



EU RegCORE Client Alert

Financial Services

September 2026



Dr. Michael Huertas

Tel: +49 160 97375760

michael.huertas@pwc.com

EU RegCORE Team

de_regcore@pwc.com

UNGA 81 and financial services: from AI agents to a new architecture of dependency

QuickTake

The 81st session of the United Nations General Assembly (UNGA 81) brings together delegations from the UN's 193 Member States for the General Debate and a wider week of high-level meetings. This Client Alert focuses on the financial-services implications of statements and events through 26 September 2026, including AI governance, payment infrastructure, sovereign and development finance, climate adaptation and pandemic preparedness.¹

UNGA 81 does not create a new financial-services rulebook. It does, however, expose a common policy problem: banks, insurers, asset managers, payment firms and financial market infrastructures increasingly depend on technology and market infrastructures that they neither own nor fully control. The emerging question is therefore not only whether an individual institution is resilient, but whether the common

¹ United Nations General Assembly, 'Meetings and Events: Eighty-first Session' (High-Level Week 2026), [available here](#).

infrastructure on which many institutions rely remains available when several firms need the same alternative at the same time.

AI is the clearest example. On 24 September, the Australian Government disclosed that an AI agent used by OpenAI's research team had gained unauthorised access to a public-facing Medicare statistics portal after encountering repeated blocks. The investigation remains ongoing; no personal Medicare information was believed to have been accessed at the time of disclosure. OpenAI separately acknowledged that its models had taken actions it did not intend.²

Five points matter most for boards, risk committees and operational-resilience teams:

1. AI can create layered dependencies that existing frameworks capture at different points, but which may not be visible as one end-to-end dependency chain;
2. the Australian incident, as presently described by the Australian Government, turns agentic-behaviour risk from a theoretical scenario into an operational risk-management question;
3. "concurrency risk" asks whether exit and substitution plans remain credible when peer institutions need the same models, cloud capacity, compute or specialist resources simultaneously;
4. the same week, the European Supervisory Authorities identified external dependencies, non-EU/EEA service providers and increasingly capable AI models as vulnerabilities for the EU financial system—giving the dependency thesis a supervisory, not merely political, anchor;³
5. and the wider UNGA 81 financial-services agenda—payments, sovereign finance, climate adaptation and crisis preparedness—is increasingly about the architecture through which capital, data and risk move.

One development merits particular attention. UK Prime Minister Andy Burnham used the post-2008 financial-crisis response as an explicit analogy for international AI coordination and said the UK intends to place AI at the centre of its 2027 G20 Presidency.⁴ The following day, the UK Foreign Secretary told the UN Security Council that resilience to AI-enabled incidents includes reinforcing financial systems and expressly listed financial stability among the issues requiring international coordination.⁵

That does not make UNGA a financial standard setter. Speeches in New York do not themselves create prudential, conduct or operational-resilience obligations. The significance lies in agenda formation—and in the fact that, during the same week, EU supervisors were already identifying external dependencies and AI

² Prime Minister of Australia, 'Press conference – New York' (24 September 2026), [available here](#); ABC News, 'Extreme concern: OpenAI agent hacked Australian public health website, prime minister says' (24 September 2026), reporting OpenAI's statement that its models took actions it did not intend.

³ European Supervisory Authorities, 'ESAs call for vigilance over external dependencies, cyber threats and private credit risks' (23 September 2026), [available here](#).

⁴ UK Prime Minister's Office, 'PM speech at UNGA: 22 September 2026' (23 September 2026), [available here](#).

⁵ UK Foreign, Commonwealth & Development Office, 'Foreign Secretary Address to the UNSC on Artificial Intelligence' (23 September 2026), [available here](#).

as financial-system vulnerabilities. This alert therefore distinguishes political signals, supervisory direction and binding legal requirements.

AI, dependency and the Australian incident: key implications for financial markets

The emerging financial-services map from UNGA 81 has several connected components: AI governance and infrastructure; payment rails and digital finance; sovereign and development finance; climate and adaptation finance; financial inclusion; and crisis preparedness. The tables below distinguish the policy signal, its legal or supervisory status and the practical significance for financial institutions as at 26 September 2026.

AI risk, agentic behaviour and operational resilience

The-matic area	Development / status as at 26 September 2026	Why it matters for financial institutions
1. Aus-tralian in-cident	Australian Prime Minister Anthony Albanese stated on 24 September that an AI agent used by OpenAI's research team during an internal evaluation gained unauthorised access to the Medicare Statistics Reporting Service portal on 18 June after encountering repeated blocks. It accessed public and non-public files and wrote files to an internal server. The Government's forensic investigation is ongoing, and no personal Medicare information was believed to have been accessed at the time of disclosure.	The incident should be treated as an investigated case, not a settled technical finding. It nevertheless illustrates an operational question that conventional model governance does not fully answer: how permissions, guardrails, monitoring and escalation should work where an agent can take multi-step actions autonomously.
2. Three AI risks	AI use risk concerns whether a use is lawful, controlled and fit for purpose. Common-provider risk concerns shared dependence on the same model, cloud or compute ecosystem. Agentic-behaviour risk concerns autonomous action that departs from the user's or provider's intended route.	A firm can manage use-case compliance and still remain exposed to a common dependency. Different risks require different controls and different owners.
3. EU su-pervisory corroboration	On 23 September, the ESAs identified external dependencies, emerging technologies and private credit as key vulnerabilities. They highlighted dependence on non-EU/EEA ICT providers and growing cyber risks linked to increasingly capable AI models.	This is the strongest bridge from the UNGA policy debate to current EU supervision. Dependency mapping and preparedness are already supervisory concerns, even though "concurrency risk" is not a defined regulatory category.
4. DORA and third-party risk	DORA applies to ICT risk, incidents, resilience testing and ICT third-party risk. Separately, the EBA published its final Guidelines on non-ICT third-party risk on 18 September; they promote a holistic approach across ICT and non-ICT	The relevant exercise is not to relabel every AI issue as DORA. It is to map the complete dependency chain and then identify which legal and supervisory requirements attach at each layer.

The-matic area	Development / status as at 26 September 2026	Why it matters for financial institutions
	services and include contracting, subcontracting, monitoring and exit strategies. ⁶	
5. Con-currency risk	Firm-level exit planning assumes an alternative exists. Sys-tem-level resilience requires a second assumption: that the alternative remains available when multiple institutions seek it simultaneously.	The useful stress test is not only whether Bank A can exit Provider X. It is whether Bank A can exit when Banks B to Z are seeking the same model, compute, cloud capacity, validation resources and specialist personnel.
6. Inter-national architec-ture	The UK intends to put AI at the centre of its 2027 G20 Presi-dency, explicitly drawing on the G20 response to the 2008 fi-nancial crisis. The EU has also framed AI regulation, finan-cial stability and resilient supply chains as problems requir-ing multilateral action. ⁷	No FSB-equivalent architecture for frontier AI ex-ists. The policy question is whether a sequence of political principles, technical standards and do-mestic implementation begins to emerge.
7. Agen-tic com-merce	Banks have warned that AI shopping agents may create fraud, scam, privacy and redress risks as agents move from product discovery towards transaction initiation.	Authentication, consent, liability, fraud monitoring and redress become more complex where the consumer's instruction is mediated or materially shaped by an autonomous agent.
8. Legal status	UNGA speeches and side-event statements are policy sig-nals. The Australian incident is under investigation. The UK G20 agenda is a stated intention, not an international stand-ard.	Do not convert a political signal into a compliance obligation. Use the signals to identify scenarios and dependencies that existing legal frameworks already require firms to govern.

Sovereign finance, climate and crisis preparedness

Thematic area	Development / status as at 26 September 2026	Why it matters for financial institutions
9. Sover-eign and develop-ment fi-nance	UNGA-week speeches and African Union initiatives repeat-edly focused on debt sustainability, the cost of capital, guar-antees, blended finance and mobilisation of domestic institu-tional savings. The AU framed the challenge as the “archi-tecture that connects capital to opportunity” and linked it to PAPSS and stronger African financial institutions. ⁸	For lenders and investors, the commercial question is less the headline volume of capital than the allocation of currency, political, credit and first-loss risk—and whether public or multi-lateral balance sheets make projects investable.
10. Cli-mate ad-aptation	At the 24 September high-level meeting on sea-level rise, the President of the General Assembly cited estimated ad-aptation needs for developing countries of USD 310–365	Physical risk can transmit through property val-ues, collateral, insurance availability, bank loss-given-default, municipal finances and ultimately sovereign balance sheets. Financing resilience

⁶ European Banking Authority, 'The EBA publishes its final Guidelines on the management of third-party risk' (18 September 2026), [available here](#).

⁷ European Council, 'Speech by President António Costa at the 81st United Nations General Assembly' (24 September 2026), [available here](#).

⁸ African Union, Selma Malika Haddadi, 'Remarks at the African Union–United Nations Top Table' (20 September 2026), [available here](#).

Thematic area	Development / status as at 26 September 2026	Why it matters for financial institutions
	billion annually by 2035, versus USD 26 billion of international public adaptation finance in 2023. ⁹	before loss crystallises becomes a balance-sheet question.
11. Payments and digital infrastructure	Brazil's General Debate address defended Pix as public payment infrastructure and linked it to financial inclusion. African Union discussions linked PAPSS, local-currency cross-border payments and financial sovereignty while stressing that sovereignty is not isolation. ¹⁰	Payment rails are increasingly treated as strategic infrastructure. AI agents add another layer by mediating instructions, authentication and consumer choice.
12. Financial inclusion	A UNGA High-Level Side Event on 21 September marked "50 Years of Micro Credit" and connected social business, responsible technology and financial inclusion. ¹¹	The policy discussion is moving beyond access alone towards whether financial infrastructure and technology improve resilience and outcomes for users. That is relevant to retail banking, payments and digital distribution.
13. Pandemic preparedness	The 25 September high-level meeting emphasised sustainable financing and responsible use of AI and AGI for epidemiological analysis, surveillance and early threat identification. ¹²	The lesson resembles financial crisis planning: capacity, information channels and contingent financing are more effective when built before a low-frequency, high-impact shock.
14. Translation into firm governance	Across the themes, the common structure is dependency on infrastructure → control → risk allocation → substitutability → system-wide capacity under stress.	The practical response is to connect AI governance, third-party risk, operational resilience, payments, crisis management and investment decisions rather than treat each as a separate policy workstream.

Making sense of the key changes

In our view, the most important conceptual change is not that UNGA has invented a new regulatory category. It is that the same dependency problem is appearing simultaneously in political, supervisory and operational-resilience discussions. That point is strengthened by the ESAs' 23 September risk update, which identifies external dependencies and AI among vulnerabilities requiring monitoring and preparedness.

The recurring analytical chain is now increasingly:

⁹ President of the UN General Assembly, 'Speech at the high-level meeting on addressing the existential threats posed by sea-level rise' (24 September 2026), [available here](#).

¹⁰ United Nations, transcript of the 81st General Debate, Day 1 (22 September 2026), [available here](#); African Union, Haddadi remarks (20 September 2026), above.

¹¹ United Nations, 'Social Business, Youth and Technology' (21 September 2026), convened under the theme '2026: 50 Years of Micro Credit, 2030: The Year of Social Business', [available here](#).

¹² President of the UN General Assembly, 'Remarks at the high-level meeting on pandemic prevention, preparedness and response' (25 September 2026), [available here](#).

critical function → external dependency → control and information rights → failure mode → substitutability
→ capacity when multiple firms act at once

For AI specifically:

use-case risk → common-provider concentration → agentic-behaviour risk → concurrency risk

These chains meet at resilience, but they should not be collapsed into one legal test. DORA, the AI Act, EBA third-party-risk expectations, contractual governance and prudential recovery or resolution frameworks have different scopes and purposes. The value of a dependency map is that it lets firms see where those regimes overlap—and where they do not.

The three AI risks: a framework for controls

The following framework distinguishes three categories of AI risk that require different controls. The Australian incident illustrates why agentic-behaviour risk is now operational rather than theoretical.

AI use risk	Core question: Is a particular use lawful, controlled and fit for purpose?
	Typical controls: AI Act compliance, model governance, data quality, validation, conduct controls
	Why it matters now: Necessary, but increasingly only the first layer of AI governance.
Common-provider risk	Core question: How many critical functions or firms depend on the same model, cloud or compute ecosystem?
	Typical controls: DORA third-party risk, concentration mapping, exit and substitution planning
	Why it matters now: Raises a system-wide concentration and concurrency question—what happens when multiple firms need the same alternative at the same time?
Agentic-behaviour risk	Core question: Can an autonomous system pursue a goal in an unintended way, including around controls?
	Typical controls: Permissions, sandboxing, human intervention points, monitoring, incident response
	Why it matters now: The Australian disclosure makes this an operational scenario. The facts remain under investigation, so controls should be based on the capability and failure mode, not on premature conclusions about causation.

Four cautions before drawing regulatory conclusions

First, UNGA speeches do not themselves create prudential, conduct or operational-resilience obligations. Second, “concurrency risk” is an analytical concept used in this alert, not a defined term in DORA or the AI Act. Third, not every model provider or AI service is automatically an ICT third-party service provider or a DORA-designated critical third-party provider; the legal analysis depends on the service and the applicable designation criteria. Fourth, the Australian facts remain subject to forensic investigation. The Government’s account and OpenAI’s acknowledgment of unintended model actions support scenario analysis, but they should not be treated as final findings on technical causation or liability.

Why the Australian incident nevertheless matters

Until this week, many agentic-AI scenarios could be treated as forward-looking. Australia’s disclosure makes it harder to assume that agent autonomy is confined to laboratory hypotheticals. That is consistent with the UN Independent International Scientific Panel on AI’s 21 September thematic brief on AI agents, misalignment and loss-of-control risks, which reviewed prior incidents while expressly avoiding predictions about the probability or timing of severe loss of control.¹³

The useful question for financial institutions is therefore practical: what permissions can an agent exercise; what does it do when a control denies the preferred route; what telemetry reveals that behaviour; who can stop it; and how quickly must the provider notify a regulated customer if unexpected activity could affect the customer’s systems or obligations? The Australian Government said it was notified on 10 September about the 18 June incident, by email to a public mailbox.

What to watch next

The following developments merit close monitoring. The table below summarises the key areas and their significance for financial institutions.

Development	Current status	Why it matters
UK G20 AI agenda	The UK intends to put AI at the centre of its 2027 G20 Presidency, using the 2008 crisis response as an analogy.	Watch whether political principles turn into concrete workstreams on transparency, access, model testing, incident reporting or cross-border crisis coordination.
UN / Security Council AI work	The Security Council heard from AI-sector leaders during UNGA week; the UK linked AI resilience expressly to financial systems and financial stability. The UN scientific panel has also published a thematic brief on AI-agent misalignment.	Watch for proposals that move from general AI safety into critical-infrastructure or financial-stability coordination.
EU supervisory treatment	The ESAs’ 23 September update identifies external dependencies and AI as current vulnerabilities; the EBA’s non-ICT third-party-risk Guidelines are final and awaiting translation.	The most immediate client relevance is supervisory: dependency mapping, non-EU provider exposure, contracting, subcontracting and exit strategy.
Australian investigation	A forensic investigation and government taskforce are ongoing.	Further facts could refine expectations on provider notification, permissions, agent monitoring and incident classification. Avoid over-generalising before the investigation concludes.
Agentic commerce	Banks have raised fraud, scam, privacy and redress concerns as AI shopping agents move closer to transaction initiation.	Payments teams should clarify authentication, consent, liability, fraud controls and consumer redress where an agent mediates the transaction.

¹³ United Nations Independent International Scientific Panel on AI, ‘Thematic Brief on AI Agents, Misalignment and the Risk of Losing Human Control’ (21 September 2026), [available here](#).

Development	Current status	Why it matters
UNGA closing statements	The General Debate continues through 26 September and concludes on 28 September.	Further developing-country statements may add to the debt, development-finance, payment-sovereignty and international-financial-architecture themes; update the policy map after the debate closes.

The UNGA 81 debate is best read as a set of policy signals, not as a substitute for current EU law. The strongest immediate validation comes from the ESAs' same-week risk update: external dependencies and AI are already part of the supervisory risk agenda. Firms therefore do not need to wait for a future international AI architecture before asking whether their present dependency maps and exit assumptions are credible.

What firms should consider doing now

The following steps translate the policy and supervisory signals into existing governance processes. They are management responses, not obligations created by UNGA.

- **Map below the application layer.** Identify the model, provider, orchestration layer, hosting, compute, data, connectivity and any downstream tools that support critical or important functions.
- **Separate the three AI risks.** Distinguish use-case compliance, common-provider concentration and agentic-behaviour risk so that controls, ownership and escalation routes are not conflated.
- **Test concurrency and substitution.** Challenge exit plans on the assumption that peers are seeking the same alternative capacity, validation resources and specialist personnel at the same time.
- **Revisit contractual information and notification rights.** Ensure contracts and operating procedures provide the information, cooperation and escalation needed for the financial entity to assess and meet its own incident-management obligations.
- **Constrain agent permissions.** Define tool access, transaction limits, human intervention points, logging, sandboxing and stop mechanisms according to the functions the agent can perform.

The non-AI themes also are likely to warrant a practical response: payment firms should assess agent-mediated authentication and redress; lenders and investors should distinguish headline development-finance commitments from funded guarantees and risk-sharing structures; and banks and insurers should test how physical climate risk can transmit through collateral, insurability and sovereign exposures.

The EU regulatory and supervisory bridge

Policy announcements and other happenings during and in the adjacent events to UNGA do not alter the EU rulebook. The practical significance is how the week's dependency themes intersect with frameworks that are already applicable or moving into implementation:

- **DORA (Digital Operational Resilience Act)** — applicable since 17 January 2025. DORA addresses ICT risk management, ICT-related incident reporting, resilience testing and ICT third-party risk. For AI, the key is to identify which parts of the model / orchestration / cloud / compute chain are ICT services

used by the financial entity and to ensure registers, contracts, concentration analysis and exit arrangements reflect the actual service chain.¹⁴

- **AI Act** — most provisions apply since 2 August 2026. Regulation (EU) 2026/1744 amended the timetable for the main high-risk-system requirements: 2 December 2027 for Annex III stand-alone high-risk systems and 2 August 2028 for Article 6(1) / Annex I high-risk systems embedded in products. The AI Act is central to use-case governance; it does not by itself answer the financial-sector aggregation and concurrency questions developed in this alert.¹⁵
- **EBA Guidelines on non-ICT third-party risk** — final, not under consultation. The EBA published the final Guidelines on 18 September 2026; they are awaiting translation and are not yet applicable. They focus on arrangements supporting critical or important functions and cover the full lifecycle, including due diligence, contracting, subcontracting, monitoring and exit, with a two-year transitional period.
- **ESAs Autumn 2026 risk update** — on 23 September, the ESAs identified external dependencies, non-EU/EEA ICT providers and increasingly capable AI models as vulnerabilities and called for stronger preparedness and monitoring. This is more directly relevant today than speculating that frontier-model providers will automatically become DORA critical third-party providers; designation depends on the services supplied and the statutory criteria.

Outlook

For much of the post-financial-crisis period, regulation focused on strengthening individual financial institutions: capital, liquidity, governance, resolvability and operational resilience. AI makes visible a related challenge—the resilience of external infrastructure shared across firms.

Payments, cloud, market data, communications, climate insurance and access to international capital raise variants of the same problem. UNGA 81 has not created a new regulatory architecture, but it has brought the dependency question into unusually sharp political focus. The ESAs' 23 September assessment shows that parts of the same issue are already embedded in the EU supervisory agenda.

For boards and senior management, the useful distinction is between signal, supervision and law. A UNGA speech does not create a regulatory requirement. A G20 analogy does not itself produce an international standard. A concurrency-risk scenario does not change capital requirements. But each can identify a failure mode that existing governance, DORA, AI Act and third-party-risk frameworks require firms to understand.

The four questions that matter are therefore:

1. *Who controls critical infrastructure?*
2. *Who bears the risk when it fails?*
3. *Can the dependency be substituted?*

¹⁴ Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA), OJ L 333/1, [available here](#).

¹⁵ Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689, including revised application dates for Chapter III, Sections 1–3, [available here](#).

4. *What happens when everyone needs the alternative at the same time?*

That is where UNGA 81 becomes useful for financial institutions: not as a source of rules, but as an early indicator of which infrastructure and dependency questions are moving towards international standard-setting and supervisory attention. The financial-services question now extends beyond an individual institution's own resilience to encompass the collective resilience of the infrastructure on which the entire financial system depends—particularly at moments when everyone needs it at the same time.



About us

PwC Legal is assisting financial services firms and market participants in forward planning for relevant developments. We have assembled a multi-disciplinary and multijurisdictional team of sector experts to support clients in addressing challenges and opportunities and to engage with their market stakeholders and regulators.

Moreover, we use a number of RegTech and SupTech tools for supervised firms, including the Rule Scanner technology within a trusted set of managed solutions from PwC Legal Business Solutions. These solutions support horizon scanning and risk mapping of legislative and regulatory developments, sanctions and fines from thousands of legislative and regulatory policymakers and other industry voices across more than 170 jurisdictions impacting financial services firms and their business.

Equally, in leveraging Rule Scanner technology within those managed solutions, we offer clients a further solution to digitise financial services firms' internal policies and procedures, create a comprehensive documentation inventory with an established hierarchy and embedded glossary, maintain version control across a defined backward- and forward-looking timeline, map critical-path dependencies and flag legislative and regulatory developments that may require corresponding changes across related policies and procedures.

The PwC Legal team's Rule Scanner-enabled work has received ALM Law.com's "2024 Disruptive Technology of the Year Award" and the "2025 Regulatory, Governance and Compliance Technology Award".

If you would like to discuss any of the developments mentioned above, or how they may affect your business more generally, please contact any of our key contacts or PwC Legal's RegCORE Team via de_regcore@pwc.com or our [website](#).



Contact

Dr. Michael Huertas

Tel: +49 160 97375760

michael.huertas@pwc.com

EU RegCORE Team

de_regcore@pwc.com

