




DATA PRIVACY | REGULATORY

AI Quality Management under the EU AI Act: Principles, Obligations and the Links to GDPR and POPIA

22 October 2026 11 min read

by Ahmore Burger-Smidt, Director and Head of Regulatory, and Boitumelo Khwene, Candidate Attorney

"The EU AI Act asks a simple question of anyone building high-risk AI systems: can your organisation reliably produce safe, lawful systems, and can it prove this? A quality management system is how you answer."

Introduction

Good outcomes stem from good processes and quality cannot be tested into a product at the end; it has to be built in from the start. This seems to be ethos behind the EU's AI Act (**the AI Act**) and how we ought to consider artificial intelligence (**AI**) systems in future.

What does this mean for providers of high-risk AI systems? Well, the organisations that develop high-risk AI systems and place them on the market must now prioritise a quality management system (**QMS**). A QMS is a written set of policies, procedures and instructions showing how the organisation designs, tests, monitors and fixes its AI systems.

None of this is entirely new. Many organisations already run quality systems under ISO 9001. Privacy law also takes the same approach. The GDPR and POPIA both rest on accountability, this demands that organisations must take responsibility for compliance and be able to prove it. An AI QMS builds on these foundations rather than replacing them.

The AI Act also reaches beyond Europe. It applies to providers outside the EU whose high-risk systems are sold in the EU, or whose outputs are used there. South African developers with European clients must therefore take note and stay alert.

The EU AI Act's QMS Requirements

The law wants organisations, not just products, to be trustworthy. One compliant AI system proves little, what matters most is whether an organisation has a dependable way of producing compliant systems every time.

The QMS duty sits in Article 17 of the AI Act. It is often confused as a duty to manage risk, but risk management is only one part of the wider system. A QMS must cover the whole life of an AI system: a plan for meeting the rules, how the system is designed and tested, how data is collected and handled, how risks are identified and reduced, and who is responsible for each step.

Through the QMS, an organisation must also show that it has satisfied the AI Act's other core requirements. These include good-quality, representative data; clear technical records; automatic logs of how the system operates; meaningful human oversight; and accuracy, robustness and cybersecurity. Once an AI system is in use, an organisation must constantly assess its performance (known as post-market monitoring) and report serious incidents to the relevant authorities.

How does the AI Act align with the GDPR and POPIA? The requirements in the AI Act mirror familiar privacy principles embedded in POPIA and the GDPR. For instance, good data governance reflects purpose limitation and data minimisation, this translate to the idea that data must only be used for a defined purpose, and should not be collected unnecessarily. From a POPIA perspective, the AI Act expresses the same ideas through its conditions of processing limitation, purpose specification and information quality. Technical records and logs provides for accountability under both laws, meaning that an organisation will be able to show, not merely say, that it complied.

What a Compliant AI QMS Looks Like

A good QMS is a living cycle, not a file on a shelf. It follows a simple loop: plan what you will do, do it, check whether it worked, and act on what you find. In this essence, the implementation of an QMS must be integrated throughout an organisation and not be treated as another compliance tick-box exercise.

ISO/IEC 42001, an international standard for managing AI across an organisation, adds an AI-specific layer and can be certified. Although certification serves as useful evidence of compliance, it does not by itself prove compliance with the AI Act.

The standard most closely tied to the AI Act is EN 18286, a European standard written specifically for the AI Act's QMS duty. The EU published its reference in the Official Journal on 19 October 2026. Following a standard published in this way (a "harmonised standard") creates a legal presumption that the requirements of the AI Act are fulfilled.

A third standard, ISO/IEC 27090:2026, published on 14 October 2026, addresses a different question: how to protect AI systems from attack. It gives guidance on threats such as data poisoning (corrupting training data), model theft and prompt injection (tricking a system with hidden instructions). At this stage, it merely provides guidance and also serves as a useful reference on how organisations can accede to the AI Act's cybersecurity requirement.

The GDPR requires data protection by design: building privacy into systems from the outset. POPIA requires responsible parties to ensure its conditions are met both when they decide how processing will work and while it takes place. It therefore important to understand that a well-built AI QMS ensures compliance with the GDPR and POPIA if done correctly.

Legal and Regulatory Implications

Rules without consequences do not change behaviour. The AI Act therefore gives regulators real powers, this shifts the spotlight on organisations as they will be forced to prove that they acted responsibly should they find themselves in hot waters with regulators.

Failing to run a proper QMS can lead to fines of up to EUR 15 million or 3% of worldwide annual turnover, whichever is higher.

Fines are not the only risk. Regulators can require changes, restrict sales, or order a system to be withdrawn or recalled, meaning that a system cannot stay on the market if it is a commercial problem, well before it becomes a financial one.

Regulators will ask three questions. Is there a written procedure for each requirement? Is it followed? Can the records prove it? After an incident, the question will focus less on whether the system made a mistake and more on whether the organisation had a way of noticing and if it was reported to the relevant authorities.

South Africa's Information Regulator follows a similar logic. It can issue enforcement notices and, where these are ignored, impose administrative fines of up to R10 million. Enforcement is increasing. The Regulator has also begun asking organisations to demonstrate compliance before any complaint is made, and is considering proposals for immediate fines similar to the GDPR model. The message is the same as under the AI Act: be able to prove what you have done.

Specific Considerations

AI rules do not replace existing rules; they sit alongside them. The aim is to fit the AI QMS into the existing governance and not build a separate structure. Therefore the AI Act must be read in harmony with the GDPR and POPIA, for local businesses.

Data Privacy and Cybersecurity

The AI Act's data and cybersecurity requirements overlap heavily with existing privacy law. Under the GDPR, organisations must assess high-risk processing in advance, build in protection by design, and protect people from decisions made solely by automated means. POPIA has a close equivalent. People may not be subject to a decision with legal or substantial effects based solely on automated profiling unless an exception applies. Where one does, they must be able to make representations and be given enough information about the logic of the processing to do so.

POPIA also requires appropriate security safeguards, notification of security compromises to the Information Regulator and affected people, and a personal information impact assessment. The EU's NIS2 cybersecurity rules add further duties, and ISO/IEC 27090:2026 offers a shared reference for AI-specific threats. A poisoned training dataset can be a security incident, a data quality failure, and possibly a personal data breach, so one well-designed assessment can often cover all three regimes.

But what are the practical implications of this AI Act?

Organisations that cope best with new regulations tend to share three habits. They start early, they build on what they already have, and they keep the evidence to maintain a record. Fortunately, the AI Act caters for all three habits.

Any programme must begin with the correct timetable. The Digital Omnibus on AI moved the QMS deadlines, which will apply from 2 December 2027 for stand-alone high-risk systems and from 2 August 2028 for high-risk AI built into regulated products. These dates are fixed and do not depend on standards being ready. It is equally important to remember what has not moved: the bans on prohibited AI practices, these have applied since February 2025, and the penalty rules have been enforce since August 2025.

With the timetable settled, the first practical question is: **what** does an organisation actually have. Very few organisations keep a complete list of the AI systems they develop or use, and fewer still have assessed which are high-risk. This exercise matters as it informs the an organisation's obligations. A provider carries the QMS duty while a user of someone else's system carries lighter obligations. However, the line can shift as an organisation undergoes significant changes or rebranding, a system can turn its user into its provider. An inventory that records both risk classification and role is therefore the foundation on which everything else rests.

The next step is to measure **existing** governance against the Act's requirements, rather than starting afresh. Most regulated organisations already run quality, risk and privacy frameworks, and the gap is often narrower than it first appears. EN 18286 provides the natural benchmark because it was written for the Act's QMS duty, and any decision to depart from it should be recorded with reasons that are well-substantiated. ISO/IEC 27090:2026 can serve as a reference point for AI-specific security controls, however, its role as guideline should not be understood as a route to presumed compliance. Privacy compliance belongs in the same exercise and AI assessments can often be combined with GDPR data protection impact assessments and POPIA personal information impact assessments. Moreover involving data protection officers or POPIA information officers from the outset avoids parallel processes that reach inconsistent conclusions.

A system is only as strong as the people accountable for it. Senior management should hold clear responsibility for the QMS, with named owners for each part, so regulators and notified bodies can see who answers for what.

The final consideration is to look outward, to incidents and to suppliers. Monitoring and incident-reporting plans are of little value if they are first tested during a real incident, so organisations should rehearse them in advance, including the timelines for notifying

regulators under both the AI Act and POPIA. Much AI is also bought rather than built, which means compliance depends on the supply chain. Contracts should therefore pass obligations down to suppliers through warranties, audit rights and access to documentation. An organisation cannot evidence a QMS for a system it cannot see inside.

Conclusion

The AI Act asks an organisational question before a technical one: can this organisation reliably produce trustworthy AI output, and can it prove this? The QMS is where that question is answered. The underlying ideas will be familiar to anyone who works with the GDPR or POPIA: take responsibility, design a compliance system and keep the records to prove it.

What this means for South African companies

The AI Act is not South African law, but it will most certainly shape South African practice. Three groups will be most affected. The first is South African developers that sell or license high-risk AI to customers in the EU. The second is companies whose AI outputs are used in the EU, even where the system itself never leaves South Africa. The third is local subsidiaries of multinational groups, who are usually expected to follow group-wide AI governance built to EU requirements.

The central lesson here is that AI governance has become an organisational discipline, not just a technical add-on. The QMS model requires written procedures, named owners, and records that prove compliance with the model. Standards give that discipline its shape and provide guidance on how to incorporate AI governance within an organisation. EN 18286 is a European instrument, but ISO/IEC 42001 and ISO/IEC 27090:2026 are international and can be adopted anywhere.

South Africa is not starting from nothing. POPIA already requires accountability, purpose specification, information quality and security safeguards, and restricts decisions based solely on automated profiling. The Information Regulator is enforcing it with growing confidence. Banks, insurers and energy operators also work within established governance and risk regimes. The gap is one of scope. However, POPIA protects personal information; it does not regulate AI's safety, robustness, or human oversight.

Nonetheless, for local businesses with European clientele, early alignment is a commercial advantage for those businesses. European customers will expect suppliers to evidence an AI QMS, and procurement contracts are likely to demand it well before any South African law does. Companies that can show alignment will find it easier to win and keep attracting business from that part of the world. They will also be better placed when domestic policy arrives, since it is likely to draw on the same international standards that are already established.



Ahmore Burger-Smidt

Head of Regulatory

[View Profile](#)