OrbusSoftware

EU AI Act Compliance Guide and Where EA is the **Critical Enabler**

EU AI Act Compliance Guide and Where EA is the Critical Enabler

The EU AI Act is raising the bar for organizations, setting new expectations for oversight, accountability, and risk management. Leaders need to know what systems exist, how they're classified, who's accountable for them, and whether the right controls are in place. This is difficult because AI shows up across the enterprise in ways that aren't always visible from the top.

Enterprise architecture is what makes this manageable. It helps leaders see how AI systems connect to the data they use, the people responsible for them, the risks they carry, and the parts of the business they impact, turning a piecemeal set of tools into something that can be governed. If your organization needs a fuller picture of its AI use, this EU AI Act compliance guide will provide clarity.

EU AI Act compliance starts with visibility.

What the EU AI Act requires organizations to understand

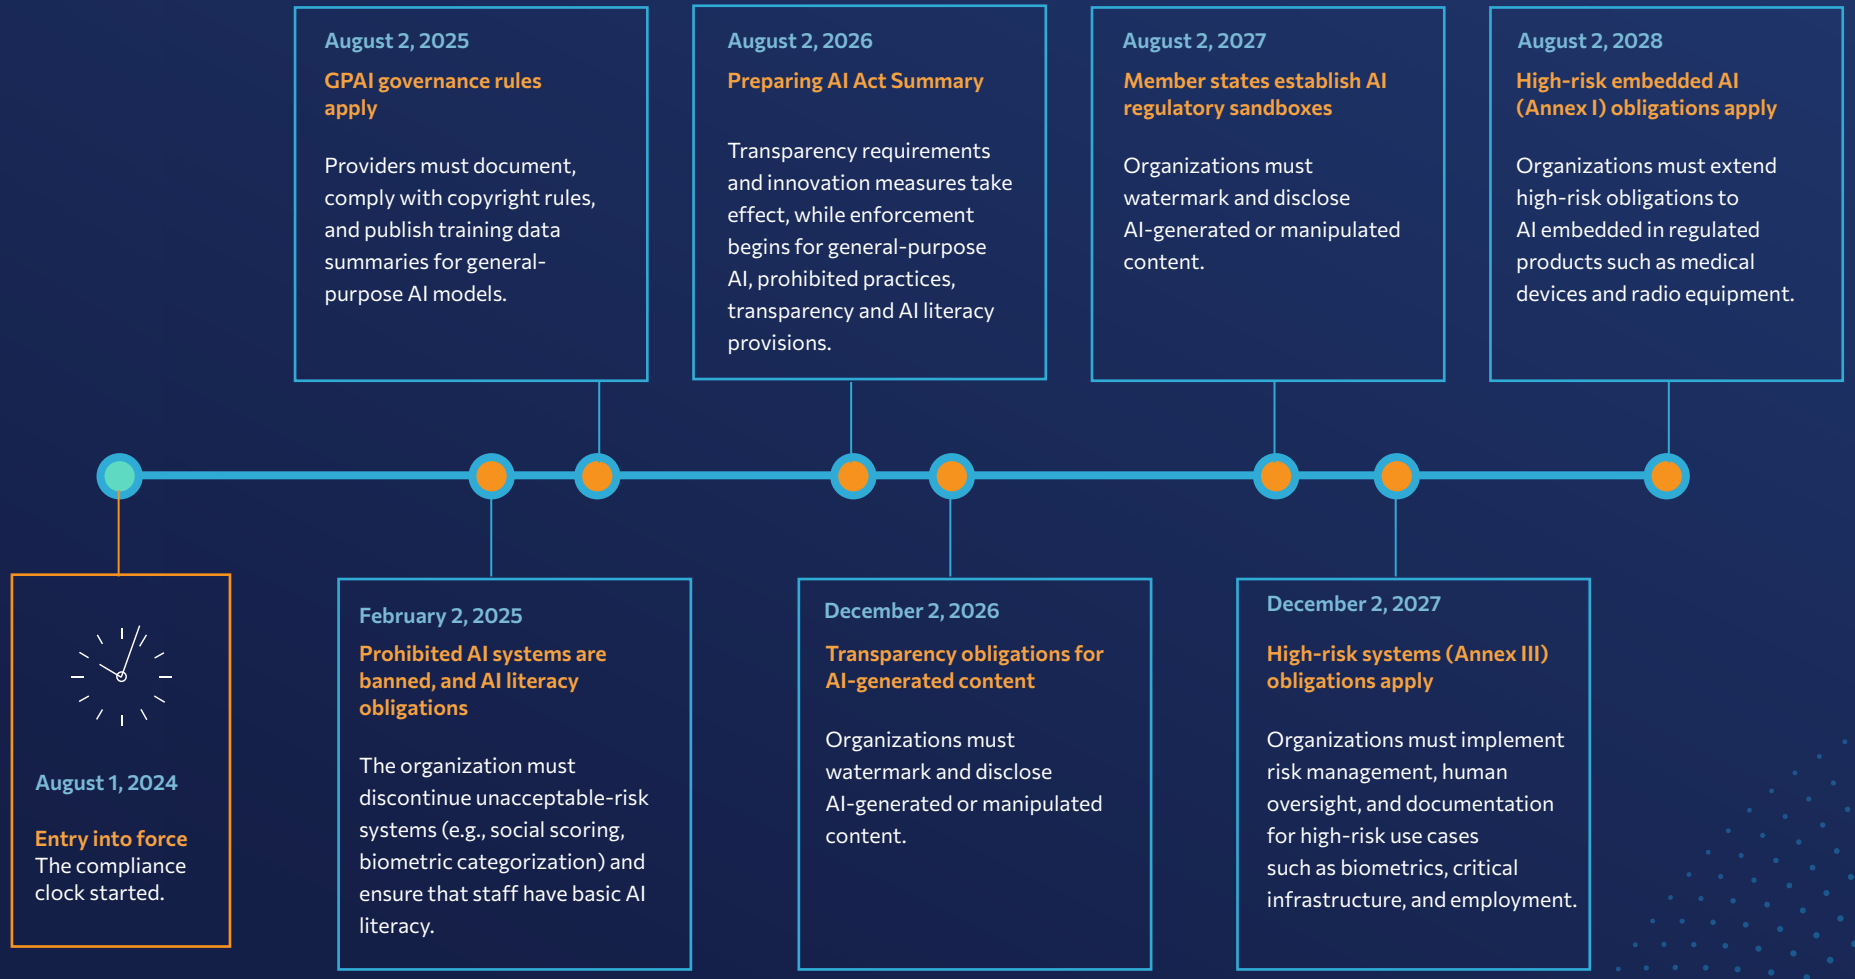

Current as of August 2026 : Regulatory dates and obligations may change. Always confirm the latest requirements with qualified legal or compliance advisers.

Your enterprise's obligation varies by role. Providers who build or supply AI are subject to documentation, risk management, and oversight requirements. Deployers, those who use AI professionally, also have obligations, but far fewer than providers. And don't make the mistake of thinking your company isn't impacted; the third-party tools embedded in everyday platforms mean that most enterprises would be classified as deployers, even if they don't see themselves that way.

Penalties can reach a significant percentage of global annual turnover, making AI a board-level concern with real business impact, rivaling data privacy.

But you can't take action on any of this if you don't know what you have: which AI systems exist, how your organization uses them, who owns them, and what risks they create. Developing that inventory is step one before beginning any risk classification or compliance work.

The EU AI Act turns AI oversight into an enterprise accountability issue.

> "98% of CIOs surveyed lack visibility into the technical and business risks of AI, creating blind spots that undermine both confidence and control."

- Steve Fulton, CEO, Orbus Software

The EU AI Act uses a risk-based approach to AI regulation, categorizing risks into four tiers. Your organization must determine where each AI system falls on this spectrum.

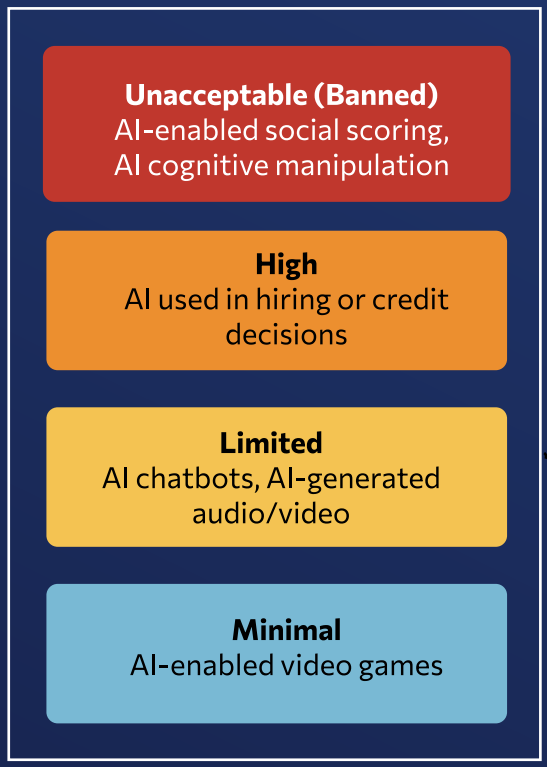

AI Act compliance starts with enterprise AI visibility

Before you can determine whether one of your systems is high-risk, limited-risk, or minimal-risk, you need a complete, accurate picture of every AI system operating across your enterprise. And if your organization is like most, you don't have one. Shadow AI tools adopted by individual teams, disconnected pilots that never made it into formal governance, and vendor-embedded AI features nobody flagged as AI all create blind spots that lead to compliance exposure. A tool no one tracked is still subject to the Act if it's in use.

Here's where enterprise architecture supports AI Act readiness. EA practices turn a scattered list of tools into a governed, classifiable AI system inventory.

> "78% of CIOs struggle with shadow AI, driven by unsanctioned tools, unclear usage, and weak enforcement mechanisms."
>
> —Orbus Software, The Global CIO Report

You can't classify, govern, or audit AI systems you can't see.

AI classification needs business and technology context

AI classification depends on more than just the model you use; it also depends on the business context in which you're using it. Two organizations could deploy the identical AI model and end up in completely different risk tiers under the EU AI Act because it isn't just the algorithm that matters. It's how they use the system, who it affects, and what it's connected to.

The problem is that these inputs rarely live in one place.

- IT knows the technical architecture
- Legal and compliance understand regulatory exposure
- Security tracks data flows and access
- Business teams understand the processes AI touches and who's impacted when it makes a decision

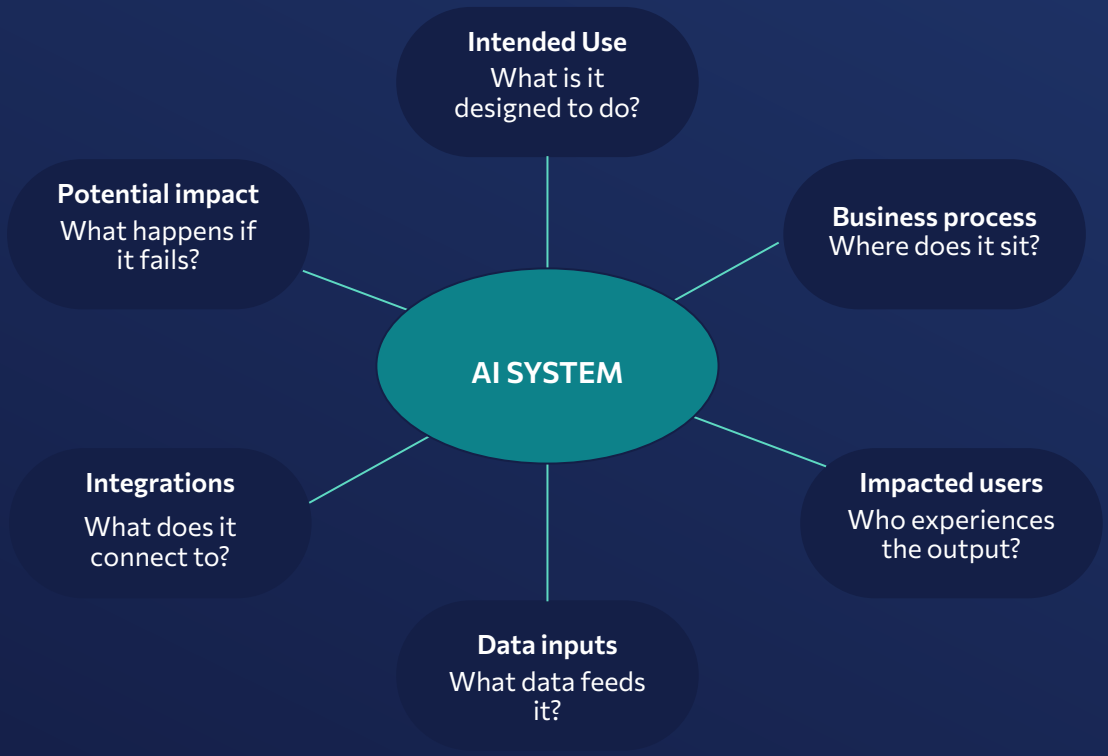

To accurately classify a system, you have to consider all of these perspectives. Enterprise architecture can help facilitate that by connecting fragmented technical, regulatory, and operational views into a single decision-making context. By mapping AI systems to business capabilities and value streams, EA makes it possible to see what a system does technically, where it sits in the business, and why that matters for risk.

Why business architecture matters.

Mapping an AI system to the business capability it supports, the value stream it belongs to, and the customer journeys it touches shows where your real risk is, not just what the system does technically. A recommendation engine used for marketing has different stakes than the same engine being used to help approve credit or guide patient care, even if the underlying model is identical.

Business architecture also aligns AI with the business value and expected ROI. This allows leaders to weigh risks against potential returns rather than addressing each AI initiative in isolation.

> “82% of CIOs rely on enterprise architects to identify and assess AI-related risk.”
>
> - Orbus Software, The Global CIO Report

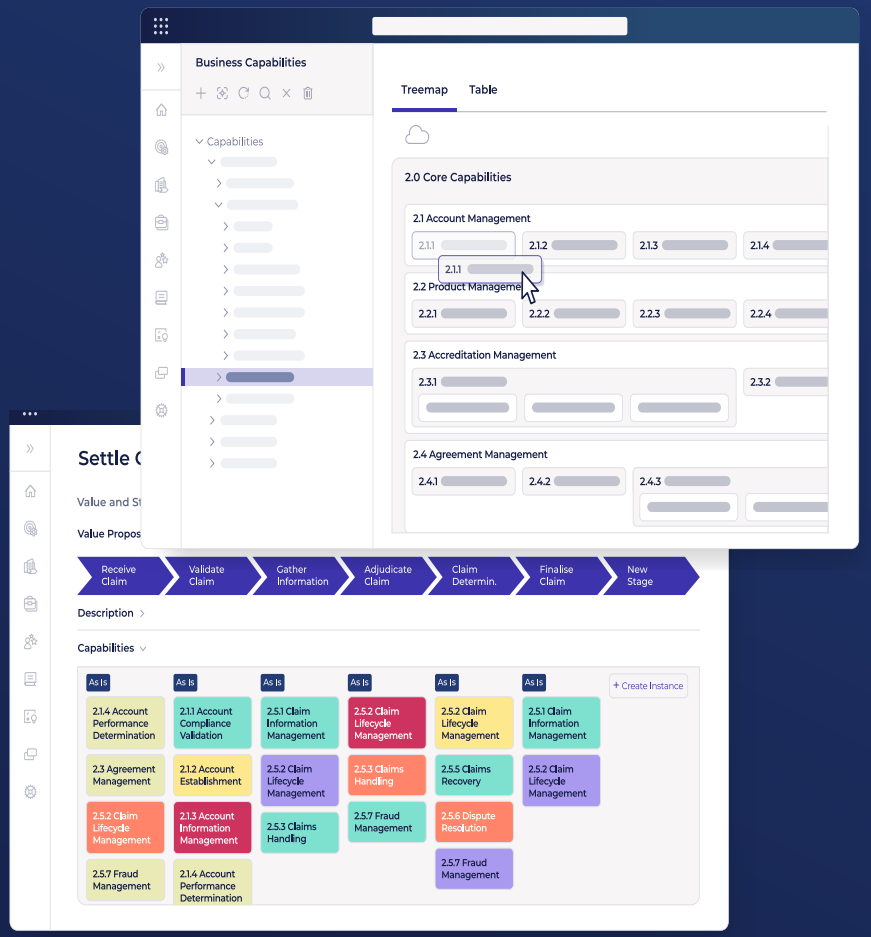

AI classification depends on business context, not the model alone.

Architectural governance turns compliance into an operating model

Policies define expectations, and governance architecture makes them repeatable. A written AI policy can state exactly what your enterprise requires. But if no one owns enforcement, review paths, or controls, the policy accomplishes little. Compliance breaks down because the organization lacks the structure to consistently translate policy into practice.

That's what AI governance architecture provides: a connective layer linking policy requirements to the systems, roles, workflows, and evidence needed to satisfy them. It's the difference between simply saying "Our high-risk systems require human oversight" and having a defined process that assigns an owner, triggers a review, and captures documentation showing that oversight occurred.

Enterprise architects can take the lead on building this layer, defining:

- **Decision rights:** Who approves what
- **Control points:** Where checks happen in a system's lifecycle
- **Escalation paths:** What happens when a system shows signs of a problem
- **Review triggers:** What prompts the organization to reassess as a system changes

Because the AI Act's obligations scale with risk, your governance layer should also apply oversight proportionally; in other words, your customer service chatbot doesn't need the same scrutiny as a hiring algorithm.

As your AI systems grow, manual review will struggle to keep pace. Automated governance that flags new systems, tracks classification status, and identifies systems that need review can help your organization stay compliance-ready.

Policies define expectations. Governance architecture makes them repeatable.

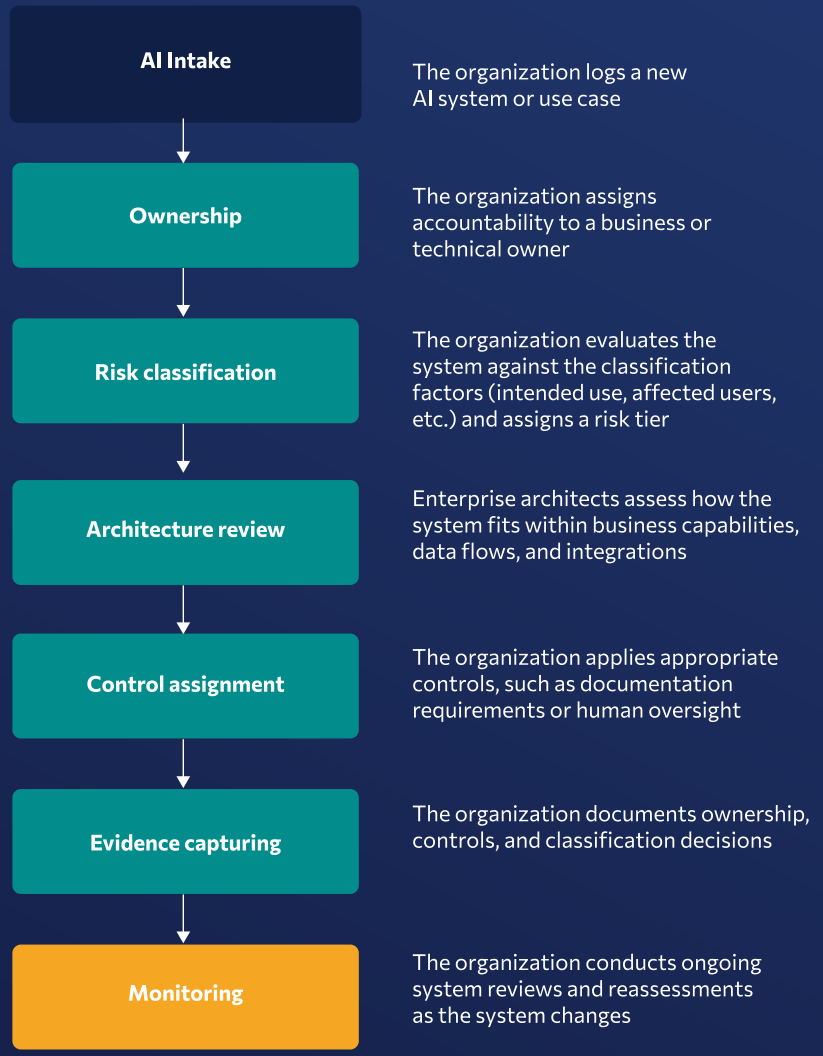

Enterprise architecture helps build compliance evidence

Meeting AI Act requirements isn't a one-time activity that you complete and move on from; regulators and auditors will want to see that your AI systems are understood, governed, and monitored. And that's difficult to do on short notice if the information is scattered in spreadsheets, old emails, and institutional memory.

An enterprise architecture repository solves the problem of disconnected information. Rather than searching multiple sources to determine who owns a system, what data it touches, and what controls sit around it, your organization's information lives in one place, kept current as part of how you're already tracking your systems.

When you link your organization's AI systems to the business capabilities and processes they support, leaders can more easily understand (and explain) how the organization's AI use aligns with business purpose, risk controls, and ROI. That's a far more credible position during an audit than handing over a static document.

> "80% of CIOs struggle to bridge technical possibilities with business priorities, and over half of projects end with unclear business value."

Orbus Software, The Global CIO Report

A Digital Twin of an Organization directly supports this. You don't want compliance evidence that's accurate today and wrong in three months. A DTO provides an exact, current-state model of the enterprise, in which AI systems, owners, dependencies, and controls update as the business changes, rather than going stale the moment something shifts.

Live operational data makes this possible. Real-time KPIs and business metrics feed directly into the model, so when a process changes or a data source shifts, the change appears immediately rather than being discovered only in an audit. That means you have compliance evidence ready whenever you need it; it isn't something you have to rebuild when an audit arises.

Compliance evidence needs to stay connected as AI systems change.

To get your enterprise ready for EU AI Act compliance

Use these six actions to move from awareness to an operating model

Inventory AI Systems

Build or update a complete list of AI tools and use cases

Assign ownership

Give every system a business/ technical owner

Map business context

Connect each system to the business capabilities, processes, and data it touches

Define review paths

Set risk-based review requirements so oversight matches the stakes

Connect controls to evidence

Document and monitor controls so you can prove compliance, not just claim it

Use EA to align compliance and business change

Keep pace with evolving AI systems and regulations so your compliance stays relevant

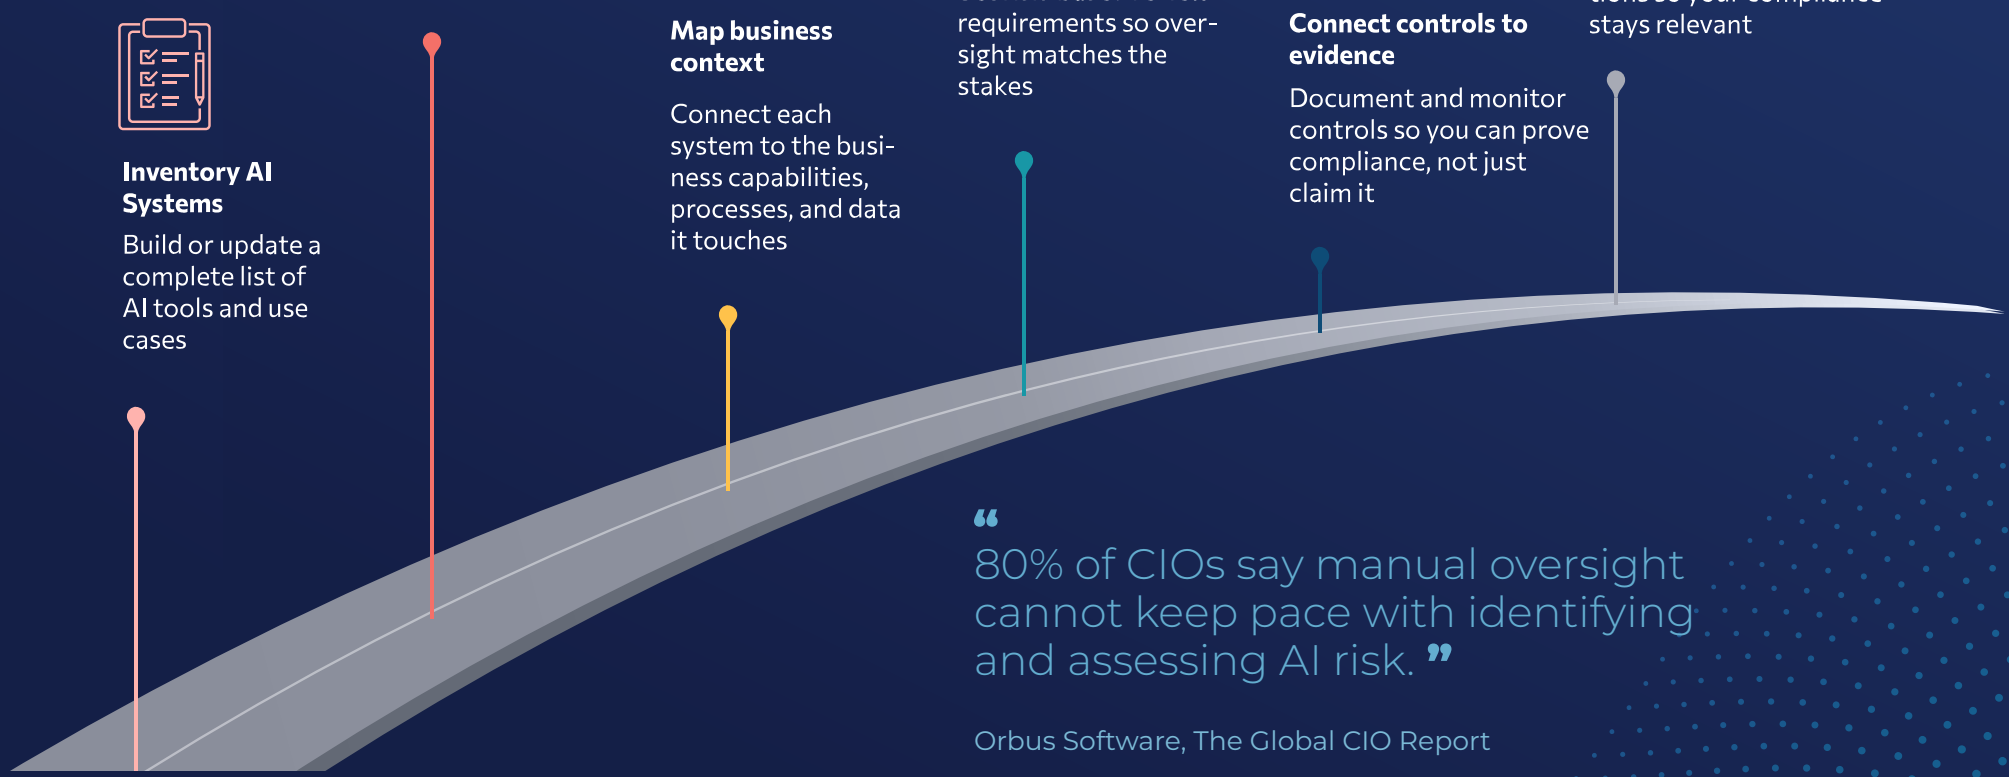

> “80% of CIOs say manual oversight cannot keep pace with identifying and assessing AI risk.”
>
> Orbus Software, The Global CIO Report

AI in Enterprise architecture makes EU AI Act compliance actionable

EU AI Act compliance comes down to three things, none of which happen through policy documents alone:

- Visibility into where AI exists
- Context to accurately classify it
- Governance that's repeatable enough to hold up over time and context
- Evidence that remains current as systems change

Enterprise architecture gives your organization the connected view it needs to effectively and responsibly classify, govern, and monitor AI, turning compliance from an ad hoc effort into a manageable ongoing process.

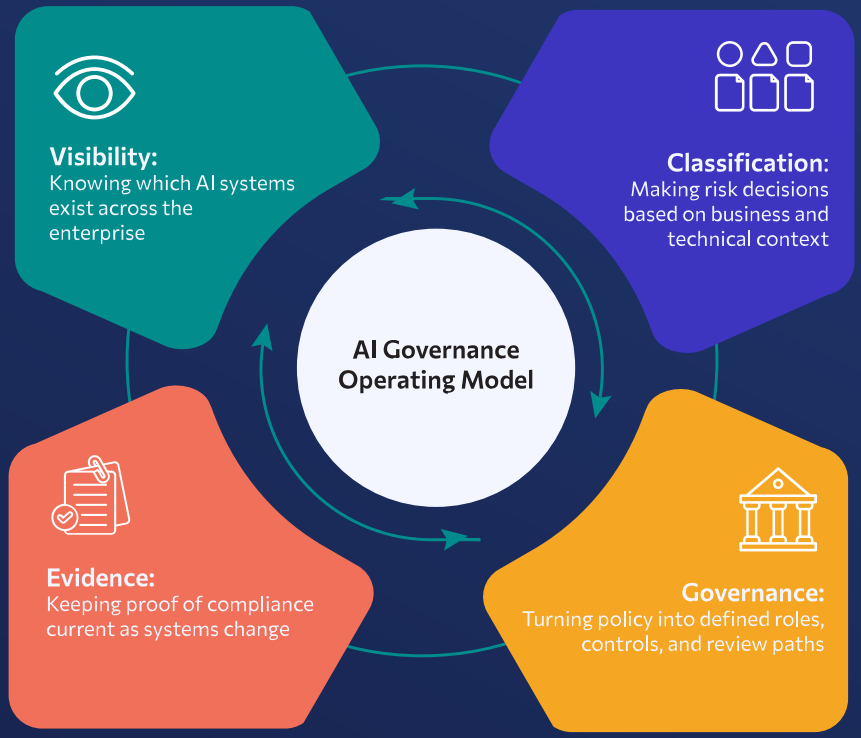

Book a demo to see how OrbusInfinity can help you build a connected AI inventory and governance model.

BOOK A DEMO

About Orbus Software

Orbus Software is a leading global provider of enterprise transformation solutions. We aim to empower customers with a strategic decision-making platform to successfully manage complex change. Our OrbusInfinity platform enables leaders to deliver business objectives, innovate faster, and ensure enterprise resiliency, while supporting them to make more informed, responsible, and sustainable business decisions.

orbussoftware.com

EBK0114/0826