The hidden architecture of AI in Omani banking.

The strategic question is no longer adoption. It is who gains capability, who retains authority, and where dependency accumulates.

A central bank does not supervise intelligence. It supervises trust — and trust survives only when failure remains visible, explainable and recoverable. Build capability. Retain authority. See the system.

HOW TO READ THIS DECK

Three kinds of statement. Each one labelled.

EVIDENCE

Supported by regulation, policy, primary sources or institutional disclosure.

OBSERVATION

Practitioner inference from how banking structures and implementation decisions behave. Reasoned, not measured.

PROPOSAL

Offered for consideration. Not an existing regulatory expectation.

Where the evidence ends and the argument begins should never be a matter of interpretation.

THE TRILOGY

SOURCE MAP · SEE P47

Three papers. Three questions. One architecture.

WHITE PAPER 01 · WHO GAINS CAPABILITY?

AI and the Omani Banking Workforce

Its contribution: the workforce effect isn't fixed in advance. What decides it is which roles are automated, which are augmented, and whether the route upward is redesigned.

WHITE PAPER 02 · WHO RETAINS AUTHORITY?

Supervising AI Before It Becomes Systemic

Its contribution: governance must rise with consequence and dependency, not with the AI label. Material AI must stay explainable, auditable, resilient and confidential.

WHITE PAPER 03 · WHO SEES THE SYSTEM?

From Digital-First to Dependency-Aware

Its contribution: an institution can manage its own dependency. Only the supervisor can see the concentration across the system.

Paper 01 followed capability. Paper 02 followed authority. Paper 03 followed dependency. Together they describe the hidden architecture of AI in banking.

ACT I — WHAT IS ACTUALLY CHANGING

OBSERVATION

AI is not the story. Dependency is.

The first question was what AI can do. That question is now too small. The real question is what banks, employees and supervisors become dependent on once AI starts doing the work.

The danger is not adoption. The danger is adoption without visibility.

ACT I — WHAT IS ACTUALLY CHANGING

FRAMEWORK

Three questions decide Oman's AI future.

CAPABILITY

Can we understand it?

Does Oman hold the human capability to operate, challenge and improve the technology — not only to be employed alongside it?

AUTHORITY

Can we overrule it?

Can the institution explain, challenge, suspend and reverse a material AI-supported decision?

DEPENDENCY

Can we withstand its failure?

Can the system understand, withstand and replace the external layers critical activity depends on?

Capability concerns people. Authority concerns decisions. Dependency concerns infrastructure.

None of the three requires owning everything. Each requires visibility, control options and credible recovery.

ACT I — WHAT IS ACTUALLY CHANGING

OBSERVATION

Three conditions, together, produce hidden dependency.

- **Scale.** More banking activity passes through technology.
- **Sameness.** Several institutions adopt the same providers, models, clouds or implementation partners.
- **Opacity.** No participant can see the full chain, or the concentration across the sector.

None of the three is a problem alone. Scale without sameness is progress. Sameness without opacity is a known trade-off. It is the three together that create exposure nobody chose.

Fragility rarely starts with one bad call. More often it's many sound decisions, all pointing the same way.

ACT I — WHAT IS ACTUALLY CHANGING

OBSERVATION

A bank can become more digital and less independently operable at the same time.

More applications launched. More processes automated. More services delivered. And underneath, fewer people who could run the thing if the provider stopped answering.

If our principal provider became unavailable tomorrow, what share of our AI roadmap would still operate?

ACT II — HOW THIS HAPPENS IN PRACTICE

OBSERVATION

Seven routes to approval. None of them meets the others.

AI investment rarely arrives through one door. It arrives through business transformation, technology renewal, fraud and compliance, customer experience, cost optimisation, digital channels, and vendor-led pilots.

Each proposal is manageable on its own terms. Each is approved by people doing their jobs properly.

The cumulative dependency stays invisible because no committee ever sees all seven together.

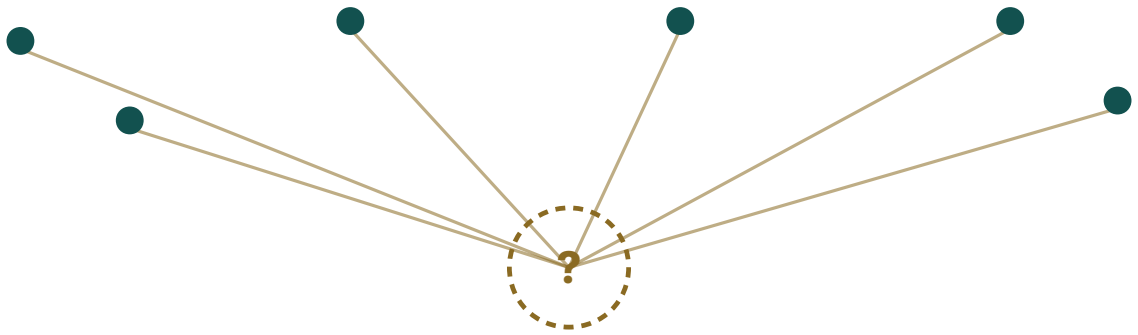
ACT II — HOW THIS HAPPENS IN PRACTICE

OBSERVATION

A bank can have governance approval without integrated governance.

Risk sees model risk	**Compliance** sees customer treatment
IT sees architecture	**Procurement** sees the contract
HR sees workforce implications	**Business** sees value

Six functions. Six correct answers. Six approvals on file.



The gap between the six — owned by none of them

Six correct approvals can still create one invisible dependency.

ACT II — HOW THIS HAPPENS IN PRACTICE

OBSERVATION

Procurement optimises for the things that are easy to compare.

WEIGHED HEAVILY

- Initial price
- Implementation speed
- Reference sites
- Feature comparison

WEIGHED LIGHTLY, IF AT ALL

- Long-term portability
- Knowledge retained inside the bank
- Provider rights to change the model
- Evidence available to internal audit
- Model substitution
- Continuity after termination

The second column decides what the bank owns in five years. The first column wins the evaluation.

ACT II — HOW THIS HAPPENS IN PRACTICE

OBSERVATION

Banks fund tools faster than they redesign roles.

Licences are issued. Training is delivered. Attendance is reported. And then:

- Decision rights remain exactly as they were
- Productivity expectations are never restated
- No specialist career path is created
- Operating roles disappear before the route upward exists
- Vendor staff retain the most valuable implementation knowledge

Which capability are we genuinely building inside the bank — and which are we renting?

ACT II — HOW THIS HAPPENS IN PRACTICE

OBSERVATION

A successful pilot proves that the pilot worked.

Pilots run with narrow data, senior attention, vendor support on call, selected users, limited customer consequence and no real production dependency.

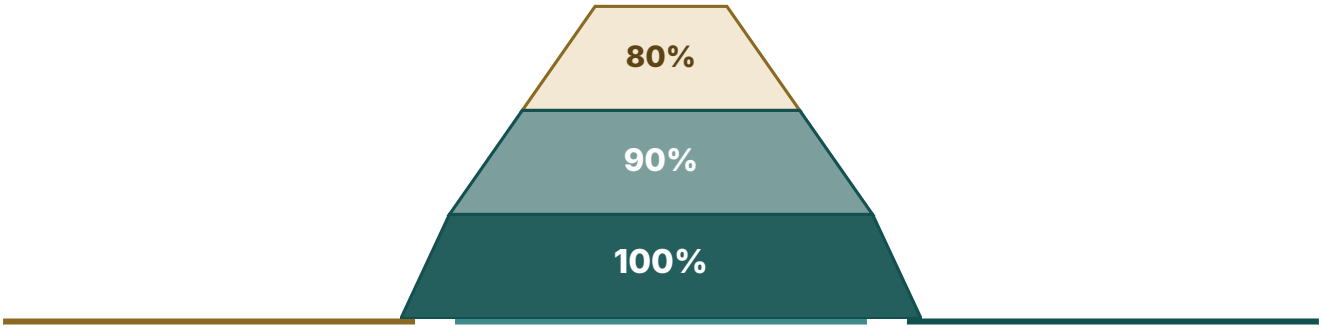
Every one of those conditions disappears at exactly the moment the model starts to matter.

The real test begins when the model becomes ordinary.

ACT III — THE OMANI CONSEQUENCE

WHITE PAPER 01

The AI Workforce Pyramid. The base is Omani. The apex is not.



Senior management · 80%
The hardest tier to fill. Where specialist dependence persists — and where augmentation may carry the greatest leverage.

Middle management · 90%
Technical and specialist roles. The contested layer AI will reshape.

Clerical & non-clerical · 100%
Mandated fully Omani. The base AI automates first if cost-logic leads.

Sector Omanisation targets by cadre, CBO Circular BM 1105 — as cited in White Paper 01 and corroborated by Chambers' 2026 Banking Regulation guide and Oxford Business Group's 2025 Oman banking overview.

The gap that persists is at the top — precisely where stronger AI-enabled judgement and oversight can matter most.

ACT III — THE OMANI CONSEQUENCE

EVIDENCE-BASED · E01

AI does not arrive in a neutral workforce structure.

It enters a sector where local participation is strongest in operational and administrative layers, where specialist expertise is slower and costlier to build, and where vendors often hold much of the implementation knowledge.

A cost-led programme therefore automates the most localised work first, leaving the deepest structural dependency intact.

The bank reports lower cost. The country receives fewer entry routes.

Evidence: cadre targets are published as 100% clerical, 90% middle and 80% senior management (CBO Circular BM 1105, as cited in White Paper 01 and corroborated by Chambers' 2026 Banking Regulation guide and Oxford Business Group's 2025 Oman banking overview). The asymmetry follows from those targets read against where AI lands.

ACT III — THE OMANI CONSEQUENCE

WHITE PAPER 01

AI does three different things to work.

Augment
Raise the capability of a human role

Substitute
Remove a task, or a role

Create
Generate work that did not exist — model oversight, validation, fairness, data governance, API and partnership roles

PAYMENTS	**Mostly augmentation** The analyst is upgraded; only first-line triage is substituted. Creates senior oversight roles.
CREDIT	**The sharpest case** Real substitution in routine underwriting, and real new senior roles in validation and explainability.
OPEN BANKING	**Job-creating** Less about automating existing work than creating new: API, partnership, product, consent and data-governance roles.

Open banking is one of the clearer places to build new Omani specialist roles around AI — APIs, consent, partnerships, product and data governance — because much of the work is additive rather than purely substitutive. Lead where technology creates a credible career path, not only a cost saving.

ACT III — THE OMANI CONSEQUENCE

PROPOSAL

Remove the lower rungs, and no one reaches the top.

If AI removes entry-level work before institutions redesign the route upward, future Omani specialists have nowhere to begin. Preserving obsolete tasks isn't the answer. Rebuilding the ladder is.

- Apprenticeships in data, product and model operations
- AI-assisted specialist roles as a destination, not a by-product
- Structured movement from operations into analytics and control
- Explicit succession paths for critical technical functions

Are we using AI to develop Omani specialists — or removing the roles specialists normally emerge from?

ACT III — THE OMANI CONSEQUENCE

WHITE PAPER 01

"Move people up" is a slogan until it has a pathway.

STEP 1 — **Free the time**
Automate the routine and ring-fence the freed hours as training capacity rather than cutting them. A budget choice, not a technical one.

STEP 2 — **Reskill toward the edge**
Train for oversight, not operation: reading model outputs, spotting failure, challenging false negatives.

STEP 3 — **Occupy the new senior roles**
Model validation, fairness and financial-crime oversight become the destination. Imported expertise is now grown.

STEP 4 — **Redeploy into growth**
Capacity beyond oversight goes to open banking, SME lending and inclusion: growth the sector could not previously staff.

Front-line staff already hold customer and market knowledge that models do not. The bridge to specialist oversight is not training alone: it needs technical learning, supervised practice and a defined role to move into. Freed capacity can help fund that bridge.

ACT III — THE OMANI CONSEQUENCE

OBSERVATION

Omanisation can rise while national capability remains externally dependent.

LOCALISATION ASKS

- How many Omanis are employed?
- In which grades?
- Against which target?

CAPABILITY ASKS

- Who designs the model?
- Who validates it?
- Who negotiates the vendor?
- Who recovers it when it fails?
- Where does institutional memory accumulate?

Omanisation measures who is in the institution. Capability ownership asks who can run it when the vendor leaves.

Capability ownership and Omanisation aren't rivals. Omanisation should increasingly produce it.

ACT III — THE OMANI CONSEQUENCE

THE HONEST BALANCE

Two things this argument does not pretend.

The incentives are misaligned. Elevation costs more in the short term. An individual bank bears most of the near-term investment to build national capability, while much of the long-term benefit accrues across the wider system. That is a collective-action problem, and one reason this is a matter for public policy rather than commercial strategy alone.

The outcome is not assured. Whether AI ultimately raises or lowers the skill threshold for advanced banking work remains uncertain. The claim is narrower: deployed deliberately, it creates a more favourable environment for local capability than previous technology waves did — not that the good outcome follows automatically.

The policy question is whether incentives reward capability building as visibly as they reward near-term cost reduction.

Structural, not personal. Reducing expat dependency means a structural skills dependency in specific functions — not the displacement of any individual or nationality.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

OBSERVATION

A human approval can preserve legal accountability after practical authority has already moved.

The approver may lack the time to examine the recommendation, the knowledge to challenge it, the standing to disagree, and may be measured against the model's own output.

A human in the loop may still be outside the decision.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

WHITE PAPER 02

Supervisability by design. Build the window in, before it closes.

Require that material AI be supervisable from the start rather than retrofitted onto systems already running.

MATERIAL AI

- Influences credit, pricing, eligibility or customer treatment
- Materially supports AML, fraud, sanctions or regulatory reporting
- Acts, recommends or communicates autonomously with material consequence, or handles sensitive personal data

LOWER MATERIALITY

- Managed under existing controls, proportionate to risk
- A model that materially influences a loan decision sits inside
- An internal drafting assistant may sit outside only if it has no sensitive data, autonomous action or material consequence

Materiality follows consequence, autonomy, data sensitivity and dependency — not the label on the tool.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

THE BOUNDARY

This doctrine is a spine — not a substitute for the controls beside it.

Capability, authority and dependency answer three structural questions. They do not replace the control systems that determine whether an AI use case is fair, valid and remediable.

FAIRNESS & CONDUCT	**Customer outcomes** Bias testing, differential outcomes, disclosure where needed, contestability and remedy.
MODEL & DATA ASSURANCE	**Validity and reliability** Data quality, baseline comparison, drift, independent validation, revalidation and decommissioning.
INCIDENT & REMEDIATION	**When harm occurs** Which customers were affected, whether earlier decisions need review, and what correction is owed.

A system can be controllable and resilient and still produce an unfair or invalid outcome. Structural control is necessary; it is not sufficient.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

WHITE PAPER 01

The minimum control floor this doctrine assumes.

Any AI model that materially influences a credit decision should face independent validation, bias testing and an explainability function answerable to governance and, where required, the supervisor. The same principle scales to other high-consequence uses.

Explainability asks why an output appeared. Validation asks whether the system is reliable enough to influence the decision at all.

A system can be explainable, auditable, resilient, locally controlled and replaceable — and still treat a customer unfairly. An AI incident is therefore not only a technology event.

This doctrine does not attempt a full rulebook for conduct, fairness or model assurance. Those disciplines need their own standards, evidence and supervisory treatment.

Clear scope is stronger than pretending one framework governs every AI risk.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

PROPOSAL

Six questions — answered with evidence, not assertion.

01 Can the human understand the basis, uncertainty and limits of the recommendation?

02 Can the human disagree without unreasonable operational friction?

03 Can the institution reconstruct the decision later?

04 Can the process continue safely without the model?

05 Can the model or the provider be replaced?

06 Is the human accountable because they exercised judgement — or because their name appears last?

Evidence means override rates, reasons recorded, time permitted for challenge, outcomes of disagreement, fallback exercises actually run.

A signature is not evidence of judgement. Without evidence, "human in the loop" is a statement about the org chart.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

WHITE PAPER 02

Four properties, required at deployment.

EXPLAINABLE

Why was this decided?

The bank can reconstruct why a material AI decision was made, on demand, at a depth appropriate to the use case.

AUDITABLE

Can it be re-examined?

A logged, reviewable trail of model versions, data and decisions a supervisor can examine after the fact.

RESILIENT

What if it is wrong, or down?

A tested fallback for when the model fails or becomes unavailable.

CONFIDENTIALITY — who can reach the data? The bank and supervisor can determine where customer data goes when the model is called, which legal routes may compel access, what is retained, and whether the bank could leave the provider.

Confidentiality is the fourth property. Without it, explainability, audit and resilience can still leave the institution exposed.

ACT IV — WHERE AUTHORITY ACTUALLY SITS

EVIDENCE-BASED · E02

Residency is where data sits. Jurisdiction is which laws may reach it.

A bank may satisfy a residency requirement and still face a lawful route through which a foreign authority — depending on the provider, contract and applicable law — can compel access. The two questions must be tested separately.

- Where do prompts and outputs go, and how long are they retained?
- Are inputs used for training, by default or by contract?
- Which subcontractors sit behind the named provider?
- Under whose law is the contract enforced?

Can the bank leave the model — not merely terminate the contract?

ACT IV — WHERE AUTHORITY ACTUALLY SITS

OBSERVATION

The bank owns the accountability. The provider may own the evidence.

Business ownership, risk and compliance, internal audit — each requires evidence. With material AI, that evidence may sit with the provider: model updates, architecture, interpretability, incident detail, exit complexity.

Where is the knowledge needed to challenge the system — inside the institution, or inside the vendor?

ACT V — THE SYSTEM

OBSERVATION

Every bank can see itself. No single bank can see the whole system.

Each institution can see its own contracts, models, systems and providers. It generally cannot see whether peers are converging on the same cloud, model, processor, region or recovery dependency.

That picture only appears above the institution.

ACT V — THE SYSTEM

WHITE PAPER 02

The AI dependency stack. Direct control usually weakens upstream.

Application	Built by the bank or supplied
Use-case model	Internal, adapted or supplied
Foundation model	A concentrated global market
Cloud	A concentrated global market
Compute & physical infrastructure	Concentrated further upstream

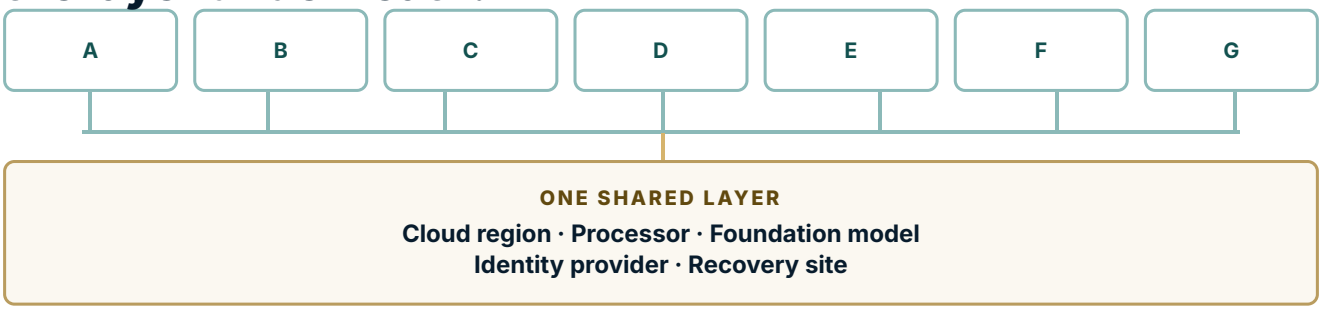
The restraint this argument keeps: most fraud and credit decisioning still relies on traditional statistical or machine-learning models, often without a shared foundation-model dependency. Concentration becomes more visible in generative-AI assistants, copilots and document workflows.

A bank can govern its model risk. Only a sector view can reveal common upstream dependency.

ACT V — THE SYSTEM

OBSERVATION

Seven banks can compete at the customer layer and share the layer underneath.



Shared infrastructure is not the problem. Unseen, untested and irreversible dependence is.

ACT V — THE SYSTEM

OBSERVATION

Four forms of commonality — and only some are concentration.

DIVERSE	**Provider layer** Several providers with credible substitutes.
SHARED BY DESIGN	**Common utility or rail** The sector intentionally relies on one shared layer; resilience depends on that layer.
COMMON	**Standard** One specification, many implementers. A common protocol is not a common provider.
UNMEASURED	**Provider chain** Actual sector concentration is not yet visible enough to quantify with confidence.

Institution-level redundancy does not automatically become sector resilience. Different front-end or core providers can still converge on the same cloud region, processor, identity service or recovery dependency.

Seven separate fallback plans can still fail together if they end at the same layer.

ACT V — THE SYSTEM

PROPOSAL

The provider name won't show you concentration. The chain will.

Bank	Named in every register
Use case	Owned, documented
Application	Usually known
Model	Sometimes known
Cloud & region	Rarely aggregated
Telecoms & recovery	Almost never mapped

Two banks with different vendors can still share a region, a processor and a recovery site.

ACT V — THE SYSTEM

EVIDENCE-BASED · E04

The next systemically critical dependency may sit outside finance.

It may take no deposits, extend no credit and appear on no bank balance sheet, yet become operationally critical to several licensed institutions at the same time.

A cloud environment. A processor. An identity platform. A foundation-model provider. In some cases, a specialist implementation partner.

The prudential perimeter can remain financial while critical dependency increasingly sits beyond it.

ACT V — THE SYSTEM

PROPOSAL

Two different jobs. One is inherently sector-level.

THE BANK

- Use-case ownership
- Model governance and validation
- Data controls and jurisdiction
- Human intervention, evidenced
- Fallback and exit
- Vendor governance
- Workforce transition

THE SUPERVISOR

- Sector concentration picture
- Cross-bank dependency visibility
- Correlated scenario testing
- Supervisory capability
- Common reporting expectations

Sector visibility doesn't hand the supervisor direct jurisdiction over every provider. What it allows is action through institutions, contracts, concentration expectations and coordinated oversight.

If several banks use one provider, who has tested that failure as a sector rather than as seven institutions?

ACT VI — THE OMANI RESPONSE

EVIDENCE + PROPOSAL · E03

Oman already has much of the foundation. The next step is to connect it.

Oman has a national AI policy. CBO frameworks already address cloud, cyber resilience, open banking and digital banks, while Vision 2040 links competitiveness to capability and skills.

This doctrine proposes one operating picture that connects AI use, decision authority, internal capability, data path, provider chain, fallback and sector concentration.

The opportunity is not another layer of policy. It is making the existing layers visible together.

THE SYSTEM BENEATH THE DOCTRINE

ORIGINAL FRAMEWORKS

The frameworks already exist. This doctrine assembles them.

WP 01	**The AI Workforce Pyramid**	Where the Omanisation gap really sits — with augment / substitute / create, the three effects AI has on work.
WP 02	**The Infrastructure Escalation**	How AI rises from tool to systemic infrastructure, and why the control must change at each layer.
WP 02	**The Four Stages of AI Maturity**	Where a supervisor stands: adoption, governance, supervision, systemic oversight.
WP 02	**Supervisability by Design**	Explainable, auditable, resilient, confidential — before critical functions depend on it.
WP 03	**The dependency chain**	Why concentration is invisible at the provider name, and why every instrument lands on the institution.

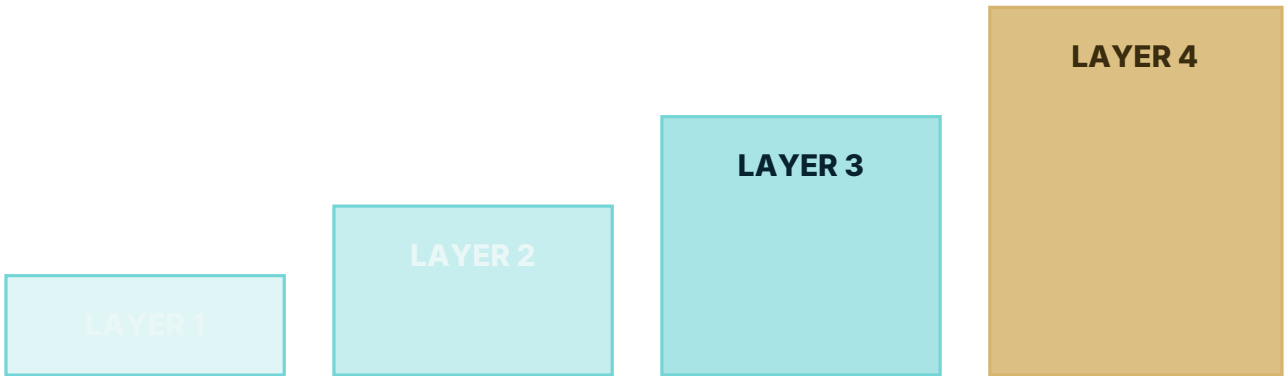
Capability. Authority. Dependency. — the spine that holds them together.

ACT VI — THE OMANI RESPONSE

WHITE PAPER 02

The Infrastructure Escalation. Four layers. Four controls.

AI climbs toward systemic status in stages, not one step, and the risk is applying tool-era oversight to infrastructure-era exposure.



Tool

Proportionate internal control. AI assists a contained task. Existing model, data, conduct and security controls may be sufficient only where consequence, autonomy, sensitivity and dependency remain low.

Critical business process

Explainability and audit. A core function now depends on the model. It must be reconstructable and reviewable.

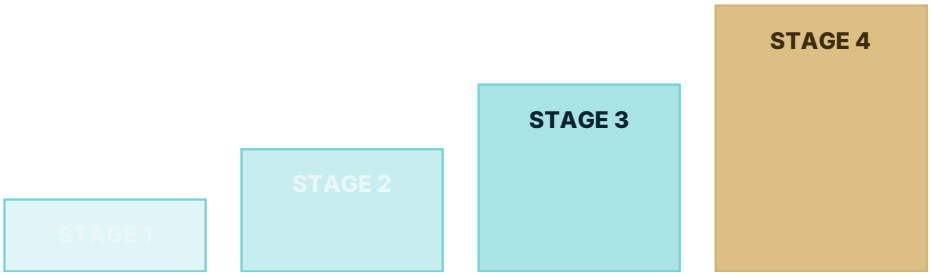
Critical infrastructure

Systemic infrastructure

ACT VI — THE OMANI RESPONSE

WHITE PAPER 02

The Four Stages of Central Bank AI Maturity.



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

AI Governance
Expectations set for how banks govern their own AI — risk appetite, accountability, documentation.

AI Supervision
The supervisor actively examines AI — requiring explainability and audit, reviewing models, using supervisory technology (SupTech) of its own.

AI Systemic Oversight
Cross-system exposure mapped and managed — concentration, correlation, resilience — that no single bank can see.

The design lesson is to avoid treating Stage 4 as a future add-on. Cross-system visibility is easier to build into reporting and approval structures before dependencies become entrenched.

ACT VI — THE OMANI RESPONSE

PROPOSAL

Two pages before approval. Make the dependencies visible.

Value Expected benefit and how it will be evidenced	**Customers** Who is affected and which decisions change
Decision rights Human authority, override and escalation	**Workforce** Roles changed and Omanisation net effect
Knowledge What capability remains inside the bank	**Data** Path, sensitivity, retention and jurisdiction
Provider chain Model, cloud, subcontractors and concentration	**Fallback & exit** Continuity, substitution and transition time

Approval becomes an institutional decision: value, customer consequence, control, capability and dependency are considered together.

ACT VI — THE OMANI RESPONSE

WHITE PAPER 02

Six questions mapped across four maturity stages.

True stage is the highest stage for which every requirement at that stage and below is met. Each gap names the next capability to build.

Q1 · S1	Is there an inventory of where AI is used across supervised institutions?
Q2 · S2	Are institutions required to govern material AI — risk appetite, accountability and documentation?
Q3 · S2	Can material AI decisions be explained, reconstructed and audited on demand?
Q4 · S3	Are material systems actively examined, and is SupTech used where it improves supervisory reach?
Q5 · S4	Is cross-institution dependency on shared models, vendors, clouds and critical services mapped?
Q6 · S4	Can correlated failure be seen — several institutions exposed to the same model, provider or dependency at once?

The model makes no claim about which stage any named supervisor currently occupies.

ACT VI — THE OMANI RESPONSE

PROPOSAL

A termination clause is not an exit plan.

A real exit test asks seven questions, and the last one is the hardest:

- Can the data be exported — in what format?
- Can decision history be reconstructed?
- Can prompts, embeddings and configurations move?
- Is there a credible alternative provider?
- How long would transition take?
- What stops operating during transition?
- **Which of our own people understands the system well enough to lead the exit?**

If the answer to the last question is a vendor name, there is no exit plan.

ACT VI — THE OMANI RESPONSE

EVIDENCE-BASED · E04

At sufficient scale, AI-intensive banking cannot be supervised effectively by manual review alone.

Stage 3 needs people, not only policy. The supervisor's own capability has to mature alongside the market:

- Machine-readable submissions rather than documents
- Use-case visibility across institutions
- Dependency and concentration analytics
- Cross-bank anomaly detection
- Supervisory AI and data capability of its own

The point is not to automate supervision. It is to give supervisory judgement enough visibility to keep pace with the system it oversees.

ACT VI — THE OMANI RESPONSE

PROPOSAL

Test degradation, not just outage.

Total failure is the easy scenario. The harder ones are quieter:

- One provider affecting several institutions at once
- Outputs degrading or becoming inconsistent rather than disappearing
- An agentic workflow taking the wrong action repeatedly before humans notice
- A cloud region, processor, identity service or recovery dependency degrading
- Vendor support, access or data movement constrained during a critical period

Banks would test individually. Only the supervisor can assess the correlated effect.

ACT VI — THE OMANI RESPONSE

OBSERVATION

A commercially successful AI programme can still be strategically negative.

A bank may show cost down and processing time improved. That alone doesn't tell us the bank got stronger. It may just as easily have removed an internal skill, transferred knowledge to a vendor, and created an exit problem that only becomes visible in three years.

- Productivity · customer outcome · risk improvement
- Error reduction · control quality · resilience
- **Employee capability · decision control · dependency and exit**

Cost reduction alone cannot tell you whether a deployment strengthened or hollowed the institution.

ACT VI — THE OMANI RESPONSE

EVIDENCE-BASED · E01

This did not begin with AI.

Market	Where value is moving
Infrastructure	How rails shape competition
Capability	Paper 01 · people
Authority	Paper 02 · decisions
Dependency	Paper 03 · infrastructure
Implementation	The operating cycle

Earlier work examined how the sector is changing and how infrastructure shapes competition. Earlier editions examined the market and its infrastructure. The three white papers then moved in sequence from capability, to authority, to dependency.

Developed from the structure and behaviour of Omani banking. Applicable wherever local capability, external technology and financial-system resilience meet.

THE DOCTRINE

Oman does not need to own every layer of artificial intelligence.

But it must understand the layers it relies upon, retain authority over the decisions they influence, and hold the capability to operate when those layers fail.

This isn't innovation versus control. It's conscious dependency versus accidental dependency.

The operating cycle is explicit: See → Classify → Control → Test → Measure → Reassess. At Measure, track value, customer outcome, override behaviour, control quality, capability retained and dependency created. Repeat whenever the model, provider, data, use case, consequence or concentration changes.

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

Build capability. Retain authority. See the system.

METHOD

SOURCE MAP

Where the evidence-marked claims come from.

Evidence-marked claims draw from four source groups. The map also separates author-originated frameworks from public sources checked directly for this release.

- **E01 · Mahmood Insight papers.** White Papers 01 and 02 were directly inspected in the controlled corpus. White Paper 03 — *From Digital-First to Dependency-Aware* — was directly inspected in both English and Arabic, standard and mobile editions for this release. Its dependency claims remain subject to the evidence tiers and limitations stated in that paper.
- **E02 · Oman data & cloud.** Personal Data Protection Law (RD 6/2022); Executive Regulation (MD 34/2024); CBO cloud rules (BM 1185); and provider-specific contractual terms.
- **E03 · Oman foundation.** MTCIT safe and ethical AI policy (2025); CBO Cyber Security & Resilience Framework (BM 1194); open-banking and digital-bank frameworks; Oman Vision 2040.
- **E04 · International supervisory context.** FSB AI financial-stability work (2024–25) and BIS/FSI work (2024–26) on AI governance, data, third parties, concentration and supervisory capability.

Everything else is marked observation or proposal, and should be read as reasoning rather than as record.

Traceability matters: the source tells you what is known; the label tells you what this doctrine is adding.

THE OPEN QUESTIONS

This is offered as a practitioner contribution — not a finished rulebook.

The next questions are practical:

- Where should Omani banks deploy AI first?
- Which capabilities must remain inside the institution?
- What should banks measure after implementation?
- Which dependencies should be visible at sector level?

I welcome challenge, evidence and practical experience from banks, regulators, providers and researchers. I am particularly interested in conversations where this thinking can be tested against a live portfolio, translated into an operating model, or applied to a real implementation.

Mahmood Al Hattali · Mahmood Insight
I work where banking strategy, product execution, financial infrastructure and institutional change meet.
mahmoodssh@gmail.com · linkedin.com/in/mahmoodinsights
For feedback, collaboration or professional opportunities.