CORTEX

— APPLIED BRIEFING 05

Assurance, Evidence and Accountability

Why AI assurance depends on structured evidence, operational traceability and visible responsibility.

Cortex

Governed AI operational infrastructure

Applied briefing document · v1.0 · 2026

This document is intended as an institutional briefing note for senior readers considering governed AI adoption across government, defence, critical infrastructure, regulated enterprise and other accountability-intensive environments.

Executive summary

AI assurance depends on evidence.

As AI moves from experimentation into operational use, organisations need more than statements of intent, supplier confidence, policy controls or retrospective review. They need evidence that can support assurance, governance, accountability and learning when AI contributes to decisions, recommendations, analysis, communications, operational activity or service outcomes.

This matters because AI-enabled activity may be difficult to review if it is not connected to operational context. A model output, usage log or approval record may show part of the picture, but assurance often needs to understand what the activity related to, what process it affected, what data was involved, which boundary applied, who reviewed it and who remained accountable.

Assurance cannot rely on isolated artefacts alone.

Organisations need structured evidence that connects AI use to the people, processes, systems, data, suppliers, controls, dependencies and outcomes around it. Without that connection, AI governance can become difficult to explain and difficult to trust.

Evidence does not automatically prove that AI use was safe, compliant, fair, lawful, effective or appropriate. It does not replace legal, audit, compliance, security, safety, regulatory or operational assurance. But it provides the material those functions need in order to ask better questions, review activity and understand accountability.

Cortex treats assurance, evidence and accountability as core requirements for governed AI operations.

The issue

AI governance is increasingly being discussed through the language of assurance.

Organisations want confidence that AI use is lawful, safe, secure, fair, robust, explainable, accountable and aligned with organisational intent. Boards, regulators, audit teams, security authorities, public bodies, risk functions, assurance teams, procurement teams and operational leaders may all ask how AI adoption will be governed.

These questions are necessary.

But assurance can become abstract if it is not supported by evidence from real operational use.

A policy may define acceptable AI use. A risk assessment may classify a use case. A supplier review may assess a provider. A model test may examine performance. A training record may show that users were briefed. A governance board may approve deployment. A technical log may record activity.

Each of these artefacts may be useful.

But none of them alone necessarily explains how AI was used inside an operational process, what context surrounded it, what evidence was created, what boundary applied or who remained accountable.

This creates a gap.

Organisations may have assurance activity around AI, but limited assurance evidence about AI-enabled work.

The central question is therefore not simply: Has AI been assessed?

It is: Can AI-enabled activity be reviewed in context, with evidence that is meaningful enough to support assurance and accountability?

Why AI assurance needs operational evidence

AI assurance needs operational evidence because AI use does not happen in isolation.

AI may support drafting, analysis, summarisation, prioritisation, classification, planning, triage, design, investigation, operational support, customer interaction or decision support. In each case, the significance of AI use depends on the environment around it.

The same AI capability may be low consequence in one setting and high consequence in another. A tool used to summarise internal meeting notes is not the same as a tool used to summarise case evidence, incident material, safety-adjacent information, regulatory submissions, operational plans or customer complaints.

The model may be similar. The assurance requirement is not.

Assurance therefore needs to understand operational context.

It needs to know which process AI supported, what data was involved, which systems or tools were connected, which user or role used it, what output was produced, whether the output influenced an action, what control boundary applied and who remained responsible.

Without that context, assurance can become dependent on assumptions.

A log may show that a model was used, but not what operational activity it affected. A supplier statement may describe platform controls, but not how those controls relate to the organisation's accountability. A policy may define rules, but not show whether activity remained within them. A human review may be recorded, but not whether the reviewer had enough context to make a meaningful judgement.

Operational evidence helps close that gap.

It connects assurance to what actually happened.

Why traceability matters

Traceability is the ability to follow the relevant path of AI-enabled activity.

That path may include the user, role, process, data, system, tool, provider, model, output, decision point, control boundary, review activity and accountability record.

Traceability matters because AI may influence work indirectly.

It may not make the final decision, but it may shape the information that a person sees. It may not approve an action, but it may recommend a next step. It may not own a customer outcome, but it may draft material that becomes part of a customer communication. It may not control infrastructure, but it may help interpret operational data. It may not decide a risk rating, but it may summarise evidence used by someone who does.

If these influences are not traceable, assurance becomes harder.

The organisation may struggle to distinguish human judgement from AI support, supplier behaviour from internal action, model output from workflow routing, or operational context from technical activity.

Traceability helps turn isolated records into reviewable evidence.

It allows assurance stakeholders to ask not only “was AI used?” but also “what did it relate to?”, “why did it matter?”, “what boundary applied?”, “who reviewed it?”, “what evidence exists?” and “who remained accountable?”

In accountability-intensive environments, those questions are essential.

Why accountability can become harder to see

AI can make accountability harder to see if organisations do not deliberately preserve it.

In most serious environments, AI should not replace organisational responsibility. A public body remains accountable for public functions. A regulated firm remains accountable for customer outcomes and control effectiveness. A defence organisation remains accountable for operational judgement. A critical infrastructure provider remains accountable for service resilience and operational safety.

But AI may complicate how responsibility is understood.

If an AI-generated summary influences a decision, who reviewed the summary? If a model recommends an action, who accepted or rejected it? If an AI tool drafts a communication, who checked it? If a supplier embeds AI into a workflow, who understands the change? If a model update alters behaviour, who knows whether the assurance position has changed?

The answer cannot be: the AI did it.

Nor can it simply be: a human was in the loop.

Accountability requires visible responsibility. The organisation must be able to show how AI contributed, how human judgement was applied, what governance conditions existed and what evidence remains for review.

This is why accountability depends on evidence.

Evidence does not remove responsibility. It helps responsibility remain visible.

Why assurance cannot be only retrospective

Retrospective review is important, but it is not enough.

Audit, assurance and investigation often look backwards. They review evidence after the event, test whether controls operated as intended, examine records, identify weaknesses and recommend improvement.

This remains essential.

But AI-enabled activity may change quickly. Prompts may change. Users may change. Workflows may evolve. Data sources may be added. Tools may be connected. Suppliers may update features.

Providers may alter model behaviour. AI may be used in adjacent ways that were not part of the original approval.

If assurance only looks back, organisations may not understand when the conditions of use have changed.

This is why assurance needs a runtime connection.

Runtime governance does not mean every decision is automated. It means governance remains connected while AI-enabled activity operates. The organisation can see enough of the activity to understand whether it remains within intended conditions, whether boundaries are being crossed, whether evidence is being captured and whether escalation or review is needed.

Assurance therefore needs both retrospective and operational dimensions.

It needs evidence after the event, and visibility during use.

Why evidence quality matters

Not all evidence is equally useful.

AI assurance may be weakened if evidence is incomplete, fragmented, over-technical, disconnected from context, inaccessible to reviewers or controlled only by suppliers.

A technical log may show events, but not organisational significance. A model card may describe a system, but not how it was used. A supplier assurance document may support procurement, but not runtime accountability. A governance approval may show intent, but not operational behaviour. A user record may show an outcome, but not AI contribution.

Evidence quality depends on whether the evidence can answer the questions assurance stakeholders need to ask.

Useful evidence should be:

- Connected to operational context.
- Structured enough to support review.
- Proportionate to consequence.
- Interpretable by authorised stakeholders.
- Linked to boundaries, controls and responsibilities.
- Available when needed.
- Protected according to sensitivity.
- Retained according to policy and obligation.
- Capable of supporting learning and improvement.

This does not mean every AI activity should be recorded in the same way. Evidence should be proportionate. Low-consequence use may require lighter records. High-consequence, regulated, public-facing, safety-adjacent, security-sensitive or operationally critical use may require stronger evidence and clearer traceability.

The important point is that evidence requirements should be designed before AI use scales.

Why it matters operationally

Assurance, evidence and accountability need to work in the real organisation.

Senior leaders need confidence that AI adoption is not creating unmanaged institutional exposure.

Boards and governance committees need evidence that AI use remains aligned with risk appetite and organisational intent.

Audit and assurance teams need reviewable material.

Risk and compliance teams need visibility of controls and boundaries.

Security teams need to understand providers, data movement, access and runtime conditions.

Legal and privacy teams need records that support purpose, accountability and data-handling review.

Operational teams need practical governance that fits the way work happens.

Procurement and supplier-management teams need evidence of supplier behaviour, contractual boundaries and continuity risks.

If these perspectives are not connected, assurance becomes fragmented.

One team may approve a use case. Another may manage the system. A third may own the operational process. A fourth may manage the supplier. A fifth may receive the audit question. Evidence may sit across logs, case records, workflow tools, supplier portals, spreadsheets, email and informal notes.

That fragmentation weakens assurance.

A governed approach should help these stakeholders reason from a shared evidence base: where AI was used, what it touched, what conditions applied, what evidence exists and who remained accountable.

The Cortex view

Cortex starts from the view that AI assurance must be connected to operational evidence.

Assurance cannot rely only on policy, approval, supplier statements, model assessments or retrospective audit. These remain important, but they need to be connected to the conditions of AI use.

For Cortex, assurance, evidence and accountability depend on six connected platform capabilities.

Cortex Atlas helps establish the operational context around AI use: the people, processes, systems, data, infrastructure, suppliers, controls and dependencies that shape consequence.

Cortex Conduit supports runtime governance while AI-enabled activity operates, helping connect governance intent to operational conditions.

Cortex Lens helps make AI-enabled activity observable and traceable, so that organisations can understand what happened and what it related to.

Cortex Ledger supports structured evidence and accountability records that can support assurance, audit, review, learning and governance conversations.

Cortex Gate helps govern the boundaries between AI, tools, systems, data, providers, suppliers and operating environments.

Cortex Bridge supports continuity as providers, models, runtime arrangements and deployment conditions change.

Together, these planes provide a foundation for assurance that is connected to AI-enabled activity in use.

This is not a claim that Cortex delivers assurance by itself. Cortex does not replace legal review, audit, compliance, security, safety assurance, regulatory engagement, public accountability, operational review or executive responsibility.

It provides structured evidence and operational visibility that can support those functions.

Assurance, evidence and accountability require

AI assurance requires more than confidence in AI capability alone.

Requirement	Public role
Operational context	Understand what AI activity related to and why it mattered.
Traceability	Follow the relevant path of AI-enabled activity across users, systems, data, tools and outputs.
Structured evidence	Capture material that can support review, assurance, audit, learning and accountability.
Runtime governance	Connect policy, control and assurance intent to activity while it operates.
Boundaries	Show which permissions, constraints, providers, tools and systems applied.
Reviewability	Make evidence interpretable by authorised stakeholders in the relevant operational context.

Continuity	Maintain evidence and governance coherence as models, providers, tools and runtimes change.
Accountability	Keep responsibility visible when AI contributes to decisions, recommendations, communications or actions.

These requirements help organisations move from AI assurance as documentation to AI assurance as an operational capability.

What organisations should consider

When designing AI assurance, organisations should ask evidence-first questions.

They should ask:

- What AI-enabled activity needs to be assured?
- Which process, service, function, customer journey, mission activity, operational workflow or regulated control does it relate to?
- What data, systems, tools, suppliers or providers are involved?
- What output does AI produce?
- Could that output influence a decision, recommendation, action, communication, escalation or record?
- What boundary or control condition applies?
- What evidence is captured while the activity operates?
- Can the evidence be interpreted in operational context?
- Who can access and review the evidence?
- How is sensitive evidence protected?
- How long should evidence be retained?
- Who remains accountable for the outcome?
- How will the organisation know if conditions of use change?
- How will assurance remain coherent if the provider, model, tool, workflow or runtime changes?

These questions are not intended to prevent AI adoption. They are intended to support AI adoption that can be explained, reviewed and improved.

The more consequential the activity, the more important it becomes to define evidence and accountability requirements before AI use scales.

How Cortex relates

Cortex supports organisations that need AI assurance connected to operational context, traceability, evidence and accountability.

It helps connect AI-enabled activity to the environment around it.

It supports governance while AI activity operates.

It helps make activity visible and traceable.

It supports structured evidence and accountability records.

It helps frame AI connections as governed boundaries.

It supports continuity across changing providers, models, tools and runtime conditions.

This makes Cortex relevant to government, defence, critical infrastructure, regulated enterprise and other environments where AI adoption must be considered through assurance, evidence and institutional responsibility.

Cortex does not claim that AI adoption becomes risk-free. It does not claim legal compliance, audit sign-off, regulatory approval, safety assurance, security accreditation or operational authorisation. It provides a disciplined foundation for more reviewable, governable and accountable AI operations.

Closing statement

AI assurance depends on evidence that can be understood.

As AI becomes part of operational work, organisations need to know more than whether a tool was approved or a model performed well in testing. They need to understand how AI was used, what it related to, what evidence exists and who remained accountable.

Assurance without evidence is fragile. Evidence without context is weak. Accountability without traceability is difficult to prove.

Governed AI adoption requires all three.

That is why Cortex treats assurance, evidence and accountability as core foundations for AI operations that can be explained, reviewed and trusted.

SUGGESTED ONWARD READING

- Explore Evidence, Traceability and Accountability in AI Operations to understand why AI-enabled activity needs to be visible, traceable and reviewable.
- Read Runtime Governance to understand how governance can remain connected while AI-enabled activity operates.
- Review Why Operational Context Matters for AI Governance to understand why evidence needs to be connected to people, processes, systems, data and dependencies.
- Read Governing AI Connections, Tools and Providers to understand why AI boundaries and provider relationships matter for assurance.
- Review Sovereign and Regulated AI Operations to understand why high-consequence environments need governance connected to context, evidence, continuity and accountability.