitimate verification and signing service then relied on poisoned blockchain data and produced a valid attestation for a forged message. Because the configuration required only one verifier, that attestation was sufficient to release assets. LayerZero's final report found no evidence of compromise beyond the RPC environment. Key protection alone is insufficient when the interface, policy data, RPC layer, or verification path has been compromised (LayerZero Labs, 2026).

The following three controls are needed:

First, high-risk transactions need an independent check of their intended operation, destination, and effect. EIP-712, created in 2017 and now final, defines how typed structured data is hashed and signed, allowing wallets to display the data for review (Bloemen, Logvinov, & Evans, 2017). Draft ERC-7730, created in 2024, proposes clearer display of smart-contract calls and typed messages (Castillo, et al., 2024). Neither standard alone protects against a compromised transaction display; verification requires trusted data and an independent control path.

Second, automated transactions reduce the time available to intervene. Tobias Adrian's April 2026 IMF Note proposes formal verification and independent audits for systemically important contracts, alongside mechanisms to pause or adjust execution under predefined emergency conditions (Adrian, 2026). These

are the author's policy recommendations, not universal legal requirements. Financial centers should consider whether firms can pause, contain, and reconcile an erroneous or malicious process before it triggers cascading transactions, liquidity pressure, or conflicting records.

Third, intervention and recovery functions create powerful administrative privileges. Stablecoin systems may include mint, burn, pause, blacklist, or freeze functions; other arrangements may rely on upgrade keys, oracle-change privileges, or governance controls. Using these functions requires proportionate authorization, segregation of duties, independent verification, monitoring, and tested procedures for compromise, revocation, and replacement. Settling the legal basis and customer consequences should happen before an incident.

Requirements differ by asset and activity. Hong Kong's August 2025 guideline for licensed stablecoin issuers covers key lifecycles, high-risk operations, smart-contract assurance, detection of key compromise, and key revocation and replacement. It also identifies transaction limits and time locks as additional measures to apply where applicable ([HKMA, 2025](#)). ESMA's July 2026 Common Supervisory Action examines governance, key and storage management, transaction controls, incident response, smart-contract risk, and third-party dependency in custody services ([ESMA, 2026](#)). The 2016 CPMI-IOSCO cyber guidance and its 2022 implementation assessment provide a baseline for financial market infrastructures, but not a complete control specification for current tokenized systems ([CPMI-IOSCO, 2016](#)).

A financial center review should extend beyond custody keys. Firms should identify the asset and settlement model, map the full transaction and recovery path, record critical shared providers, and show how they would reconcile on-chain, off-chain, and legal records after a disruption. A common baseline can support licensing and supervision, but should not replace either.

Zanarc's 2026 advisory on real-world asset tokenization outlines the four-record, five-stage models used here ([Zanarc, 2026](#)).

For real-world assets, Zanarc's four-record model checks consistency between the legal claim, off-chain asset record, on-chain token record, and settlement record. These records can diverge even when the smart contract works as intended, so normal controls and recovery procedures should include four-way reconciliation.

Test minting, burning, freezing, changing an oracle, or moving reserves as a five-stage chain: request, verify and authorize, execute, settle, then reconcile or recover. No person, system, or provider should control the entire chain without an independent check. Evidence should identify who performed each stage and the data used.

The signing path may include the user's device, front end, API, RPC service, policy engine, wallet software, hardware or multiparty-computation service, bridge verifier, and approver display. Testing should identify shared providers and single points of trust across that path. Threshold signing offers little protection when every signer receives the same false representation of the transaction.

Recovery plans should identify who can pause, freeze, burn, reissue, or upgrade; the legal authority for each action; how to reconcile conflicting records; and how to inform customers and counterparties. These decisions require advance agreement and testing with affected providers.

2.4. The cryptographic floor is moving

Post-quantum migration is the one item on this list with published deadlines. The G7 Cyber Expert Group's roadmap for the financial sector sets 2030–32 for migrating critical systems and functions, with 2035 as the overall completion target. The EU's coordinated implementation roadmap requires cryptographic inventories, dependency maps, and supply chain assessments by the end of 2026, with high-risk use cases made quantum-safe by the end of 2030. The UK's NCSC expects migration plans by 2028 and completion by 2035. NIST published its core standards in August 2024.

Two features make this a center-level issue rather than purely a firm-level one. It is a "harvest now, decrypt later" problem, so exposure accrues today from data captured today. And it is an inventory problem before it is a cryptography problem: an institution cannot migrate what it cannot enumerate, and that enumeration extends to suppliers.

The perception gap is wide. Only 37 percent of World Economic Forum respondents expect quantum technologies to affect cybersecurity within 12 months (WEF; Accenture, 2026), while the EU inventory deadline is three months away as of this writing. Closing that gap through member guidance, a coordinated inventory drive, and supplier questionnaires is inexpensive and unusually well-defined work for a center to own.

2.5. What the evidence does not yet show

A working paper is more useful if it is candid about the limits of its own evidence base.

The canonical deepfake fraud case, the January 2024 loss of HKD 200 million by Arup in Hong Kong ¹, involved an engineering firm, not a financial institution, and remains the only large, victim-confirmed case of its kind on the public record more than two years later. FinCEN reports a rise in suspicious activity reporting involving deepfake media but provides no quantification (FinCEN, 2024). The FBI's 2025 Internet Crime Report, in its first year in which AI appears as a discrete category, records approximately USD 893 million in AI-related fraud losses out of USD 17.7 billion in total reported fraud: a significant, but smaller proportion than the discourse implies (FBI, 2026). OWASP's review of agentic AI exploitation in the first quarter of 2026 documented real-world incidents across several sectors and none in financial services (Clinton, 2026).

None of this argues for complacency; supervisory action of the kind described above does not wait for loss data, and should not. It argues for a particular register. A paper that argues from trajectory, structural analysis and supervisory response, and that says plainly where the loss evidence runs out, will hold up better than one built on vendor surge statistics, an ecosystem in which figures circulate widely with untraceable provenance.

¹ <https://www.theguardian.com/technology/article/2024/may/17/uk-engineering-arup-deepfake-scam-hong-kong-ai-video>

3. The concern this paper digs into: third-party dependency

3.1. Where incidents actually originate

Consider the incidents over the last three years that most severely disrupted market functioning or public confidence, and ask what was breached.

ION Cleared Derivatives, January 2023: a post-trade processing provider. Clearing members reverted to manual trade confirmation, margining, and reporting at a moment of high volatility, with knock-on effects on EMIR and CFTC reporting.

ICBC Financial Services, November 2023: a broker-dealer subsidiary, not the parent bank. Its settlement agent was temporarily owed approximately USD 9 billion against the subsidiary's net capital of roughly USD 481 million; settlement details were physically transported so that trades could be cleared manually.

CrowdStrike, July 2024: not an attack at all, but a faulty content update to a security product, which, by Microsoft's estimate, crashed 8.5 million Windows devices and disrupted financial institutions across more than a dozen jurisdictions. Parametrix estimated direct losses to Fortune 500 companies at approximately USD 5.4 billion, only a small fraction of which was insured.

Bybit, February 2025: a wallet infrastructure provider's developer machine.

Toppan Next Tech, April 2025: a printing vendor whose ransomware incident exposed customer data of DBS and Bank of China's Singapore branch.

C&M Software, June–July 2025: a technology provider connecting smaller institutions to Brazil's instant payment system, breached via purchased employee credentials. Reported loss estimates vary widely across sources and should be treated with caution, but at least six institutions were affected, and dependent institutions could not transact for several days. Two further Brazilian incidents followed the same year through providers and providers' providers.

SolarWinds, disclosed December 2020, provides an earlier example: a network management software provider. Attackers compromised its Orion software build process and inserted the SUNBURST backdoor into updates released between March and June 2020. Customers installed the malicious code through the provider's normal software update process. Thousands of organizations received affected software, although a much smaller group suffered further intrusion, including US government agencies. Microsoft confirmed malicious SolarWinds code in its environment and unauthorized access to some source code, but found no evidence of access to customer data or production services.

The pattern is not incidental. It is the natural consequence of a sector that has outsourced most of its technology surface while retaining all trust obligations. The World Economic Forum finds 65 percent of large organizations now identify third-party and supply chain vulnerabilities as their single greatest cyber challenge, up from 54 percent a year earlier, with inability to assure the integrity of third-party software ranked first among supply chain risks. (WEF; Accenture, 2026)

Two data points sharpen it further. The IMF's April 2024 *Global Financial Stability Report* found that over half of the IT providers serving global systemically important banks serve two or more major banks (IMF, 2024). And the IMF's November 2025 FSAP technical note on Swiss cyber risk supervision records that

third-party-related incidents accounted for 50 percent of reported incidents in 2023 and 30 percent in 2024 (IMF, 2025b).

3.2. Why the regulatory perimeter cannot close the gap

Supervisors have recognized the problem and built real machinery for it. That machinery has structural limits rather than accidental ones.

Under DORA, the European Supervisory Authorities designated a first cohort of around twenty critical ICT third-party providers in November 2025. Cloud providers, data and infrastructure firms, and major systems integrators: each assigned a Lead Overseer with powers to assess, investigate, inspect, issue recommendations, publicize non-compliance and, as a last resort, require financial entities to suspend or terminate use of the service. In the United Kingdom, HM Treasury designated four critical third parties in July 2026: Amazon Web Services, Google Cloud, Microsoft and Oracle, with regulatory oversight beginning immediately.

Three limits follow.

The regimes reach the top of the distribution, not the tail. Four designations in the UK and around twenty in the EU capture the providers whose failure would be systemic. None of the entities in Section 3.1 would plausibly appear on such a list. Incidents that damage a center's reputation are concentrated in the very population the designation regimes do not reach.

Enforcement is incomplete. The IMF's July 2025 euro area FSAP found that the ESAs lack the authority to fine critical providers for non-compliance, leaving oversight dependent on voluntary cooperation. It also found that incident reporting infrastructure was decentralized across national competent authorities, that no on-site inspections were conducted for TARGET services despite their systemic importance, and that only three staff were dedicated to financial market infrastructure cyber oversight (IMF, 2025a).

Supervisory capacity is thin relative to the surface. The Swiss FSAP is the clearest available picture, and it concerns a highly developed center. In 2023, external audit firms performed 211 regulatory audits on cyber-related matters, compared with 17 FINMA on-site examinations. The supervisory unit, covering cyber, ICT, data, and outsourcing risk across the sector, had 9 staff, excluding leadership. Cyber and ICT rules applied comprehensively to banks but not explicitly to insurers and asset managers. No published testing framework comparable to CBEST or TIBER-EU existed, and the incident reporting template did not capture loss data. The note also identifies concentration in SIX Group as a systemic exposure warranting attention.

This should be read as a diagnosis of a general condition rather than of one jurisdiction. Supervisory resources scale with mandates and budgets; the attack surface scales with digitalization and outsourcing. The two are diverging, and no plausible increase in supervisory headcount closes the gap for the long tail of suppliers.

3.3. Concentration is now reappearing in the AI layer

The third-party problem is not static. It is reproducing itself one layer up.

The designation regimes were designed around cloud and infrastructure dependency. Model provision is now developing the same characteristics: a small number of providers, deep integration into core

processes, opacity about model changes, and correlated behavior across institutions that use the same underlying systems. The Bank of England has explicitly flagged reliance on a handful of frontier AI providers as a potential new channel of concentration risk. The FSB's June 2026 consultation identifies the same concern, adding that unannounced model changes and homogeneous behavior drive correlated outcomes. The Cambridge survey shows 43 percent of regulators are concerned about critical third-party dependencies, compared with 23 percent of vendors.

A center that has learned the cloud concentration lesson has a short window to apply it before AI dependency is as entrenched.

3.4. What firms are not doing, and what supervision will not ask them to do

The most striking evidence here is about practice, not policy. Across organizations globally, the World Economic Forum finds that 66 percent evaluate supplier security maturity and 65 percent involve security in procurement. Still, only 33 percent comprehensively map their supply chain ecosystem, and only 27 percent run cyber incident simulations with recovery exercises involving those partners. (WEF; Accenture, 2026)

That last figure is the heart of the matter. Contractual assurances and questionnaires are widespread; joint rehearsals are rare. Yet every incident in Section 3.1 was, operationally, a coordination failure between a provider and its dependent institutions: how quickly the provider communicated, how dependent firms resorted to manual processes, and how counterparties decided whether it was safe to reconnect.

Two forms of false comfort compound the problem, and both qualify the dependency map proposed in Section 0.

The first is sequential. When a later control reuses the profile an earlier one created, an onboarding mistake can reappear in monitoring as though it were independent confirmation. A check that repeats a prior conclusion, rather than testing it against source records, adds no assurance; it launders the original error.

The second is structural. Using two providers does not establish independence if both draw on the same underlying data. (Krisjanous, 2026). A firm may have diversified its suppliers but not its sources, and a dependency map built on supplier names will record diversity that doesn't exist. Useful mapping must reach past the contracting entity to the data, models, and infrastructure behind it, which is why the map proposed in Section 6.1 starts with critical services rather than supplier names.

No supervisor will convene that rehearsal, because it spans entities that answer to different authorities or to none. No individual firm will convene it, because the benefit accrues to the market. It is a textbook collective action gap, and it is the single most valuable thing a financial center can build.

4. Why this is a financial center concern

4.1. Evidence on the effects of cyber incidents on confidence

The European Systemic Risk Board's framework distinguishes three amplification channels through which a cyber incident can become systemic: operational contagion, financial contagion, and confidence erosion. On the third, it observes that "uncertainty and the loss of confidence are critical catalysts in triggering financial instability" and notes that coherent, timely communication is correspondingly important ([ESRB, 2022](#)).

The empirical support is partial and should be presented as such. The IMF's April 2024 GFSR chapter documents modest but real deposit outflows (retail and wholesale deposits at smaller US banks falling by around 5 percent cumulatively over six quarters following an incident) and finds equity price declines of 0.1 to 0.2 percentage points due to malicious attacks, larger for small firms. No significant "cyber run" has been observed yet. The transmission channel is well argued; it is not an established empirical regularity, and claiming otherwise weakens the paper ([IMF, 2024](#)).

The digital-asset sector, less protected by deposit insurance and central bank backstops, offers the closest thing to an observed run. After the February 2025 theft, Bybit experienced net outflows of more than USD 4 billion (roughly a third of the platform's assets) within days. It survived through pre-positioned liquidity and continuous public disclosure. The episode illustrates the mechanism without proving that it operates the same way in insured banking.

In his paper "Measuring the Performance of International Financial Centers," the second one in this series, Biedermann identifies confidence as crucial to an international financial center's long-term performance, with reputation, trust, and resilience as major drivers. He concludes that "reputation and trust are cumulative, slow-building, and easily damaged assets" ([Biedermann, 2026](#)). A serious cyber incident with system-wide repercussions can quickly and profoundly undermine trust.

4.2. The operational channel: better measured, and more alarming

Where confidence effects are theorized, operational effects have been quantified.

Eisenbach, Kovner, and Lee at the Federal Reserve Bank of New York find that impairment of any of the five largest US banks affects, on average, 38 percent of the Fedwire network ([Eisenbach, Kovner, & Lee, 2021](#)). Where institutions respond to uncertainty by hoarding liquidity, forgone payment activity exceeds 2.5 times daily GDP. And—the finding most relevant here—banks with under USD 10 billion in assets can trigger significant system-wide impairment through third-party linkages.

Kotidis and Schreft's study of a real multi-day attack on a technology service provider used by more than 50 banks found that affected banks sent 16.4 percent fewer payments by count and 26.5 percent less by value, with day-one declines of 36 percent and 50.7 percent, respectively. They also document an innocent bystander effect: banks merely exposed to disrupted senders saw their incoming payments fall. Their counterfactual, had a top-five bank been affected, implies an effect roughly seventeen times larger ([Kotidis & Schreft, 2025](#)).

The conclusion for a center is direct: **system-level damage does not require a systemically important firm.** It requires a shared dependency. The population that creates shared dependencies is precisely the population outside the supervisory perimeter.

4.3. The jurisdictional dimension

Two lines of evidence connect this to a financial center's standing.

The first is by analogy. FATF grey-listing is a reputational signal rather than an operational failure, and IMF research links it to reductions in total capital inflows of about 7.6 percent of GDP, driven by bank de-risking and investors using the listing as a risk heuristic ([Kida & Paetzold, 2021](#)). That is the order of magnitude of a jurisdiction-level trust shock when it is measurable.

The second is direct. The IMF's Swiss FSAP technical note places Switzerland in the "Tier 2 — Advancing" category of an international cyber security index (a measure of national commitment, not a ranking of financial centers) against the United Kingdom, Germany and Singapore in Tier 1, and frames the resulting regulatory gaps as bearing on the attractiveness of the Swiss financial center ([IMF, 2025b](#)). Whatever one makes of the underlying index, the significant fact is that an international financial institution now discusses cyber supervisory maturity and center attractiveness in the same paragraph.

This is why cyber resilience is a financial center concern, not only a supervisory one. The consequences of an incident spread across a jurisdiction's institutions and, increasingly, its standing.

5. Is promoting standards enough?

A reasonable position would be that financial centers should encourage adoption of established frameworks (e.g., ISO/IEC 27001, the NIST Cybersecurity Framework, DORA compliance in the European Union) and leave the rest to supervisors. Three lines of evidence suggest this is a floor, not a strategy.

The centers themselves have identified the gap. At the EU Roundtable of Financial Centers meeting in Frankfurt on 1 October 2025, members treated cyber resilience as a systemic risk, noted emerging threats from AI-enabled attacks and quantum computing, and called for harmonized implementation of DORA, along with common metrics for cyber preparedness ([EU-RFC, 2025](#)). This is a precise diagnosis by the constituency best placed to make it. The problem is not a lack of rules, but inconsistent implementation and a lack of comparable measurement. Neither is remedied by adding standards.

Certification coexists with basic control failure. The 2025 CBEST thematic report, drawn from threat-led penetration testing of systemically important UK firms and infrastructures, found persistent patching gaps in operating systems and application configuration, credential management failures including weak password enforcement, missing multi-factor authentication and inadequate network segmentation, detection deficiencies in monitoring and alerting, and inadequate oversight of remediation by risk management and internal audit ([Bank of England, PRA, FCA, 2025](#)). These findings were substantially the same as the preceding year for firms holding current certifications.

Flagship requirements are not always met. The CPMI-IOSCO Level 3 assessment of financial market infrastructures' cyber resilience found that some FMIs had no adequate response and recovery plan capable of meeting the two-hour recovery time objective that is the flagship expectation of the 2016 guidance, and identified further shortcomings in scenario planning, retesting after major system changes, and stakeholder involvement in response testing ([CPMI-IOSCO, 2022](#)).

The obligation often already exists. FATF Recommendation 15 requires financial institutions to identify and assess the money laundering and terrorist financing risks of new products and business practices, and of the use of new or developing technologies for both new and existing products — and to do so before launch rather than after. ([FATF, 2026a](#))

Implemented through national law, this is already an enforceable requirement to assess an AI tool before relying on it. Recommendation 11 requires retaining customer due diligence records, including the results of any analysis undertaken; in an automated process, that means the record must show what the tool received, what it returned, and who reviewed it.

The missing piece is implementation, and the FATF says so. Its July 2026 targeted update found significant gaps in translating risk assessments into effective mitigation and in making licensing and registration work in practice. ([FATF, 2026b](#))

That is the CBEST finding restated by a different standard-setter in a different domain: the requirement is in force, the assessment is on file, and the control is not demonstrably working. It supports this section's conclusion rather than qualifying it. The gap financial centers can usefully address is the distance between a documented assessment and evidence that the control performs.


Practitioner perception fits the same pattern. The World Economic Forum finds 74 percent holding positive views of cyber regulation's effectiveness, with 58 percent citing increased board awareness. Still, only

around 30 percent in Europe report actual security improvements, a gap the report partly attributes to regulatory complexity.

This doesn't mean standards are unimportant. Certification is a floor and a procurement signal, and both matter, particularly for the unregulated edge, where a common baseline is often absent altogether. The implication is that a center adds least value where the standards ecosystem is already crowded, and most value where it is empty: implementation consistency, comparable measurement, tested response, shared intelligence, trained people, and a rehearsed voice.

6. The practical role: seven functions for a financial center organization

Each function sits above the individual firm and below the regulator. None requires statutory authority, and none asks a center to become either of the two things it should not be: a supervisor, or an operational cyber-response body.



Centers do not set requirements, examine firms, run security operations, or lead an incident. They can close the gaps that open up between firms, service providers, authorities, and jurisdictions: gaps in

coordination, testing, and learning. Each function has working precedent. The examples name centers and bodies that have acted; the absence of a center from an example should not be read as a judgment about it, since much center-level work is deliberately unpublicized.

6.1. Map dependence on common providers and services

Section 3 established that incidents that damage a center regularly arise from shared dependencies, and that the designation regimes apply only to the largest providers. The first practical function follows directly: someone has to know what the center actually depends on.

No individual firm can see this. A bank knows its own suppliers and, with effort, its material fourth parties. It cannot see that eleven other institutions in the same center rely on the same reconciliation platform, the same statement printer, or the same connectivity provider into the national payment system. Supervisors see more through information registers and outsourcing notifications. Still, that view is organized by supervised entity, constrained by confidentiality, and, in most jurisdictions, is not shared back with the market.

A center-level dependency map is a different instrument, and a modest one. It does not require firms to disclose contracts or security posture. It requires only that participating firms indicate, against a common list of service categories, which providers support functions they could not operate without for a defined period. Aggregated and anonymized, the result answers questions no single participant can currently answer: which providers appear across a material share of the center; where one provider sits behind several apparently independent institutions; which of those dependencies fall outside every designation regime; and where the center's concentration differs from that of its peers.

The output is neither a supervisory judgment nor a scorecard. It is a shared factual base for the other functions in this section: it tells the center which providers belong inside the information-sharing arrangement, which should be in the room for the next exercise, and which subjects deserve a thematic review.

Two cautions. The map is sensitive, and its value depends on aggregation: participating firms should see the aggregate picture and their own position within it, not each other's returns. And it decays. A dependency map is worth maintaining on a fixed annual cycle or not at all.

A center-level map should start with critical services, rather than supplier names alone. Each participating firm would record the service, maximum tolerable outage, direct provider, known material third-party providers, hosting or processing location, data dependency, fallback, estimated substitution time, last exercise, and incident-notification route. A common taxonomy allows aggregation, while each firm's register remains the system of record.

The map should complement, not replace, DORA registers of information or national outsourcing returns. Where authorities participate, the center should agree in advance which aggregated questions they can answer and what information they cannot share back. The exercise is most useful when it identifies concentrations across regulated firms and adjacent participants, dependencies beyond directly designated providers, or services whose contractual alternatives would take too long to operationalize.

An independent function or trusted third party should manage collection under a written protocol. Participation should be voluntary unless an authority supplies a different legal basis. Report provider names and concentration data only above a minimum aggregation threshold, with competition-law review

and a clear separation between resilience analysis and the center's promotional work. Participants should receive their own position against the aggregate, not other firms' returns.

The map can identify candidates for information-sharing, exercises, contingency work and thematic review; it cannot establish that a provider is safe or compliant. The proposed refresh cycle is annual and after material provider changes, subject to participants' agreement. Include hidden service chains, including services operating through accounts or infrastructure provided by another firm, where they support a critical function.

Mapping should lead to at least one tested action for each material concentration. A joint exercise can test loss or compromise of the shared service, the time required to invoke fallback arrangements, access to the data and configuration needed for substitution, and the provider's ability to support several affected firms at once. Assign, date, and retest actions; do not present an unexercised contractual exit as an operational fallback.

A map that identifies a shared provider raises an obvious question: what can members do about it? Individually, a small licensee has no leverage over a provider serving thousands of customers. Collectively, firms in a center that depend on the same provider have more leverage.

The most transferable use of a dependency map is therefore to support convergence on a small number of common contractual asks. Three recur: access to the records of decisions the service makes on the firm's behalf; notice of material changes to the service before they reach production, in enough detail to decide whether re-testing is needed; and a defined position on supervisory access where national law provides for it. (Krisjanous, 2026)

Where a provider will not grant access, the firm should be able to explain how it obtains sufficient evidence by other means.

None of this requires the center to negotiate on members' behalf, or to endorse or exclude any provider. Publishing an agreed form of words that members may use in their own negotiations is enough, and it turns an observation about concentration into something a member can act on.

6.2. Connect existing information-sharing arrangements

Regulators bind licensees. Financial centers can reach the entities on which licensees depend, but they should not start by building something new.

Most centers already sit within a dense set of arrangements: a sectoral Information Sharing and Analysis Center (ISAC), a national Computer Emergency Response Team (CERT) or cyber agency, a bankers' association committee, a supervisory information-sharing forum, and the firms' own bilateral relationships with their providers. The problem is rarely that none of these exist. It is that they do not connect. A national agency holds indicators the sector never sees; a sectoral ISAC's membership stops at regulated firms; a supervisor's incident data is not fed back in any usable form; and none of them has a view across jurisdictions. An ISAC-style arrangement may be part of what a center needs, but it should not be the default proposal. The higher-value work is connective: identifying which arrangements exist, where the gaps lie, who is excluded, and how the center can bridge them.

The precedents below are best read as connective moves rather than as new institutions.

FS-ISAC's Critical Providers Program, launched in 2022, connects member financial firms directly with key infrastructure and service providers through secure channels: an explicit attempt to bring the supplier population inside the sharing arrangement rather than merely contracting with it.

In the United Kingdom, the Financial Sector Cyber Collaboration Center operates as a public-private incident-coordination service, with membership open to all UK-regulated financial organizations. During a crisis, it convenes a fusion cell drawn from member firms alongside government partners and provides situation reports to authorities without attribution. Free membership and non-attribution are not administrative details; they remove the two standard obstacles to sharing: cost and exposure. FS-ISAC UK became the steward of the center in November 2025, at a time when the government's own general-purpose sharing platform was being retired, a notable shift of load from state infrastructure to sector infrastructure.

Luxembourg's CIRCL has addressed the exposure problem through technical measures. Its MISP platform supports financial-sector attributes and, importantly, delegation of publication, allowing a firm to share indicators in a pseudo-anonymous manner through a trusted partner (CIRCL, 2016).

In Hong Kong, the Hong Kong Association of Banks and the HKMA operate the Cyber Intelligence Sharing Platform, and in December 2024 extended it to enable cross-sector sharing across banking, insurance, and capital markets (HKMA, 2024).

Singapore has taken the baseline route. Following the Toppan Next Tech incident, MAS and the Cyber Security Agency jointly proposed in April 2025 that vendors serving financial institutions obtain the Cyber Essentials or Cyber Trust marks, with certification potentially becoming a precondition in some contexts.² This converts a voluntary standard into an ecosystem requirement without extending the supervisory perimeter: a model other centers can copy.

The scale of the opportunity is evident in the fact that only 34 percent of organizations globally provide any formal, structured support to smaller or less-resourced partners in their ecosystems.

None of these bodies conducts incident response itself. They route information, convene the people who do so, and give authorities a consolidated view; the response remains with the firms, their providers, and the national agencies.

For the World Alliance of International Financial Centers, as the association of the leading financial centers, the same logic scales upward. The connective gap between centers is wider than the gap inside any one of them: an indicator, a control failure, or a supplier weakness identified in one center rarely reaches the others. Section 8 proposes a mechanism for that.

Gaps open not only between firms, providers, and jurisdictions. They also open between authorities within a single jurisdiction.

A tool that operates inside a regulated firm rarely falls under a single supervisory remit. An agent that moves client funds or issues tokens creates financial-crime risk alongside operational, conduct, and prudential risk, and different teams typically supervise these—or, in smaller authorities, one team wears several hats. (Krisjanous, 2026)

Each authority sees the part of the tool that falls to it, and no one necessarily sees the whole.

² <https://www.mas.gov.sg/news/letters-to-editor/2025/response-to-st-forum-letter-on-14-april-2025>

A single inventory of a firm's AI uses, and of what each tool is permitted to do, is the cheapest way to show where joint review is needed. It is the instrument proposed in Section 0 applied to a second purpose: the dependency map tells a center what its firms collectively rely on, while an inventory of permitted actions tells the relevant authorities where their interests overlap. A center cannot assign supervisory responsibility and should not try. It can convene the conversation in which authorities establish who is looking at what, which is exactly the kind of gap this paper is about.

6.3. Coordinate cross-firm exercises

Supervisory testing assesses institutions one at a time. Threat-led testing regimes such as CBEST, TIBER-EU and the ABS Adversarial Attack Simulation Exercise guidelines are valuable and increasingly mandatory. Still, they answer the question "can this firm withstand an attack," not "can this market keep functioning."

SIFMA's Quantum Dawn VIII, held from 4–6 November 2025, is the most instructive example currently available. It engaged roughly 1,000 participants from more than 100 financial institutions, critical market utilities, public agencies, central banks, and law enforcement, representing more than 20 countries. The scenario was deliberately a polycrisis: a Category 5 hurricane, a severed transatlantic communications cable, a market infrastructure outage, and a state-actor zero-day. The after-action report, published in March 2026, named five priorities: 1. third-party infrastructure risk, 2. access credentialing for essential personnel, 3. integrating disconnection and reconnection protocols, 4. geographic dispersion, and 5. planning beyond prior crisis scenarios (Sifma, protiviti, 2026).

The program originated with a government agency and was subsequently handed to an industry body, which has run it biennially since 2011.

In the United Kingdom, the Cross Market Operational Resilience Group delivers market-wide SIMEX exercises, and UK Finance and the Bank of England jointly resource its program office—a genuinely shared ownership model. SIMEX 24 tested the sector's ability to shut down and restart.³

In Japan, the FSA's annual Delta Wall exercise spans banks, securities firms, insurers, credit unions and payment providers, with a public summary published after each round (FSA, 2023). FS-ISAC's CAPS exercise⁴ offers a free, confidential payment-systems tabletop that deliberately includes communications, legal and executive functions alongside technology teams.

Three design choices distinguish the strongest programs: third parties are in the room; the scenario is compound rather than single-cause; and findings are published. The third is what converts an exercise from a private benefit into a public good.

Given that only 27 percent of organizations run recovery exercises with their partners, a center that convenes one credible market exercise a year is supplying something the market demonstrably will not supply for itself.

The center's role here is convening and design, not command. It agrees on the scenario with members, brings third parties into the room, and ensures the lessons are documented and circulated. Participants run the response.

³ <https://www.bankofengland.co.uk/news/2024/october/simex-24-testing-the-uk-financial-sector-resilience>

⁴ <https://www.fsisac.com/what-we-do/resilience>

6.4. Help smaller firms improve

Small organizations are roughly twice as likely as large ones to fall short of the minimum level of resilience. FinTechs and small licensees typically fail on the same short list of controls: administrative account security, patching, multi-factor authentication, backup and recovery. The cost of fixing that list is small relative to the cost of an incident that brings the center's name into disrepute.

Singapore's Cyber Essentials and Cyber Trust marks, extended in April 2025 to cover cloud, AI, and operational technology security, provide a graduated baseline delivered through accredited third parties and supported by co-funding of up to 70 percent of security consultancy costs for eligible smaller enterprises. The United Kingdom's Cyber Essentials performs a similar function with procurement leverage attached. The UK's sector collaboration center's free membership model is arguably the strongest small-firm measure of all, since it removes cost as a barrier to intelligence rather than adding a requirement.

The financial center's contribution here is rarely to invent a scheme. It is to make an existing national scheme financially and practically reachable for the smallest regulated and adjacent firms, and to give it procurement meaning within the sector.

A second small-firm problem, distinct from the baseline controls above and less often addressed, remains.

A small firm that buys an off-the-shelf automated service (identity verification, document checking, screening, or transaction monitoring) inherits a control it cannot meaningfully evaluate. It can read the vendor's own performance figures and run a handful of known genuine and fraudulent cases through the service as configured, which may expose obvious weaknesses. From a sample of that size, it cannot establish reliable error rates or show how the service performs across different customer groups. NIST offers the same caution about generative AI, warning against extrapolating performance from narrow, non-systematic, anecdotal assessments. ([NIST, 2024](#)) ([Krisjanous, 2026](#))

This is an unusually clean case for collective action. The cost of assembling a credible test set (current forgery techniques, synthetic identities, names in several scripts, poor-quality documents) is substantial for one small firm and modest when shared. A center that commissions such a set once and makes it available for members to run against their own configurations gives every small member something none could produce alone, without assessing any firm or endorsing any vendor. When several members use the same service, the aggregate result also answers a question the dependency map raises but cannot settle: whether a shared provider performs as assumed.

A center operating at level 3 or 4 on the small-firm dimension in Section 7 might reasonably be expected to offer this alongside a baseline certification scheme. ([Krisjanous, 2026](#))

6.5. Support communication during major incidents

This is the most transferable of the seven functions, and the one least often institutionalized. It is also the function most easily misread: the point is not that the center speaks for the firms or the authorities, but that what each says is coherent with the rest.

The clearest model is the Incident Communications Group within the UK's Sector Response Framework, housed at UK Finance. Its stated purpose is to provide a single unified communication voice for the finance sector during an incident. Its membership is heads of media, communications, or corporate affairs at member firms. It designs coordinated external communication plans, aligns messaging with the

authorities' response framework, and prepares FAQs and materials in advance. It sits alongside a Business and Operational Continuity Committee at UK Finance and a Cross-Market Business Continuity Group hosted by the Bank of England. It has been exercised, including in a communications-focused scenario.

Internationally, the G7 Cyber Expert Group's fundamental elements of collective cyber incident response and recovery require communications experts to be directly involved in the response arrangement itself, and for crisis communication strategies to define in advance the extent to which messaging is harmonized across members. The FSB's incident response and recovery toolkit requires a cross-functional communication team that delivers consistent messaging and maintains trusted, redundant communication channels, with contact lists regularly validated. This was demonstrated when ICBC Financial Services staff reverted to public email during their incident.

The empirical lessons from recent incidents are consistent and worth stating plainly because they run counter to instinct.

The dominant failure mode is over-reassurance, not over-disclosure. ICBC Financial Services publicly stated that the affected trades had been successfully cleared, even as its settlement agent was owed billions and settlement details were being processed manually. UK NCSC guidance makes the same point in general terms: statements that require later retraction cost more credibility than honestly stated uncertainty.

Speed on the nature of the event beats completeness. CrowdStrike's confirmation within roughly six hours that a faulty update, not an attack, caused the incident is the best-documented example of an early, narrow, high-confidence statement halting a speculative spiral.

A single voice must be pre-constituted. The Incident Communications Group exists and is staffed before incidents. No one can assemble these arrangements during a live event.

The transparency question is not binary. The G7 resolves it through tiered disclosure—Traffic Light Protocol, Chatham House Rule, defined audiences, and a confidentiality classification—rather than choosing between silence and full publication. This is the practical answer to the objection that disclosure aids attackers.

Sustained transparency can hold a franchise together. Bybit's chief executive went public within hours, appeared in person repeatedly, and maintained a public timeline of the incident during a withdrawal run that took roughly a third of the platform's assets. By contrast, ION Trading issued a single short statement for an incident with market-wide reporting consequences.

A center does not need statutory authority to build this, and it displaces no one. The affected firm speaks for itself, the authorities speak for themselves, and the center adds coherence between them: a standing group, an agreed protocol, an annual rehearsal, and a working relationship with the authorities' communications function established before it is needed.

6.6. Build talent as center infrastructure

ISC2's 2025 workforce study found 95 percent of teams reporting at least one skills gap, with 59 percent describing gaps as critical or significant (up from 44 percent a year earlier) and 41 percent naming AI skills as the most pressing need (ISC2, 2025).

The distribution is unequal in a way that matters for ecosystems: the World Economic Forum finds 46 percent of small organizations cite a lack of cybersecurity expertise, compared with 29 percent of large

ones, and 85 percent of organizations with insufficient resilience report skills gaps, compared with 22 percent of highly resilient ones. (WEF; Accenture, 2026)

Toronto Finance International (TFI) offered ⁵ the clearest example of a center body leading rather than signposting. It convened a working group of major financial institutions (including BMO, CIBC, OMERS, RBC, Scotiabank, Sun Life, TD Bank Group and TMX Group) and co-developed a cybersecurity and threat management graduate certificate with Seneca Polytechnic to financial services specifications.

Hong Kong shows a complementary division of labor: the HKMA designed the Enhanced Competency Framework on Cybersecurity, and the Hong Kong Institute of Bankers, an industry body, administers it and issues a recognized professional credential. Singapore's Institute of Banking and Finance maintains the sector skills framework and has run attach-and-train conversion programs with multi-agency support.

Luxembourg demonstrates the cluster variant, and markets it: Luxembourg for Finance, its financial center promotion agency, publishes cybersecurity as part of the center's proposition, citing 317 cybersecurity companies in the jurisdiction, of which 86 have cybersecurity as their core business, alongside the national cybersecurity house, its incident response and prevention units, and an immersive crisis simulation facility.⁶

6.7. Speak collectively, and carry lessons between centers

The EU Roundtable of Financial Centers demonstrates the seventh function. A single financial center asking for harmonized implementation and common preparedness metrics is a lobby. A standing roundtable of financial centers asking for it is a constituency, and its diagnosis carries evidential weight precisely because it is collective.

Carnegie's proposal for a "G7+" arrangement on the FATF model (extending the G7 Cyber Expert Group's work to include major financial centers beyond the G7) describes an open institutional gap. It is a convening gap, not an expertise gap. The bodies that would need to be in such a room already exist; what is missing is the room.

For the World Alliance of International Financial Centers, this function has the clearest comparative advantage because it is the only one on this list that requires global rather than national membership.

The same membership supports the less visible half of this function. Lessons do not currently travel between centers. The after-action findings of one center's exercise, the control weaknesses a review surfaced, the supplier problem one jurisdiction discovered—all of it tends to stay where it was found. Circulating that material in anonymized, comparable form costs little, and it is the one thing a global association can do that no national body can. Section 8 sets out how. Both parts of this function are assessed under the Collective representation and learning dimension in Section 7.

⁵ TFI closed in 2022.

⁶ <https://www.luxembourgforfinance.com/portfolio/luxembourg-on-alert-against-cyberthreats/>

7. A resilience maturity model

The model below is offered as an evidence-supported self-assessment instrument. It is deliberately left unpopulated with member placements and is not a league table. World Alliance members differ in mandate, statutory basis, funding, and access to firms; a center without convening power over supervised entities isn't failing when it records a lower level than one that hosts the sector's crisis machinery. The model's purpose is therefore longitudinal rather than comparative. It lets a center see what the next step looks like from wherever it currently stands, and measure itself against its own previous position. The final column makes the destination's shape visible: no center known to the authors operates at level 4 across all nine assessment dimensions listed in the first column of the table below.

The model uses nine dimensions to assess the maturity of the seven functions described in Section 6 and the arrangements supporting their delivery. Each dimension is assessed separately against four maturity levels: Advocacy, Convening, Operating and Leading, with documentary evidence supporting the level recorded. Seven dimensions correspond to the seven functions, including collective representation and shared learning under Section 6.7. The two additional dimensions are Mandate, which assesses the organizational responsibility and resources supporting these functions, and Forward risks, which assesses preparation for developments in AI, digital assets and post-quantum cryptography. The thematic reviews described in Section 8 help centers identify common weaknesses, share lessons and assess progress over time.

Dimension	1 — Advocacy	2 — Convening	3 — Operating	4 — Leading
Mandate	Cyber referenced in center strategy	Standing cyber or resilience committee	Named owner and funded program	Resilience forms part of the center's external proposition
Dependency mapping	Third-party risk raised with members	Common service categories agreed with members	Aggregated dependency map maintained and shared back to members in anonymized form	Map refreshed on a fixed cycle, reconciled against supervisory registers, and compared with peer centers
Information-sharing connectivity	Signposts members to the national CERT or cyber agency	Members encouraged to join a sectoral ISAC	Existing arrangements mapped and connected; gaps identified; vendors and FinTechs included	Anonymized sharing; standing exchange of indicators and lessons with peer centers

Dimension	1 — Advocacy	2 — Convening	3 — Operating	4 — Leading
Collective exercising	Members exercise individually	Center participates in a national exercise	Center co-designs a sector exercise	Regular market-wide exercise including third parties, with published findings
Incident communications	No sector arrangement	Media contact list maintained	Standing communications group with a pre-agreed protocol	Group exercised annually; tiered-disclosure protocol agreed with authorities
Talent	Signposts to available training	Sector skills gap surveyed	Co-designed curriculum or professional credential	Cyber capability is an inward-investment asset for the center
Small-firm baseline	Awareness material issued	Baseline standard promoted	Certification or support scheme reachable for small firms; supplier expectations articulated	Ecosystem-wide baseline with procurement meaning; suppliers inside the sharing and exercising arrangements
Collective representation and learning	Common concerns and lessons noted informally	Members discuss common resilience concerns and contribute to proposed collective positions	Thematic review completed and anonymized findings circulated to members; agreed collective positions presented to relevant policymakers or authorities	Subjects re-reviewed on an agreed cycle, with progress reported and lessons shared between centers; review findings used to update collective positions
Forward risks	Awareness of AI, digital assets and post-quantum	Position papers published	Member guidance issued	Cryptographic inventory drive; agent authorization and

Dimension	1 — Advocacy	2 — Convening	3 — Operating	4 — Leading
				key-governance work with authorities

A few notes on use:

- The dimensions are not equally costly. Incident communications and forward-risk guidance are largely organizational and can be built quickly. Collective exercising and ecosystem baselines require funding and sustained convening power.
- The order in which they are built matters less than the fact that level 3 in any dimension requires an owner.
- Throughout, the distinction between level 2 and level 3 is the difference between encouraging and operating.
- Positive exemplars exist for level 3 and level 4 in every row, and several are named in Section 6. They come from centers of very different sizes, suggesting the binding constraint is mandate and focus rather than scale.

7.1. Applying the model: assessment method

A maturity level is a claim, and a claim needs a method. What follows is a starting point for participating centers to agree on and refine, not a finished procedure.

Assess each dimension independently, and record the highest level for which every descriptor is met, not the highest level for which it can produce an example. Record partial achievement as the level below, with a note on what is missing. That convention makes the model useful for planning rather than presentation.

The center conducts the assessment itself, because only the center knows its mandate and its members. Three practices keep it honest. The person accountable for the function signs off on the assessment, not the communications team. It is dated and retained, so that the next assessment is a comparison rather than a fresh start. A peer center reviews it reciprocally, which costs nothing and is the cheapest corrective to optimism.

Levels are not aspirations. A center that intends to reach level 3 on a dimension within the year records level 2 now, with the intended step noted, rather than recording level 3 conditionally.

A standard assessment record should contain, for each dimension: the level claimed; the descriptors judged applicable; a short rationale; the evidence references; any exceptions; the target level; and the action owner and review date. 'Not applicable' should require an explanation linked to the center's mandate rather than being treated as a neutral score.

For Collective representation and learning, record evidence for representation and learning separately. The maturity level recorded must be supported by all applicable descriptors for that level in both parts. Where collective representation falls outside the centre's mandate, record that part as not applicable and explain why.

Mixed achievement should not be averaged into a decimal maturity score. A Level 2 capability with two Level 3 features is recorded as 'Level 2 - progressing to Level 3', with the unmet Level 3 conditions listed. The format preserves the model's ordinal meaning and gives management a usable work plan.

Peer review should rotate among participating centers and use a short common checklist. The reviewer checks consistency between the descriptor, rationale, and evidence; challenges unsupported exceptions; and records questions without changing the center's score. Keep identifiable firm or provider material with the originating center unless a separate lawful sharing arrangement applies.

For the proposed pilot, a coordinating body could select two or three willing financial centres to assess the same three functions from Section 6, using their corresponding dimensions in the maturity model. Each selected function would be assessed separately against the four maturity levels, with documentary evidence recorded to support the level. Participants would compare the time required, available evidence and points of interpretation. The purpose would be to refine the assessment method and descriptors for those dimensions before wider adoption, rather than to publish participants' maturity levels.

Record capability and assurance separately. The capability level records the highest level supported by evidence against all applicable descriptors; the assessment also records whether the team reviewed that evidence internally, peer-reviewed it, tested it in an exercise, or independently validated it. Record unsupported statements as unverified and exclude them from achieved levels. Stronger assurance does not by itself raise the capability level, and peer review is not certification.

7.2. Demonstrating progress through evidence

The difference between a maturity model that drives improvement and one that becomes a brochure is whether evidence supports the levels. The test proposed here is deliberately simple: for every level claimed, a center should be able to name a dated artifact that a peer could inspect.

The artifacts are ordinary. A budget line and a named owner evidence a funded program. A standing group is evidenced by terms of reference and dated minutes or attendance. A scenario document, a participant list showing which third parties attended, and an after-action report evidence an exercise. A dependency map is evidenced by its aggregated output and the date of the last refresh. A small-firm scheme is evidenced by uptake, not by the scheme's existence. A thematic review is evidenced by the circulated findings and, on the second pass, by the comparison with the first.

Two consequences follow. The absence of an artifact means the level is not met, however confident the center is about the underlying capability. The artifact is what makes that capability transferable to a successor and inspectable by a peer. And progress becomes measurable in a way that self-declaration never is: the same dimension, assessed against the same evidence standard two years apart, either moved or it did not.

Evidence falls into three classes. Design evidence shows that the activity is authorized and owned, such as a mandate, terms of reference, budget, or protocol. Operating evidence shows that it happened, such as dated minutes, submissions, logs, participant records, or a published output. Effectiveness evidence shows what changed, such as closed and retested exercise actions, broader participation, reduced recovery time, or an improved result at re-review.

A simple evidence register should record the artifact, owner, date or period covered, applicable dimension and level, access classification, retention period, reviewer, and any follow-up action. The proposed retention period is at least two assessment cycles, subject to applicable retention requirements. You can share a reference with a peer even when the underlying artifact must remain internal.

Sharing should follow three tiers. The originating center retains full artifacts, available only to authorized internal reviewers. A peer may receive the reference, selected extracts, or an anonymized evidence pack under agreed confidentiality. Public reporting should use aggregated results and minimum cell sizes, with sensitive operational details removed.

Report progress as a change from a dated baseline, with the action and evidence that explain the change. Activity counts, such as meetings held or guidance downloaded, are useful operating measures but do not prove effectiveness. Where possible, pair them with an outcome measure or retest; where that is not feasible, state the limitation.

Each evidence reference should record its scope, age, source, and sampling limitations. Exercise and operating records should identify coverage, exceptions, and whether corrective actions were retested. A successful test supports conclusions within its scope; neither that result nor the absence of incidents proves overall effectiveness. The record should include adverse evidence, such as overdue actions and failed retests.

8. Continuous review and improvement

A thematic review examines one risk or dependency across a defined group of participating organizations, using common questions and an agreed collection period. Findings should guide improvements to the functions described in Section 6. For example, a cloud concentration review could identify gaps in a center's dependency map and inform the scenario for a joint exercise.

The four subjects below are topics to investigate. The nine dimensions in Section 7 assess how well established the center's functions and supporting arrangements are. Review findings and evidence of follow-up actions can inform those assessments, but completing a review does not automatically justify a higher maturity level. A change in level requires evidence that the relevant descriptors have been met.

The analysis in Sections 2 and 3 suggests the following four subjects for thematic review.

1. Cloud concentration. Which cloud services support functions that participating firms could not operate without; where those services are common across firms in the center; and what exit and fallback arrangements actually exist, as opposed to being contracted for.

2. Wallet and signing infrastructure. For firms holding client assets in tokenized or crypto form: how cryptographic keys are generated, held, and used; whether what an approver sees is what the transaction does; and which third parties sit in the signing path. Section 2.3 explains why the review should examine wallet and signing infrastructure alongside smart-contract security.

3. AI used within the platform. Which models and providers are embedded in customer-facing or decision-making processes; what authority, if any, an agent holds over a transaction; and how model changes reach production. Section 3.3 explains why mapping this dependency matters before it entrenches.

Whether firms re-test after a vendor updates a model or service; how they learn that an update has occurred at all; and which controls across the participating firms depend on the same vendor's settings. A review that established nothing else would still be worth running if it established that. ([Krisjanous, 2026](#))

4. Preparation for post-quantum cryptography. Whether a cryptographic inventory exists, whether it extends to suppliers, and how it compares to the published 2030 and 2035 milestones.

Anonymizing and aggregating findings can reduce the risk of identifying participating organizations and support candid participation. Check reports for distinctive details, especially when few participants are involved. Comparing findings across centers can help distinguish recurring weaknesses from local problems and identify useful practices that other centers could adopt. A repeat review should assess whether agreed actions were implemented and whether the evidence shows improvement.

The review should complement existing supervisory work, with its scope agreed in advance to avoid unnecessary duplication. The center's role is to coordinate the collection of responses and supporting evidence, independently aggregate the findings, and help participants agree on follow-up actions. The review protocol should define how findings requiring escalation will be handled, taking account of applicable reporting obligations and confidentiality arrangements.

The proposed review cycle has seven steps: select the subject against agreed criteria; define scope and terms; approve a common question set and evidence sample; collect responses within a fixed window; independently aggregate and quality-check the results; agree actions and owners; and re-run the core

review after 12 to 18 months. Report changes to scope or population so apparent improvement is not mistaken for a change in sample.

The output should distinguish baseline measures, common weaknesses, useful practices and agreed next actions. It should state the number and types of participants, the collection period, exclusions, aggregation threshold and limitations. Findings can be shared in tiers: full internal results for the steering group, a confidential comparative pack for participating centers, and a public summary containing no identifiable firm or provider data.

Subject selection should consider cross-firm relevance, likely decision value, overlap with current supervisory work, feasibility of obtaining comparable evidence, and the availability of an owner for follow-up. A small steering group of centers, participating firms and relevant authorities can approve the scope without giving the center a supervisory role.

The thematic review pilot would involve two or three willing centers examining one agreed subject, with a four- to six-month collection and analysis period and a pre-agreed re-review date. It could involve the same centers as the maturity assessment pilot in Section 7, which tests the assessment of three selected functions. The maturity pilot should produce feedback on the assessment method and descriptors; the thematic pilot should produce findings and agreed actions on the selected subject.

A short survey of prospective participating organizations and selected interviews can refine the question set, while case studies explain how controls operate in context. The thematic report should distinguish evidence supplied by participating organizations from desk research and record any questions it could not answer.

The review team should combine subject expertise with independence from the activities being assessed, record conflicts of interest and use more than questionnaire responses where proportionate. Before publication, participants should be allowed to correct factual errors and identify confidentiality concerns, but not to remove supported findings. Follow-up should report whether agreed actions were completed and retested, and whether unresolved weaknesses were referred through the appropriate firm or authority.

9. Conclusion

The regulatory architecture for cyber resilience in finance is mature, extensive and improving. By design, it is also an architecture that acts on firms. The risks that most threaten a financial center's standing arrive through entities the architecture does not reach and produce consequences—loss of confidence, loss of market function, loss of reputation—that no single firm internalizes.

Artificial intelligence sharpens this rather than changing it. Lowering the cost of attack while the sector converges on a narrower set of shared dependencies makes correlated, multi-institution events more likely. Agentic AI extends the same logic into the transaction itself. Digital assets demonstrate, in an environment with fewer backstops, what happens when a supplier's privileged access is compromised. The post-quantum transition adds a deadline.

The answer isn't more standards, and Europe's financial centers have already said as much: they need harmonized implementation and comparable measurement. Nor is the answer for centers to imitate supervisors, or to take on operational response. It is for centers to close the gaps that no one currently owns — to know what the center collectively depends on, to connect the information-sharing arrangements that already exist, to exercise across firms and their providers, to make a baseline reachable for the smallest firms, to rehearse a communications capability before it is needed, to build a talent pipeline to sector specification, and to carry lessons between centers.

Each of these exists somewhere today. None exists everywhere. Assembling them is not a technical problem, and it was never primarily a resourcing problem. It is a collective-action problem—which is to say, exactly the kind of problem financial centers exist to solve.

Coordination needs to continue through regular assessment and review so that participating centres can check whether their arrangements are improving. A center that assesses itself against a common frame, supports each claim with something a peer could inspect, and reviews one subject at a time until it can say whether conditions improved, is doing something more useful than declaring itself resilient.

During the preparation of this work, the author used AI tools for research and to improve language and readability. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

Bibliography

- Adrian, T. (2026, April). Tokenized Finance. *IMF Note 2026/001*. Retrieved from <https://www.imf.org/en/-/media/files/publications/imf-notes/2026/english/insea2026001.pdf>
- Bank of England, PRA, FCA. (2025). *2025 CBEST thematic*. London: Bank of England, Prudential Regulation Authority, Financial Conduct Authority. Retrieved from <https://www.bankofengland.co.uk/financial-stability/operational-resilience-of-the-financial-sector/2025-cbest-thematic>
- Biedermann, J. (2026). *Measuring International Finance Part II: Measuring the Performance of International Financial Centers*. Frankfurt. Retrieved from <https://financial-centers.org/publications/financial-center-performance/>
- Bloemen, R., Logvinov, L., & Evans, J. (2017). EIP-712: Typed structured data hashing and signing. *Ethereum Improvement Proposals*. Retrieved from <https://eips.ethereum.org/EIPS/eip-712>
- Castillo, L., Rein, D., Aoun, P., Galansky, A., Rozwarski, B., Uzdogan, K., . . . Forshtat, A. (2024). ERC-7730: Structured Data Clear Signing Format. *Ethereum Improvement Proposals*. Retrieved from <https://eips.ethereum.org/EIPS/eip-7730>
- CIRCL. (2016). *MISP - Malware Information Sharing Platform & Threat Sharing for the financial sector*. Luxembourg: Computer Incident Response Center Luxembourg. Retrieved from <https://www.circl.lu/assets/files/MISP-financial-sector-infosharing.pdf>
- Clinton, S. (2026). *OWASP GenAI Exploit Round-up Report Q1 2026*. OWASP Foundation. Retrieved from <https://genai.owasp.org/2026/04/14/owasp-genai-exploit-round-up-report-q1-2026/>
- CPMI-IOSCO. (2016). *Guidance on cyber resilience for financial market infrastructures*. Retrieved from <https://www.bis.org/cpmi/publ/d146.htm>
- CPMI-IOSCO. (2022). *Implementation monitoring of the PFMI: Level 3 assessment on Financial Market Infrastructures' Cyber Resilience*. Committee on Payments and Market Infrastructures and International Organization of Securities Commissions. Retrieved from <https://www.bis.org/cpmi/publ/d212.htm>
- Debenedetti, E., Zhang, J., Balunovic, M., Beurer-Kellner, L., Fischer, M., & Tramèr, F. (2024). AgentDojo: A dynamic environment to evaluate prompt injection attacks and defenses for LLM agents. *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*. Retrieved from https://proceedings.neurips.cc/paper_files/paper/2024/hash/97091a5177d8dc64b1da8bf3e1f6fb54-Abstract-Datasets_and_Benchmarks_Track.html
- Eisenbach, T. M., Kovner, A., & Lee, M. J. (2021). *Cyber Risk and the U.S. Financial System: A Pre-Mortem Analysis*. New York: Federal Reserve Bank of New York. Retrieved from https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr909.pdf
- ESMA. (2026). *ESMA launches Common Supervisory Action on CASPs' digital operational resilience for custody*. European Securities and Markets Authority. Retrieved from

- <https://www.esma.europa.eu/press-news/esma-news/esma-launches-common-supervisory-action-casps-digital-operational-resilience>
- ESRB. (2022). *Mitigating systemic cyber risk*. Frankfurt am Main: European Systemic Risk Board. Retrieved from <https://www.esrb.europa.eu/pub/pdf/reports/esrb.SystemiCyberRisk.220127~b6655fa027.en.pdf>
- EU-RFC. (2025). *Communiqué*. EU Roundtable of Financial Centers. Frankfurt am Main: World Alliance of International Financial Centers. Retrieved from <https://downloads.waifc.finance/regional/EU-RFC%20Communique%CC%81%20Frankfurt%2001.01.25.pdf>
- European Commission. (2024). *Commission Implementing Regulation (EU) 2024/2956. Standard templates for the register of information under DORA*. Retrieved from https://eur-lex.europa.eu/eli/reg_impl/2024/2956/oj/eng
- FATF. (2026a). *International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation: The FATF Recommendations (updated June 2026)*. Paris: Financial Action Task Force. Retrieved from <https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/fatf-recommendations-2012.pdf>
- FATF. (2026b). *Virtual Assets: Targeted Update on Implementation of the FATF Standards on Virtual Assets and Virtual Asset Service Providers*. Paris: Financial Action Task Force. Retrieved from <https://www.fatf-gafi.org/en/publications/Fatfrecommendations/targeted-updated-virtualassets-vasps-2026.html>
- FBI. (2026). *FBI's 2025 Internet Crime Report*. Washington, D.C.: Federal Bureau of Investigation. Retrieved from https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf
- FinCEN. (2024). *FinCEN Alert on Fraud Schemes Involving Deepfake Media*. U.S. Department of the Treasury's Financial Crimes Enforcement Network. Retrieved from <https://www.fincen.gov/system/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf>
- FSA. (2023). *Financial Industry-Wide Cybersecurity Exercise (Delta Wall VIII)*. Tokyo: Financial Services Agency. Retrieved from <https://www.fsa.go.jp/en/news/2023/20231218/deltawall8.pdf>
- FSB. (2024). *The Financial Stability Implications of Tokenisation*. Financial Stability Board. Retrieved from <https://www.fsb.org/uploads/P221024-2.pdf>
- Garratt, R., & Shin, H. S. (2023). Stablecoins versus tokenised deposits: implications for the singleness of money. *BIS Bulletin 73*. Retrieved from <https://www.bis.org/publ/bisbull73.htm>
- HKMA. (2024). *Cyber Intelligence Sharing Platform Enhancements*. Hong Kong: Hong Kong Monetary Authority. Retrieved from <https://brdr.hkma.gov.hk/eng/doc-ldg/docld/getPdf/20241227-4-EN/20241227-4-EN.pdf>
- HKMA. (2025). *Guideline on Supervision of Licensed Stablecoin Issuers*. Hong Kong Monetary Authority. Retrieved from https://www.hkma.gov.hk/media/eng/doc/key-functions/ifc/stablecoin-issuers/Guideline_on_supervision_of_licensed_stablecoin_issuers_eng.pdf

- IMF. (2024). *IMF Global Financial Stability Report: The Last Mile: Financial Vulnerabilities and Risks*. Washington D.C.: International Monetary Fund. Retrieved from <https://www.elibrary.imf.org/display/book/9798400257704/9798400257704.xml>
- IMF. (2025a). *Euro Area Policies: Financial System Stability Assessment*. Washington D.C.: International Monetary Fund. Retrieved from <https://www.imf.org/-/media/files/publications/cr/2025/english/1eurea2025002.pdf>
- IMF. (2025b). *Switzerland: Financial Sector Assessment Program-Technical Note on Cyber Risk Supervision*. Washington D.C.: International Monetary Fund. Retrieved from <https://www.imf.org/-/media/files/publications/cr/2025/english/1cheea2025005-source-pdf.pdf>
- ISC2. (2025). *Cybersecurity Professionals Navigate Evolving Workplaces While Seizing New Opportunities*. Alexandria: International Information System Security Certification Consortium. Retrieved from https://edge.sitecorecloud.io/internationalf173-xmc4e73-prodbc0f-9660/media/Project/ISC2/Main/Media/insights/Features/2025/12/COMMS_2025_Cybersecurity_WFS_Report_Whitepaper.pdf
- Kida, M., & Paetzold, S. (2021). *The Impact of Gray-Listing on Capital Flows: An Analysis Using Machine Learning*. Washington D.C.: International Monetary Fund. doi:<https://doi.org/10.5089/9781513582436.001.A001>
- Kotidis, A., & Schreft, S. L. (2025). The Propagation of Cyberattacks through the Financial System: Evidence from an Actual Event. (T. A. Association, Ed.) *The Journal of Finance*, 80(6), 3313-3358. doi:<https://doi.org/10.1111/jofi.13475>
- Krisjanous, M. (2026). AI in VASP Customer Checks and the Implications for AML Supervision. (Z. Limited, Ed.)
- LayerZero Labs. (2026). *LayerZero Labs KelpDAO Incident Report*. Retrieved from <https://layerzero.network/publications/kelpdao-incident-report.pdf>
- Maurer, T., & Nelson, A. (2021, March). The Global Cyber Threat. (IMF, Ed.) *IMF F&D Magazine*. Retrieved from <https://www.imf.org/-/media/files/publications/fandd/article/2021/march/en/global-cyber-threat-to-financial-systems-maurer-en.pdf>
- NIST. (2024, July). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. *NIST Trustworthy and Responsible AI (NIST AI 600-1)*. Retrieved from <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- OWASP Gen AI Security Project. (2025). *LLM06:2025 Excessive Agency. OWASP Top 10 for Large Language Model Applications*. Retrieved from <https://genai.owasp.org/llmrisk/llm062025-excessive-agency/>
- Sifma, protiviti. (2026, March). Quantum Dawn VIII: After-Action Report. Retrieved from <https://live-sifma-new.pantheonsite.io/wp-content/uploads/2026/03/SIFMA-Quantum-Dawn-VIII.pdf>

Sygnia. (2025). *Sygnia's Investigation into the Bybit Hack: What We Know So Far*. Retrieved from <https://www.sygnia.co/blog/sygnia-investigation-bybit-hack/>

TRM Labs. (2026). *H1 2026 Crypto Hacks Reach Record High as Losses Fall Below USD 1 Billion*. Retrieved from <https://www.trmlabs.com/resources/blog/h1-2026-crypto-hacks-reach-record-high-as-losses-fall-below-usd-1-billion>

WEF; Accenture. (2026). *Global Cybersecurity Outlook 2026*. Geneva: World Economic Forum. Retrieved from https://reports.weforum.org/docs/WEF_Global_Cybersecurity_Outlook_2026.pdf

Zanarc. (2026). *AI cybersecurity risk in Web3. Development, operations, and third-party exposure. Advisory*. Retrieved from <https://resources.zanarc.com/ai-web3-cybersecurity-risk-advisory>

Zanarc. (2026). *RWA Tokenisation Cybersecurity Advisory*. Retrieved from <https://resources.zanarc.com/rwa-tokenisation-cybersecurity-advisory>