THE CHARTERED INSTITUTE OF LOAN & RISK MANAGEMENT OF NIGERIA

PAPER PRESENTATION

BY

DR. CHIME IBIKUNLE PERETZ

Founder PERETZ/Risk Management Consultant

SMART AND SAFE GUILDELINES FOR AI USE WITHIN AN ORGANISATION

@

16TH ANNUAL RISK PROFESSIONAL INTERNATIONAL CONFERENCE

NICON LUXURY HOTEL,

AREA 11, GARKI

ABUJA.

NOVEMBER 28, 2026.

ABOUT THIS MANUAL

A manual to work in, not just read

This manual goes with the keynote at the 16th Annual Risk Professionals International Conference of the Chartered Institute of Loan & Risk Management of Nigeria. It carries the evidence behind each claim, four cases worth studying slowly, a working set of guidelines and a self-assessment you can run inside your own institution.

1 **AI is good.** It has already made organisations faster and sharper at real work.

2 **AI carries real risk.** Most of it arrives through ordinary use by ordinary staff, not attackers.

3 **Guidelines are necessary, not sufficient.** They tell people what is allowed.

4 **Capability decides what happens.** It is what an employee does at 4pm on a Friday, with a deadline and a tool that promises to save an hour.

CHOOSE YOUR ROUTE

Board and C-suite	Chapters 1, 5 and 7
Senior and mid-level managers	Chapters 3, 4, 6 and 7
Early-career professionals and students	Chapters 3, 4 and 6

☐ **Pause and reflect.** A prompt, with room to write.

☐ **Try it.** A short exercise on a real task.

Source. Where each figure comes from.

Dr. Chime Ibikunle Peretz
MSc (Aberdeen), FCILRM, MAPM, APMP, MRICS

NICON Luxury Hotel, Abuja · Saturday 28 November 2026

01

CHAPTER 1

Guidelines won't save you

...and why that is not the last word.

Necessary. Not sufficient.

Most organisations answer AI risk with a policy. A document is written, circulated, signed off and filed. It answers one narrow question well: what is permitted. It rarely answers the question that decides what happens under pressure.

> *What will an employee do at 4pm on a Friday, with a deadline and a tool that promises to save them an hour?*

Policies are necessary, and a working set appears in Chapter 6. But a guideline that lives on a server has never once stopped an employee, mid-deadline, from doing the reasonable-sounding thing in front of them. What stops an AI failure from becoming an institutional failure is capability.

GUIDELINE

A document, written once and read rarely, that tells people what is permitted.

CAPABILITY

A trained reflex, built, tested and owned by a named person, that tells people what to do next.

PAUSE AND REFLECT

Which half of that do you find easier to say out loud in your institution: "AI is good" or "AI is risky"? What does that tell you about where conversations stall?

02

CHAPTER 2

The state of play

Four positions. None of them settled.

Four positions, visible at once

It is worth being precise about what each one is, rather than treating them as equal opinions.

BRUSSELS

The strictest rulebook, running late

The EU AI Act remains the toughest AI law in force anywhere: binding, with fines up to 7% of global turnover. In July 2026 the Digital Omnibus on AI pushed back the high-risk deadlines, citing competitiveness. Delayed, not dropped.

WASHINGTON

Move fast, and rename it

The US administration has moved against state AI laws and publicly played down safety concerns. On 19 September 2026 it ran a poll on X and Truth Social to rename “Artificial Intelligence”. “Superior Intelligence” won. Renaming took less debate than governing.

THE LABS

Agents that break their own cage

In July 2026 OpenAI disclosed that two of its models, during an internal cyber-capability test, escaped their sandbox and breached Hugging Face’s production systems. They were hunting for benchmark answer keys.

THE BUILDERS

Now saying it out loud

On 12 September 2026 Anthropic’s Dario Amodei published “We Must Pace the Frontier”, urging the industry to slow down. Sam Altman agreed the same day, Elon Musk reshared it with “Dario is right”, and Satya Nadella backed deliberate pacing the day after.

SO WHAT. Governments disagree, researchers are alarmed and builders are asking for brakes. Your organisation cannot wait for a consensus. It has to decide its own position.

SOURCE Lewis Silkin, 27 July 2026; Cloud Security Alliance research note, 30 July 2026; Amodei essay and coverage, 12–13 September 2026; press reports of the 19 September 2026 poll.

CHAPTER 2

What actually changed in Brussels

A widely repeated claim says Europe is loosening AI regulation. The record says something narrower.

24 Jul 2026	Regulation (EU) 2026/1744, the Digital Omnibus on AI, is published in the Official Journal.
27 Jul 2026	It enters into force.
2 Aug 2026	**The original date for high-risk (Annex III) obligations. Now deferred.**
2 Dec 2026	Grace period ends for labelling synthetic media on systems already on the market.
2 Aug 2027	National authorities must have at least one AI regulatory sandbox.
2 Dec 2027	**New date for high-risk (Annex III) obligations.**
2 Aug 2028	**New date for AI embedded in regulated products (Annex I).**

Read it correctly. The Act is still the strictest AI law in force anywhere. What moved is the calendar for the high-risk rules, plus some simplified reporting. Deferred, not cancelled. Even the toughest regulator on earth feels the commercial pull, and that is the point.

SOURCE Lewis Silkin, “The Digital Omnibus on AI enters into force today”, 27 July 2026; Cloud Security Alliance, “EU AI Act’s High-Risk Deadline: Deferred, Not Cancelled”.

CHAPTER 2

Where we actually stand

There is hardly an organisation left that is not already using AI. Banning it was never a realistic control. In many organisations the picture looks less like an AI programme and more like this: no approved platform, no budget set aside for one, and staff quietly using whichever free tool gets the work done fastest.

Strategy papers, staff records and customer data go straight into the prompt, on a platform nobody in the organisation reviewed.

63%

of Nigerian organisations are still only exploring or piloting AI.

17%

have actually scaled it into how the business runs.

SOURCE Boston Consulting Group, Nigeria Leadership Survey / Nigeria Executive Outlook, 2026.

> *If the people building this technology say it is moving faster than anyone can govern, an organisation with no AI position at all is already behind.*

Anthropic's CEO made that argument on 12 September 2026, in the essay that cited the July sandbox breach as one reason the industry should slow down.

PAUSE AND REFLECT

If you asked ten colleagues tomorrow which AI tools they used last week, how many different answers would you get, and how many would you have approved?

03

CHAPTER 3

What failure looks like

Eight failure modes, named plainly.

The eight failure modes

1 Data and confidentiality leakage

Company or customer data pasted into a public AI tool, outside any agreement governing how it is stored or used.

ASK What has been pasted into a free tool this week?

2 Confident, wrong outputs treated as fact

A fabricated figure, a non-existent regulation, an invented precedent, accepted because it sounded right.

ASK Who checks before it leaves the building?

3 Biased or discriminatory outcomes

An AI-assisted credit, hiring or risk-scoring decision repeats bias already present in its training data.

ASK Can we explain why it decided that?

4 Regulatory and compliance breach

AI use crosses an already regulated line, such as data protection or AML/KYC, before anyone flags it as an AI question.

ASK Which rules apply the moment AI touches this process?

5 Third-party and vendor exposure

AI built into a vendor's software creates exposure the organisation never approved and often cannot see.

ASK Which suppliers now use AI on our data?

6 Erosion of professional judgement

Staff defer to AI output without independent checking, so errors a trained professional once caught now pass through.

ASK When did anyone last override the tool?

7 Manipulated media and fraud

AI-generated audio, video and images are now good enough to defeat normal verification. See the Arup case, page 10.

ASK What would we do if the CFO appeared on video asking for a transfer?

8 Reputational exposure

One visible failure, a wrong chatbot answer or a public error, becomes a reputational event out of proportion to its scale.

ASK Who speaks, and who approves it, when it happens?

CHAPTER 3

This is already happening here

None of this is projected from somewhere else. In Nigeria, workers themselves say so.

72%

say they have relied on AI output at work without checking whether it was accurate.

61%

say AI has already caused them to make a mistake on the job.

SOURCE KPMG and University of Melbourne, *Trust, Attitudes and Use of Artificial Intelligence: Nigeria Insights*, 2025.

That is direct local evidence for failure mode 2 (confident, wrong outputs) and failure mode 6 (erosion of professional judgement). The workers surveyed are not careless. They are doing their jobs, under pressure, with a tool that makes the work quicker.

PAUSE AND REFLECT

If those two figures held in your institution, which processes would carry an unchecked AI output into a decision? Name one.

04

CHAPTER 4

Four cases

What stopped it, and what did not.

CASE 1

A chatbot said the ship carried nuclear weapons

WHEN	WHERE	STAKES
18 September 2026	United States / China	War

WHAT HAPPENED

A US Special Operations Command analyst asked an AI chatbot to assess a Chinese vessel's cargo. The chatbot blended open-source material with classified signals intelligence into a confident, fabricated report: the ship carried nuclear weapons components.

The analyst used AI again to package the finding as a standard intelligence report and distributed it. Fighter jets were in the air. Armed personnel were preparing to board.

WHAT STOPPED IT

Officials caught the error just before the operation launched. The report was, in one source's words, "entirely false", yet it "almost started a war" between the United States and China. A trained person reviewed the output before anyone acted on it.

WHAT IT TEACHES

It was not a guideline that stopped it. It was capability, proven in the highest-stakes setting imaginable. Note also how the error compounded: AI output fed into a second AI task, and the confidence travelled with it.

CAPABILITY QUESTION

Where in your own workflow does one AI output become the input to another, with no human check between?

SOURCE CNN exclusive, 18 September 2026, corroborated by Rolling Stone, Engadget and other outlets.

CASE 2

Verify before you transfer

WHEN	WHERE	LOSS
2024	Hong Kong	About $25 million

WHAT HAPPENED

A finance employee at the engineering firm Arup joined what looked like a routine video call with the company's CFO and several colleagues. Everyone looked and sounded right. During the call he was instructed to make transfers, and he made fifteen, totalling roughly $25 million.

Every other person on that call was an AI-generated deepfake. He was the only real participant.

WHAT FAILED

The normal check, seeing and hearing the person, no longer proved anything. The employee followed the process he knew, with the evidence he trusted, and still lost the money.

WHAT IT TEACHES

Manipulated media defeats verification by sight and sound. The control has to move to something that cannot be faked on a call: a separate channel, a call-back to a known number, a second approver who is not on the call.

CAPABILITY QUESTION

If the CFO appeared on video today and asked for an urgent transfer, what exactly would your team do, and has anyone tested it?

SOURCE Case background widely reported after Arup's confirmation of the incident, 2024.

CHAPTER 4

Case 3: verify what you rely on

The caution applies to what AI tools *tell* you, not only what they show you. A source that looks authoritative inside an AI-generated answer can trace back to something as unreliable as an unverified social media post. This is increasingly common in AI search summaries and "overview" features, which give an answer with confidence but not always with a credible source behind it.

1 **Ask**

Ask the AI tool where the information came from.

2 **Open**

Open that source yourself. Do not accept a summary of it.

3 **Decide**

If it cannot produce one, treat the answer as a guess, not a fact.

TRY IT

Take one AI answer you relied on this week. Run the three steps on it now. Record what you found.

Confident phrasing is not the same as a credible source. This is guideline 6 in Chapter 6, and the habit behind failure mode 2.

CASE 4

One employee, one prompt, one filing

WHEN	WHERE	TRIGGER
May 2026	Pennsylvania, USA	One pasted prompt

WHAT HAPPENED

An employee at a Pennsylvania community bank pasted non-public customer information, names, Social Security numbers and dates of birth, into an AI tool that IT had never approved. The employee had a valid login and legitimate access to the data. There was no hacker, no attack and no outage.

WHAT IT TRIGGERED

The bank caught the upload early and reached the AI vendor in time to stop the data being used to train its model. It still triggered two obligations within days: a material-event disclosure to the SEC under Item 1.05, and a 36-hour notice to the bank's prudential regulator. It was reported as the first public "shadow AI" cybersecurity filing of its kind.

WHAT IT TEACHES

The employee was not careless. They were doing the job, under normal pressure, with a tool that made the work faster. Nobody had made the boundary real enough, in a form the frontline could act on, for that moment to go differently. That is a capability gap, not a compliance failure.

CAPABILITY QUESTION

Could every person with data access in your institution state, in their own words, what must never go into an unapproved tool?

SOURCE American Banker; Wilson Sonsini Goodrich & Rosati; BankInfoSecurity, 2026.

05

CHAPTER 5

The evidence

Boards and the frontline are not looking at the same picture.

WHAT BOARDS ARE DOING

44%

of boards have no agreed position on AI at all.

2%

have a formal governance framework with clear accountability and regular reporting.

Board-level AI literacy also trails executive AI competence by a statistically significant margin in the first controlled academic study to measure the two directly. The people setting the rules understand the technology less well than the people running it.

> **SOURCE** BoardPro AI Governance Report, 2026; *Frontiers in Artificial Intelligence*, “Board-level AI literacy as a missing governance capability”, 2026.

CHAPTER 5

What the frontline is doing, and what it costs

THE FRONTLINE

84%

of organisations already use AI across the business.

33%

keep a formal, written employee AI-use policy.

67%

of executives believe their organisation has suffered a data breach or leak from an unapproved AI tool.

36%

have no formal plan for supervising what AI agents do inside the business.

SOURCE ISACA 2026 AI Pulse Poll (released 7 September 2026□; Writer Enterprise AI Adoption Survey, 2026.

THE COST OF NOBODY WATCHING THE GAP

43%

of breaches studied involved shadow AI, up from 20% a year earlier.

$5.39M

average cost where it was involved, against a global average of $4.99M.

68%

of breached organisations lacked the AI governance to manage or detect it.

19%

reported coordination between their governance and security teams.

SOURCE IBM *Cost of a Data Breach Report 2026* □602 breached organisations, Ponemon Institute).

CHAPTER 5

Nigeria in focus

The regulator is moving faster than most policies are reviewed. Nigeria ranks 38th of 135 on the Global Index on Responsible AI 2026, first in Africa: policy on paper, not yet safeguards in practice.

17

CBN regulatory actions in 14 months, covering cybersecurity, anti-money laundering and data protection.

6

compliance deadlines falling between March 2026 and March 2028.

₦15.42bn

sanction on a commercial bank in 2025: breaches now carry heavy costs.

SOURCE Adhere and TechCabal, “The Compliance Reckoning: Regulating Financial Services in the Age of AI”, 2026, as reported by ThisDay, 25 July 2026.

Confidence is not capability. KPMG found leader confidence in AI governance rose from 57% to 73% in one quarter, yet only 49% have defined where AI may not decide alone. A static policy reviewed yearly (Guideline 8) cannot keep pace.

PAUSE AND REFLECT

Which of your institution’s AI uses would a regulator ask about first? Who in your organisation could answer on the day?

06

CHAPTER 6

The floor: eight guidelines

A baseline every organisation needs. Not the finish line.

Guidelines 1 to 5

1

Know where AI is being used

Keep a live inventory of approved AI tools. Nothing enters production or touches customer or company data without a review first.

TEST Could you hand me the inventory today, including tools nobody approved?

2

Draw a hard line around sensitive data

Customer personal data, financial records and credentials never enter a public or consumer AI tool without an enterprise data-processing agreement in place.

> The line staff can remember: *nothing with a customer's name, a number, or a staff record goes into an AI tool this organisation has not approved.*

TEST Ask three junior staff for the rule. Do they give the same answer?

3

Keep a human in the loop for consequential decisions

Credit decisions, fraud flags, disciplinary and personnel matters: AI can assist, but a named person approves.

TEST Pick one consequential decision from this month. Who approved it, by name?

4

Name an owner, not a committee

Every AI system in use has one accountable individual who can be asked directly what is and is not approved.

TEST Can staff name the owner without looking it up?

5

Disclose AI use where it touches customers

Particularly in regulated financial services, which is the direction current CBN guidance is already moving.

TEST Where, exactly, would a customer be told that AI was involved?

CHAPTER 6

Guidelines 6 to 8

6

Verify before you trust an AI-sourced fact

Check what an AI answer, or the material behind a prompt, is actually based on. Confident phrasing is not a credible source. Ask where it came from, open the source yourself, and treat an answer with no source as a guess.

TEST Run the three-step check on one AI answer used this week.

7

Build the escalation path before you need it

A known channel, a defined response time, and a regulatory-notification checklist, ready in advance rather than improvised during an incident.

TEST Run a 30-minute drill: a prompt leaked customer data. Who is called first?

8

Revisit the guideline itself

Review at least annually. A regulator issuing new requirements every few weeks makes a static policy dangerous on its own.

TEST When was the guideline last changed, and what triggered the change?

Allowed is not the same as done

Guidelines	Capability
A document	**A trained reflex**
Written once, read rarely	**Built, tested, owned by name**
Tells people what is permitted	**Tells people what to do next**
Silent under real pressure	**Holds up under real pressure**

Grounded in the NIST AI Risk Management Framework, ISO/IEC 42001, and current Central Bank of Nigeria cyber and data-protection requirements. This is a baseline, not an exhaustive compliance programme. Legal and regulatory counsel should confirm what applies to your institution and licence class.

07

CHAPTER 7

The capability check

The boardroom-to-frontline AI check. Twenty minutes, inside your own institution.

This is not a scoring exercise. It is a way of finding out where your organisation's gap sits. Tick one box per question. Be honest, not generous.

A · BOARD AND GOVERNANCE

Has the board formally discussed AI risk in the past twelve months, with a minute recorded?	☐ YES	☐ UNSURE	☐ NO
Is there a named individual, not a committee, who owns AI risk and can be asked directly what is and is not approved?	☐ YES	☐ UNSURE	☐ NO
Would that individual know within a week if an unapproved AI tool were being used somewhere in the business?	☐ YES	☐ UNSURE	☐ NO

Run it three times. Ask a board member, a manager and a frontline employee to complete all nine questions separately, then compare. The differences between their answers are the finding.

CHAPTER 7

B · FRONTLINE AND DAILY USE

Can staff in customer-facing or data-handling roles state, in their own words, the one thing they must never put into an AI tool?	☐ YES	☐ UNSURE	☐ NO
Is there a clear, known channel for a staff member to ask “is this allowed?” and get an answer within the same working day?	☐ YES	☐ UNSURE	☐ NO
Has anyone tested, not assumed, what a frontline employee would do under a real deadline with a tool available that could help?	☐ YES	☐ UNSURE	☐ NO

C · ESCALATION AND RESPONSE

If an AI-related data exposure happened today, is it clear within the first hour who is told, and in what order?	☐ YES	☐ UNSURE	☐ NO
Has the organisation identified which regulatory notifications would be triggered, and within what timeframe, before an incident forces that research?	☐ YES	☐ UNSURE	☐ NO
Is there a mechanism for lessons from a near-miss to reach the board, not just IT or compliance?	☐ YES	☐ UNSURE	☐ NO

CHAPTER 7

Reading your pattern

Count the “No” and “Unsure” answers in each section.

MOSTLY IN A

The gap is at the top. Nobody owns AI risk by name, and the board is working from assurance, not evidence.

MOSTLY IN B

The gap is at the desk. The rules exist, but the people they are meant to protect cannot state them or have not been tested.

MOSTLY IN C

The gap is the connective tissue. Both ends may be sound, but when something goes wrong the path between them is improvised.

Most institutions find their gap is not where they expected.

Ask these on Monday

If nothing else from this session goes back into the organisation, these three questions are the minimum viable start.

1. Do we know where AI is already being used inside this organisation without our sign-off?
2. Can our frontline staff explain, in their own words, the one thing they must never put into an AI tool?
3. Does our board have someone who owns AI risk by name, not by committee?

A NOTE FOR EVERY LEVEL IN THE ROOM

The argument holds from every seat

If you sit on the board

Close the gap by naming an owner and asking for evidence, not assurances. The World Economic Forum is plain: a named executive must own the consequence, and agents are software to audit, not staff to trust.

If you manage a team

Close it by making the guideline concrete for your own people, in your own words.

If you are early in your career

Close it by asking, before using a new tool on real data: has anyone actually said this is allowed? Asked honestly and often, that question is most of what capability looks like.

SOURCES

- Regulation (EU) 2026/1744 ("Digital Omnibus on AI"), Official Journal, 24 July 2026; Lewis Silkin, 27 July 2026; Cloud Security Alliance research note on the high-risk deadline.
- Cloud Security Alliance Labs, research note on the OpenAI and Hugging Face sandbox incident, 30 July 2026.
- Amodei, "We Must Pace the Frontier", 12 September 2026, and coverage of responses, 12–13 September 2026.
- Boston Consulting Group, Nigeria Leadership Survey / Nigeria Executive Outlook, 2026.
- KPMG and University of Melbourne, *Trust, Attitudes and Use of AI: Nigeria Insights*, 2025.
- CNN exclusive, 18 September 2026 (US Special Operations Command incident), with Rolling Stone and Engadget.
- Arup incident, Hong Kong, 2024, as confirmed by the company and widely reported.
- American Banker; Wilson Sonsini Goodrich & Rosati; BankInfoSecurity, 2026 (Pennsylvania community bank).
- BoardPro AI Governance Report, 2026; *Frontiers in Artificial Intelligence*, board-level AI literacy study, 2026.
- ISACA 2026 AI Pulse Poll, 7 September 2026; Writer Enterprise AI Adoption Survey, 2026.
- IBM *Cost of a Data Breach Report 2026*, Ponemon Institute.
- PwC, *AI Performance Study*, April 2026; KPMG, *Q3 2026 AI Pulse Survey*, 24 September 2026.
- Global Index on Responsible AI 2026, via BCG Nigeria; World Economic Forum, April 2026; *Harvard Business Review*, 6 May 2026.
- Adhere and TechCabal, "The Compliance Reckoning", 2026, via ThisDay, 25 July 2026.
- NIST AI Risk Management Framework; ISO/IEC 42001.

ABOUT THE SPEAKER

Dr. Chime Ibikunle Peretz

FCILRM, MAPM, APMP, MRICS

Founder and Chief Performance Officer of PERETZ, an Abuja-based organisational capability and performance consultancy. His work centres on closing the gap between an organisation's stated strategy and the capabilities it actually has to deliver on it.

234 810 422 5521

PERETZ

Prepared for the 16th Annual Risk Professionals International Conference of the Chartered Institute of Loan & Risk Management of Nigeria. © 2026 The Peretz Advisory Limited, trading as PERETZ