Anduril's Lattice Lacks Configurable Harm Ordering

by Nick Clark | Published April 25, 2026

What Lattice Provides

Anduril's Lattice is a software platform unifying autonomous defense systems across sensors (Sentry towers, Wisp, ALTIUS), processing (CCM, mission planning), and effectors (Roadrunner, Pulsar, Bolt). The platform handles target detection, classification, tracking, and engagement coordination. It produces what Anduril describes as a 'mission autonomy' layer that operates faster than human-in-the-loop systems while maintaining authorized engagement parameters.

What the platform does not externalize is the harm-ordering policy that determines, in ambiguous engagement scenarios, how the autonomous decision-making weights potential outcomes. Civilian-presence weighting, allied-unit risk, infrastructure damage, mission-critical-target priority — all are computed inside Lattice with the resulting weighting visible to operators but not directly configurable by the governing authority.

Why LAWS Governance Demands Externalized Harm Ordering

Lethal autonomous weapons systems face a governance challenge that no other technology category has. The international debate over LAWS — at the UN

Convention on Certain Conventional Weapons, in academic ethics, in defense policy — converges on a single architectural requirement: meaningful human control over the use of force. The disagreement is over what 'meaningful' means.

The architectural answer is that the human control must include configuring the harm ordering, not merely approving the result. A system where the manufacturer hardcodes ordering and the operator reviews outcomes does not meet the standard. A system where the governing authority (national command, theater commander, mission ROE authority) credentials the ordering policy and the system executes it under audit-grade lineage does meet the standard. The architectural distinction is determinative.

How Externalization Would Sit on Top of Lattice

Confidence-governed actuation accepts harm-ordering policies as credentialed observations from the governing authority. For defense autonomy, the credentialing chain runs through the national command authority (NCA), through theater command, through mission-specific ROE issuance — each level signing within its scope, with the operating system consuming the composite policy through composite admissibility.

Lattice's autonomous engagement logic remains Anduril's. The harm-ordering policy moves to where the actual authority lives: in the credentialed governance chain that descends from the NCA through to the specific mission. ROE updates issue as credentialed observations propagating through the mesh; the platform consumes them; every engagement is recorded in audit-grade lineage with the policy under which it was evaluated.

What This Enables for Defense Autonomy Acceptance

DOD's Joint AI Center, DARPA's autonomy programs, and the broader defense-autonomy procurement environment increasingly require auditable governance. The LAWS debate is converging on architecture rather than capability restrictions: the question is not whether to deploy autonomous weapons, but under what auditable governance structure.

Externalized harm ordering is the architectural primitive that meets the audit requirement structurally. Anduril remains the integrator; the harm-ordering layer above Lattice becomes the procurement-relevant differentiator. The patent positions the primitive that defense autonomy will need as the LAWS governance debate resolves into actual procurement requirements.