OCTAVE

EBOOK

Building an AI governance framework for quality management

A practical guide to implementing AI responsibly in regulated environments

Executive brief

Let's start with the stats: 33% of manufacturers are already using AI. 49% say they plan to use it within the next two years. But enthusiasm doesn't equal readiness.

Rushing to deploy AI without the governance foundation to support it can result in hallucinated compliance data, audit trails that can't be explained and decisions made based on recommendations that take time, or nobody, can trace back to their source.

AI governance in quality management is the foundation that enables responsible AI adoption in regulated environments.

The four essential components:

1. **Assessment first:** You can't govern what you don't see. Make a list of where you use AI (including shadow AI), classify each by risk and identify regulatory requirements before building governance structures.
2. **Ecosystem, not hierarchy:** Governance works when it's a connected system—cross-functional teams, clear ownership and feedback loops from frontline users to executive leadership.
3. **Embedded, not bolted-on:** AI that's embedded in your QMS inherits existing workflows, approvals and audit trails. AI that's separate from your QMS requires parallel governance, which adds complexity and risk.
4. **Continuous, not project-based:** AI models degrade. Regulations change. Governance needs continuous monitoring, regular audits and a culture of improvement.

Technology considerations for AI-ready QMS platforms:

When evaluating whether your quality management system can support governed AI, look for:

- Flexibility to accommodate embedded AI, partner integrations and custom models
- Built-in workflow automation and approval structures AI can leverage
- Robust data governance and audit trail capabilities
- Role-based permissions that extend to AI functionality
- Integration capabilities with specialized AI tools and data platforms

Introduction

What This Framework Is and Isn't

This isn't a theoretical governance model designed by consultants who've never dealt with a CAPA backlog. It's a practical roadmap built from the ground up for quality professionals operating in regulated environments.

What you'll find: Specific steps for assessing your AI readiness, building governance infrastructure that actually works, implementing AI with oversight that enables speed (not just control) and evolving your approach as AI capabilities mature. Generic "establish a committee" advice or compliance checklists divorced from operational reality or frameworks that assume unlimited resources and infinite executive patience.

Who this is for: Quality leaders who need to build something that works without bringing quality operations to a halt.

How to use it: You can implement all four parts sequentially or jump to the section most relevant to your current situation. Each part stands alone while building toward a comprehensive governance ecosystem.

One critical principle underlies this entire framework: AI governance works best when it's embedded in your quality management system, not bolted on as an afterthought. Governance shouldn't feel like a separate compliance burden. It should be woven into the workflows where AI is actually being used.

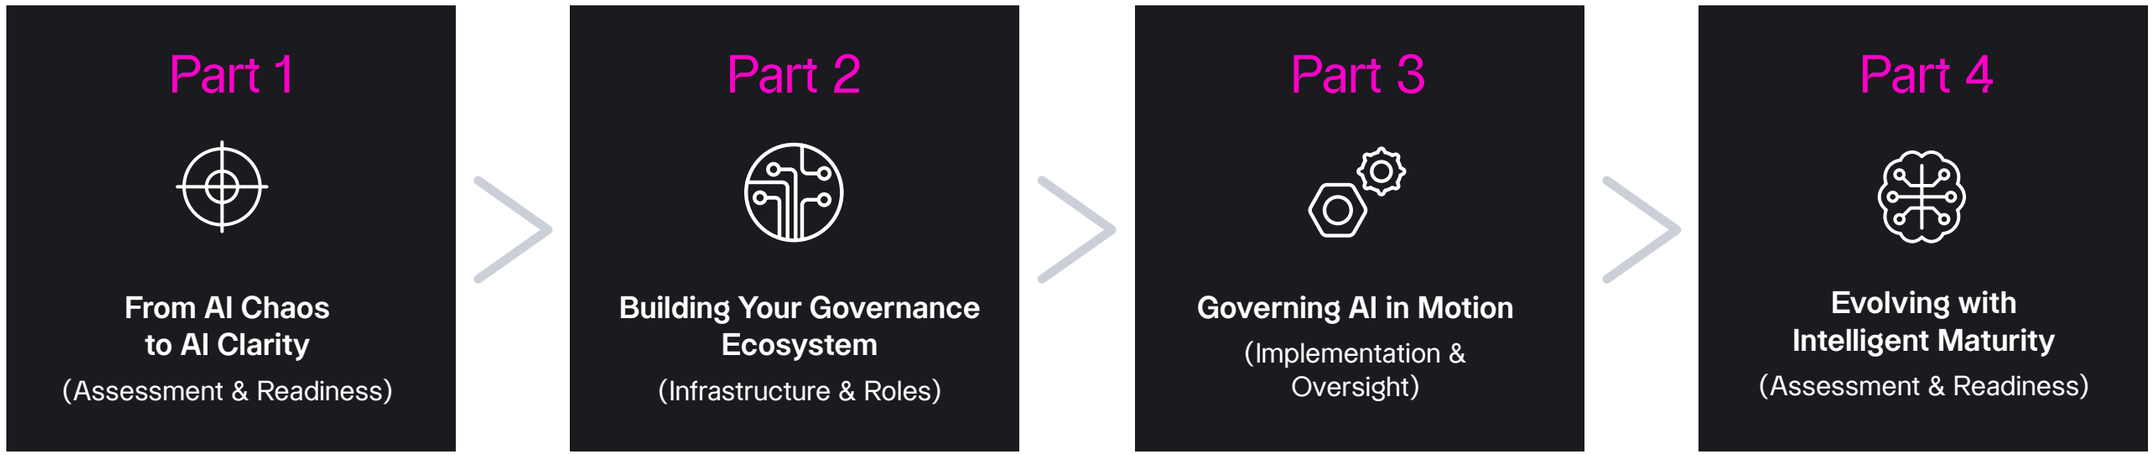

01

PART 1

From AI chaos to AI clarity

If we asked you where AI is being used in your organization, would you know?

It might be someone in manufacturing using ChatGPT to write root cause analysis descriptions. It might be the quality team testing a predictive analytics tool. Or it might be IT piloting an AI-powered document classifier.

When nobody's coordinating, and nobody's governed, governance failures begin. None of this is with malicious intent, either. Instead, it's with disconnected experiments that never get documented or properly evaluated.

The first step isn't building governance structures. It's achieving clarity about your current state.

- Where is AI being used (officially or unofficially)?
- What types of AI are in play (tools embedded in your QMS, specialized partners, or custom models built internally)?
- What regulations apply?
- What data is being used?

Without this diagnostic step, you're building governance blind. You might over-regulate low-risk applications while missing high-risk AI operating in the shadows.

This assessment phase answers three critical questions: (1) What AI do we have? (2) What risks does it create? (3) What governance do we actually need?

Key actions

- **Inventory current AI usage:** Survey departments to uncover both sanctioned and shadow AI tools
- **Classify by risk level:** High-risk (decision-making in regulated processes) vs. low-risk (administrative tasks)
- **Identify data sources:** What data is training these models? Customer data? Quality records? Production data?
- **Map regulatory requirements:** FDA, ISO, EU AI Act — which apply to your AI use cases?
- **Assess organizational readiness:** Do you have the roles, skills, and infrastructure to govern AI effectively?

AI governance readiness assessment

1	**AI Visibility & Control**	No visibility of where AI is used	1	2	3	4	5	Complete inventory and tracking
2	**Risk Classification**	No risk framework exists	1	2	3	4	5	Mature risk assessment process
3	**Data Governance**	Ad hoc data usage	1	2	3	4	5	Comprehensive data governance
4	**Regulatory Alignment**	Compliance requirements unknown	1	2	3	4	5	Fully aligned with all regulations
5	**Roles & Ownership**	No defined AI ownership	1	2	3	4	5	Clear RACI and accountability
6	**Change Management Capacity**	Organization resistant to change	1	2	3	4	5	Highly adaptive culture

0-10 points: Start with AI inventory and discovery
11-20 points: Build governance foundation (focus on Part 2)
21-25 points: Ready for implementation (move to Part 3)
26-30 points: Focus on continuous maturity (Part 4)

02

PART 2

Building your governance ecosystem

Governance isn't a person or a policy document. It's a connected ecosystem of roles, responsibilities, standards and feedback loops that span from the shop floor to the C-suite.

The challenge in quality management is that AI doesn't stay in one department. It touches everyone, from quality engineers using AI-assisted form completion and compliance officers monitoring AI-driven audit readiness to plant managers relying on predictive analytics and executives making decisions based on AI-generated insights.

Your governance framework needs to cover cross-functional ownership and clear escalation paths and standards. The goal is to create clarity, so you know the answers to the following questions when, not if, they arise:

- Who decides if an AI application is high-risk?
- Who reviews AI recommendations before they impact compliance?
- Who owns the data?
- Who responds when AI produces an unexpected output?

Key actions

- **Establish a Center of AI Excellence (CoE):** Cross-functional team (Quality, IT, Compliance, Legal, Operations) that sets standards and adjudicates edge cases
- **Define clear ownership roles:** Chief AI Officer (strategy), AI governance lead per function (implementation), frontline "AI translators" (day-to-day usage)
- **Create AI standards & policies:** Data usage rules, transparency requirements, explainability standards, model validation protocols
- **Build decision frameworks:** Risk classification criteria, approval workflows by AI type, escalation procedures
- **Enable governance:** Design systems that make governed AI easier than ungoverned AI

AI governance ecosystem map

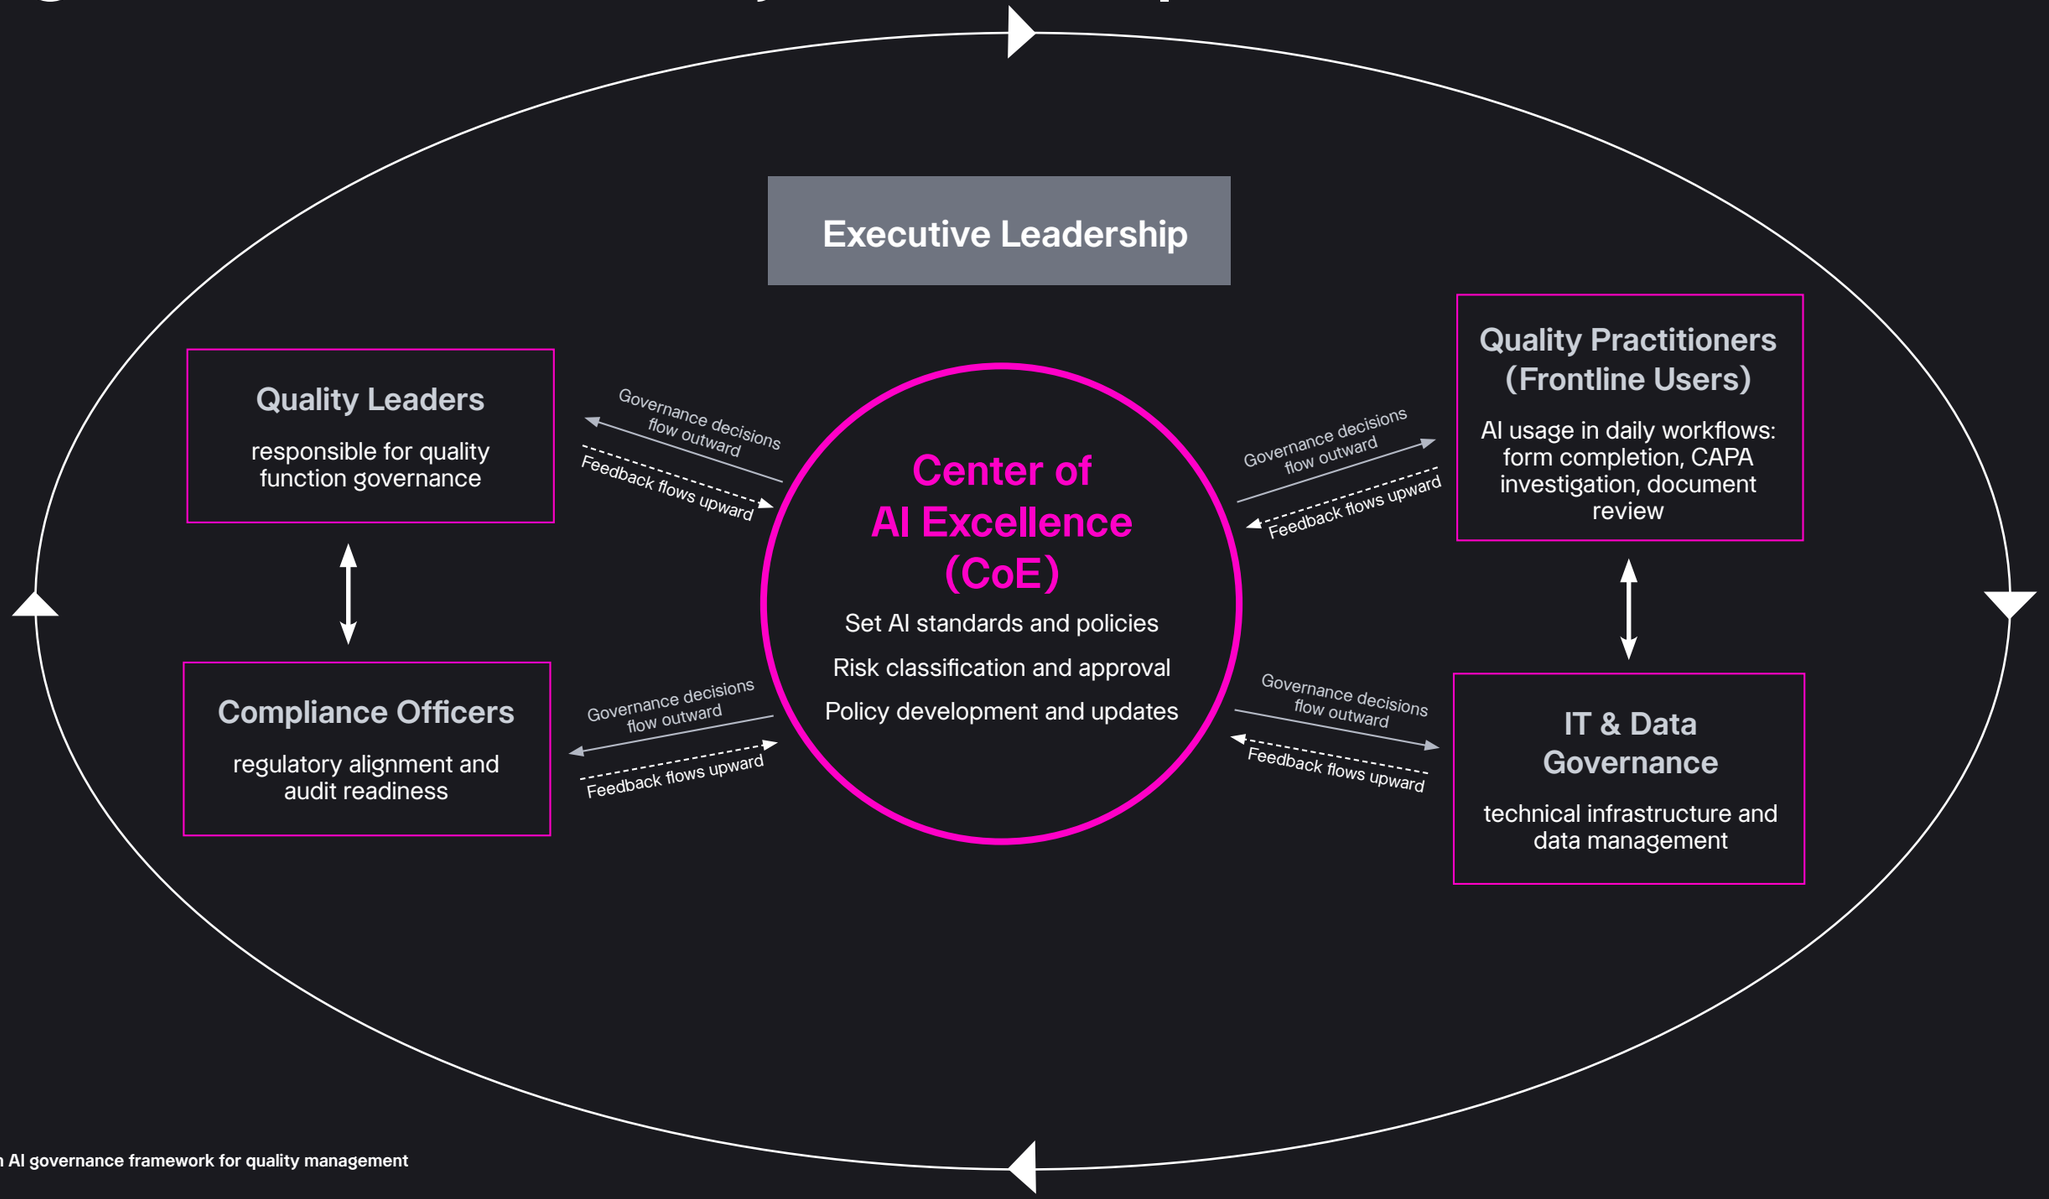

03

PART 3

Governing AI in motion

The best governance framework on paper is worthless if it slows quality operations to a crawl. The challenge is to implement AI capabilities quickly while maintaining appropriate oversight.

To do this, embed governance into your workflows instead of layering it on top. When a quality engineer uses AI to recommend form field values, governance should be invisible, built into the tool itself through confidence scoring, explainability and user control. When a compliance officer reviews AI-generated audit readiness reports, governance should feel like quality checks, not bureaucratic hurdles.

This is where the “embedded vs. integrated” distinction matters — and where the same governance considerations apply to any AI solution, whether native or integrated. AI that operates within your QMS inherits the governance, workflows and compliance structures already in place. AI that operates through an integrated solution should be evaluated against those same standards.

If your QMS already has approval workflows, electronic signatures, audit trails and role-based permissions, ask whether your AI operates within that system. When it does, governance becomes automatic. When it doesn’t, you risk managing two separate compliance frameworks where gaps can emerge. Applying consistent governance standards across all your AI solutions, regardless of how they’re deployed, is what keeps your compliance posture intact.

Key actions

- **Start with single, high-value use cases:** Document automation or predictive quality analytics and prove value before scaling
- **Implement transparency by design:** Every AI recommendation includes a confidence score, data sources and an explanation
- **Maintain human-in-the-loop control:** Users can accept, reject or modify AI suggestions (AI augments, not replaces)
- **Build compliance into workflows:** AI validation happens at point of use, not in separate review cycles
- **Monitor AI performance continuously:** Track accuracy, user acceptance rates and downstream quality impacts

AI Implementation journey

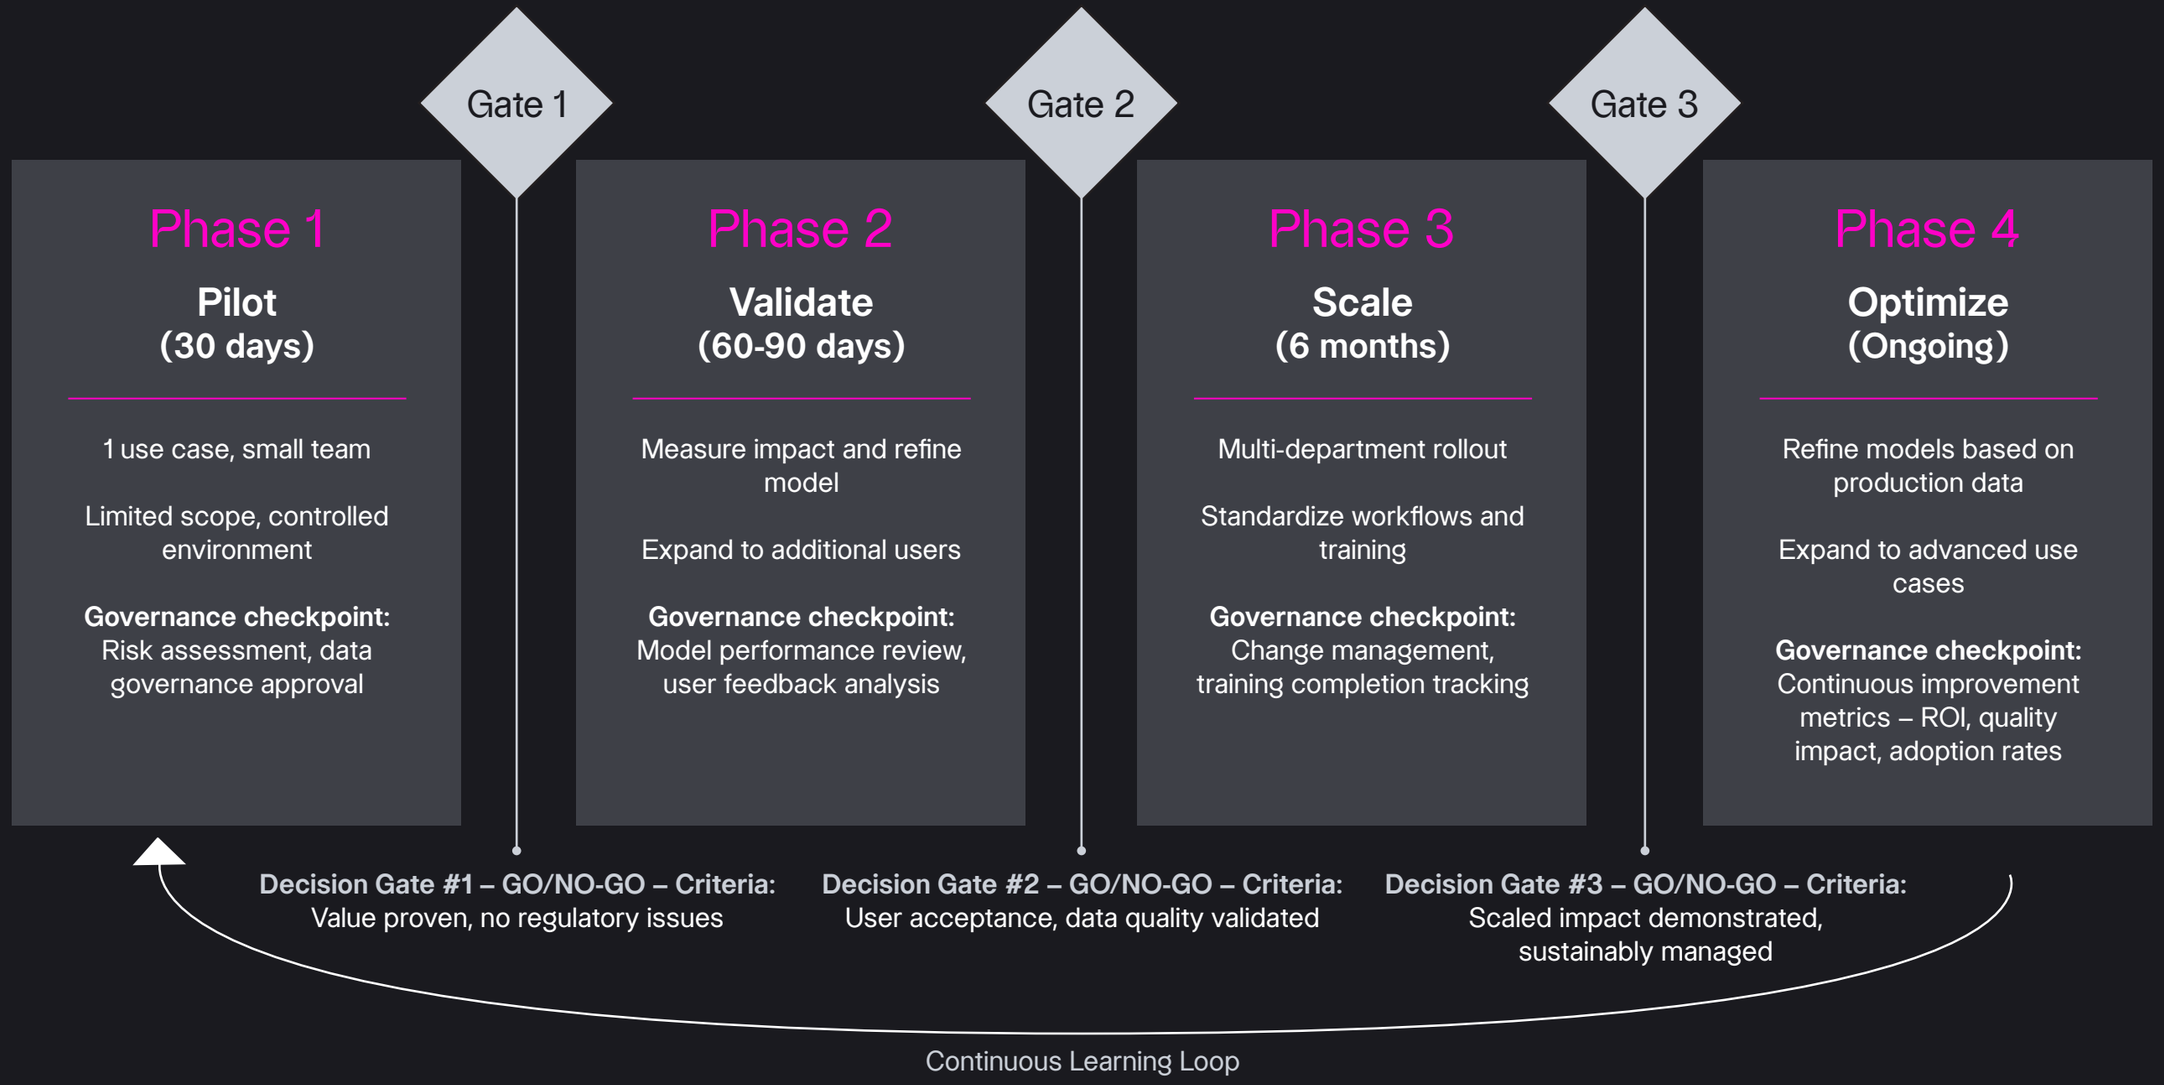

04

PART 4

Evolving with intelligent maturity

One of the challenges with AI models is that they degrade over time. They're trained on historical data, but business conditions change, regulations evolve and new quality issues emerge that weren't in the training dataset.

A model that performs well at launch might give you garbage recommendations six months later. If you're not monitoring continuously, you won't know until audit findings or customer complaints force the issue.

Mature AI governance means continuous auditing, such as:

- Tracking model accuracy in production
- Monitoring user acceptance rates (are people rejecting AI recommendations more often?)
- Detecting data drift (is the incoming data significantly different from training data?)
- Validating that AI outputs still align with current regulations

This is where most governance frameworks fail. They're great at launch but terrible at sustainability. Which makes the final component of governance a culture of continuous improvement where AI is treated like any other quality system and is regularly audited, continuously improved and never assumed to be "good enough."

Key actions

- **Implement continuous model monitoring:** Track prediction accuracy, user override rates and downstream quality metrics in real-time
- **Schedule regular governance reviews:** Quarterly assessment of AI performance, risks and alignment with business/regulatory changes
- **Maintain an AI assumptions ledger:** Document what assumptions models are built on and revisit when conditions change
- **Build feedback loops from frontline users:** Quality practitioners often spot issues first so create easy channels to report concerns
- **Plan for model retirement:** Not all AI experiments work long-term, so define criteria for when to sunset underperforming models

Continuous AI Governance Cycle

Plan

Set governance KPIs and metrics

Define acceptable thresholds for model performance

Do

Monitor AI performance continuously

Collect user feedback

Track key metrics in production

Act

Update and retrain models

Refine governance policies

Train teams on improvements

Check

Audit results against established criteria

Identify gaps and risks

Assess compliance alignment

About Octave

Octave is a leader in enterprise software, turning data into decisive action and intelligence into your edge. Our software solves for and simplifies complexity, from the design and build to operations and protection of people, property and assets – for any scope, at any scale. For decades, we've partnered with customers to sharpen performance, elevate efficiency and amplify results. From factory floors to entire cities, our solutions are tuned to scale up what's possible from day one onward.