The Agentic Coalition

AI Agents and Decision Advantage in NATO and the EU

scale

Scale AI's Mission:

Develop reliable AI systems for the world's most important decisions

ABOUT SCALE AI

Founded in 2016 by Alexandr Wang, Scale's mission is to develop reliable AI systems for the world's most important decisions. Growing along with the evolving AI industry, Scale has worked with leading autonomous vehicle companies, frontier AI labs, Fortune 500 companies, and governments around the world. Today, Scale builds data and model training solutions as well as complex, multi-step AI applications for enterprise and public sector clients.

Since its early years, Scale has developed deep experience supporting the national security community, fielding proven and trusted capabilities across all levels of classification to tackle the most consequential operational problem sets. Our unique capabilities in training, testing, and evaluating AI models, coupled with our sophistication in building applications, make us the first-choice partner to help the U.S. and its allies and partners navigate the opportunities and threats posed by AI and the emerging "agentic battlefield". Scale is committed to delivering robust, mission-ready capabilities to ensure we secure our collective strategic advantage into the future.

AGENTIC WARFARE:

THE OPPORTUNITY

July 14th, 2026

We stand at a watershed moment of the most unfortunate kind.

The Pax Europa era is over, shattered by the return of high-intensity war to the continent and now in the greater Middle East. Europe's NATO members are rightly racing to rearm. But we must be clear-eyed and intentional: rearmament without making use of advanced AI tools and the decision advantage they provide will simply give us inventory—impressive hardware platforms on paper, but brittle in crisis.

Mass is essential, but the decisive variable is whether the Alliance can convert that mass into synchronised action at speed.

If we invest in one new capability, it should be the decisional architecture to make maximal use of our evolving defense assets. Agentic Warfare is how we seize that advantage. AI agents—dynamic systems capable of proactively planning, testing, and orchestrating multi-step actions under human authority—can now decisively compress the gap between insight and execution. By equipping allied commanders with a tireless digital staff, we can master the velocity and complexity that overwhelm traditional headquarters, enabling NATO forces to anticipate, adapt, and act faster than our most capable adversaries.

This discussion draft, building on a longer paper on Agentic Warfare released by Scale, explores how agentic systems piloted today in the U.S. Department of War can help to tackle critical NATO problem sets—from the GIUK Gap to Baltic escalation. Crucially, European nations need not make impossible choices. If architected properly, this new era of command can be built on federation and interoperability, ensuring seamless transatlantic integration while preserving sovereign governance. Together, the coalition can fight at the same tempo, from the same shared understanding, without compromising national freedom of action.

The strategic imperative is clear: we must build the agentic coalition infrastructure of tomorrow, today.

Dan Tadros
Head of Public Sector
Scale AI

Billy Long
Head of Global Defense
Scale AI

Executive Summary: The Promise of Agentic AI

European defence and security stand at a crossroads. As the post-Cold War security architecture collapses, European nations are racing to regenerate mass, stockpiles, and industrial capacity to restore deterrence and blunt aggression by Russia and in the near-abroad. The strategic problem is that Europe risks rearming for a battlespace that has already moved on. In modern conflict, the decisive question is not simply how much capability NATO fields, but how quickly the Alliance can sense, decide, and act.

Against opponents that weaponise ambiguity, NATO's slow and staff-intensive decision-making becomes a critical vulnerability. NATO is built for symmetric shared defence, while adversaries wage asymmetric campaigns carefully calibrated below an obvious Article 5 threshold, stripping the Alliance of optionality. Russia does not need to outfight NATO in a protracted war if it can outpace it in the first 72 hours—seizing ground, fracturing political cohesion, and daring the Alliance to reverse a fait accompli. In this context, decision advantage promises to be Europe's most immediate and most essential deterrence multiplier. New systems of AI agents already proven in pilot command and control systems in the U.S. Department of War must now quickly be mobilised in the Alliance.

For decades, software has accelerated traditional C2 workflows, enabling precision strike and network-centric warfare without resolving the human bottlenecks around sense-making and coordination. As outlined in Scale's foundational *Agentic Warfare* whitepaper[1] (https://scale.com/agentic-warfare), that paradigm is now breaking down with the emergence of agentic AI. These systems are defined not just by machine-speed correlation, but by agency: the ability to plan, test, and execute multi-step actions under human authority. Such systems promise a step-change in military power, equipping commanders with persistent machine-assisted staff support that can triage sensor data, stress-test operational plans, and

[1] *This paper provides UK, NATO and EU readers with a strategic overview of how agentic technologies can help tackle structural challenges in existing command and control and thereby unlock new levels of military decision advantage. It focuses on the operational logic for adoption and the institutional principles that should shape it, building on the foundational concepts set out in Scale's Agentic Warfare whitepaper, published in January 2026. Follow-on papers will examine specific governance, technical and assurance concepts in greater depth. This series reflects our commitment to advancing coalition understanding and delivering robust, mission-ready capabilities to secure our collective strategic advantage and security into the future.*

compress the gap between insight and execution. Deterrence strengthens when an adversary knows that NATO can see the battlespace more clearly, adapt more rapidly, and synchronise thousands of assets more precisely. That is the essence of Agentic Warfare: deterrence by decision advantage.

This transformation is already underway. Through the Defense Innovation Unit, the United States is moving rapidly to integrate agentic capabilities to maintain global overmatch. If U.S. forces fight at machine speed while European allies remain constrained by manual workflows, the result will be a widening capability gap that strains coalition cohesion. Yet it is precisely here, in the decision infrastructure that makes mass decisive, that European rearmament is dangerously splintered. Sovereignty instincts and domestic industrial policies risk driving architectures apart, creating vendor-locked 'AI islands' that cannot be trusted—or connected—when a crisis hits. If state-of-the-art AI is not taken up in NATO's command and control architecture, NATO's European pillar could spend historic sums rebuilding force structure and still lack the capacity to coordinate at the tempo modern war demands.

Sovereignty instincts and domestic industrial policies risk driving architectures apart, creating vendor-locked 'AI islands' that cannot be trusted—or connected—when a crisis hits.

This discussion draft argues that NATO's European members must treat decision advantage as the central component of combat power. The agentic systems already working in U.S. combatant commands offer a template. NATO's European pillar must now build the technical and trust architectures they require to adopt such systems—forging a shared, operational capability at the heart of NATO command and collective deterrence.

Crucially, building this connective tissue does not mean centralising around a single, monolithic system. Transatlantic integration at the decision layer and the close protection of sensitive national assets are not mutually exclusive. Rather, the core design principles must be federation and sovereign interoperability. In practice, a shared decision advantage stack places the agentic layer—the digital staff—between sovereign data inputs and coalition-readable outputs.

National control should be conceived not as a binary choice, but rather as a ‘Sovereignty Rheostat’: a mechanism to calibrate integration according to operational need, data sensitivity, and political risk. Just as NATO ballistic missile defence combines a shared C2 backbone with nationally owned interceptors, an agentic coalition preserves national control where required while enforcing integration where it produces a decisive collective effect.

To move the Agentic Coalition from concept to procurement, this paper outlines how agentic capabilities can be deployed inside real headquarters workflows. Doing so requires navigating several tensions: the imperative for rigorous, through-life assurance; acquisition frictions across 27 separate European procurement markets shaped by EU industrial incentives (EDF/PESCO); and Responsible AI mandates that provide the evidence—and permission framework—that commanders require to use agentic systems lawfully and confidently. This paper focuses on the strategic case for adoption of agentic decision systems and the institutional principles that should shape that adoption. Follow-on papers in this series will address operational evidence, governance, and technical and assurance concepts in greater depth

Europe is rearming. The question is whether it will modernise for the war it will face—one where ambiguity is weaponised, warning times shrink, and coalition speed matters as much as platform count. In this war, the side that decides and acts fastest shapes escalation and keeps the initiative. The Alliance must establish the standards, technical architectures, and assurance artefacts necessary for headquarters adoption of agentic systems at scale. With the right sequencing of (relatively modest) investment, experimentation, and governance, Europe’s NATO nations can strengthen decision advantage while preserving the political legitimacy on which collective defence depends.

PART I: A VISION OF ALLIANCE DECISION ADVANTAGE

The town of Narva on Estonia's northeastern border is often cited as a key test of NATO resolve and unity, one of many potential flashpoints across the Nordic-Baltic region. The terrain itself is not uniquely valuable. But its location and its large ethnic-Russian population makes it an ideal stage for a calculated Russian gamble: contest a small piece of NATO territory, keep escalation manageable, and expose the Alliance as divided and procedurally constrained.

The logic is both imperial and instrumental: intimidate its neighbours, undermine confidence in NATO's Article 5 guarantee, and widen Russian freedom of action across the continent.

Russia would not need a dramatic armoured thrust to generate a strategic shock. For years it has waged a continuous campaign against Europe using sabotage, cyber operations, electronic attack, disinformation, election interference, and deniable incidents at sea and in the air. These activities probe for seams in Western decision-making, exploiting ambiguity to delay attribution and fracture political consensus. In a Narva-style crisis, such tools would converge to cloud understanding while a fait accompli is imposed under a fog of competing narratives. These tactics exploit a structural asymmetry: NATO is designed for symmetric defence under Article 5, while adversaries increasingly pursue asymmetric actions below the threshold that would trigger collective response.

This is NATO's deterrence failure mode: not shortage of sensors or warnings, but a coalition that struggles to convert partial, differently classified truths into a shared, coalition-actionable decision at speed. Brittle interoperability, classification seams, and staff-heavy headquarters processes turn what should be machine-speed correlation into manual reconciliation across national chains of command. Beyond a political vulnerability, it is a failure of imagination and optionality. The slow pace of NATO's mandraulic staff processes fundamentally limits its ability to generate timely, proportionate counters to sub-threshold aggression. Meanwhile, Russia's strategy exploits this weakness, sustaining ambiguity long enough to narrow NATO's options and present the Alliance with a stark choice between escalation and accommodation.

The importance of NATO creating decisional advantage—to keep ahead of likely Russian assaults—extends beyond NATO's eastern flank. European NATO faces a complex, 360-degree threat landscape that includes persistent grey-zone pressure, instability across the Middle East and Mediterranean, and the growing alignment of authoritarian powers. The outbreak of conflict with Iran in early 2026 highlighted how quickly allied strategic bandwidth can be

stretched across multiple theatres. With U.S. attention increasingly divided, and supply chains for munitions and strategic assets under concurrent stress, Europe must assume greater responsibility for its own continental defence.

In this environment, slowness in decision-making—and the lack of readily available options to employ the full range of alliance power—is not just tactical friction, it is a strategic vulnerability. The question now is how to meet the moment. Russia's invasion of Ukraine triggered European rearmament after decades of underinvestment. But rearmament without using AI technologies to achieve decision advantage is inventory: impressive on paper, brittle in crisis. Military advantage in future conflicts will not come from stealthier jets, faster missiles, or larger drone swarms alone—though Europe needs more of all three.

What matters most is speed to action: the ability to identify, understand, and orchestrate effects faster than an adversary can exploit uncertainty. Strategic integration at the decision-layer therefore becomes the most immediate and effective route to strengthen deterrence—enabling NATO to change posture quickly, reducing friction across multinational command, and maximising the impact of existing forces without waiting years for new platforms to arrive.

AGENTIC WARFARE

Recent advances in agentic AI offer a direct path to realise this decision-layer integration, thereby unlocking a genuine step-change in military command tempo and optionality. While earlier generations of AI showed value in narrow analytic tasks like summarising information, classifying imagery, or detecting anomalies in data, they remained fundamentally reactive and operated within tightly bound workflows. Agentic systems represent a definitive break from that model. They are always-on, goal-driven systems capable of continuously monitoring data, reasoning across multiple sources, and proactively initiating tasks.

Whereas today's narrow and brittle AI capabilities can perform, at best, like a clever junior staffer, agentic systems unlock the full potential of AI to act as a genuine mission partner.

Whereas today's narrow and brittle AI capabilities can perform, at best, like a clever junior staffer, agentic systems unlock the full potential of AI to act as a genuine mission partner. Imagine a network of specialised AI agents working alongside commanders across the Observe, Orient, Decide, Act (OODA) loop. Together, they can tackle levels of complexity that would swamp the best human staff: filtering noise and deception, stress-testing thousands of scenarios in parallel, continuously updating plans, and pushing validated recommendations forward fast enough to matter. Crucially, these systems augment rather than replace human judgement. The commander still decides, but with far greater clarity and confidence.

That is the essence of Agentic Warfare: decision advantage across every echelon of command, enabled by human-machine teaming that accelerates coalition sense-making, planning, and coordination.

ADOPTION IN NATO AND EUROPE

Realising the potential of agentic systems is fundamentally harder in a multinational alliance than within the unified apparatus of a single nation-state. For NATO, the challenge is rarely a lack of sensors or capable staffs, but a problem of composition. Information across the Alliance exists as partial truths, held at varying classification levels and governed by strict national caveats. Headquarters must constantly reconcile these differences, often relying on slow, manual processes to sanitise data for release across coalition networks. As modernisation accelerates across Europe, the decision infrastructure required to make mass decisive is splintering, threatening to fracture the

Information across the Alliance exists as partial truths, held at varying classification levels and governed by strict national caveats. Headquarters must constantly reconcile these differences

shared operational picture exactly when a crisis demands machine-speed coordination.

Political and industrial crosswinds add additional friction. NATO demands seamless interoperability, while the EU, through mechanisms like Permanent Structured Cooperation (PESCO) and the European Defence Fund (EDF), heavily incentivises European supply security and industrial capacity. This drive for 'strategic autonomy' is rooted in acute operational risk. Fearful of vendor lock-in, opaque update cycles, and supply risk under extra-territorial pressure, European militaries simply cannot accept single points of failure in mission-critical systems.

If nations pursue AI-enabled decision tools as national or EU-only stacks, the result will be a fractured ecosystem of national stovepipes, incompatible data standards, and brittle interfaces.

Yet, while these sovereignty instincts are rational, they are self-defeating in a coalition context. If nations pursue AI-enabled decision tools as national or EU-only stacks, the result will be a fractured ecosystem of national stovepipes, incompatible data standards, and brittle interfaces. These vendor-locked AI islands would replicate and magnify the very interoperability problems NATO has spent decades trying to solve. In the agentic era, where operational advantage hinges on the decision layer's ability to orchestrate action inside an adversary's decision cycle, failing at the digital seam could prove disastrous.

The challenge, therefore, is not simply adopting new technologies, but doing so in a way that strengthens coalition integration rather than undermining it. Agentic capabilities must operate within NATO's federated structure, allowing nations to protect sensitive data and maintain political control while still enabling rapid sharing of the information and decisions that matter most. Large language models already demonstrate how this can work in practice: systems can translate complex operational data into structured summaries, risk assessments, or recommended actions that can be shared across systems without exposing the underlying sensitive inputs. Achieving that balance securely is the central task facing NATO and European defence ministries as they modernise. The Alliance must build a decision infrastructure capable of operating at machine speed without sacrificing the sovereignty, accountability, and political legitimacy that underpin collective defence.

The new class of agentic capabilities in development today cannot simply be layered onto fragmented digital foundations. Their success depends on the data architectures that NATO and European militaries are already beginning to build. Initiatives such as Federated Mission Networking (FMN) provide a framework for connecting national command networks into coalition mission environments, while emerging programmes like Digital Ocean seek to fuse sensors and maritime surveillance data across allied forces.

These efforts point toward a federated model of coalition warfare, but they remain focused primarily on networking and sensing. Realising the full potential of agentic systems will require extending this logic to the alliance decision layer itself. This begins with disciplined data foundations: common standards, machine-readable formats, and shared semantic models that allow systems across the Alliance to exchange and interpret operational information. Ultimately, AI is only as capable as the data environment in which it operates.

The Alliance must build a decision infrastructure capable of operating at machine speed without sacrificing the sovereignty, accountability, and political legitimacy that underpin collective defence.

PART II: FROM LLMs TO AGENTIC SYSTEMS

Over recent years, Ministries of Defence across Europe have experimented heavily with large language models and other generative AI tools across a wide range of missions and domains. Yet despite promising demonstrations, few have transitioned into operational military systems.

The reasons are structural. Traditional generative models are fundamentally reactive. They respond to prompts and produce outputs, but they do not independently track evolving situations, reconcile conflicting information, or orchestrate sequences of action. In practice, they help solve writing problems. However, writing is rarely the central bottleneck in military operations; it is sense-making, updating assumptions as conditions change, and translating decisions into coherent action across complex military systems.

Agentic AI systems address precisely this challenge. Rather than passive tools responding to queries, this new class of AI systems coordinate sequences of tasks: gathering information, evaluating alternatives, and updating conclusions as new data emerges. This extends AI beyond isolated analytical tools toward integrated systems capable of supporting decision-making across complex operational environments.

WHAT MAKES AI "AGENTIC"

What makes a system 'agentic' rather than simply 'generative'? It requires an architecture designed for operational execution:

- ***Continuous Loops:*** *Agents do not wait for a prompt. They operate in continuous monitor-reason-act loops, dynamically adapting to new data.*
- ***Tool Use:*** *Agents actively query databases, call specialised external models, and cue physical action.*
- ***Memory and Orchestration:*** *Agents track shifting assumptions over days or weeks, and hand off work between specialised sub-agents for optimised outcomes.*
- ***Bounded Delegation:*** *Agents operate under strict, engineered controls. In the near-term, authority is restricted to time-sensitive but reversible actions—such as retasking ISR coverage, prioritising alerts or drafting COAs—while retaining human approval for operational tasking and requiring explicit human authorisation as an immutable requirement for tasks related to escalatory action or kinetic targeting.*

The potential defence applications for agentic systems are wide-ranging, spanning logistics optimisation, cyber defence, intelligence fusion, targeting support, and operational planning. Much of the political and public debate regarding military AI often focuses on Lethal Autonomous Weapons Systems (LAWS). While important, lethal autonomy involves complex technical, ethical, legal, and alliance-wide governance challenges that will take years to resolve. It is a downstream challenge.

For NATO and Europe, the most immediate opportunity lies in the staff work that underpins command and control. Modern headquarters must process enormous volumes of information while coordinating complex multinational forces under tight time pressure. The friction in this process is concentrated at two points in the decision cycle: the Observe/Orient phase, where commanders struggle to extract signal from data noise, and the Decide phase, where planning is slowed by manual analysis and coordination. Here, agentic systems can support headquarters workflows at scale, compressing the gap between insight and execution and enabling coalition operations at the tempo modern warfare demands.

AGENTIC ALERTING: RESTORING THE SIGNAL-TO-NOISE RATIO

Agentic Alerting revolutionises how commands sense the battlefield. In today's saturated operational environments, the Observe/Orient phase is defined by a deluge of sensor data that overwhelms human cognition. Agentic Alerting restores the signal-to-noise ratio, empowering commanders to make rapid, high-stakes decisions, turning a reactive force into a proactive one. Its near-term promise is not a costly, brand-new sensor architecture, but a dynamic synthesis layer that sits on top of what the force already has: the existing sensors, data pipes, and analytic tools that work well in their lanes but remain siloed and relatively static. By ingesting massive volumes of multi-modal feeds—telemetry, imagery, traditional alerts, and operational context—it executes decision chains at machine speed and combines disparate inputs into a coherent, mission-relevant picture.

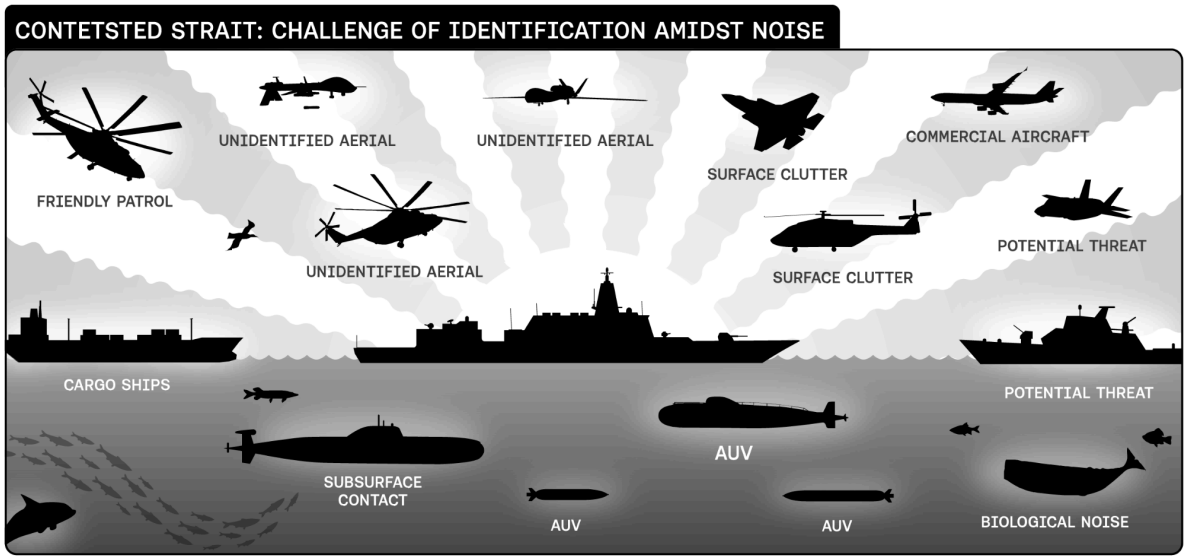

The key differentiator between Agentic Alerting and legacy fusion or correlation engines is judgement. Existing deployed systems operate at machine speed but are inflexible; they cannot reason across modalities or draw upon situational context outside their discrete rule sets. Agentic Alerting introduces dynamic intelligence into the military's real-time nervous system through:

- **Contextual reasoning:** Whereas traditional alert systems rely on fixed rules, often producing floods of false alerts, agentic systems interpret signals in context, comparing them with other data and recent activity. Instead of flagging every unusual sonar contact, for example, the system might weigh shipping traffic, prior detections, and intelligence reporting to prioritise the alerts that matter.
- **Autonomous investigation:** When anomalies appear, agents gather additional information automatically. A suspicious seabed vibration might cue a nearby drone or satellite sensor to confirm whether it reflects routine activity or a potential threat.
- **Network prioritisation:** Sensor networks generate more data than commanders can absorb. Agentic systems continuously identify which signals matter most, summarising routine activity while focusing sensing and analysis on emerging risks.

The result is a force that sees sooner, understands faster, and buys time for the right decisions.

CASE STUDY: DEFENDING THE ATLANTIC BASTION

The GIUK (Greenland-Iceland-UK) Gap is NATO's most critical maritime chokepoint, dense with vulnerable Critical National Infrastructure (CNI) such as subsea data cables and energy interconnectors. As the UK and allied nations deploy the hybrid Atlantic Bastion network to counter aggressive, grey-zone undersea activity, task forces face a dual bottleneck. First, the acoustic harshness of the High North generates paralysing false positives. Second, strict classification barriers—sovereign seams—prevent the rapid sharing of national data, forcing vital intelligence to be manually sanitised and granting hostile actors a critical time advantage.

Agentic Alerting can help to transform this vulnerable maritime flank into a more proactive deterrent. If a UK seabed array detects an anomaly, an agent could rapidly synthesise available multi-coalition feeds and recommend the cross-cueing of relevant sensors (surface warships, uncrewed underwater vehicles, surveillance aircraft) to verify the target. Because the AI summarises data locally, it can act as an automated security gateway, algorithmically abstracting the highly classified parameters of the detecting sensor, outputting only a sanitised, provenance-rich alert—what the threat is, where it is heading, and the mathematical certainty projection—directly to NATO Maritime Command. By translating raw, sovereign data into actionable, shareable intelligence at machine speed, this architecture would allow a multi-national force to instantly converge on a shared ground truth and neutralise the threat.

AGENTIC PLANNING: FROM STATIC BINDERS TO "LIVING PLANS"

The latest generation of agentic systems under development and testing within the U.S. Department of War points to yet another way NATO's European Nations can integrate AI—what Scale calls Agentic Planning. Military planning is core to coalition dominance, yet it relies on a methodical, linear process optimised for deliberation rather than speed. Producing theater-level operational plans can take up to two years, resulting in static binders of Courses of Action (COAs) that routinely fall victim to time lag and technological drift as the battlespace evolves. More critically, these static plans suffer from crisis inadaptability—once an adversary makes an offensive move, significant portions of standing plans become irrelevant, and manual updates are simply too slow to validate before a commander must act.

Agentic Planning systems that are being developed through the Defense Innovation Unit today address these vulnerabilities, rewiring how military staffs develop plans and allowing them to rapidly iterate as circumstances evolve. The near-term advantage of Agentic Planning is not that it uses AI tools to complete the staffwork of planning faster. It lies in revolutionising wargaming and simulation, which is frequently the first step abandoned during a crisis due to latency despite being a comparative alliance strength. Currently, rigorous modeling of combat scenarios requires a specialised cadre of analysts to operate fragmented systems over days or weeks. Agentic tools automate calls to simulators through a natural language interface, democratising access for broader planning staffs without requiring niche experts. Working in continuous loops, agents retrieve simulation data, analyse findings, and automatically initiate secondary runs with adjusted variables. This generates operational 'multiverses', running thousands of permutations to present planners with optimal, pre-validated COAs with the highest likelihood of mission success.

Ultimately, Agentic Planning systems will evolve these static products into 'living plans' by integrating live, multimodal data directly into planning feeds. By using agents to ingest real-time intelligence—from space-based imagery to logistics telemetry—while simultaneously validating new COAs through simulators, plans will essentially self-heal in response to battlefield events. Commanders will be provided validated, confidence-bound COAs adapted to current conditions on the battlefield. This ensures that when they make the decision to act, they are acting on tested, probability-weighted analysis rather than stale assumptions, transforming episodic planning into dynamic living plans. Instead of staffs pulling and synthesising raw intelligence, the system pushes only the most relevant, empirically validated details at the exact moment required, creating an entirely new paradigm of decision advantage. Crucially, that includes optionality against asymmetric threats: by testing a broad spectrum of approaches, from

non-kinetic and cyber to localised deterrence options, commanders can be rapidly provided with proportionate options to deny the adversary grey-zone advantage without triggering escalation.

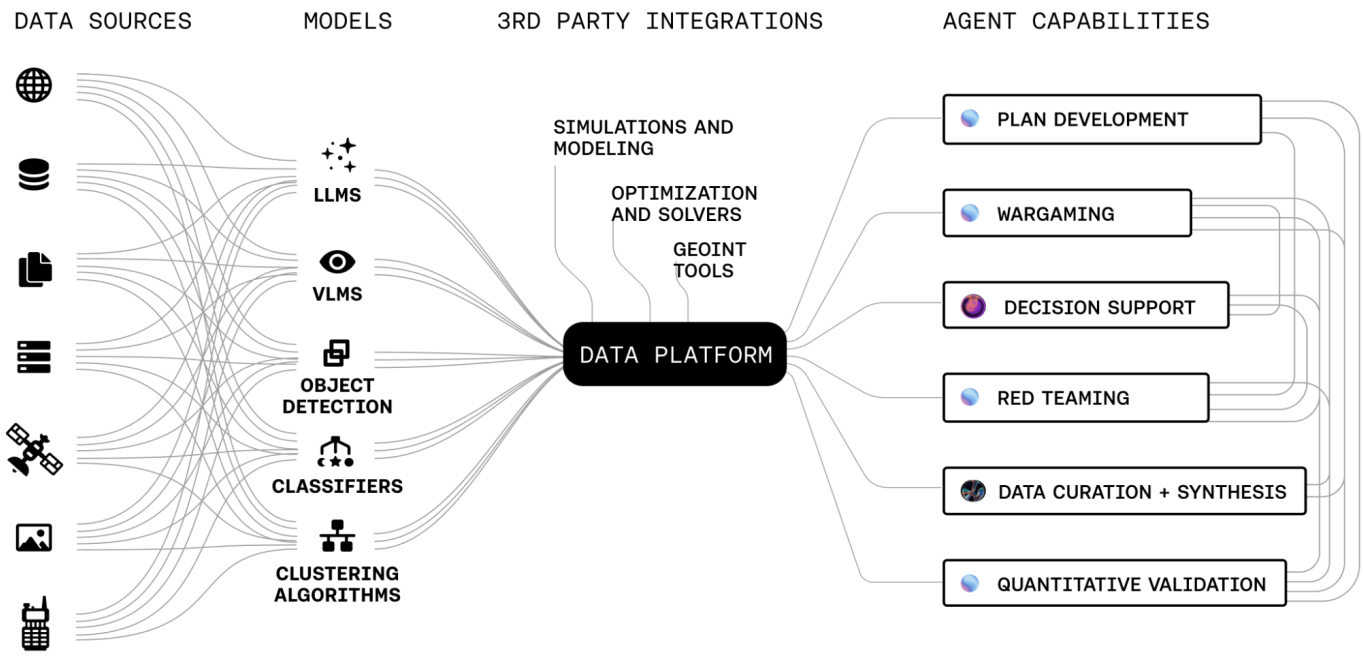

CASE STUDY: BALTIC ESCALATION MANAGEMENT

In the Baltic theatre, a sophisticated grey-zone crisis can render deliberate operational plans obsolete within hours. Consider a scenario in which cyber disruptions hit Lithuanian logistics nodes while localised GPS spoofing interferes with navigation systems across parts of the region. At the same time, a coordinated disinformation campaign floods social media with claims of NATO escalation and fabricated civilian casualties, designed to fracture political consensus and delay a collective response.

In this contested environment, the UK-led Joint Expeditionary Force and forward NATO elements face a rapidly shifting operational picture. Key variables like port accessibility, national rules of engagement, the availability of ISR assets, and air and missile defence coverage change by the hour. Attribution remains uncertain and intelligence reports conflict. Traditional planning cycles are poorly suited to such conditions; the time required to reconcile information and coordinate action across national caveats risks ceding the initiative to an adversary seeking to consolidate gains around the Suwalki corridor.

Agentic planning helps restore operational tempo. As conditions evolve, an agentic system continuously evaluates the environment against the existing operational plan, testing alternative reinforcement routes and assessing the impact of shifting constraints. It can recommend the repositioning of ISR and air defence assets to protect critical corridors and sustain reinforcement flows. Instead of waiting for manual planning to catch up, commanders receive updated, probability-weighted options in near real time, allowing the coalition to maintain decision advantage and respond before a grey-zone crisis hardens into a strategic setback.

THE SOVEREIGNTY RHEOSTAT: HOW NATO'S EUROPEAN NATIONS RETAIN CONTROL

If agentic systems are to deliver decision advantage across the Alliance, they must be integrated in a way that respects national control while enabling coalition action at speed. The solution is not to centralise NATO's digital infrastructure, but to architect it differently.

Conceptually, a shared decision advantage stack places the agentic layer between sovereign data inputs and coalition-readable outputs.

Conceptually, a shared decision advantage stack places the agentic layer between sovereign data inputs and coalition-readable outputs. Nations retain control of sensitive data and capabilities while sharing the decision products needed to coordinate action. NATO already applies this principle in its Federated Mission Networking (FMN), which enables national command networks to interconnect for coalition missions. An agentic decision layer would extend this model to decision-making itself. The key design principle is federation. Rather than pooling raw data, allies share derived outputs—alerts, assessments, and continuously updated plans—while allowing agentic systems to query sovereign environments. The result is a shared operational picture without requiring nations to surrender control over sensitive systems and data.

...sovereignty should be understood not as a binary choice but as a set of adjustable controls.

In this context, sovereignty should be understood not as a binary choice but as a set of adjustable controls. A useful way to think about this balance is the "Sovereignty Rheostat": a conceptual model for calibrating how data, access, and agent authorities are shared across the coalition according to operational need and political risk.

THE SOVEREIGNTY RHEOSTAT:

Just as a mechanical rheostat can be turned up or down to regulate electrical flow, sovereignty in coalition AI systems should not be treated as a binary choice but as a set of configurable controls, dialed by mission, data sensitivity, and political risk. The whole agentic stack does not need to be sovereign. What matters is that sovereignty is encoded where national control is essential, and relaxed where coalition effect is created. This can be imagined as three rheostats:

- ***Dial one—Data custody:*** *Keep raw data within sovereign enclaves, allow federated query across boundaries, and share derived products (alerts, tracks, confidence-weighted assessments etc) to enable coalition sense-making.*
- ***Dial two—Releasability:*** *Configure partner access depending on mission and sensitivity (national; bilateral; minilateral; NATO) so the same agentic workflow can generate products at different classifications without manual workarounds.*
- ***Dial three—Agent authority:*** *Set authorised boundaries (assistance; recommendations; non-kinetic orchestration) in line with policy, assurance levels and trust.*

Advanced, privacy-preserving technical enablers can underpin these conceptual dials. Technologies such as federated learning, confidential computing, and zero-trust architectures, allow agents to train on or query highly classified national intelligence enclaves without exposing the underlying data. The agent extracts insights and operational context, not the sovereign intelligence, solving the classification seam at a cryptographic level.

Settings are enforced through technical and policy controls: data rights, model governance, auditability, operational and update control, and strategic off ramps. Interoperability is created through interfaces, shared semantics, workflow standards, and common assurance artefacts.

The 'settings' of the Sovereignty Rheostat would depend on the scenario. For example, for a purely national capability, such as a UK-only counter-terrorism operation, the dials might be configured as follows:

- Dial one—Data custody: lock data custody to national enclaves
- Dial two—Releasability: restrict information sharing to 'UK Eyes Only'.
- Dial three—Agent authority: permit higher agent authority given the singular chain of command.

Conversely, for a NATO coalition engagement (e.g. Baltic Air Policing), the dials might be configured as follows:

- Dial one—Data custody: allow federated queries across allied sensors.
- Dial two—Releasability: dynamically share alerts to NATO networks.
- Dial three—Agent authority: limit agents to 'recommendations' given the need to navigate multiple national legal frameworks and ensure command consensus before action is taken.

NATO has successfully adopted similar governance models before. Ballistic missile defence, for example, combines a shared command architecture with nationally owned sensors and interceptors. The precedent is not the technology, it is the governance pattern: national control retained where necessary, allied integration enforced where it creates collective effect.

Agentic systems can adopt the same pattern—if designed for it. Nations can adopt cutting-edge agentic capabilities in high-consequence domains while preserving the ability to govern, validate, and intervene in the operation of systems to preserve mission assurance, operational integrity, and lawful command. But sovereignty must be treated as an engineering requirement from day one, including through clear data ownership, transparent model governance, jurisdictional clarity and common interoperability standards.

...sovereignty must be treated as an engineering requirement from day one...

PART III – FROM CONCEPT TO COALITION REALITY

The emergence of agentic systems offers a credible route to improve NATO's decision advantage. However, technological possibility does not automatically translate into operational capability. The Alliance has long experience with innovations that stall at the stage of pilots.

In NATO and Europe, three factors will determine whether agentic capabilities move beyond experimentation: assurance, acquisition, and organisational adaptation.

For agentic systems, the challenge is not simply technical maturity, but the institutional conditions required for militaries to trust and adopt decision-support technologies at scale. This paper focuses on the strategic case for action and the institutional principles that should shape it. The next questions—how to define trusted performance envelopes, structure procurement and interoperability, calibrate delegation boundaries, and generate the evidence needed for large-scale adoption—require more detailed treatment and will be developed in follow-on papers. In NATO and Europe, three factors will determine whether agentic capabilities move beyond experimentation: assurance, acquisition, and organisational adaptation.

ASSURANCE: THE FOUNDATION OF DECISION ADVANTAGE

Commanders will only delegate decisions to AI-enabled systems they trust. In an agentic force, assurance is therefore the gateway to operational tempo.

Agentic decision advantage capabilities differ from traditional software in an important respect: they influence decisions. By prioritising information, proposing options, and shaping information flows within HQ workflows, these systems can accelerate command, but they can also subtly steer judgement. The risk is not simply "wrong answers", but wrong influence: recommendations that compress deliberation or redirect attention in ways that affect operational outcomes.

In Agentic Warfare, test and evaluation (T&E) is not a back-office function—it is the discipline that enables commanders to delegate under pressure. Testing is rarely pass-fail; its purpose is

to discover operational limits and ensure reliable use, while generating the evaluation data needed to continuously improve safety and performance.

Testing also requires more than simple performance benchmarks. Because agents operate through multi-step reasoning processes, they must be evaluated across entire workflows and stress-tested against adversarial conditions, such as spoofed inputs, incomplete information, and manipulated data. The objective is not to prove a system infallible—no complex system ever will be. Instead, assurance seeks to define the conditions under which the system performs reliably, identify failure modes, and ensure awareness when limits are reached.

Equally important, assurance must include human feedback loops. Operator corrections, overrides, and judgements during live use of the tools must be captured and fed back into the update process to improve the effectiveness of models over time, and to sharpen the policy and operational workflow wrappers.

This allows commanders to understand the trusted performance envelope. Within those boundaries, agentic systems can safely accelerate the tempo of staff work. Outside them, human oversight and intervention remain essential.

Degrees of system autonomy will vary across mission types. Time-critical or reversible actions, such as sensor tasking or logistics coordination, may justify higher levels of delegation. High-consequence activities such as targeting or escalation will continue to demand direct human decision. The design challenge is therefore to ensure that machine processes remain traceable, bounded, and contestable, allowing commanders to intervene whenever required.

Building this assurance infrastructure is the central operational challenge of agentic warfare. Deploying an initial model is only the beginning; the larger task is establishing the evaluation, governance, and monitoring systems that allow these tools to operate safely inside real headquarters workflows. Without that foundation, agentic capabilities will remain confined to demonstrations. With it, they can become trusted infrastructure for coalition decision-making.

RESPONSIBLE AI AS OPERATIONAL DOCTRINE:

Responsible AI is the permission framework for scaling agentic systems in European defence, linking technical behaviour to command responsibility and the lawful employment of force.

On both sides of the Atlantic, military users demand systems that are lawful, secure, reliable under stress, traceable, governable, and accountable. Within NATO, the Principles of Responsible Use of AI (PRUs) dictate engineering requirements and essentially function as the operating doctrine for human-machine teaming.

They set the conditions under which commanders can delegate machine-speed work without surrendering judgement, lawyers can defend how decisions were supported, and coalitions can share workflows without importing unbounded risk.

Accountability must be engineered (who authorises use and owns configuration). Governability must be real (bounded permissions, approval thresholds, disengagement). Traceability must be reconstructable (inputs, sources, assumptions, uncertainty). And trust must be evidence-based (behaviour under contest, not persuasive explanations).

Handled rigorously, Responsible AI therefore becomes an enabler of operational tempo: reducing approval friction, enabling delegation, and allowing coalition integration at speed.

THE ACQUISITION REALITY

Even when the operational value of a technology is clear, defence adoption is shaped by procurement realities. European ministries evaluating digital systems ask familiar questions: Where does the data reside? Who controls updates? What happens if the supplier relationship changes—or if access to critical technology is restricted? And how can governments avoid becoming dependent on proprietary ecosystems?

These concerns are particularly acute for agentic systems because they sit at the intersection of operational data, national sovereignty, and industrial policy. Without careful design, advanced decision-support systems could become single points of failure—vulnerable to vendor lock-in, external policy decisions, or disruptions to commercial supply chains.

Addressing these risks requires procuring agentic systems as a capability ecosystem—combining models, data infrastructure, and continuous evaluation—not a static software product. Three principles follow:

- **Keep operational data inside national or NATO enclaves** while sharing derived insights across coalition networks.
- **Build composable systems using open interfaces and interoperable standards** so components can be swapped without breaking the architecture.
- **Hard-wire through-life testing into contracts**, including ongoing evaluation and red-teaming.

These principles don't eliminate dependence on external technology providers. But they reduce the risk that critical decision-support capabilities become brittle or inaccessible during crises.

PATHFINDERS AND MINI-LATERALS

Reaching consensus across the entire NATO alliance for common systems is historically slow. A more effective approach is operational pathfinders: bounded pilots among small groups of allies, with NATO institutions explicitly involved. Allied Command Transformation (ACT) can prove concepts and translate them into doctrine and training. The NATO Communications and Information Agency (NCIA) can harden reference architectures, security patterns, and accreditation. Joint / Component Commands can operationalise these systems in live headquarters conditions.

Done properly, these pilots produce more than demos; they generate the interfaces, releasability rules, evaluation evidence, and governance templates required to scale across the Alliance.

Existing multinational frameworks provide natural engines for such experimentation. Allies could use PESCO and the EDF to co-finance FMN-compatible data foundations; the NATO Innovation Fund (NIF) and the Defence Innovation Accelerator for the North Atlantic (DIANA) to validate and accelerate solutions; and the Joint Expeditionary Force (JEF) or NATO operational headquarters as hosts for early deployments.

Potential pathfinders should be tightly scoped and operationally hosted. In the High North, a JEF-based Alerting Cell could fuse multi-domain feeds into auditable alerts and automated tasking recommendations, sustaining tempo in sparse and contested Arctic conditions. In the Baltic, a 'Living Plan' could continuously generate reinforcement options as constraints shift—degraded logistics, ISR gaps, changing air-defence coverage—giving commanders traceable, data-driven assumptions at crisis speed. Finally, a Sovereignty Rheostat demonstration between a small set of allies could prove that agentic workflows can turn highly classified national data

into shareable coalition decision products, enabling coordination without exposing underlying intelligence sources or sovereign algorithms.

ALLIANCE DOCTRINE, TRAINING AND PROFESSIONAL MILITARY EDUCATION

Technology alone will not deliver decision advantage. History shows that the true impact of military innovations emerges only when institutions adapt their doctrine, training, and organisational culture to exploit new capabilities.

Agentic systems require a shift in how militaries think about command and control, challenging long-standing assumptions about how Headquarters function. The traditional general staff model assumes human analysts perform most tasks of sensing, interpreting, and planning, with software in support. In an agentic environment, that balance shifts. Agents perform the first pass of many analytical tasks—triaging information, testing assumptions, and proposing options—while humans set intent, apply judgement, and own the risk. Commanders move from hands-on controllers to mission directors, supervising machine-speed processes while retaining authority over consequential decisions.

Adapting to this model requires new competencies. Officers will need greater familiarity with data and AI systems, including an understanding of where tools are reliable and where they are brittle. Staffs will need to learn how to supervise digital agents, challenge machine-generated recommendations, and integrate automated workflows into operational planning.

Professional military education must evolve accordingly. War colleges, staff programmes, and NATO institutions will need to incorporate AI literacy, human-machine teaming, and the critical evaluation of automated decision-support tools into their curricula. Just as earlier generations learned to command radar networks or nuclear forces, future commanders must understand the capabilities and limitations of agentic systems.

The underlying principle remains unchanged: technology enhances the art of command, it does not replace it. Agentic systems can accelerate sensing, planning, and coordination, but judgement and decision remain human responsibilities.

SENIOR LEADER SUMMARY: BUILDING THE AGENTIC COALITION

Europe is rearming. Budgets are rising, production lines restarting, and forces rebuilding mass. The decisive question is whether NATO can turn inventory into coordinated action at the speed modern warfare demands. The Alliance must pivot from platform-centric C2 towards decision advantage.

Strategic Priorities for Leaders:

1. **Treat decision advantage as a core element of combat power.** Mass remains essential, but its value depends on coalition tempo. Modernisation must be measured not only in platforms, but in time-to-detect, time-to-decide, time-to-task, and time-to-synchronise.

2. **Build agentic capabilities as coalition infrastructure.** Decision-support systems must be designed from the outset for interoperability across NATO's federated architecture. Open interfaces, shared data standards, and coalition-ready workflows should be treated as non-negotiable procurement requirements.

3. **Adopt sovereignty-by-design and assurance-by-design.** Systems must be architected so that national control, auditability, and responsible AI governance are built into their technical foundations. These are not optional add-ons, they are the conditions that allow commanders and governments to trust and deploy systems at scale.

Core principles for modernisation:

- **Decision advantage is Europe's fastest deterrence multiplier.** Mass matters; coordination at speed unlocks optionality and, where necessary, makes mass decisive.

- **Agentic planning and alerting are the practical route.** Implementing digital staff at scale tackles the most acute friction points in the coalition OODA loop.

- **Sovereignty is adjustable, not binary.** The "Sovereignty Rheostat" offers a way to calibrate data sensitivity and releasability without compromising operational security.

- **Assurance is the bridge from pilots to trust.** In NATO, rigorous, through-life test and evaluation is what makes delegation, legality, and operational tempo possible at scale.

- **Pathfinders prove the model.** The Alliance should back operational pathfinders to prove sovereign interoperability in real NATO workflows while generating the evidence, interoperability standards, and governance templates for alliance-wide scaling.

The 2026 NATO Summit presents a timely opportunity to lock in this direction through a new strategic framework on decision advantage. This would mark a critical step forward, committing allies to a common, interoperable fabric for agentic AI and accelerating NATO's evolution into an agentic coalition projecting modern, collective deterrence.

scale