CORTEX

— APPLIED BRIEFING 02

AI Governance in Defence

Why defence environments need AI governance connected to operational context, mission consequence, evidence and control.

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 defence, national-security-adjacent and high-assurance operational environments.

Executive summary

AI adoption in defence carries operational consequence.

AI may support analysis, planning, intelligence preparation, logistics, engineering, simulation, training, cyber operations, procurement, assurance, operational support, mission planning or decision support. These uses may improve speed, coordination and situational understanding, but they also raise questions of control, assurance, security, accountability, resilience, dependency and operational consequence.

For defence, AI governance cannot be treated only as a digital innovation issue, a model-performance issue or a supplier-selection issue. It must be connected to the operational environment in which AI is used.

Defence organisations need to understand where AI is being introduced, what mission or support process it touches, what data and systems are involved, which security or assurance conditions apply, what evidence is created, which boundaries are crossed and who remains accountable for decisions, recommendations or actions.

This is especially important because defence environments are complex, federated and sensitive. Activity may span commands, agencies, suppliers, coalition partners, security domains, legacy systems, operational networks, deployed environments, information-sharing arrangements and mission-specific constraints.

AI introduced into this environment can inherit that complexity.

AI governance in defence therefore requires more than policy, principles or experimentation controls. It requires governance that is connected to operational context, applied during use, supported by evidence and resilient to provider, runtime and mission change.

Cortex treats these requirements as central to governed AI adoption in defence and high-assurance environments.

The issue

Defence organisations are under pressure to understand and adopt AI.

AI may be seen as a way to improve operational tempo, reduce analytical burden, support decision-making, strengthen logistics, accelerate planning, improve engineering insight, enhance cyber defence or improve the management of complex information environments.

These opportunities are significant.

But defence adoption carries a different governance burden from ordinary enterprise use.

A defence organisation may need to consider security classification, operational consequence, mission assurance, sovereign control, coalition interoperability, data sensitivity, supplier dependency, continuity, legal review, rules of engagement, safety-adjacent activity, command accountability and the operational impact of degraded or incorrect information.

In these settings, an AI capability cannot be judged only by whether it produces useful outputs.

The organisation must understand where the capability is operating, what it can influence, what systems or data it touches, what assumptions it relies on, what boundary applies, what evidence is retained and how accountability remains visible.

This matters even where AI is not the final decision-maker.

AI may influence how information is interpreted, which options are considered, which risks are highlighted, how quickly a response is prepared or how confidence is placed in an assessment. In defence environments, those influences may carry operational significance.

The central question is therefore not simply: Can AI support defence activity?

It is: Can AI be adopted in a way that remains governable, reviewable and accountable within the mission, security and operational conditions in which it is used?

Why defence AI governance is different

Defence environments are different because the relationship between information and action can be consequential.

Information may shape operational planning, force protection, logistics, prioritisation, threat assessment, engineering decisions, cyber response, operational readiness or mission support. Even when a human remains responsible, AI can affect the information available to that human and the way that information is understood.

Defence also operates across multiple layers of constraint.

Some activity is highly sensitive. Some depends on legacy systems. Some involves suppliers or sovereign industrial partners. Some depends on coalition or cross-domain interoperability. Some must operate in degraded, denied, contested or disconnected environments. Some may involve safety, legal, ethical or operational-assurance conditions. Some must be reviewed without exposing sensitive detail.

AI governance needs to reflect those realities.

A general policy may say that AI should be used responsibly. A procurement review may approve a tool. A security assessment may examine a provider. A model test may examine output quality. These controls are important, but they do not fully answer how AI is being used inside defence activity.

Defence needs governance that understands operational context.

It needs to know which process AI supports, which operational or support function may be affected, what information is involved, which boundary applies, what evidence exists, what system or provider dependency has been introduced and who remains accountable.

Without that understanding, AI adoption can create blind spots.

A low-risk tool may become more consequential if connected to sensitive information. A productivity use may become operationally important if embedded in planning. A supplier feature may alter assurance conditions. A model update may change behaviour. A runtime change may affect where evidence is captured or whether governance remains visible.

In defence, the conditions of use matter as much as the capability itself.

Why operational context matters

AI cannot be governed properly in defence if the surrounding operational context is unclear.

Operational context includes the people, roles, units, processes, systems, data, infrastructure, suppliers, security domains, operational dependencies and decision points around AI-enabled activity.

That context determines consequence.

The same AI capability may be low consequence in one setting and significant in another. A summarisation tool used to prepare an internal administrative note is not the same as a summarisation tool used to support operational planning, intelligence preparation, incident response, logistics prioritisation or mission-assurance work.

The model may be similar. The operational context is not.

Defence organisations therefore need a shared picture of where AI is being introduced and what it touches.

They need to understand whether AI is operating in strategic, operational, tactical, support, engineering, training, procurement, intelligence, cyber or administrative contexts. They need to understand whether activity is connected to classified data, operational systems, deployed networks, mission partners, suppliers, coalition environments or decision-support processes.

If that context is missing, governance becomes abstract.

A governance board may approve a use case without seeing the dependencies around it. A team may use AI for analysis without creating evidence for review. A supplier may add AI capability into a platform without fully exposing the operational implications. A human may remain accountable, but the organisation may not be able to reconstruct what AI contributed.

Operational context helps defence organisations connect AI governance to reality.

It makes it possible to reason about where governance should apply, what level of evidence is needed, what boundaries are required and how assurance should be proportionate to consequence.

Why evidence and assurance matter

Defence AI governance needs evidence that can support review.

Evidence is not the same as compliance. It is not the same as certification. It does not remove the need for legal, security, safety, operational or mission assurance.

But without structured evidence, defence organisations may struggle to understand how AI-enabled activity occurred, what conditions applied and whether use remained within intended boundaries.

They may need to review:

- What AI was used for.
- Which mission, support or operational process it related to.
- What information, data or systems were involved.

- Which provider, model or runtime was used.
- What output was generated.
- Whether the output informed an assessment, recommendation or action.
- What governance, security, legal, operational or assurance boundary applied.
- Who reviewed the activity.
- Who remained accountable.
- What changed as the activity moved through systems, teams or operational environments.

In defence contexts, evidence may also need to be handled carefully. It may need to preserve sensitive information, respect classification boundaries, avoid unnecessary exposure and remain useful to authorised reviewers.

This creates a practical challenge.

Evidence must be structured enough to support review, but controlled enough to fit the security environment.

Traceability is central to this. Defence organisations need to understand not only that AI was used, but what it related to, which operational context surrounded it, which boundary applied and what evidence is available for authorised review.

Accountability also needs to remain visible.

A statement that a human is responsible is not enough if the organisation cannot see what information the human relied on, what AI contributed, how the activity was governed and what evidence remains.

In defence, accountability must be preserved through the operational chain, not only asserted at the end of it.

Why boundaries and control matter

AI governance in defence must pay close attention to boundaries.

A boundary is where conditions matter.

It may define which users can access a capability. It may separate security domains. It may control what data can be used. It may govern whether AI can connect to tools or systems. It may determine which provider or runtime is permitted. It may define what evidence must be captured. It may require escalation, review or human authorisation.

AI connections can change consequence.

A standalone drafting tool may carry one level of risk. The same capability connected to classified repositories, operational planning tools, engineering systems, logistics platforms, cyber tooling or coalition information environments may carry a different level of risk.

The connection changes what AI can touch.

It may expand access to sensitive data. It may introduce a provider dependency. It may create a new evidence requirement. It may affect reviewability. It may alter the assurance boundary. It may change who needs to approve, monitor or be accountable for use.

Defence organisations therefore need to govern AI connections as operational boundaries, not simple technical integrations.

Provider and runtime change also matter.

AI supply is dynamic. Models change. Providers change. Deployment options change. Capabilities change. Terms change. Costs change. Availability changes. Embedded AI features appear inside existing tools. Defence organisations may need to restrict, replace, combine or abstract providers depending on mission, sovereignty, security or continuity requirements.

If governance is locked to a single provider, implementation or model, it may become fragile.

Defence needs governance that can remain coherent as the AI supply environment changes.

Why it matters operationally

AI governance in defence needs to be practical for the way defence actually operates.

Senior leaders need to understand where AI creates advantage and where it creates operational exposure.

Command and operational sponsors need to understand how AI affects mission support, planning, readiness or decision support.

Security authorities need to understand data, providers, runtime environments and boundaries.

Architecture and engineering teams need to understand systems, dependencies, interfaces and resilience.

Assurance stakeholders need evidence that can support review without overexposing sensitive detail.

Commercial and procurement teams need to understand supplier, continuity and sovereignty implications.

Operational users need to know how AI fits into real work and what constraints apply.

If these perspectives are not connected, AI adoption can fragment.

One team may approve the use case. Another may own the system. A third may manage the supplier. A fourth may hold operational responsibility. Evidence may sit elsewhere. The result may be a capability that appears useful but is difficult to govern, review or assure.

A governed approach should help defence organisations reason from a shared operational picture: where AI is being introduced, what it touches, which boundary applies, what evidence is available and who remains accountable.

The Cortex view

Cortex starts from the view that defence AI governance must be connected to operational reality.

AI governance in defence should not remain a static policy layer around dynamic operational activity. It should connect intent, context, boundaries, evidence, runtime governance and accountability.

For Cortex, governed AI adoption in defence requires six connected capabilities.

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

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

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

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

Cortex Gate helps govern the boundaries between AI, tools, systems, data, suppliers, security domains and operational environments.

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

Together, these planes provide a foundation for AI governance in defence and high-assurance environments.

This is not a claim that Cortex provides mission assurance, security accreditation, sovereign approval, legal compliance, safety certification or operational authorisation. Cortex does not replace defence governance, security authority, legal review, operational command, safety assurance, procurement scrutiny or mission accountability.

It provides a structured way to make AI-enabled activity more understandable, governable, reviewable and accountable.

AI governance in defence requires

Defence AI adoption requires more than AI capability alone.

Requirement	Public role
Operational context	Understand where AI is introduced into defence activity and what mission or support process it touches.
Runtime governance	Apply governance while AI-enabled activity operates, not only before or after use.
Traceability	Make AI activity visible enough for authorised review and operational understanding.

Evidence	Support assurance, governance, learning and accountability conversations.
Boundaries	Govern connections between AI, systems, tools, data, providers and security domains.
Continuity	Maintain governance coherence as models, providers, runtime arrangements and operating conditions change.
Accountability	Keep responsibility visible when AI contributes to assessments, recommendations, planning or action.

These requirements help defence organisations move from AI experimentation to governed AI operations.

What defence organisations should consider

When considering AI adoption, defence organisations should ask operational governance questions early.

They should ask:

- Where is AI being introduced into defence activity?
- Which mission, support, operational, engineering, logistics, cyber, intelligence, training or administrative process does it relate to?
- What users, roles, units, commands, suppliers or partners are involved?
- What data, information, records or operational material does the AI activity access, generate, transform or expose?
- Which systems, tools, infrastructure, networks or providers does it depend on?
- What classification, security, legal, operational, assurance or policy boundary applies?
- How is governance applied while the activity operates?
- What evidence is captured?
- Can that evidence be interpreted by authorised reviewers in operational context?
- Who can review the activity?
- Who remains accountable for decisions, recommendations or actions?
- What happens if the provider, model, runtime or operating condition changes?

- How will governance remain coherent across commands, domains, suppliers, partners or deployed environments?

These questions are not intended to prevent useful AI adoption. They are intended to support adoption that can be governed, assured and trusted in operationally consequential settings.

The more consequential the activity, the more important it becomes to answer these questions before AI use scales.

How Cortex relates

Cortex supports defence and high-assurance organisations considering AI adoption in operationally complex environments.

It helps connect AI-enabled activity to the operational context 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, runtimes and deployment conditions.

This makes Cortex relevant where AI adoption must be considered through mission consequence, security, assurance, operational resilience, evidence and accountability.

Cortex does not claim that AI adoption becomes risk-free. It does not claim mission assurance, sovereign certification, security accreditation, safety assurance or legal compliance. It provides a disciplined foundation for more governable, reviewable and accountable AI operations.

Closing statement

Defence AI adoption cannot be treated as ordinary technology adoption.

When AI enters defence environments, organisations need to understand more than whether the capability works. They need to understand where it operates, what it touches, what conditions apply, what evidence exists and who remains accountable.

AI governance in defence must be connected to operational context, applied during use, supported by traceability and evidence, bounded across systems and providers, and resilient to change.

That is how defence organisations can pursue AI advantage without losing sight of control, assurance and accountability.

That is why Cortex treats defence AI governance as a question of context, runtime governance, evidence, boundaries, continuity and accountable operational control.

SUGGESTED ONWARD READING

- Explore AI as Operational Infrastructure to understand why AI governance must connect operational context, runtime governance, evidence and accountability.
- Read Why Operational Context Matters for AI Governance to understand why AI use needs a shared picture of people, processes, systems, data and dependencies.
- Review Runtime Governance to understand how governance can remain connected while AI-enabled activity operates.
- Read Evidence, Traceability and Accountability in AI Operations to understand how structured evidence supports review and accountability.
- Review Sovereign and Regulated AI Operations to understand why government, defence, critical infrastructure and regulated environments need AI governance connected to context, evidence, continuity and accountability.