MAHMOOD.INSIGHT

A POLICY WHITE PAPER · THE FUTURE OF FINANCIAL INFRASTRUCTURE

Supervising AI Before It Becomes Systemic

Banking supervisors are taking different approaches to AI. Oman has established national principles for its safe and ethical use; a banking-specific supervisory posture is still emerging. A study of international practice, and the supervisability that can be built before AI dependencies become difficult to unwind.

White Paper 02
2026

Audience
CBO · Boards · Risk & Compliance

Reference work
Bank for International Settlements

Mahmood Al-Hattali · Independent Strategic Advisor · Banking & Payments

IN THIS PAPER

Contents

— Executive summary 02

01 AI is becoming critical banking infrastructure 04

02 The Infrastructure Escalation — four layers, four controls 05

PART I How major jurisdictions are approaching AI 06

03 Three supervisory postures & the international reference layer 06

04 Three risks that recur across the frameworks reviewed 08

PART II A maturity model for AI supervision 09

05 The Four Stages of Central Bank AI Maturity 09

06 Supervisability by design — and the material-AI perimeter 11

07 Foundation-model concentration risk 12

08 Where the data goes — the fourth property 13

09 What the CBO could require 14

— Methodology & definitions 15

EXECUTIVE SUMMARY

The argument, in one page.

AI is moving from a tool banks use to infrastructure banks depend on. When that shift completes, AI risk becomes systemic risk — and the time to build supervisability is before, not after.

1	**AI is becoming critical infrastructure, not a feature.**	Once fraud defence, credit and operations depend on models, a model failure stops being one bank's problem. *That threshold is what this paper is about.*
2	**Jurisdictions have gone different ways. There is no settled model to copy.**	EU: high-risk obligations from 2 Dec 2027 and 2 Aug 2028. US: generative and agentic AI outside the scope of revised model-risk *guidance* — novel and rapidly evolving, not to speed adoption. Germany: folded into existing ICT controls.
3	**An international reference layer sits alongside them.**	The BIS, the FSI and the Innovation Hub have built a body of work on AI governance, financial stability and supervisory technology. *It corroborates. It confers no authority.*
4	**Three risks recur across the frameworks reviewed.**	Explainability. Third-party concentration. Operational resilience. They sit alongside — not instead of — data governance, privacy, cybersecurity and human oversight. *They are the three that decide whether a supervisor can see what is happening.*
5	**Sequencing is an advantage, used well.**	A compact supervisor can require supervisability by design — explainable, auditable, resilient, confidential — before dependencies deepen, instead of retrofitting oversight afterwards.
6	**The ask is principle, not prescription.**	A posture for consideration. The rulebook is the CBO's to write.

ONE IDEA, THREE SUPPORTS

If this paper is remembered for one principle, let it be **supervisability by design**: that material AI be explainable, auditable, resilient and confidential before critical functions depend on it. The other frameworks serve it — the *Escalation* explains *why*, the *Four Stages* show *where* a supervisor stands, and *concentration* names *what* to watch most closely.

01 · AI AS CRITICAL INFRASTRUCTURE

The threshold that changes everything.

There is a moment when a technology stops being something a bank chooses to use and becomes something it cannot operate without. AI in banking is approaching that line. Supervision is easier to build before it is crossed.

A bank that uses AI to screen fraud faster is using a tool. A bank that can no longer screen fraud without the model is depending on infrastructure. Volume, speed and sophistication have outrun the manual fallback. The same shift is underway in credit decisioning, AML monitoring and customer operations.

This matters to a supervisor for one reason. **Infrastructure dependence can create systemic consequences where exposure is shared, correlated, or embedded in critical functions.** Tool failures are generally more containable — unless the tool is widely shared or critical to service continuity. The transmission mechanisms are specific: shared vendors, shared cloud or model infrastructure, correlated outputs, common data or design weaknesses, and the inability to revert to an effective manual process. The BIS has examined the financial-stability implications of exactly this.[4]

> **THE PREMISE OF THIS PAPER**
>
> Broad prohibition is neither realistic nor consistent with current adoption trajectories. The practical supervisory question is how material uses are governed and kept visible — whether the systems AI runs on remain supervisable by the time critical functions depend on them. That window is open now, and it narrows as dependencies deepen.

Part I maps how major jurisdictions are approaching that window — three supervisory postures, the international reference layer alongside them, and three risks that recur across the frameworks reviewed. Part II sets out where a supervisor stands today, and the posture that sequencing makes available: the opportunity to observe where other regimes have needed to clarify, delay, or retrofit controls before defining a sector-specific posture of one's own.

02 · THE INFRASTRUCTURE ESCALATION — WHY SUPERVISABILITY MATTERS

As AI rises through four layers, the control must change.

AI does not become systemic in one step. It climbs. Each layer calls for a different control, and the risk is that supervisors apply tool-era oversight to infrastructure-era exposure.

LAYER 1
Tool

Control: internal governance. AI assists a task; failure is contained. The bank's own model-risk policy suffices.

↓

LAYER 2
Critical business process

Control: explainability & audit. A core function now depends on the model. It must be reconstructable and reviewable.

↓

LAYER 3
Critical infrastructure

Control: operational resilience. The bank cannot operate without it. Fallback, override, and tested contingency become mandatory.

↓

LAYER 4
Systemic infrastructure

Control: systemic oversight. Many banks depend on the same AI. Only the supervisor can see and manage the cross-system exposure.

THE POINT OF THE LADDER

Most supervisory debate sits at Layers 1–2 while AI climbs toward Layer 4. Match the control to the layer. Build the higher control before the system gets there. By Layer 4 retrofitting is more expensive and less reliable.

PART I

How major jurisdictions are approaching AI

The international landscape available to study before committing — three supervisory postures, one international reference layer, three recurring risks.

03 · THREE SUPERVISORY POSTURES

No single model to copy — which is the opportunity.

Major jurisdictions have gone different ways. That is useful. A supervisor still defining its own posture can watch where each has had to clarify, delay or retrofit, and borrow selectively — within its own legal context.

THE INTERNATIONAL SUPERVISORY LANDSCAPE

	Posture 1 · EU	Posture 2 · Germany / ECB	Posture 3 · United States	Reference layer · BIS / FSI
Instrument	A dedicated horizontal legal regime — the AI Act	Existing prudential, ICT and outsourcing controls	Revised model-risk *guidance* (SR 26-2)	Research, convening and supervisory-enablement work
Legal status	Binding legislation	Supervisory practice, alongside the AI Act — not replacing it	Non-binding supervisory guidance	No binding force. A reference, not a regime
Where it stands	High-risk obligations apply 2 Dec 2027 (stand-alone) and 2 Aug 2028 (embedded). Council approval 29 Jun 2026[5]	BaFin addresses AI-related ICT and third-party risk; ECB adds cloud, outsourcing and resilience expectations[6]	Generative and agentic AI outside scope — *novel and rapidly evolving*, not to accelerate adoption[7]	Governance, financial-stability analysis, SupTech — alongside, not above, any jurisdiction
Reusable here?	The classification logic, not the legislative machinery	**Yes — closest fit.** The CBO already supervises ICT and outsourcing	The materiality test, not the scope exclusion	Yes — as the anchor for terminology and expectations

THE BIS REFERENCE FRAME

Where the evidence backs the frameworks.

The frameworks here are the author's own. They are not asserted in a vacuum. The BIS, the FSI and the Innovation Hub have published work that tests each of their load-bearing claims — and it holds. Used as a reference point, not as borrowed authority. No endorsement is implied.

Across 2024–25 the BIS published on the governance of AI adoption in central banks, regulating AI in the financial sector, the financial-stability implications of AI, generative AI in supervision, and Project AISE.[9] Read against the frameworks here, it corroborates. It is not their source.

SUPPORTS · THE MATURITY MODEL

Governance begins with risk appetite

The BIS frames AI governance as starting from what an institution will and will not accept, adapting existing structures — exactly the Stage 1→2 transition this paper's maturity model describes.

SUPPORTS · SUPERVISABILITY BY DESIGN

Unexplained means untrusted

The BIS reference work is direct on this point: supervisors expect firms to explain AI models, and results that cannot be understood are unlikely to be trusted — the explainability pillar of supervisability by design, corroborated in the BIS/FSI reference work.

SUPPORTS · STAGE 3 (AI SUPERVISION)

The regulator should use AI too

The BIS documents central banks using AI for supervision itself — the SupTech capability that defines reaching Stage 3 on this paper's ladder.

WHY THIS ORDERING MATTERS

The BIS/FSI work corroborates the frameworks without authoring them. Ownership and reasoning remain the paper's own, and no endorsement is implied.

04 · THREE RISKS THAT RECUR ACROSS FRAMEWORKS

Three risks that recur across the frameworks reviewed.

Three risks recur across the frameworks reviewed. They are not the whole of any regime — data governance, privacy, cybersecurity, human oversight, bias and lifecycle validation sit alongside them. But these three decide whether a supervisor can see what is happening. A defensible place to start, not a claim of universal consensus.

RISK 1 · EXPLAINABILITY

Models that cannot be explained cannot be governed

A model that declines a loan or flags a transfer should be reconstructable — by the bank, and on demand by the supervisor. Without sufficient traceability and interpretability, prudential review, customer-outcome testing, and accountability are materially impaired. Of the three, this is the expectation that appears most consistently across the frameworks reviewed.

RISK 2 · THIRD-PARTY CONCENTRATION

When everyone uses the same few providers

Where most banks rely on a handful of AI vendors and cloud providers, a single failure or weakness can propagate beyond one institution. Concentration in the AI supply chain is a stability risk distinct from any one bank's model risk, and the BIS has examined it directly.[4]

RISK 3 · OPERATIONAL RESILIENCE

What happens when the model is wrong, or down

Where AI becomes infrastructure, it can become a single point of failure. Resilience — fallback processes, human override, tested contingency — is what keeps a model failure from becoming a service outage. It appears consistently across the frameworks reviewed.

THE STRATEGIC SIMPLIFICATION

You do not need to settle every open AI-governance debate to start well. Require explainability. Map third-party concentration. Test operational resilience. The contested questions can wait for the international picture to settle.

PART II

A maturity model for AI supervision

An original framework for where a supervisor stands today, and where it may choose to reach — and why sequencing is a strong hand for climbing it.

05 · THE FOUR STAGES OF CENTRAL BANK AI MATURITY — WHERE A SUPERVISOR STANDS

From adopting AI to overseeing a system that runs on it.

Supervisors are not all in the same place. The first step is knowing which stage you are actually at.

1

STAGE 1

AI Adoption

The supervisor permits and observes AI in banks, relying on existing rules. Awareness, not yet a posture.

2

STAGE 2

AI Governance

The supervisor sets expectations for how banks govern their own AI — risk appetite, accountability, documentation.

3

STAGE 3

AI Supervision

The supervisor actively examines AI — requiring explainability and audit, reviewing models, using SupTech of its own.

4

STAGE 4

AI Systemic Oversight

The supervisor maps and manages cross-system AI exposure — concentration, correlation, resilience — that no single bank can see.

THE SEQUENCING ADVANTAGE ON THE LADDER

Early movers climbed Stages 1–3 by trial and error and are now retrofitting Stage 4. A supervisor building now can design for Stage 4 visibility from the start — while still building each necessary capability in turn, and learning from the costly missteps. Sequencing is not a lower rung. It is the chance of a better-planned climb.

05 · PUTTING THE MODEL TO WORK

Which stage is your central bank at?

Six questions place any supervisor on the ladder. True stage = the highest where **all** boxes tick. Each gap names the next capability to build. The model makes no claim about which stage any named supervisor currently occupies.

THE MATURITY SELF-ASSESSMENT

☐	Stage 1	Does the supervisor maintain an inventory of where AI is used across its banks?
☐	Stage 2	Does it require banks to govern their own AI — risk appetite, accountability, documentation?
☐	Stage 2	Are AI model decisions required to be explainable and auditable on demand?
☐	Stage 3	Does it actively examine banks' AI models — and use SupTech of its own?
☐	Stage 4	Does it map cross-bank dependency on shared foundation models, vendors, and clouds?
☐	Stage 4	Can it see correlated model risk — where many banks could fail the same way at once?

06 · SUPERVISABILITY BY DESIGN — THE CENTRAL IDEA

Build the window in, before it closes.

The idea is straightforward. Require that material AI be supervisable from the start — explainable, auditable, resilient, confidential — instead of retrofitting oversight onto systems already running. Security by design, applied to supervision.

MATERIAL AI — THE DECISION FLOW

Does the AI influence credit, pricing, eligibility or customer treatment — or support AML, fraud, sanctions or regulatory reporting?

YES →

Material AI. Supervisability required at deployment. If it also supports a critical function: infrastructure-grade controls — fallback, override, tested contingency.

NO — but does it act autonomously, or handle sensitive or high-volume personal data?

If yes: material AI.
If no: existing governance is sufficient.

A model that declines loans sits inside the perimeter. A chatbot that drafts internal emails does not.

SUPERVISABILITY BY DESIGN — WHAT IT MEANS IN PRACTICE

Three properties, required at deployment, not after.

Explainable: the bank can reconstruct why a material AI decision was made, on demand, at a depth appropriate to the use case. **Auditable:** a logged, reviewable trail of model versions, data and decisions a supervisor can examine after the fact. **Resilient:** a tested fallback for when the model is wrong or unavailable. **Confidential:** the supervisor can say where customer data goes when the model is called, who can lawfully compel access to it, and whether the bank could leave the provider if it had to.

Requiring this now shapes how banks build. A bank told at the design stage that its AI must be explainable and auditable builds those properties in cheaply. A bank told afterwards rebuilds. That is the cost the early movers are absorbing.

THE FORENSIC CONNECTION

Supervisability is the ability to reconstruct a decision after the fact. It is the discipline of a forensic audit.

07 · FOUNDATION-MODEL CONCENTRATION RISK — WHAT TO WATCH MOST CLOSELY

The risk no single bank can see.

Model risk lives inside a bank. The bank can manage it. Concentration risk lives between banks.

Outsourcing risk used to ask which cloud a bank runs on. There is a layer beneath that. Most fraud and credit decisioning still runs on traditional statistical and machine-learning models — bank-specific, no shared upstream provider. Foundation-model dependency concentrates elsewhere: assistants, copilots, document analysis, generative support to alert investigation.[8]

THE AI DEPENDENCY STACK



The bank controls the top two layers and depends on the bottom three. Cloud, vendor and foundation-model concentration are distinct exposures.

THE SUPERVISORY VISIBILITY PYRAMID

A single bank sees	its own applications, its own models, its own vendor contracts	●
It cannot see	which foundation models its peers depend on	○
It cannot see	whether the system is converging on the same few providers	○
It cannot see	correlated failure — many banks failing the same way at once	○
Only the supervisor sees	the cross-system exposure	●

08 · WHERE THE DATA GOES — THE FOURTH PROPERTY

Supervisable, but not confidential, is not supervisable.

The stack shows what a bank depends on. It does not yet show what crosses it. When a model is called, customer data moves — and that movement is a supervisory event, not a procurement detail.

THE FOUR QUESTIONS THE STACK RAISES

When the model is called, does customer data leave Oman?	**PDPL Art. 23**	Permitted with consent, subject to national interest — but the bank stays accountable for protection at the far end.[2]
Who else can lawfully compel access to it there?	**Sovereignty**	Encryption does not answer foreign legal process. Location is one question; jurisdiction is another.
Does the prompt itself become training data?	**Confidentiality**	A query sent to a model *is* a transfer. Retention terms are a supervisory matter, not a procurement footnote.
If the provider withdrew tomorrow, could the bank leave?	**Reversibility**	Model portability is harder than cloud portability. Dependency that cannot be exited is the one that matters.

None of this needs new law. The PDPL governs cross-border transfer and has been fully enforceable since February 2026. The CBO already supervises outsourcing, cloud, encryption, key control and data segregation. What is missing is not authority. It is the supervisory picture of where the data actually goes when the models are called — and only the supervisor can assemble it.

WHY THIS BELONGS IN THE SUPERVISORY PERIMETER

A model whose decisions are explainable and auditable, but which sends customer data to a jurisdiction the supervisor cannot see into, is not supervisable in any sense a central bank would accept. Confidentiality is not a fourth concern bolted onto the other three. It is the condition that makes the other three worth having.

09 · WHAT THE CBO COULD REQUIRE

A posture, offered — not a rulebook.

What follows is offered for consideration, not as prescription. The rulebook is the CBO's to write. This is one shape it could take.

One Oman-specific observation. Oman is not starting from a blank page. The Public Policy for the Safe and Ethical Use of Artificial Intelligence Systems entered into force in April 2025, setting national principles of transparency, human oversight, accountability, and privacy, and inviting regulators to issue sector-specific provisions consistent with it.[1] The Personal Data Protection Law and its Executive Regulations already impose controller and processor obligations.[2] What is still emerging is a consolidated *banking-specific* supervisory posture. Meanwhile the CBO has advanced both an open-banking framework and a regulatory framework for digital banks.[3] These deepen data sharing, API interdependence, and reliance on common technology providers. They do not, by themselves, create systemic AI risk — but they make early visibility over AI, cloud, and third-party dependencies more valuable, not less. The supervisory window is therefore current rather than theoretical, and the task is operationalisation rather than invention.

1	REFERENCE	**Draw on the BIS/FSI reference work.** Align with international reference material rather than importing any single jurisdiction's regime wholesale.	*Costs nothing. Start here.*
2	REQUIRE	**Supervisability by design, for material AI.** Explainable, auditable, resilient — built in at deployment, not retrofitted.	*Cheapest now. Expensive later.*
3	MAP	**Build the concentration picture.** Which banks depend on which foundation models, vendors and clouds.	*Only the supervisor can see this.*
4	EQUIP	**Use AI in supervision itself.** Grow the supervisor's capability alongside the banks it oversees.	*Stage 3 needs people, not only policy.*

THE LINE TO REMEMBER

A central bank does not supervise intelligence. It supervises trust — and trust survives only when failure remains visible, explainable, and recoverable.

SOURCES & METHODOLOGY

Grounding.

A regulator-facing paper must be transparent about its evidence. P1 = binding or formally issued primary material. P2 = authoritative institutional research or interpretation. Author = the paper's own inference, labelled as such. Evidence verified July 2026.

#	Source — institution, title, date	Tier
1	MTCIT (Oman), *Public Policy for the Safe and Ethical Use of Artificial Intelligence Systems* — in force 9 April 2025	**P1**
2	Personal Data Protection Law, RD 6/2022, Art. 23; Executive Regulations, MD 34/2024, Arts. 37–38 — cross-border transfer; controller accountable for protection abroad. Fully enforceable 5 Feb 2026	**P1**
3	Central Bank of Oman — Open Banking Framework; Digital Banks Framework (eff. 1 Jun 2025); outsourcing and cloud rules (encryption, key control, segregation, audit trail)	**P1**
4	BIS — governance of AI adoption in central banks (Jan 2025); regulating AI in the financial sector (Dec 2024); financial-stability implications of AI (Jun 2025)	**P2**
5	Council of the EU — *Digital Omnibus on AI*, final adoption 29 June 2026. High-risk application: 2 Dec 2027 (stand-alone, Annex III); 2 Aug 2028 (embedded in regulated products, Annex I)	**P1**
6	BaFin — guidance on ICT risks arising from AI use (Dec 2025); ECB Banking Supervision — *Guide on outsourcing cloud services* (2025)	**P1**
7	Federal Reserve / OCC / FDIC — SR 26-2, *Revised Guidance on Model Risk Management*, 17 April 2026. Non-binding supervisory guidance. Gen/agentic AI outside scope as "novel and rapidly evolving" (fn. 3)	**P1**
8	Author's analysis. Foundation-model concentration is an inference drawn from vendor-market structure, not a measured finding for Oman	**Author**
9	BIS Innovation Hub — Project AISE (AI Supervisory Enablement)	**P2**

Method & stance

- **Anchored on international reference work.** The paper draws on the body of work produced by the BIS, the Financial Stability Institute, and the Innovation Hub, with peer-regulator material as comparative evidence. That work corroborates the paper's frameworks; it does not author them, and no endorsement is implied.
- **Posture, not prescription.** Section 08 offers a strategic posture for consideration. Supervisory rules are the CBO's to make; nothing here represents the position of any institution.
- **Evidence vs argument.** Regulatory facts carry numbered endnotes. The frameworks are the author's reasoned analysis — proposals, not established standards.
- **What is not claimed.** The paper does not assert that Oman lacks an AI policy. It recognises the April 2025 national policy and proposes a banking-specific operationalisation of it. Unverifiable claims were removed, not qualified.

DEFINITIONS

The terms, in one line each.

So the paper reads cleanly for prudential, board, and technical audiences alike.

Term	Definition
Systemic risk	Risk that a failure spreads across the financial system rather than staying contained in one institution — the threshold at which AI failure becomes a supervisory concern.
Model risk	The risk that a model is wrong, misused, or misunderstood, and that decisions based on it cause loss or harm.
Explainability	The ability to reconstruct and articulate why a model produced a given output — the precondition for supervisory trust.
Auditability	The presence of a logged, reviewable trail of model versions, data, and decisions that a supervisor can examine after the fact.
Operational resilience	The capacity to keep functioning when a model is wrong or unavailable — through fallback processes, human override, and tested contingency.
Concentration risk	System-level risk arising when many institutions depend on the same AI vendors, cloud providers, or model approaches.
SupTech	Supervisory technology — AI and data tools used by the regulator itself to monitor banks and the system.
Supervisability by design	The principle, proposed in this paper, that AI be explainable, auditable, resilient and confidential from deployment rather than retrofitted later.
The Infrastructure Escalation	This paper's framework describing AI's rise through four layers — tool, critical process, critical infrastructure, systemic infrastructure — each requiring a different supervisory control.
Four Stages of AI Maturity	This paper's model of where a central bank stands in governing AI: Adoption, Governance, Supervision, Systemic Oversight.

MAHMOOD INSIGHT · ORIGINAL FRAMEWORKS

The AI Workforce Pyramid · White Paper 01 — where the Omanisation gap really sits
The Infrastructure Escalation · White Paper 02 — how AI rises from tool to systemic infrastructure
The Four Stages of Central Bank AI Maturity · White Paper 02 — where a supervisor stands
Supervisability by Design · White Paper 02 — the principle: explainable, auditable, recoverable before systemic

A growing body of frameworks developed across the series, intended to be cited independently of any single paper.

Mahmood Insight — White Paper No. 02 · First Edition — July 2026
© Mahmood Insight. All rights reserved.
This paper synthesises publicly available information for policy discussion. It is not investment, legal, or regulatory advice, and does not represent the position of any institution. Figures are accurate to the author's knowledge at publication.