CORTEX

EXECUTIVE BRIEFING 03

# Runtime Governance: Beyond Policy and Retrospective Audit

Why AI governance must move beyond policy, approval and retrospective review into operational use.

**Cortex**

Governed AI operational infrastructure

Public briefing document · v0.1 · 2026

This document is intended as an institutional briefing note for senior readers considering governed AI adoption in operationally complex environments.

## Executive summary

**AI governance cannot rely on policy alone.**

As AI moves from experimentation into operational use, organisations need governance that can be connected to activity while it is taking place. Policy, procurement review, risk assessment and audit remain important, but they are not sufficient on their own when AI begins to influence real work, operational decisions, service delivery or accountable outcomes.

Runtime governance is the discipline of governing AI-enabled activity in use. It connects organisational intent to operational conditions: what is permitted, what is visible, what is evidenced, what boundaries apply and who remains accountable.

This matters because AI use can change quickly. Models, providers, prompts, data sources, workflows, users and connected tools may all alter the conditions under which AI operates. Governance therefore needs to be more than a static approval event. It needs to remain connected to the activity it is intended to govern.

Cortex treats runtime governance as a core requirement for governed AI operations.

## The issue

**Many organisations approach AI governance through familiar control points.**

They create AI policies. They assess suppliers. They review proposed use cases. They define acceptable-use rules. They examine model risk. They set approval routes. They establish review boards, principles, documentation and assurance checkpoints.

These are necessary steps. They help organisations set intent and create boundaries around adoption.

But a gap can appear once AI moves into operational use.

The organisation may know that a use case was approved. It may know which supplier was selected. It may know which policy applies. It may know which team owns the process. But it may not be able to see clearly how AI-enabled activity is behaving during use, whether the original conditions still apply, what evidence is being created or how accountability remains visible in practice.

That gap becomes more important as AI moves closer to operational work.

A policy can say what should happen. A risk assessment can describe what was expected. A supplier review can examine the provider. An audit can look back at selected evidence. But operational governance needs to understand what is happening while AI-enabled activity is actually being used.

This is where runtime governance becomes important.

## Why policy and audit are not enough

**Policy is essential, but it is not the same as operational control.**

A policy may define acceptable use, prohibited use, escalation routes, approval requirements or accountability expectations. But unless policy intent is connected to operational activity, the organisation may struggle to know whether those expectations are being met during use.

Retrospective audit is also essential, but it is not the same as live governance.

An audit can review evidence after the event. It can test whether controls operated as intended. It can examine whether governance processes were followed. But audit depends on evidence being available, meaningful and connected to the relevant operational context.

AI complicates this because use can be fluid.

A user may change a prompt. A workflow may change. A connected tool may be added. A model may be updated. A provider may alter capability. A team may start using AI for adjacent work. A low-consequence use may become more significant because it becomes embedded in a real process.

None of this means AI should be treated as ungovernable. It means governance has to stay connected to the conditions of use.

If governance only operates before or after AI activity, organisations may miss what happens in between.

## Why it matters operationally

**Runtime governance matters because AI-enabled activity can influence work before traditional assurance processes notice.**

In many organisations, operational activity is made up of small steps: information is gathered, interpreted, summarised, recommended, reviewed, routed, approved, acted upon or escalated. AI may influence one or more of those steps without being the final decision-maker.

That influence still matters.

It may shape what a staff member sees. It may affect how quickly a case is handled. It may change the framing of a recommendation. It may prioritise one option over another. It may trigger an action in a connected tool. It may create a record that later becomes part of an assurance, legal, service or operational process.

If runtime governance is weak, organisations may not have a clear view of:

- where AI is being used;
- what policy boundary applies;
- which operational context surrounds the activity;
- what tools, systems or data are involved;
- what evidence is captured;
- whether use remains within the intended conditions;
- who can review what happened;
- who remains accountable.

In high-consequence environments, those questions cannot be left until after adoption has scaled.

The issue is not whether governance exists on paper. The issue is whether governance can remain meaningful during operational use.

## The Cortex view

Cortex starts from the position that AI governance must become operational.

That does not mean every governance decision should be automated. It does not mean human accountability should be replaced. It does not mean technology can guarantee compliance, safety or assurance.

It means the organisation needs a way to connect governance intent to runtime activity.

For Cortex, runtime governance is about making the conditions of AI use visible, governable and reviewable while activity is operating. It connects policy intent, operational context, evidence, boundaries and accountability.

Cortex Conduit is the platform plane associated with this need. It supports runtime governance by helping AI-enabled activity operate within defined organisational conditions.

The Runtime Control Plane describes where this governance is coordinated, applied and evidenced. It is not a claim that every operational risk can be eliminated. It is a way of understanding where governance needs to become active, observable and accountable as AI enters operational use.

This matters because AI governance cannot remain a static layer around dynamic activity. As AI use changes, governance needs to remain connected to what is happening.

## Runtime governance connects

Runtime governance helps connect governance intent to operational use.

| RUNTIME GOVERNANCE ELEMENT | PUBLIC ROLE |
|---|---|
| **Policy intent** | Connects organisational rules, permissions and expectations to operational activity. |
| **Operational context** | Relates AI use to the people, processes, systems, data and dependencies around it. |
| **Runtime activity** | Focuses governance on what happens while AI-enabled work is being carried out. |
| **Boundaries** | Helps define where use is permitted, constrained, escalated or reviewed. |
| **Evidence** | Supports review by capturing structured material about activity and conditions of use. |
| **Accountability** | Keeps responsibility visible when AI contributes to decisions, recommendations or actions. |

Runtime governance gives organisations a practical frame for asking whether AI use remains within the conditions under which it should operate.

## What organisations should consider

**When moving AI into operational use, organisations should ask runtime governance questions early.**

They should ask:

- What governance intent applies to this AI use?
- Where is that intent expressed: policy, risk control, approval, operating procedure or assurance requirement?
- How is that intent connected to activity during use?
- What operational context surrounds the activity?
- Which users, systems, tools, data sources or providers are involved?
- What boundaries define permitted, constrained or escalated use?
- What evidence is captured while the activity operates?
- Who can review that evidence?
- How will the organisation know if conditions change?
- Who remains accountable for the outcome?

These questions help organisations move from static governance artefacts to operational governance capability.

They are especially important where AI use may affect public services, regulated activity, security-sensitive work, infrastructure operations, mission environments, customer outcomes, legal processes, financial decisions or other high-consequence contexts.

A runtime governance approach does not remove the need for policy, risk management, audit or assurance. It strengthens their practical connection to the way AI is actually used.

## How Cortex relates

Cortex supports runtime governance by connecting AI-enabled activity to operational context, visibility, evidence, boundaries and accountability.

Cortex Atlas helps establish the operational context around AI use.

Cortex Conduit supports runtime governance of AI-enabled activity.

Cortex Lens helps make activity observable and traceable.

Cortex Ledger supports structured evidence and accountability records.

Cortex Gate helps govern interoperability boundaries.

Cortex Bridge supports provider and runtime abstraction as AI supply and deployment conditions change.

Together, these planes help organisations reason about AI governance as something that operates through the lifecycle of use, not only before or after it.

Cortex does not replace an organisation's legal, audit, compliance, security or assurance responsibilities. It does not claim to automate accountability or guarantee compliant outcomes. It provides a structured foundation for connecting governance to operational AI use.

## Closing statement

**AI governance must move closer to the point of use.**

Policies, approval routes and retrospective audits remain important, but they cannot carry the full burden of governing AI once it becomes part of operational work.

Organisations need to understand not only whether AI was approved, but how it is being used, what context surrounds it, what evidence exists and who remains accountable.

Runtime governance is how AI governance becomes operational.

That is why Cortex places runtime governance at the centre of governed AI adoption.

SUGGESTED ONWARD READING

- Explore **Runtime Governance** to understand how Cortex frames governance while AI-enabled activity operates.
- Read **Runtime Control Plane** to see where runtime governance is coordinated, applied and evidenced.
- Review **Cortex Conduit** to understand the platform plane associated with runtime governance.
- Read **Governance Evidence** to understand how structured evidence supports review and accountability.
- For broader context, read **Briefing 01: AI as Operational Infrastructure** and **Briefing 02: Why Operational Context Matters for AI Governance**.