HTech / National Enterprise AI Program™

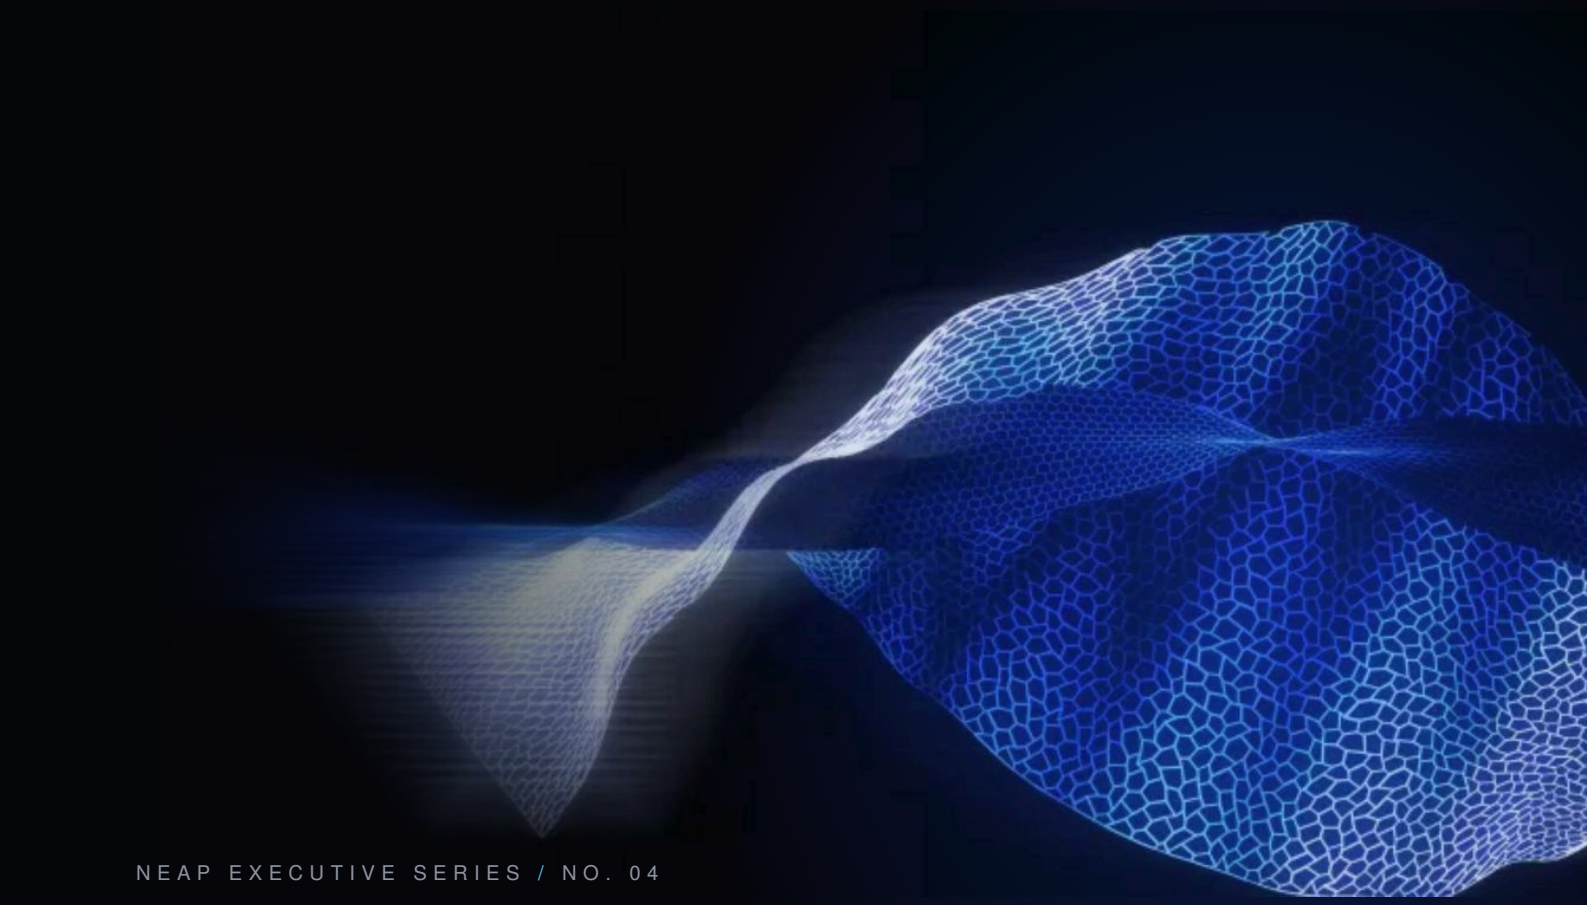

NEAP EXECUTIVE SERIES / NO. 04

Governing AI: rules that let you move faster

Responsible AI, data readiness and regulation for Lebanese enterprises, with a governance starter kit you can put in place in 90 days.

ABOUT THIS GUIDE

Good governance is not a brake on AI. It is what lets you press the accelerator.

Companies without AI rules do not avoid AI. Their staff use it anyway, with company data, in tools nobody approved. This guide explains the risks in plain terms, sets out what Lebanese and European rules mean for a Lebanese company, and gives a starter kit that a mid-sized firm can put in place within a quarter.

01	Why governance is a speed issue	Shadow AI and liability
02	The risk map	What can go wrong, and the controls
03	The governance starter kit	Seven components, sized for mid-market
04	Data governance for AI	Classify, then decide what goes where
05	The Lebanese legal landscape	Law 81/2018 in practice
06	The EU AI Act from Beirut	When it applies and what changed in 2026
07	Frameworks worth borrowing	NIST, ISO/IEC 42001, OWASP
08	When something goes wrong	An AI incident playbook
09	Your 90-day governance plan	And a one-page policy outline

READ THIS IF

You are accountable for risk, compliance, legal, IT or data, or you chair the committee that will approve AI cases.

TIME TO READ

About 45 minutes. Chapters 03, 08 and 09 are designed to be used directly.

NOT LEGAL ADVICE

This guide explains the landscape. Confirm how any law applies to your company with qualified counsel.

CHAPTER 01

Why governance is a speed issue

The choice is not between governed AI and no AI. It is between governed AI and shadow AI, and shadow AI is the more expensive of the two.

Employees adopt AI faster than their employers. MIT's 2025 research found that over 90% of employees in its sample used personal AI tools for work, while only about 40% of companies had bought official subscriptions.[1] When the approved route is slow or absent, company data goes into personal accounts with unknown terms.

63%

of breached organisations had no AI governance policy or were still developing one.

IBM, 2025

1 in 5

organisations reported a breach involving shadow AI, the use of unapproved AI tools.

IBM, 2025

$670K

Extra average breach cost at organisations with high levels of shadow AI.

IBM, 2025

IBM's 2025 study of 600 breached organisations also found that 13% had suffered breaches of AI models or applications, and that 97% of those lacked proper AI access controls. Among organisations that did have AI policies, only 34% audited regularly for unapproved AI use.[2]

You are responsible for what your AI says

In *Moffatt v. Air Canada* (2024), a Canadian tribunal held the airline liable after its website chatbot gave a customer wrong information about bereavement fares. The airline argued the chatbot was responsible for its own statements; the tribunal called this "a remarkable submission" and found the company had not taken reasonable care to ensure its chatbot was accurate.[3] The sums were small. The principle is not: an AI system that speaks to customers speaks for the company.

The governance gap is about to widen

BCG's 2026 research found that 42% of companies expect to give AI agents genuine decision-making authority by 2030, while only 5% currently have the full set of controls that would require.[4] As AI moves from drafting to acting, the cost of missing controls rises sharply.

THE CASE FOR ACTING NOW

A short policy, an approved tool and a simple review process take weeks to put in place. They remove the main reason staff use shadow AI, give you a record of what AI is doing, and make it possible to approve more ambitious cases with confidence.

CHAPTER 02

The risk map

AI risks fall into a handful of families. Each has known, proportionate controls. Name them, then size the controls to the use case.

RISK	WHAT IT LOOKS LIKE	PROPORTIONATE CONTROLS
Wrong answers	Confident but false output ("hallucination"); outdated policy quoted to a customer.	Ground answers in approved sources with citations; human review where errors are costly; evaluation sets; clear disclaimers.
Data leakage	Client or personal data pasted into public tools; an assistant revealing documents a user should not see.	Approved tools with no-training terms; data classification; access that mirrors existing permissions; logging.
Manipulation	Hidden instructions in an e-mail or web page hijack an AI assistant ("prompt injection").	Limit what AI can act on; separate reading from acting; approval for sensitive actions; security testing.
Unfair outcomes	Screening, pricing or credit decisions that disadvantage groups of people.	Avoid fully automated decisions about people; test outcomes by group; documented human review and appeal.
Excessive autonomy	An agent sends, pays, deletes or commits beyond what was intended.	Least-privilege permissions; amount and scope limits; approval steps; reversible actions; full audit trail.
Intellectual property	Confidential know-how exposed; generated content that reproduces others' work.	Contract terms on data use and ownership; review of public-facing content; internal guidance.
Cost runaway	Usage-based bills grow faster than value; loops in automated agents.	Budgets and alerts by tool and case; cost per case tracked; rate limits.
Dependence	Staff stop checking; knowledge lives in a vendor's system; a model change breaks a process.	Training on when not to trust AI; exportable configurations; notice of model changes; re-testing.

The OWASP security community's list of the top risks for applications built on large language models places prompt injection first and sensitive information disclosure second; its 2025 edition also names excessive agency, misinformation and unbounded consumption, the last covering runaway usage and cost.[5]

CHAPTER 03

The governance starter kit

Seven components give a mid-sized company proportionate AI governance. Most can be drafted in days; the value comes from using them consistently.

1 **An AI policy and acceptable-use rules.**
One page for all staff: which tools are approved, what data must never be entered, when output must be checked, who to ask. A longer policy for those who build or buy AI.

2 **An AI inventory.**
A living list of every AI tool and use case: owner, purpose, data used, users, vendor, risk tier and status. You cannot govern what you cannot see.

3 **Risk tiers.**
Three levels, each with required controls, so low-risk uses move fast and high-risk uses get proper scrutiny. See the table below.

4 **Clear roles.**
A business owner and a risk owner for each case, and a small AI review group (operations, IT or security, legal or compliance, and data) that meets monthly.

5 **A review process for new cases.**
A short intake form, a tier decision within a week, and a go-live checklist matched to the tier.

6 **Human oversight rules.**
Where a person must approve, what they check, and how they can stop or reverse an AI action.

7 **Monitoring and incidents.**
Quality sampling, usage and cost reports, and a simple way to report and learn from AI incidents and near misses.

Risk tiers and the controls they require

TIER	TYPICAL CASES	MINIMUM CONTROLS
Low	Drafting internal e-mails, summarising public information, coding help, brainstorming.	Approved tool; acceptable-use rules; no confidential or personal data unless the tool is approved for it.
Medium	Internal knowledge assistants, document extraction, drafting customer replies that a person sends, internal reporting.	All of the above, plus a named owner, inventory entry, access controls, evaluation before launch, human review of outputs that leave the company, periodic sampling.

High	Anything that decides about people (hiring, credit, pricing, claims), speaks to customers unsupervised, or acts in financial or core systems.	All of the above, plus formal review by the AI group, documented testing including fairness where relevant, legal and data protection sign-off, approval points, full audit logs, incident plan, periodic re-approval.

Proportion matters. Treating every use as high risk drives staff back to shadow AI. Treating every use as low risk invites the incidents in Chapter 01. Tiers let you be fast and careful at the same time.

INTAKE FORM FOR A NEW AI CASE: TEN QUESTIONS

1	What business problem does it solve, and who is the business owner?
2	Which tool or vendor, and under which contract terms?
3	What data will it read, and in which data class (Chapter 04)?
4	Does it process personal data? If so, is the processing covered by our declarations?
5	Who will use it, and who will see its outputs: staff only, or clients and the public?
6	Does it make or shape decisions about people, money or legal rights?
7	Can it take actions in other systems, and with what limits?
8	What is the worst realistic error, and how would we notice it?
9	Where does a person review or approve, and how is that recorded?
10	How will we measure quality, value and cost once it is live?

CHAPTER 04

Data governance for AI

The single most useful data decision is a simple classification that tells everyone which data may go into which kind of AI tool.

Data problems are the leading reason AI projects are abandoned: Gartner predicts that through 2026 organisations will abandon 60% of AI projects unsupported by AI-ready data.[6] For governance, the questions are narrower but just as important: what data do we hold, how sensitive is it, and where may it be processed?

Four classes, three kinds of tool

DATA CLASS	PUBLIC CONSUMER AI TOOLS	COMPANY-APPROVED AI UNDER BUSINESS TERMS	ON YOUR OWN INFRASTRUCTURE
Public Published material, marketing	Allowed	Allowed	Allowed
Internal Procedures, internal reports	Not allowed	Allowed	Allowed
Confidential Client files, contracts, finances	Not allowed	Allowed for approved cases, with access controls	Allowed
Restricted Personal data, health, banking secrecy, security	Not allowed	Only with legal review and specific approval	Preferred route, with legal review

An example policy. Adjust the classes and rules to your sector and contracts. "Business terms" means a contract that excludes training on your data and states where data is processed and how long it is kept.

Five practices that make data AI-ready and safe

- **Inventory by use case.** For each approved case, list the data it uses, the owner and the class.
- **Access that follows people.** AI assistants should only retrieve what the asking user is already allowed to see.
- **Quality at the source.** Fix capture where data is created, rather than cleaning it repeatedly downstream.
- **Retention and deletion.** Keep personal data only for the period you declared and need; delete AI logs on a schedule.
- **Digitise selectively.** Scan and structure the document sets your priority cases need, not the entire archive at once.

CHAPTER 05

The Lebanese legal landscape

Lebanon has no AI-specific law yet, but Law No. 81 of 2018 already governs much of what AI does with personal data, and a national AI agenda is taking shape.

Law No. 81 of 10 October 2018 on Electronic Transactions and Personal Data recognises electronic documents and signatures and, in its fifth title, sets rules for the collection and processing of personal data.[7] It does not create an independent data protection authority; instead the Ministry of Economy and Trade receives declarations of processing and issues licences where required.[7,8] Commentators note that the law contains no data breach notification provision and that its implementing regulations have not yet been issued, which leaves room for interpretation.[8,9]

WHAT THE LAW REQUIRES	WHAT IT MEANS FOR AN AI CASE
Purpose (Article 87)	Personal data must be collected in good faith for legitimate, specific and explicit purposes. Using customer data to train or prompt AI for a new purpose needs to be checked against why it was collected.
Information (Article 88)	People have a right to be informed about the processing of their data. Update privacy notices when AI processes customer or employee data.
Storage limits (Article 90)	Data may be kept only for the period stated in the declaration or authorisation. Apply this to AI logs, prompts and stored conversations too.
Security (Article 93)	Controllers must protect data against loss, damage and unauthorised access. This covers AI tools, their vendors and their logs.
Declaration (Articles 94–95)	Processing that is not exempt must be declared to the Ministry of Economy and Trade; certain categories require a licence. New AI processing of personal data may need a declaration or an update.

Practical steps for Lebanese companies

- Map which AI cases touch personal data, and check each against your existing declarations to the Ministry of Economy and Trade.
- Prefer AI vendors that commit in writing on data location, retention and no training on your data.
- Be especially careful with data covered by professional or banking secrecy and with health data; take legal advice before using it with any external AI service.
- Check sector-specific rules from your regulator, which may add requirements on outsourcing, cloud and data location.
- Keep records of your reasoning. Where the law is unclear, documented good faith and proportionate controls are your best position.

THE NATIONAL AGENDA

Lebanon's Ministry of State for Technology and AI launched LEAP, a five-year strategy whose goals include placing Lebanon among the world's top 50 AI users by 2030.[10] NEAP is designed to work alongside the national framework, focusing on the established private sector. Expect AI-specific guidance to follow; companies with basic governance in place will adapt most easily.

A personal data check for every AI case

- [ ] Which personal data does the case use, about whom (customers, employees, patients, suppliers), and from which systems?
- [ ] Is this use consistent with the purpose for which the data was originally collected, and with what people were told?
- [ ] Is the processing covered by an existing declaration or licence, or does it need a new one?
- [ ] Where will the data be processed and stored, including by the AI vendor and its subcontractors?
- [ ] How long will prompts, outputs and logs that contain personal data be kept, and who can see them?
- [ ] How would we respond if a person asked what data about them the AI used, or asked for it to be corrected?

CHAPTER 06

The EU AI Act from Beirut

Europe's AI law can reach Lebanese companies, through their European clients, their products and their recruitment tools. Its timetable changed in 2026.

The EU AI Act (Regulation (EU) 2024/1689) applies not only to companies in the EU. It also covers providers that place AI systems on the EU market, and providers and deployers outside the EU where the output of their AI system is used in the Union.[11] A Lebanese software company selling to European clients, a service provider whose AI-generated work is delivered to EU customers, or a group with European subsidiaries may therefore be in scope.

The timetable after the 2026 amendments

In July 2026 the EU adopted the Digital Omnibus on AI (Regulation (EU) 2026/1744), which entered into force on 27 July 2026 and moved several deadlines.[12]

DATE	WHAT APPLIES
February 2025	Bans on prohibited AI practices, such as manipulative techniques and social scoring.
August 2025	Obligations for providers of general-purpose AI models.
27 July 2026	Amended AI literacy duty: providers and deployers must take measures to support the AI literacy of staff who operate or use AI on their behalf, without guaranteeing any individual's level.[13]
2 December 2026	New bans on AI that generates non-consensual intimate imagery and child sexual abuse material.[12]
2 December 2027	Obligations for high-risk AI systems listed in Annex III, such as AI used in recruitment, creditworthiness assessment and some insurance pricing (moved from August 2026).[12]
2 August 2028	Obligations for high-risk AI embedded in products covered by EU product-safety law (Annex I).[12]

What a Lebanese company should do now

- Check whether any AI you provide or use produces output used in the EU, or is sold into the EU.
- If so, list those systems and check them against the prohibited practices and the high-risk categories.
- Use the time before December 2027 to build the documentation, testing and oversight high-risk systems will need.
- Treat AI literacy as good practice regardless: it is now an explicit European expectation and a core NEAP principle.

CHAPTER 07

Frameworks worth borrowing

You do not need to invent AI governance. Three widely used frameworks cover most of what a company needs, at whatever depth suits it.

RISK MANAGEMENT

NIST AI RMF

A free, voluntary framework from the US National Institute of Standards and Technology, organised in four functions: Govern, Map, Measure and Manage. Govern runs across the other three.[14]

Use it to structure your risk process and your review questions.

MANAGEMENT SYSTEM

ISO/IEC 42001

The first international standard for an AI management system, published in December 2023 and certifiable by independent bodies. It follows the same structure as ISO 9001 and ISO/IEC 27001.[15]

Use it to build on existing quality or security systems, or to show clients certified governance.

SECURITY

OWASP Top 10 for LLM applications

A community list of the most critical security risks in applications built on large language models, with mitigations. A separate list for agentic applications followed in December 2025.[5]

Use it to brief developers and test AI systems before launch.

How the frameworks fit the starter kit

STARTER KIT COMPONENT	NIST AI RMF	ISO/IEC 42001
Policy and roles	Govern	Leadership, AI policy, roles and responsibilities
Inventory and risk tiers	Map	AI risk assessment and system impact assessment
Testing and evaluation	Measure	Operational controls across the AI system life cycle
Oversight, monitoring, incidents	Manage	Performance evaluation, nonconformity and continual improvement

> **Our recommendation.** Start with the starter kit in Chapter 03 and use NIST's four functions as its backbone. Consider ISO/IEC 42001 certification when clients, regulators or tenders begin to ask for evidence, particularly if you already hold ISO 27001.

Questions a board should ask once a quarter

- [ ] How many AI systems are in our inventory, in which risk tiers, and which were added this quarter?
- [ ] What value has each medium- and high-tier case delivered against its target?
- [ ] Were there any AI incidents or near misses, and what changed as a result?
- [ ] What share of staff have completed AI training, and is shadow AI use falling?

CHAPTER 08

When something goes wrong

Every company using AI at scale will have incidents. What matters is noticing quickly, limiting the damage and learning from each one.

An AI incident is any event where an AI system causes, or nearly causes, harm: wrong information given to a customer, confidential data exposed, an unintended action in a business system, an unfair outcome, or costs that suddenly spike. Near misses count too; they are the cheapest lessons you will get. A culture where staff report mistakes without fear of blame is the most important control of all.

Six steps

1 **Detect.**
Reports from staff and customers, quality sampling, cost alerts and log monitoring. Make reporting a one-click action with a named contact.

2 **Contain.**
Pause the AI feature, switch to the manual fallback, revoke access or block the tool. Every medium- and high-tier system needs a tested way to stop it quickly.

3 **Assess.**
What happened, how many people or records were affected, whether personal or confidential data was involved, and whether it is still happening.

4 **Inform.**
Tell the business owner, risk owner and leadership. Decide with legal counsel whether clients, partners or authorities must be told under contracts, sector rules or applicable law.

5 **Fix.**
Correct the cause, whether data, prompt, permissions, model or process. Add the failing case to the evaluation set so it is tested every time.

6 **Learn.**
A short written review within two weeks: what happened, why, and what changes. Share the lesson, not the blame.

Severity levels

LEVEL	EXAMPLES	RESPONSE
Minor	Internal draft with errors caught by review; small cost overrun.	Log it; fix within the normal cycle.
Significant	Wrong information reached a client; internal data seen by the wrong team; repeated failures.	Contain the same day; owner and risk owner informed; review within two weeks.
Serious	Personal or client-confidential data exposed outside the company; an unintended financial or legal action; harm to a person.	Contain immediately; leadership and legal counsel engaged; external obligations assessed at once.

WHAT TO RECORD FOR EVERY INCIDENT

When and how detected	Date, time, who reported it, through which channel.
System and owner	Which AI system, its inventory entry, business and risk owners.
What happened	A plain description, with examples of the inputs and outputs involved.
Impact	People, records, clients and amounts affected; data classes involved.
Containment	What was done to stop it, and when.
Root cause	Data, prompt, permissions, model change, process or human factor.
External obligations	Contracts, regulators or individuals considered, and the decision taken with counsel.
Lessons and changes	What changed in the system, the controls or the training, and who checks it worked.

CHAPTER 09

Your 90-day governance plan

In one quarter, a mid-sized company can move from no rules to a working, proportionate governance system.

DAYS 1–30

Contain

Publish the one-page acceptable-use policy. Provide an approved AI tool under business terms. Start the AI inventory. Name the AI review group.

DAYS 31–60

Structure

Agree data classes and the tool matrix. Define risk tiers and their controls. Map personal data in AI cases against Law 81/2018 declarations.

DAYS 61–90

Operate

Run the first cases through the review process. Train staff. Start monthly reporting on inventory, incidents, usage and cost.

OUTLINE OF A ONE-PAGE ACCEPTABLE-USE POLICY

Purpose	We use AI to work better and faster, safely and lawfully.
Approved tools	The list of approved AI tools and how to request a new one.
Never enter	Restricted data in any tool that is not approved for it: personal, health, banking secrecy, security and client-confidential information.
Always check	Facts, figures, names and references before anything leaves the company. You remain responsible for the output.
Be open	Say when content sent to clients was substantially produced by AI, where it matters to them.
Decisions about people	AI may support, but not make, decisions about hiring, credit, pricing or claims without approved controls.
Report	Any mistake, leak or misuse involving AI, to a named contact, without blame for honest reports.
Owner and review	The policy owner, and the date of the next review.

REMEMBER

The aim is not to slow AI down. It is to make the safe route the easiest route, so your people use AI in the open, on your terms, and you can say yes to more ambitious cases.

SOURCES AND NOTES

Where the numbers come from

1. Challapally, A., Pease, C., Raskar, R. and Chari, P., *The GenAI Divide: State of AI in Business 2025*, MIT Project NANDA, July 2025; as reported by BankInfoSecurity and others.
2. IBM and Ponemon Institute, *Cost of a Data Breach Report 2025: The AI Oversight Gap*, and IBM press release, 30 July 2025 (600 organisations).
3. *Moffatt v. Air Canada*, 2024 BCCRT 149, British Columbia Civil Resolution Tribunal, 14 February 2024; summarised in ABA Business Law Today, February 2024.
4. Boston Consulting Group, *AI Is Starting to Pay Off* (Applied AI Index 2026), press release, 30 September 2026.
5. OWASP Gen AI Security Project, *OWASP Top 10 for LLM Applications 2025*, November 2024; OWASP Top 10 for Agentic Applications, December 2025.
6. Gartner, *Lack of AI-Ready Data Puts AI Projects at Risk*, 26 February 2025.
7. Lebanon, Law No. 81 of 10 October 2018 on Electronic Transactions and Personal Data, Title V; DataGuidance, *Lebanon – Data Protection Overview*.
8. DLA Piper, *Data Protection Laws of the World: Lebanon* (last reviewed 2022); Council of Europe, study on the conformity of personal data provisions with Convention 108+ (Lebanon), July 2020.
9. Legal Business, *The Electronic Transactions and Data Protection Law in Lebanon*; DataGuidance, *Lebanon: Data protection in the financial sector*.
10. L'Orient Today, reporting on the LEAP strategy of the Ministry of State for Technology and AI (Launch, Apply, Advance, Promote), 2025.
11. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Article 2 (scope).
12. Regulation (EU) 2026/1744 (Digital Omnibus on AI), published in the Official Journal on 24 July 2026, in force 27 July 2026; summaries by Orrick (29 July 2026), Gibson Dunn and CMS.
13. Article 4 of the AI Act as replaced by Regulation (EU) 2026/1744; summaries by IAPP, Akin and DLA Piper (2026).
14. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, 26 January 2023.
15. ISO/IEC 42001:2023, *Artificial intelligence — Management system*, December 2023; ISO, *ISO/IEC 42001 explained*.

This guide is published by HTech S.A.R.L. as part of the National Enterprise AI Program (NEAP), an HTech program. It is general information for business leaders, not legal, financial or regulatory advice; check how any law or standard applies to your company with qualified counsel. Figures are quoted from the sources listed, as published by them, and were checked in November 2026. You are welcome to share this guide unchanged and free of charge. © 2026 HTech S.A.R.L.

YOUR NEXT STEP

Put the guardrails in before the first case goes live.

Governance is cheapest when it is designed alongside the first use case, not bolted on after an incident. In a free debrief, HTech can review where your company stands on the starter kit in this guide and what to put in place first.

FREE · 3 MINUTES

Take the Readiness Check

Ten questions across the five pillars of the NEAP Readiness Index. You get your level, your weakest pillar and the one next step that fits.

FREE · 30 MINUTES

Book a debrief with HTech

Walk through your result with an HTech consultant and leave with a short list of where AI can create value in your company first.

Book a debrief

NEXT NEAP COHORT

Join the Foundation Workshop

One day for your leadership team. You leave with a full readiness profile and a commitment card for your first AI case.

joinneap.com

HTech S.A.R.L. · Berytech, Mathaf, Beirut, Lebanon
info@htechco.com · +961 81 733 831 · htechco.com · joinneap.com

HTech / National Enterprise AI Program™