THE AI OPERATING SYSTEM

AIOS

The Complete AI Operating System Framework

LEVEL 1 — FOUNDATION & MANIFESTO

A Successor Framework to Zachman, TOGAF,
and the Gartner Enterprise Architecture Approach

FREE EDITION — AIOS.INSTITUTE

FLORIAN KRUEGER • AIOS INSTITUTE • EUROPE-FIRST EDITION • 2026

Contents

Preface: Who This Is For

This document is Level 1 of the AIOS Framework — the Foundation and Manifesto upon which every subsequent level is built. It is written for a specific reader: the strategic planner, the CEO, the COO, the Chief Digital Officer, the Board member who must make consequential decisions about AI adoption before architects, consultants, or technologists are commissioned to execute.

If you are a consultant or enterprise architect, this document is for you too — but read it through the eyes of the person who will commission your work. Understanding what they are learning, and what language they now share, is how you serve them better and deliver more precise, higher-value work.

Level 1 does three things. First, it makes the intellectual case for a new enterprise architecture framework — one that succeeds Zachman, TOGAF, and the Gartner EA approach rather than extending them incrementally. Second, it establishes the European regulatory landscape as the architectural spine of any serious AI operating model. Third, it states the AIOS design principles clearly and without compromise, so that every subsequent level of the framework rests on a firm, agreed foundation.

> "*The goal is not more AI. The goal is a better operating model.*"

Level 1 closes with the AIOS Diagnostic — a structured self-assessment that any strategic planner can complete in approximately 90 minutes. It produces a Layer Priority Map: a clear, honest picture of where your organisation currently stands, and where to direct attention first.

PRINCIPLE 1 BUSINESS STRATEGY LEADS. ALWAYS.

The direction of every architectural decision flows from business strategy. Technology choices, AI deployment decisions, and automation priorities are derived from strategic intent — they do not define it. Where AI capability creates new strategic options, those options are evaluated and decided at the strategic layer before technical execution begins. The dog wags the tail. Where this principle is violated, the organisation has begun to be led by its tools.

Before you proceed to Level 2, you should be able to articulate the case for this framework to a board audience, understand why the EU regulatory landscape is architecturally load-bearing rather than a legal afterthought, and have your Layer Priority Map in hand. Everything that follows builds on that foundation.

PRINCIPLE 2 INTELLIGENCE IS DESIGNED, NOT ACCUMULATED.

AI capability is introduced into the operating model through deliberate architectural decisions, not through opportunistic tool adoption. Every AI system deployed in an AIOS-compliant organisation has a defined purpose, a governance owner, a clear authority boundary, and an explicit place in the operating model. Ad hoc AI adoption — tool by tool, team by team, without architectural oversight — is explicitly excluded from AIOS compliance. The inventory of AI in use is a governed document, not a guess.

Chapter 1: The Architecture of Intelligence

The most important shift in enterprise architecture since the advent of the internet is not the arrival of AI tools. It is the emergence of intelligence as a designable, deployable, and governable organisational resource.

For most of the history of enterprise, intelligence resided exclusively in people. Organisations designed processes around human capability, human judgement, and human bandwidth. Technology was used to record, transmit, store, and process information — but the decisions, the interpretations, the adaptive responses — those remained human. Enterprise architecture frameworks were built on this assumption, whether explicitly or not.

That assumption no longer holds.

AI systems can now perform meaningful cognitive work: pattern recognition, language processing, decision support, content generation, process orchestration, and increasingly, autonomous action. This is not a future possibility — it is a current operational reality for thousands of organisations globally. The question is not whether to design for it. The question is whether to design for it deliberately, with governance and clarity, or to accumulate it chaotically, tool by tool, use case by use case, until the enterprise loses its own legibility.

The Scarcity Has Moved

The bottleneck in enterprise performance has shifted. For most of the last four decades, the scarcity was software production capacity: the ability to build, integrate, and maintain technology systems. Organisations that could produce working software faster than their competitors gained structural advantage. Enterprise architecture frameworks evolved to manage this production bottleneck — to ensure that the right systems were built, in the right order, with the right integration.

That scarcity has moved. AI dramatically reduces the cost and time of software production. Code can be generated, tested, and deployed at a pace that renders the old production bottleneck almost irrelevant for many classes of problem. The new bottleneck is clarity: the ability to define what the enterprise should do, in what sequence, with what governance, and with what accountability. Organisations that cannot answer these questions with precision will find themselves surrounded by AI capability they cannot govern.

> "*The bottleneck is no longer software production. It is clarity, governance, and the discipline to act on both.*"

Intelligence as a Designable Resource

The architecture of intelligence is the discipline of designing deliberately. It means treating AI not as a technology add-on, but as a first-class actor in the operating model — an entity that can hold tasks, execute decisions within defined authority boundaries, produce outputs that inform human judgement, and operate at a scale and speed that human teams cannot match.

It means designing the relationship between human intelligence and artificial intelligence as carefully as any organisation designs its reporting structures, governance frameworks, or financial controls. The organisations that will define the next decade of enterprise performance are not those with the most AI tools. They are those with the clearest operating models: where they know precisely what they want AI to do, what they never want it to do, who is accountable when it acts, and how they will know if it stops working well.

AIOS is the framework that makes that clarity possible — and that keeps it under the control of the people who are accountable for the enterprise.

Chapter 2: Standing on the Shoulders of Giants

No serious framework is built without acknowledging what came before. Zachman, TOGAF, and the Gartner Enterprise Architecture approach each made genuine contributions to how organisations think about architecture. Understanding what they got right — and what they were never designed to address — is essential to understanding why AIOS succeeds rather than simply supplements them.

John Zachman and the Enterprise Ontology

John Zachman published his framework for information systems architecture in 1987. It remains one of the most intellectually coherent contributions to enterprise architecture thinking. The Zachman Framework is an ontology — a classification system for architectural artefacts. It defines two dimensions: the interrogatives (What, How, Where, Who, When, Why) and the perspectives (the different viewpoints of the people who interact with the enterprise, from executive to technician).

What Zachman got profoundly right: He established that a complete architectural description requires both dimensions to be addressed at every intersection. An enterprise that can answer "What?" at the executive level but cannot answer "Who?" or "When?" at the operational level has gaps that will cause real problems. The framework is a discipline of completeness — and that discipline remains valuable.

What Zachman was never designed to address: Intelligence as an actor. In Zachman's matrix, the "Who" column describes people and organisations. It does not describe agents — non-human entities that can hold and execute responsibilities within the enterprise. There is no row for AI systems in the Zachman framework, because in 1987 there was no need for one. Today, there is. AIOS extends the Zachman ontology by adding three AI-native interrogatives: Which intelligence, Under what authority, and With what provenance.

TOGAF and the Architecture Development Method

The Open Group Architecture Framework (TOGAF) is the most widely adopted enterprise architecture framework in the world. Its core contribution is the Architecture Development Method (ADM): a structured, iterative process for developing, managing, and governing enterprise architecture across eight interconnected phases. TOGAF gives organisations a repeatable process with defined deliverables, governance checkpoints, and artefact templates.

What TOGAF got profoundly right: The recognition that architecture is not a document but a process. Its iterative ADM builds in structured review and evolution phases. Its governance model — the Architecture Board, the Architecture Contract, the Compliance Review — provides a credible and tested governance vocabulary that many organisations have built real capability around.

What TOGAF was never designed to address: The speed of AI-era change, and the governance of autonomous agents. TOGAF's ADM operates on timescales of months to years. AI capability evolves on timescales of weeks to months. TOGAF's governance model assumes human actors in all governed roles. It has no mechanism for governing an AI system that makes decisions within a TOGAF-defined architecture. AIOS updates the ADM for the AI era: each phase carries explicit AI governance checkpoints, timescales are compressed and made continuous, and the governance model explicitly accounts for non-human architectural actors.

Gartner and Business-Outcome Architecture

Gartner's contribution to enterprise architecture is less a framework than a philosophy: architecture exists to produce business outcomes, not to produce architecture. Gartner consistently challenged the tendency of EA practices to become academic exercises — producing beautiful diagrams and comprehensive documentation while the business continued to operate as it always had. Their emphasis on business alignment and the measurability of architectural value is at the philosophical core of AIOS.

What Gartner got profoundly right: The insistence that architecture must be justified in business terms, not technical terms. The willingness to say, bluntly, that EA practices that cannot demonstrate their contribution to business outcomes should be restructured or ended. This philosophy runs through every level of AIOS.

What Gartner was never designed to address: The operational integration of AI as a strategic capability, and the new regulatory landscape that governs it. Gartner's EA frameworks predate both the EU AI Act and the operational maturity of generative AI. Their guidance is valuable — but requires significant extension to be fit for purpose in the current environment.

The Common Gap: The Primary Reader Problem

All three frameworks share a structural assumption: that the primary reader of an enterprise architecture framework is a technical practitioner — an architect, a consultant, or an IT leader — and that business leaders will engage with architecture through briefings, summaries, and governance committees. This assumption has produced a generation of EA practices where the business does not truly understand what it has commissioned, and the technology does not truly understand what the business needs.

> "*Too often, the tail wags the dog. The consultant defines the architecture. The IT leader approves the budget. The CEO ratifies what they do not fully understand. AIOS exists to reverse this.*"

AIOS is built on a different assumption: that the strategic planner — the CEO, COO, CDO, or serious operator — must be the primary reader and the primary beneficiary of the framework. Every concept in AIOS is designed to be understandable and actionable by a business leader without a specialist background in enterprise architecture or AI. Every tool produces outputs that a business leader can use directly to make decisions, prioritise investment, and commission work with precision.

The dog wags the tail. This is not a metaphor. It is a structural design principle of the entire AIOS framework — and you will find it expressed, in practical form, at every level.

Chapter 3: The European Regulatory Spine

For organisations operating in or doing business with the European Union, the regulatory landscape is no longer a compliance consideration to be addressed at the end of an architecture process. It is load-bearing infrastructure. The EU has produced, and continues to produce, the most comprehensive and consequential regulatory framework for AI and digital operations in the world.

Organisations that build their AI operating models to the EU standard are, in almost every case, over-compliant for every other jurisdiction. This is not a constraint — it is a structural competitive advantage. AIOS is built Europe-first. The regulatory obligations described in this chapter are not treated as constraints on the framework; they are part of the framework. An architecture designed to meet these obligations is a better architecture: more governable, more transparent, more resilient, and more trustworthy.

The EU AI Act

The EU AI Act entered into force in August 2024 and is being implemented in phases through 2027. It is the world's first comprehensive legal framework specifically governing artificial intelligence systems. Its central mechanism is risk-based classification.

Unacceptable risk: Prohibited entirely. Social scoring by governments, real-time remote biometric surveillance in public spaces (with limited exceptions), AI that exploits psychological vulnerabilities, and manipulation of behaviour below conscious awareness. No AIOS-compliant organisation deploys systems in this category.

High risk: The most stringent obligations. AI used in critical infrastructure, education, employment decisions, essential services, law enforcement, border management, and the administration of justice. Providers and deployers must conduct conformity assessments, maintain technical documentation, implement risk management systems, ensure data governance, provide transparency to users, enable human oversight, and register in the EU database. For most enterprises, AI in HR, credit scoring, or access-to-services decisions will fall here.

Limited risk: Transparency obligations. Chatbots, emotion recognition systems, and deepfake-generating AI must inform users they are interacting with an AI system.

Minimal risk: No specific obligations beyond the general duty not to cause harm. Most AI tools in use today fall here.

The practical implication for AIOS is clear: every automation decision — every point at which a human task is delegated to an AI system — must be classified against the EU AI Act risk taxonomy before implementation. This classification happens at Layer 2 (Architecture & Ontology). The governance obligations that flow from each classification are managed at Layer 3 (Governance & Operating Model).

GDPR and Automated Decision-Making

The General Data Protection Regulation remains the most significant data governance instrument in the world. For AI systems, the critical provision is Article 22, which prohibits solely automated decisions that produce legal or similarly significant effects on individuals, unless specific conditions are met: explicit consent, necessity for the performance of a contract, or authorisation under EU or Member State law.

This has direct implications for AI-driven employment decisions, credit assessments, insurance pricing, and access-to-service determinations. AIOS treats Article 22 compliance as an architectural constraint: any workflow involving individual-affecting automated decisions must include a human decision gate unless one of the three exemptions is demonstrably and documentably met.

The data minimisation and purpose limitation principles constrain what data AI systems can be trained on and used with. AIOS data architecture is designed around these constraints: data provenance, consent records, and purpose documentation are structural components, not retrospective additions.

NIS2: Network and Information Security

The NIS2 Directive, effective from October 2024, significantly expands cybersecurity obligations for organisations operating in the EU. It covers a wider range of sectors and organisation sizes than its predecessor and imposes stronger requirements on risk management, incident reporting, supply chain security, and executive accountability. Under NIS2, senior management can be held personally liable for cybersecurity failures.

For AIOS, NIS2 establishes the baseline security requirements for Layer 5 (Trust, Security, and Infrastructure). The obligations around AI system security — protecting against adversarial inputs, ensuring integrity of AI outputs, and managing AI-related supply chain risk — are addressed specifically in the Layer 5 implementation playbooks (Level 4 of the framework).

DORA: Digital Operational Resilience Act

DORA applies to financial services entities operating in the EU and became applicable in January 2025. It establishes comprehensive requirements for ICT risk management, incident classification and reporting, digital operational resilience testing, and third-party ICT risk management. For financial

services organisations deploying AI systems, DORA creates specific obligations around AI-related operational risk: the resilience of AI models, management of AI vendor dependencies, and stress testing of AI systems.

AIOS includes a dedicated DORA Compliance Module within Level 3 (Governance & Operating Model) for financial services implementations.

EU Data Act

The EU Data Act, applicable from September 2025, establishes rights and obligations around data generated by connected devices and related services. It creates data portability rights, obligations on data holders to share data under defined conditions, and explicit protections against vendor lock-in in cloud and edge services. For AI systems that generate or process operational data, the Data Act creates obligations around data access and sharing that must be designed into the architecture from the outset — not retrofitted when a contract dispute or regulatory inquiry arises.

The AI Liability Directive

The proposed AI Liability Directive introduces civil liability rules for AI-caused harm, including a presumption of causality in certain circumstances and an evidence preservation obligation. Its direction is clear: when an AI system causes harm, the burden of demonstrating that the harm was not caused by the system's fault will increasingly fall on the deployer.

This makes audit trail design — a central component of AIOS Layer 3 — not merely a governance best practice but a practical legal necessity. An organisation that cannot produce a clear, traceable record of how an AI system made a decision, under what authority, and who reviewed it, is an organisation that is exposed.

The Regulatory Implication for Architecture

The combined effect of this landscape is that architectural decisions in the EU now carry direct legal consequences. The choice of how to design a workflow, configure a human oversight gate, document a decision boundary, or maintain an audit trail is no longer purely a technical or governance matter. It is a legal matter.

> "*If your architecture cannot be explained to a regulator, it cannot be explained to your board. If it cannot be explained to your board, it is not under control.*"

AIOS does not treat this as a burden. It treats it as a clarifying principle. The EU regulatory framework, for all its complexity, pushes organisations toward the kind of architectural clarity that produces better operating models regardless of whether a regulator is ever watching. Organisations that build to this standard do so not only because they must, but because it makes them better.

Chapter 4: The AIOS Manifesto

These ten principles govern every decision in the AIOS framework. They are not aspirational statements — they are structural constraints. Any architectural choice that violates them is not AIOS-compliant, regardless of its technical sophistication or its apparent business rationale. They should be read, understood, and actively tested against at every stage of framework adoption.

Chapter 5: The AIOS Diagnostic

PRINCIPLE 3 GOVERNANCE PRECEDES AUTOMATION.

Before any task or decision is automated, the governance architecture for that task or decision must be in place. This means the authority boundary is defined, the audit trail is designed, the escalation path is clear, and the regulatory classification has been assessed. Automating first and governing later is the most common and most expensive structural failure mode in AI adoption. AIOS prevents it by design, not by hope.

The AIOS Diagnostic is a structured self-assessment designed to give any organisation a clear, honest picture of where it currently stands across the five architectural layers of the AIOS operating model. It consists of 25 questions — five per layer. It is designed to be completed by a strategic planner in approximately 90 minutes, with input from two or three colleagues who have direct operational knowledge.

PRINCIPLE 4 THE EUROPEAN REGULATORY SPINE IS LOAD-BEARING.

EU regulatory obligations — the AI Act, GDPR, NIS2, DORA, the Data Act, and the AI Liability Directive — are not compliance additions to the architecture. They are structural requirements that shape it. Every AIOS layer is designed to meet these obligations. An organisation that treats regulatory compliance as a retrospective exercise will rebuild its architecture repeatedly, at increasing cost. An organisation that builds to the regulatory spine from the outset builds once.

Score each statement using the following scale:

1 — Not in place: This does not exist or has not been considered.

2 — Partially in place: Some elements exist but are incomplete, informal, or inconsistently applied.

3 — Mostly in place: The majority of this is functioning, but gaps or inconsistencies remain.

4 — Fully in place and functioning: This is operational, documented, and actively maintained.

The diagnostic produces a Layer Score (out of 20) for each of the five layers. Use the scoring bands in the Priority Map at the end of this chapter to determine your most urgent architectural priorities.

PRINCIPLE 5 HUMAN AUTHORITY IS EXPLICIT AND PROTECTED.

There are categories of decision that must remain human-led — not because AI is incapable of making them, but because human accountability for those decisions is ethically, legally, or commercially non-negotiable. AIOS requires every organisation to define, document, and govern these categories explicitly. Human authority is not the residue that remains after AI takes over everything else. It is a designed, protected, and actively maintained component of the operating model.

Reading Your Results: The Layer Priority Map

PRINCIPLE 6 TRUST IS ARCHITECTURAL, NOT REPUTATIONAL.

An organisation's trustworthiness in the AI era is not a matter of brand or communication strategy. It is a function of how the architecture works: whether audit trails are maintained, whether human oversight gates are functioning, whether data is handled lawfully, whether AI outputs can be explained and challenged. In AIOS, trust is built through architectural design. It cannot be claimed — only demonstrated. An organisation that claims to be trustworthy but cannot show the architecture has nothing.

Total your scores for each layer independently. A maximum score of 20 indicates the layer is fully operational and functioning. Use the bands below to determine your priority sequence.

The layer with your lowest score is your most urgent priority. Before commissioning any AI automation work, ensure the layer immediately above it in the operating model scores at least 11. Automating into an ungoverned operating model is the most common and most expensive structural mistake in AI adoption.

> "*Your lowest-scoring layer is not a weakness to be embarrassed by. It is the most precise possible answer to the question: where do I start?*"

Field Practitioner Checklist: Level 1

Complete this checklist before proceeding to Level 2. Each item represents a capability or commitment that Level 2 assumes you have in place. If you cannot check an item, return to the relevant chapter before continuing.

I can articulate the case for AIOS as a successor framework to a board-level audience — not as a technical matter, but as a strategic and governance matter.

I understand the difference between the Zachman ontology, the TOGAF method, and the Gartner business pragmatism — and can explain why all three are pre-intelligence frameworks that require succession.

I have reviewed the five key EU regulatory instruments (EU AI Act, GDPR Article 22, NIS2, DORA, EU Data Act) and have assigned ownership of each to a named person in my organisation.

I can state the ten AIOS design principles, or can locate them immediately when needed to evaluate an architectural or operational decision.

I have completed the AIOS Diagnostic with input from at least two colleagues, and I have my Layer Priority Map.

PRINCIPLE 7 COMPLEXITY IS MANAGED THROUGH LAYERS, NOT SUPPRESSED.

The AIOS operating model uses a layered architecture because the alternative — treating the enterprise as a flat system — produces a form of illegibility that prevents governance. Each layer has clear inputs, outputs, governance obligations, and interfaces with adjacent layers. Complexity is not eliminated; it is located. When something goes wrong, the layered architecture makes it possible to identify exactly where, whose responsibility it is, and what the correct response is. This is the difference between a governable system and a complicated one.

I have identified which layer requires my most urgent attention and have a clear next step for addressing it.

I have scheduled the executive briefing at which the AIOS framework will be formally adopted as the governing approach for my organisation's AI operating model.

I understand that the strategic planner — not the consultant, not the architect, not the IT function — is the primary owner of this framework, and I have accepted that responsibility.

Proceed to Level 2: Architecture & Ontology

Level 2 introduces the AIOS Architecture Canvas, the full five-layer operating model in detail,

and the extended ontology that gives your organisation a complete architectural language for the AI era.

AIOS Institute • aios.institute • 2026

PRINCIPLE 8 MATURITY IS MEASURED IN OPERATIONAL TERMS.

AIOS does not measure architectural maturity by the completeness of documentation, the sophistication of tooling, or the number of AI systems deployed. It measures maturity by the operational effectiveness of the AI operating model: whether AI capability is delivering business outcomes, whether governance is functioning as designed, whether the organisation can absorb new AI capability without structural disruption, and whether the regulatory posture is demonstrably sound. Documentation without operational effect is not maturity — it is theatre.

PRINCIPLE 9 THE FRAMEWORK EVOLVES WITH THE ENVIRONMENT.

AI capability, regulatory requirements, and business context all change faster than any static framework can accommodate. AIOS is designed to evolve: it uses structured review cycles, defined update triggers (regulatory change, capability change, strategic change), and a versioning model that allows the framework to be updated without invalidating existing implementations. Organisations should expect to revisit their AIOS architecture at minimum annually, and whenever a material change in AI capability, regulation, or business strategy occurs.

PRINCIPLE 10 THE STRATEGIC PLANNER IS THE PRIMARY BENEFICIARY.

Every tool, template, canvas, and diagnostic in AIOS is designed to be used directly by the strategic planner — the executive who must make decisions before work is commissioned. If a tool cannot be used by a competent leader who is not a specialist in enterprise architecture or AI, it has failed its design intent. The consultant, the architect, and the technologist are well served by AIOS because the strategic planner is well served first. Not the other way around.

LAYER 1 — LEADERSHIP, INTENT & GOVERNANCE	SCORE (1–4)
1. The organisation has a documented AI strategy connected to business strategy, with named executive ownership.	____
2. There are clear, communicated principles governing what AI will and will not be used for in this organisation.	____
3. There is a governance body or designated function responsible for AI decisions at the enterprise level.	____
4. The organisation can describe, without hesitation, which decisions must remain human-led and why.	____
5. Regulatory obligations under the EU AI Act and GDPR have been formally assessed and assigned to named owners.	____
Layer Total (max 20)	____ / 20

LAYER 2 — CAPABILITY & OPERATING MODEL	SCORE (1–4)
1. The organisation has a current capability map that includes AI capability as a distinct, governed category.	____
2. AI-related capabilities are clearly owned, with named individuals accountable for each.	____
3. There is a defined, documented process for evaluating whether a new AI capability should be adopted.	____
4. The organisation has a clear view of which operational activities have been assessed for AI suitability or risk.	____
5. The operating model is documented at sufficient detail to integrate new AI capabilities without ad hoc redesign.	____
Layer Total (max 20)	____ / 20

LAYER 3 — WORKFLOW & EXECUTION	SCORE (1–4)
1. Core workflows are documented in sufficient detail to identify where AI intervention adds value or introduces risk.	____
2. There is clear routing logic defining what goes to AI, what goes to a human, and what requires escalation.	____
3. Escalation paths from AI systems to human decision-makers are explicitly defined and have been tested.	____
4. Outputs of AI systems are reviewed on a regular schedule by accountable humans.	____
5. Any AI-assisted decision can be traced back to its inputs, logic, and authorising governance documentation.	____
Layer Total (max 20)	____ / 20

LAYER 4 — AI, AGENTS & AUTOMATION	SCORE (1–4)
1. The organisation maintains a current inventory of all AI systems in use, including team-level and shadow AI.	____
2. Each AI system in the inventory has been classified against the EU AI Act risk categories.	____
3. Authority boundaries for each AI system are documented: autonomous decisions vs. those requiring human ratification.	____
4. There is a process for retiring or modifying AI systems when they cease performing to their designed purpose.	____
5. AI tooling choices are documented against privacy and data sovereignty criteria, with EU regulatory requirements explicit.	____
Layer Total (max 20)	____ / 20

LAYER 5 — TRUST, PRIVACY, SECURITY & INFRASTRUCTURE	SCORE (1–4)
1. Data provenance — origin, consent basis, and processing history — is documented and auditable for AI systems.	____
2. Access controls for each AI system are explicit: who can configure, modify, or override.	____
3. NIS2 obligations have been formally assessed and implementation accountability has been assigned.	____
4. AI-related security risks (adversarial inputs, model integrity, supply chain) are included in risk management.	____
5. A communication plan for AI-related incidents exists, covering both internal governance and external regulatory notification.	____
Layer Total (max 20)	____ / 20

SCORE	PRIORITY STATUS	RECOMMENDED ACTION
5	Begin immediately	This layer is non-operational. Begin architectural work here before any AI work is commissioned or expanded.
6–10	Critical priority	Significant gaps. Include in your 30–90 day architectural roadmap. Do not expand AI operations in this domain until addressed.
11–15	Scheduled attention	Needs structured attention. Include in your 6-month roadmap. Current operations can continue with heightened governance oversight.
16–20	Maintain and monitor	This layer is functioning. Review quarterly. Update when AI capability, regulatory requirements, or business strategy changes materially.