

ENTERPRISE AI EXECUTION

The Kriyanta AI Deployment
Standard for Industrial Workflows

Governed. Measurable. Client-controlled.

The Deployment Gap

Industrial companies do not lack AI ideas. They lack a reliable way to put them into operation.

The practical questions begin once the demonstration ends.

Who owns the decision? What proprietary data must the model see? How does the recommendation reach SAP or another enterprise system? And who retains the knowledge created when employees correct and improve the workflow?

The enterprise AI debate is moving beyond traditional data security. Companies are increasingly asking who controls the prompts, decisions, corrections, workflow history and institutional knowledge created through use – and whether that operating capability remains intact if the underlying model or provider changes.

These choices must be built into the architecture, workflow and contractual boundary from the outset.

The Kriyanta AI Deployment Standard translates these concerns into a practical model for governed industrial workflows.

It defines how every Kriyanta AI project moves from application to live workflow – with transparent decision logic, governed execution, controlled AI use, measurable

outcomes and client-controlled learning.

The model is not the operating system. It is one component inside a governed execution workflow.

Six Disciplines of the Kriyanta AI Deployment Standard

1. Start with the business outcome

Every workflow begins with the result it must deliver: cash, margin, service, cost or productivity.

Technology choices follow the value at stake – not the other way around.

2. Define what good looks like

The client sets the targets, thresholds and limits against which recommendations are judged. AI does not decide what is commercially or operationally acceptable for the business.

3. Keep the decision logic transparent

Core calculations, scoring and financial-impact logic remain explainable and separate from the language model wherever possible. The user should be able to understand why a recommendation was generated.

4. Share proprietary data only for a defined task

Enterprise data is provided to an approved AI model only when required for a specific purpose – such as explaining one recommendation or interpreting one exception.

The full underlying dataset is not handed to the model wholesale.

5. Separate the workflow from the model

Industrial logic, approval rights and workflow history should remain intact if the underlying model changes.

Different tasks may use different approved models – or no language model at all.

6. Measure outcomes and retain the learning

Decisions, exceptions and realised results are captured as client-controlled traces, quality checks and approved rule changes.

This allows the next recommendation to become more specific to the client without relying on uncontrolled model learning.

Over time, each validated outcome strengthens the next decision – improving execution speed, protecting operational performance and compounding financial value within the client's business.

From AI Application to Enterprise Execution System

Kriyanta treats enterprise AI as a three-layer execution system.

Industrial decision logic determines what should happen through transparent operational rules, financial logic and the client's definition of what good looks like.


Governed workflow and approvals determine who reviews the recommendation, who has authority to act and how exceptions are escalated.

Controlled AI execution uses approved models only where they add value and only with the task-specific context required.

The separation matters. A pricing floor, service threshold or sourcing guardrail is not generic model knowledge. It is part of how the client runs its business. The model can change. The operating logic and authority structure must remain.

The Three Layers of an Enterprise Execution System

How an AI application becomes a governed industrial workflow



Decision logic determines what should happen. The workflow determines who decides and how the action is governed. AI supports selected tasks such as explanation,

classification, drafting and exception handling. Keeping these roles separate keeps commercial rules and operational guardrails

visible, explainable and controlled by the business.

AI should not define what is acceptable for the business. The business sets the targets, thresholds and guardrails.

The Learning Stays With the Client

When operators accept, reject or adjust a recommendation, they add judgement that may never exist in the source system.

A planner may know that a stock reduction creates unacceptable service risk. A salesperson may understand why a strategically

important quotation requires an exception. An engineer may reject a cost reduction because of a constraint that is not captured in the structured data.

These decisions are not noise around the application. They are institutional knowledge – and they should remain with the client.

The workflow should therefore capture the recommendation, the operator's decision, the reason, the action taken and the measured result.

That knowledge should remain inside the agreed client boundary and should not be silently pooled across unrelated companies.

Improvement is governed rather than automatic. Only reviewed and approved changes enter the next version of the workflow.

Defining What Good Looks Like

A learning loop is useful only if the company knows what result the workflow is expected to achieve. For a pricing application, that may include contribution margin, quotation win rate, discount leakage and below-floor approvals. For an inventory application, it may include realised cash, service stability and stock-out risk.

How the Client Learning Loop Works

Governed workflow improvement under client control

Not automatic self-learning. The client decides what becomes institutional knowledge.



For procurement, it may include implemented savings, qualification success, supplier quality and resilience of supply.

The test is not whether an AI output sounds intelligent.

The test is whether the workflow produces the business outcome it was designed to deliver.

From Demonstration to Live Deployment

A prototype proves the analytical logic, financial opportunity and workflow concept.

A live implementation adds the enterprise capabilities required to operate safely and reliably:

- approved data connections;
- identity and access rights;
- task-specific model routes;
- approval and escalation logic;
- decision and outcome records;
- audit trails;
- impact monitoring;
- version control;
- retention and deletion rules.

The exact architecture depends on the sensitivity of the workflow and the client's technology environment.

The standard does not change: transparent logic, governed execution, controlled AI use and measurable impact.

Enterprise AI Execution

Enterprise AI should not leave an industrial company less transparent, more dependent or less in control of its own operating knowledge.

The objective is not simply to deploy an AI application.

It is to build an execution system in which industrial logic, human judgement and AI work together – under client control – to deliver measurable value.

That is the Kriyanta AI Deployment Standard.

**A
recommendation
is not impact.
Impact begins
when the
decision is
executed,
measured and
sustained.**