



Myria Consulting

— TECHNOLOGY SYSTEMS

AI Governance & *Architecture*

Production-grade orchestration layers, control chains, and risk mitigations.

CONTENTS

What's inside.

01	Executive summary	p. 03
02	The challenge	p. 04
03	Orchestration layers	p. 05
04	Validation chains	p. 06
05	Audit trails	p. 07
06	Multi-tier hazard isolation	p. 08
07	The governance operating model	p. 09
08	The framework at a glance	p. 10
09	Where to start	p. 11
10	Working with Myria	p. 12

MYRIA PERSPECTIVE

Every automated decision should be explainable, traceable and contained.

EXECUTIVE SUMMARY

Governance is what lets AI move fast in production.

Enterprises are moving AI from pilots into processes that touch customers, money and regulated decisions. At that point, the question changes from 'does the model work?' to 'can we prove it behaves, and contain it when it does not?'

This framework sets out how to structure AI systems so they can be governed by design. It covers the orchestration layers that keep models replaceable, the validation chains that check every input and action, the audit trails that make decisions explainable, and the isolation tiers that limit damage.

Each chapter combines architecture patterns with the governance practices that make them work, and closes with an example from production.

4

orchestration layers that keep models swappable and governable

3

tiers of hazard isolation, matched to risk

100%

of automated decisions traceable end to end

WHO IT IS FOR

Technology, risk and compliance leaders deploying automated models in production who need control without slowing delivery.

THE CHALLENGE

Why AI stalls between pilot and production.

Most organisations have more successful AI pilots than production systems. The models usually work. What holds them back is the lack of a structure that risk, legal and operations teams can trust.

- **Models wired straight into apps**
Each team calls a vendor model directly, so swapping providers or adding controls means rewriting every application.
- **Checks added after the fact**
Validation is bolted on once something goes wrong, rather than built into the path every request takes.
- **No record of why**
When a customer or regulator asks why a decision was made, nobody can reconstruct the inputs, prompts and outputs.
- **One failure, wide impact**
Agents with broad access to tools and data can turn a single bad output into a costly incident.

WHY IT MATTERS

Without governance, every new AI use case adds risk faster than it adds value.

CHAPTER 1

01

Orchestration layers

Separate what the business wants from which model happens to deliver it.

We organise AI systems in four layers: the experience layer where people interact, the application layer that holds business rules, the intelligence layer of agents, tools and memory, and the provider layer where vendor models plug in.

Applications never call a model directly. They call a defined capability — 'summarise a contract', 'draft a reply' — and the orchestration layer decides which model, prompt and controls to use.

KEY POINTS

Capability contracts

Each AI task has a defined input, output and owner.

Provider adapters

Vendor models plug in behind a common interface.

Policy routing

Requests are routed by sensitivity, cost and quality.

Model portability

Switch or combine providers without touching applications.

IN PRACTICE

A bank moved three customer-service assistants to a new model provider in two weeks — without changing a single screen.

CHAPTER 2

02

Validation chains

Every request passes the same checkpoints before and after the model.

A validation chain is a fixed sequence of checks around each model call. Before the call, inputs are checked for personal data, policy breaches and prompt injection. After it, outputs are checked for accuracy against sources, tone, format and permitted actions.

High-risk actions — payments, account changes, external messages — require a human approval step, built into the flow rather than handled by email.

KEY POINTS

Input screening

Personal data, policy and injection checks on every request.

Output checks

Grounding, format and safety verified before release.

Human approval

Built-in sign-off for actions above a risk threshold.

Continuous evaluation

Test suites run on every model or prompt change.

IN PRACTICE

An insurer's claims assistant routes about 8% of drafts to a human reviewer — the rest are released automatically with full checks recorded.

CHAPTER 3

03

Audit trails

If you cannot explain a decision, you cannot defend it.

Every AI interaction produces a record: who asked, which data was retrieved, which model and prompt version answered, what checks ran and what action followed. Records are append-only and linked to the business transaction they affected.

Humans and agents are logged the same way, so an agent acting through a tool carries the same permissions and leaves the same trail as an employee doing the task by hand.

KEY POINTS

Complete lineage

Inputs, sources, model, prompt and output for every decision.

Tamper-proof logs

Append-only records with retention matched to regulation.

Same rules for agents

Agents act with user-level permissions and audit.

Explainability on request

Reconstruct any decision in minutes, not weeks.

IN PRACTICE

A lender answered a regulator's request on 400 automated credit decisions in two days, using records already in place.

CHAPTER 4

04

Multi-tier hazard isolation

Assume something will go wrong, and decide in advance how far it can spread.

Use cases are sorted into three tiers by potential impact. Tier 1 covers internal assistance with no external effect. Tier 2 covers customer-facing content under review. Tier 3 covers actions that move money, change records or affect rights.

Each tier has its own limits on data access, tool permissions, spending and rate of action, plus circuit breakers that pause automation and hand over to people when signals look wrong.

KEY POINTS

Risk tiers

Controls scale with the impact of each use case.

Least privilege

Agents see only the data and tools their task needs.

Circuit breakers

Automatic pause and human handover on anomalies.

Fallback paths

A safe manual route for every automated process.

IN PRACTICE

When a model update degraded accuracy, a retailer's circuit breaker paused automated refunds within minutes and switched to manual review.

CHAPTER 5

05

The governance operating model

Architecture makes control possible. People and routines make it real.

We define who owns each AI capability, who approves new use cases and how risk is reviewed. A small AI council sets policy; product teams own delivery; risk and compliance hold a clear veto on Tier 3.

Governance runs on a regular rhythm — intake, risk assessment, launch review and quarterly performance review — so new use cases move quickly through a predictable path.

KEY POINTS

Clear ownership

Every capability has a business and a technical owner.

Use-case intake

One front door with a risk-based fast track.

Regulatory readiness

Mapped to the EU AI Act, GDPR and sector rules.

Performance reviews

Quality, cost and incidents reviewed every quarter.


IN PRACTICE

With a fast track for Tier 1 use cases, one client cut approval time for new internal assistants from three months to ten days.

THE FRAMEWORK AT A GLANCE

Control at every layer, from request to record.

The architecture and the governance model work together. Each layer has a defined job, and every request flows through the same controls in the same order.



HOW TO READ IT

The same five stages apply to every use case. What changes between tiers is the strictness of each check and whether a human approval is required before release.

WHERE TO START

Your first 90 days.

Governance does not have to wait for a full platform. A focused 90-day programme puts the essentials in place around the use cases you already run.

DAYS 1–30

- Inventory every AI use case and model in use, and assign a risk tier.
- Name owners and agree a simple intake process for new use cases.

DAYS 31–60

- Introduce a shared orchestration layer for the highest-risk use cases.
- Add input and output validation and start logging full decision records.

DAYS 61–90

- Set tier-based limits and circuit breakers.
- Run the first quarterly governance review with risk and compliance.

QUESTIONS TO ASK YOUR LEADERSHIP TEAM

- *Do we know every place AI is making or shaping decisions today?*
- *Could we explain any one of those decisions to a regulator this week?*
- *If a model misbehaved tomorrow, how far could the damage spread?*

WORKING WITH MYRIA

Human judgment and specialist AI, in one shared space.

Myria pairs senior management consultants with a virtual team of specialist AI advisors. You get the judgment of people who have sat in the room when hard calls were made, and the speed of analysis that used to take weeks. Engagements are shaped around the decision you need to make.

01

AI governance assessment

Four weeks to map your AI use cases, risks and gaps against the framework.

02

Architecture blueprint

A target design for orchestration, validation, audit and isolation, fitted to your stack.

03

Myria OS

Governed AI advisors working alongside your team, with every action recorded.

04

Advisory retainer

Senior guidance on new use cases, vendor choices and regulatory change.

NEXT STEP

Every engagement starts with a 45-minute conversation about the decision in front of you — not a sales pitch.



The clarity to move from *ambition to action.*

Talk to a Myria advisor about putting these ideas to work.

info@myriaconsulting.com

<https://myriaconsulting.com/>