Allied Autonomy Is Creating a New Seam in Extended Deterrence

By
Burak Oktenli

Published: July 23, 2026

In February 2026, more than 200 personnel from Australia, the United Kingdom, and the United States gathered on Australia's east coast for the latest AUKUS Maritime Big Play exercise, where operators tested roughly thirty next-generation autonomous capabilities, including a live strike by a one-way attack munition. Earlier iterations of the series had already demonstrated something more consequential than any single platform: the ability of the three navies to control autonomous assets belonging to one another. Allied autonomy has moved out of the laboratory and into operational formations, and it is moving faster than the alliance frameworks that govern who may authorize a machine to act.

The Seam in the Umbrella

Extended deterrence has always rested on shared expectations about decision-making. The United States and Japan formalized this in December 2024, when the two governments issued their first Guidelines for Extended Deterrence, reinforcing consultation and communication procedures for the gravest decisions the alliance could face. Consultation procedures of this kind assume that the decisions in question move at human speed, through ministers and commanders who can meet, confer, and align. Autonomous systems introduce a class of alliance decisions that will not wait for a meeting.

Every ally that fields autonomous systems encodes national rules about when a machine may engage, when it must ask a human, and when it must stop. The United States codifies its approach in Department of Defense Directive 3000.09, which sets rigorous verification, testing, and senior-review requirements designed to minimize unintended engagements. Allies maintain their own rules, shaped by their own laws, doctrines, and political constraints.

Those national rule sets do not automatically align, and in a combined force the differences become seams. A coalition formation may contain systems that would engage a given contact under one national authority profile and hold fire under another. Adversaries study seams.

The probing will not be hypothetical. China has spent more than a decade practicing pressure beneath the threshold of armed conflict, from maritime militia swarms to coercive air intercepts, precisely because ambiguity about an alliance response is exploitable. Russia has run the same playbook against NATO members in the air and undersea domains. A coalition whose autonomous systems answer that pressure under visibly different national rules hands its adversaries a map of where the alliance hesitates.

The consequences for deterrence are direct. The most immediate risk is entrapment at machine speed: when an allied autonomous system engages inside the American security umbrella, the United States can inherit an escalation decision it never made, at a pace that forecloses the consultation the alliance promised itself. Mismatched authority thresholds also invite probing,

because an adversary can calibrate provocations that one member's systems would answer and another member's system would not, and can harvest knowledge of alliance rules from the pattern of responses.

Assurance suffers as well. An alliance that cannot state clearly how its machines hold fire loses precision in its guarantees at exactly the moment allies under threat are asking for more precision, not less. Credibility in the machine age includes the credibility of restraint.

Authority Interoperability

The building blocks for closing the seam already exist. Directive 3000.09 gives the United States a mature national template. The Five Eyes cybersecurity agencies showed in May 2026 that allied governments can speak with one voice on machine autonomy when they jointly published guidance on agentic artificial intelligence, urging organizations to prioritize "resilience, reversibility and risk containment over efficiency gains." The Maritime Big Play program has likewise turned toward shared command-and-control software and a common autonomous baseline across the three AUKUS navies. Each of these efforts is necessary. What none of them yet delivers is the specific artifact extended deterrence requires: a legible statement of each ally's engagement authority rules that coalition systems and coalition commanders can rely on before they operate together.

Allied autonomy needs authority interoperability, built as deliberately as data interoperability was built over the past three decades. Three practical steps would supplement the work already underway. Allies could publish national authority profiles, meaning standardized declarations of the engagement tiers their systems may exercise, the conditions that trigger human review, and the recall mechanisms that stop a pending action. Combined forces could adopt shared reversion standards, so that any coalition system can be halted through a common, tested procedure regardless of which flag it flies. Extended deterrence consultations, on the model the United States and Japan set in 2024, could add an autonomy annex that establishes in advance how allied autonomous engagements inside the umbrella will be communicated, reviewed, and, where necessary, walked back.

These steps ask allies to formalize what alliance staffs already negotiate informally under time pressure. Deterrence theory has long held that credibility rests on capability and will. Combined autonomy adds a third term: legibility. An alliance whose machines carry legible, compatible, and reversible authority is an alliance an adversary cannot divide at the seam. Extended deterrence in the machine age will belong to the alliances that make their authority architecture as interoperable as their ammunition.

Burak Oktenli is an independent researcher in Washington, DC, conducting research on human-machine authority architecture and the governance of autonomous military systems. He is the author of the ten-volume Authority and Architecture series on the governance of autonomous decision authority, and his analysis has appeared in RealClearDefense, Task and Purpose, and The Space Review. Views expressed in this article are the author's own.