EU AI ACT · REGULATORY BRIEFING · Regulatory Briefing · July 2026

High-risk classification: what MENA deployers must document

The EU AI Act does not stop at Europe's borders — but it does not blanket the Gulf either. For MENA organizations, the practical questions are two: am I actually in scope? — a question worth answering precisely, because the honest answer is often no — and, if so, what evidence must I be able to produce when a European counterparty, notified body or market surveillance authority asks?

This briefing answers both — updated for the new compliance timeline adopted in June 2026.

Who is actually in scope — and who is not

Let's be precise, because overstating the Act's reach serves no one. A Saudi organization serving the Saudi market, with AI systems whose outputs are used in the Kingdom, is in most cases simply not subject to the EU AI Act. Unlike the GDPR, whose extraterritorial reach is broad, the AI Act catches non-EU actors through two narrower doors: placing an AI system on the EU market, or operating a system whose output is used in the EU.

In practice, that translates into three profiles of MENA organizations with real exposure. First, regional groups with European operations — where an HR screening tool, a credit model or a routing system deployed group-wide also serves an EU entity. Second, and most commonly, Saudi subsidiaries of international groups: here the subsidiary is often not itself the regulated party, but must align with group-wide AI governance built for the Act — which means living with its documentation logic regardless. Third, MENA tech companies exporting to Europe — a small segment today, but the one growing fastest.

If your organization fits none of these profiles, the AI Act is a horizon issue, not a current obligation — and your compliance energy belongs with SDAIA and PDPL first. If it fits one of them, read on: the question becomes what evidence you must hold, and the answer just gained a new deadline.

First: the timeline just changed

In June 2026, the EU institutions formally adopted the Digital Omnibus on AI, the first amendment package to the AI Act. Its headline change is a deferral of the high-risk obligations: stand-alone high-risk systems under Annex III (recruitment, credit scoring, education, essential services, among others) must now comply by 2 December 2027, and high-risk AI embedded in regulated products under Annex I (medical devices, machinery, vehicles) by 2 August 2028.

Two things have not moved. The transparency obligations of Article 50 — informing people they are interacting with an AI system, disclosing AI-generated content — still apply from 2 August 2026, with the watermarking duties for systems already on the market following on 2 December 2026. And the Act's core architecture — its risk-based classification, its prohibited practices, its documentation logic

— is unchanged. The deferral buys time; it does not reduce what will be required.

The temptation for deployers is to read December 2027 as permission to pause. That reading fails for a simple reason: the hardest part of high-risk compliance is not filling documentation templates in 2027. It is knowing, today, which of your systems are high-risk at all — and that inventory work gets harder, not easier, as systems multiply.

What Annex III classification actually turns on

Annex III classifies systems by use, not by technology. The same underlying model can be out of scope in one deployment and high-risk in another. For MENA organizations, the categories that most often bite are employment (CV screening, candidate ranking, performance evaluation), access to essential private services (creditworthiness, insurance pricing), and education (admission, assessment). If your system influences decisions in these areas for people in the EU, assume you are in scope until a documented analysis says otherwise.

That documented analysis is itself the first artefact auditors expect: a classification assessment per system, recording the use case, the affected persons, the Annex III category considered, and the reasoning for the conclusion — including where the conclusion is “not high-risk.” Under the amended Act, providers who consider their Annex III-adjacent systems exempt must still register them in the EU database; an undocumented “we decided it doesn't apply” is not a position, it is an admission.

The evidence chain auditors look for

For systems that are high-risk, the Act's obligations translate into a chain of documents. In our audit practice, these are the artefacts whose absence is most consequential:

- **Technical documentation** describing the system's purpose, architecture at a functional level, training approach, and performance metrics — sufficient for a third party to assess conformity. A marketing datasheet does not meet this bar; an internal model card, properly maintained, usually can.
- **Data governance records:** where training and testing data came from, how it was assessed for relevance and representativeness, and what was done about identified bias. The amended Act clarifies the conditions under which sensitive personal data may be used for bias detection — but only within strict necessity, and only if the work is documented.
- **Risk management file:** identified risks, mitigation measures, residual risk, and the reasoning for accepting it — maintained as a living document, not a one-time exercise.
- **Logging and post-market monitoring:** the system must record events sufficient to reconstruct decisions, and the deployer must be able to show those logs are retained and reviewed. This is the artefact most often missing in practice, because it requires engineering work, not paperwork.
- **Human oversight design:** who can intervene, how, and what training they received. “A human reviews the output” is a claim; a documented oversight procedure with named roles is evidence.

What MENA deployers should do now

Between now and December 2027, three workstreams matter, in order. First, inventory and classify every AI system that touches EU persons or markets — everything else depends on this, and it is unaffected by the deferral. Second, close the documentation gaps on systems already classified high-risk, starting with technical documentation and logging, which take longest to retrofit. Third, fix the contract layer: procurement teams in Europe are already writing AI Act representations into vendor agreements, which means the commercial deadline arrives before the legal one.

One regional nuance deserves emphasis. Organizations operating in Saudi Arabia are not documenting for the AI Act in a vacuum: SDAIA's accountability expectations and PDPL's records-of-processing duties overlap substantially with the Act's documentation logic. A documentation program designed once, mapped to all applicable frameworks, costs marginally more than an EU-only exercise — and produces a compliance position instead of a compliance file.

The bottom line

The deferral to December 2027 changes when high-risk obligations bite, not what they require. For MENA deployers, the rational response is not relief but sequencing: classify now, document deliberately, and treat the extra time as what it is — the difference between a defensible position and a rushed one.

Oxon X audits AI systems against EU AI Act, SDAIA, PDPL and NCA requirements from a single methodology.

Schedule a compliance audit → contact@oxonx.sa