Policy Brief

Series Information:
This policy brief is part of the EPINOVA Policy Brief Series on Strategic Competition, AI-Enabled Warfare, and Information Conflict.

Recommended Citation:
Wu, S. (2026). From AI-enabled weapons to AI-orchestrated warfare: The emerging global military AI stack in 2026 (Policy Brief No. EPINOVA–2026–PB–69). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.pb.2026.069

GLOBAL AI GOVERNANCE RESEARCH CENTER

EPINOVA

From AI-Enabled Weapons to AI-Orchestrated Warfare:

The Emerging Global Military AI Stack in 2026

Author: Shaoyuan Wu

Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC

Date: August 17, 2026

Key Judgments

- **The unit of military AI competition is shifting from the model or weapon to the stack.** Model quality matters, but operational advantage depends on whether models can access trusted data, operate on secure infrastructure, connect to command workflows and weapons, survive communications disruption, and improve through operational feedback.
- **The most consequential effect of military AI is compression, not necessarily autonomy.** AI can alter the tempo and structure of lethal decision-making without receiving formal command authority. Systems that filter information, rank targets, generate options, or sequence fires can exercise substantial agenda-setting power while a human retains the final authorization.
- **The United States and China are developing different pathways toward orchestration.** The U.S. model combines commercial platforms, private capital, classified data, alliance networks, and combat feedback. China's model emphasizes domestic technological control, system-of-systems planning, and autonomous mass. Neither pathway is complete, and both contain dependencies and organizational tensions.
- **Alliance interoperability is becoming a software and data-governance issue.** NATO's adoption of MSS creates the possibility of a common digital warfighting environment. It also raises questions about platform concentration, national data control, model substitution, and whether interoperability can persist if access to a vendor or model is disrupted.
- **Battlefield adaptation can generate significant military value despite limited frontier compute.** Ukraine and Russia demonstrate that small machine-vision models, onboard guidance, and swarm coordination can produce immediate military value under electronic warfare. Military AI power cannot be inferred from civilian model rankings or data-center capacity alone.
- **Governance is broadening but remains fragmented and behind deployment.** International processes now address lethal autonomous weapons, military AI more broadly, and AI-related nuclear command risks, but they remain divided across forums and largely nonbinding. Meanwhile, operational systems already shape target production, planning, prioritization, and fire sequencing. Existing initiatives do not yet provide integrated oversight of the full military AI stack.

Executive Summary

Military artificial intelligence is entering a new phase. For most of the past decade, public debate focused on AI-enabled weapons: autonomous drones, machine-vision targeting, loitering munitions, and robotic platforms. Those systems remain important, but they no longer capture the principal direction of change. The emerging competition is over **AI-orchestrated warfare**—the integration of compute, military data, models, command platforms, sensors, weapons, autonomous systems, and allied networks into a continuously operating warfighting architecture.

An AI-enabled weapon applies machine intelligence to a platform or bounded mission. An AI-orchestrated war system uses AI to structure relationships among many platforms, units, data sources, headquarters, and effects. Its operational value lies not only in recognizing a target, but also in determining which information receives attention, generating courses of action, matching weapons to targets, sequencing strike waves, coordinating autonomous systems, and reassessing results before an adversary can recover.

Developments through August 2026 indicate that this transition is no longer theoretical. The United States has directed that Palantir's Maven Smart System become a formal program of record by the end of fiscal year 2026 after its reported use in high-tempo operations against Iran. The U.S. Army is also negotiating with private investors to build commercial hyperscale data centers on military installations, treating land, power, and computing capacity as defense-relevant infrastructure. NATO's version of Maven reached full operational capability in June 2026 and now serves as a foundation for an expanding allied software ecosystem. Japan is combining consideration of U.S. platforms with a domestic foundation model, a tactical AI satellite, and a real-time digital twin. China has publicly disclosed an intelligent strike-planning system coordinating more than 100 tactical units during live-fire validation, although the technical details and operational effects remain independently unverified.

These developments point to a seven-layer military AI stack: energy and compute; data and models; integration and digital environments; command decision support; firepower orchestration; edge autonomy; and alliance governance. No state possesses all seven layers equally. Public evidence indicates particular U.S. strengths in platform scale, private-sector integration, and access to operational data. China is pursuing a more sovereign architecture centered on system-of-systems coordination and autonomous mass. NATO members and U.S. allies are developing hybrid arrangements that combine interoperability with varying degrees of national control. Ukraine and Russia are producing a distributed form of military AI shaped by battlefield adaptation, low-cost hardware, and electronic warfare.

The central strategic effect of this stack is **decision compression**. AI compresses information, decision time, organizational processes, operational sequencing, and the range of options presented to commanders. Consequently, keeping a human "in the loop" is necessary but insufficient. A human may retain formal authority while an AI system determines what the human sees, how threats are ranked, which options appear feasible, and how quickly a decision must be made.

Military AI governance should therefore move beyond a narrow focus on fully autonomous weapons. Policy must also cover decision-support platforms, target-generation systems, model updates, data provenance, supplier dependence, review time, human override, and degraded operations. The appropriate standard is not simply whether a human approves a strike. It is whether the full system remains **auditable, contestable, replaceable, resilient, sovereign in its critical functions, interoperable, and capable of failing safely**.

Why This Matters

The transition to AI-orchestrated warfare affects deterrence, escalation, alliance politics, defense acquisition, and civilian protection. Militaries that can interpret the battlefield and coordinate effects faster may impose operational dilemmas on slower adversaries. Yet the same compression that creates advantage can shorten the time available to verify intelligence, consult allies, conduct legal review, or distinguish an attack from preparation for a larger escalation.

The transition also reallocates power inside the defense enterprise. Software providers, cloud operators, data-center developers, chip suppliers, and model companies are becoming components of operational readiness. The dispute over Anthropic's role in U.S. defense workflows illustrated how restrictions on a model provider can create a military supply-chain problem after the model has already been embedded in mission software. The relevant sovereignty question is therefore not whether every component is domestically produced. It is whether a military can inspect, replace, and continue operating critical functions without the consent of a single external provider.

Finally, AI orchestration changes the meaning of human control. If a system can produce targets faster than personnel can review them, or if it presents only a narrow set of machine-ranked options, formal human approval may become procedural rather than substantive. Governance that focuses only on the final trigger decision risks overlooking the upstream systems that define the target, recommend the weapon, set the tempo, and frame the commander's choices.

Analytical Note

This brief uses AI-enabled weapons to describe platforms or mission systems in which AI performs a bounded function such as classification, navigation, tracking, or terminal guidance. AI-orchestrated warfare refers to the use of AI across multiple organizational and technical layers to integrate information, generate or rank operational options, allocate resources, coordinate units, sequence effects, and support reassessment. Orchestration does not require fully autonomous command or autonomous use of lethal force. AI-orchestrated warfare overlaps with JADC2, mosaic warfare, and kill-web concepts but emphasizes the full socio-technical stack—including compute, commercial suppliers, model governance, alliance permissions, and degraded operations.

The military AI stack is a socio-technical analytical framework rather than a claim that every country is building an identical architecture. It separates three questions that are often conflated: how AI capacity is formed, how it is converted into military effects, and how it reshapes alliance and governance relationships. The layers are interdependent, but weaknesses in one layer do not prevent a state from developing substantial capabilities in another.

Evidence in this brief is described according to three standards. Confirmed developments are documented in government releases, contracts, formal policies, or institutional announcements. Reported developments are supported by credible independent journalism, sometimes relying on anonymous sources or documents not publicly released. Claimed capabilities are disclosed primarily by a government, military, state media outlet, or vendor and lack sufficient independent performance assessment. A claimed test may confirm that an organization demonstrated a system; it does not establish that the system performs as described under combat conditions.

1. The Seven-Layer Military AI Stack

Table 1. The Seven-Layer Military AI Stack

Layer	Operational Function	Representative 2025–2026 Indicators	Strategic Test
Energy And Compute	Electricity, chips, data centers, military cloud, cooling, communications, and offline compute	U.S. base data-center negotiations; South Korean defense AI data-center plan	Can priority workloads continue through grid, network, supplier, or cooling disruption?
Data And Models	Military datasets, foundation models, computer vision, model routing, evaluation, and updating	Maven model integration; Japan's PLaMo; Chinese military-linked use of domestic models	Are data provenance, model behavior, and update effects known and controllable?
Integration And Digital Environments	Common operational pictures, data ontologies, digital twins, and simulation	MSS NATO; Japan's real-time digital twin; Ukraine's DELTA ecosystem	Can heterogeneous systems exchange data without creating a single point of dependence?
Command Decision Support	Intelligence fusion, threat ranking, planning, wargaming, and course-of-action generation	Maven; Japan's planning prototype; German decision-support initiative	Can commanders understand, challenge, and reject system recommendations?
Firepower Orchestration	Sensor-to-shooter connection, weapon-target matching, strike sequencing, and reassessment	Reported U.S. use of Maven in operations against Iran; China's disclosed strike-planning test	Does faster coordination preserve verification, proportionality review, and escalation control?
Edge Autonomy	Onboard perception, navigation, target tracking, terminal guidance, and swarm coordination	Chinese, Russian, and Ukrainian unmanned systems	Can systems operate reliably under deception, communications loss, and electronic warfare?
Alliance And Governance	Technical standards, permissions, supplier relationships, legal review, audit, and human control	MSS NATO; allied sovereignty debates; UN military AI, LAWS, and nuclear-risk processes	Can allies interoperate without surrendering control over data, mission rules, or force authorization?

Sources: Author's framework and synthesis. Illustrative cases draw on Brave1 (2025), Global Times (2026b), Janes (2026), Japan Ministry of Defense (2026a, 2026b), Lee (2026), Ministry of Defence of Ukraine (2025, 2026), North Atlantic Treaty Organization (2026), Preferred Networks (2026), United Nations General Assembly (2025b), U.S. Army Communications and Outreach Office (2026), and reporting by Baptista and Potkin (2025), Copp et al. (2026), and Jeans (2026). The table describes functions and policy tests; it does not rank national capability.

The seven layers are analytically distinct but operationally interdependent. Advanced models have limited military value without reliable data, contested-network connectivity, and integration with command and weapons systems. Conversely, inexpensive autonomous platforms may generate significant tactical effects without frontier models or centralized cloud infrastructure.

Military AI power therefore cannot be inferred from civilian AI leadership alone. It depends on conversion, the capacity to translate technical resources into repeatable operational effects, and resilience. A stack that performs well in exercises but fails when electricity, communications, cloud access, or contractor support is disrupted may create strategic vulnerability rather than advantage.

2. The United States: Platform, Capital, and Combat Data

The United States is constructing one of the most comprehensive publicly documented military AI stacks, combining commercial technology, global intelligence collection, private capital, classified computing, and operational deployment. Maven is becoming the principal platform through which many of these elements are connected.

Originally established in 2017 as a computer-vision initiative, Maven evolved into a broader command-and-control environment capable of analyzing information from satellites, aircraft, drones, radar, sensors, and intelligence reporting. In March 2026, a Pentagon memorandum directed that Maven become a program of record by the end of the fiscal year. The memorandum also shifted oversight to the Chief Digital and Artificial Intelligence Office and assigned future contracting responsibilities to the Army (Jeans, 2026). Program-of-record status provides a pathway to stable funding, institutional ownership, and broader joint-force adoption; it should not be described as evidence that the transition had already been completed at the time of the memorandum.

Maven and Anthropic's Claude were reportedly used during U.S. operations against Iran to support intelligence analysis, target proposals, prioritization, target coordinates, and post-strike assessment. The Washington Post reported that the United States struck 1,000 targets in the first 24 hours and that Maven drew on extensive classified data flows (Copp et al., 2026). Those figures describe operational tempo, not independently verified AI effectiveness. Public evidence does not isolate Maven's contribution to accuracy, civilian harm, ammunition efficiency, or strategic outcomes. The defensible conclusion is narrower: AI-supported workflows entered high-tempo targeting operations at substantial scale.

The Pentagon–Anthropic dispute suggests a second lesson. A foundation model embedded in mission workflows may not be removable as easily as an office productivity application. Replacement may require software changes, security accreditation, testing, and user retraining (Copp et al., 2026). Model providers have therefore become part of the defense supply chain. A contractual or policy dispute can create operational consequences even when the government retains formal control of the surrounding platform.

At the infrastructure layer, the U.S. Army conditionally selected Carlyle to negotiate a data-center project on approximately 1,384 acres at Fort Bliss and CyrusOne for a project on approximately 1,201 acres at Dugway Proving Ground. The developers would finance, build, operate, maintain, and decommission the facilities with no upfront construction cost to taxpayers. Initial operating capability is projected for fiscal year 2027 at Fort Bliss and fiscal year 2029 at Dugway, but final lease terms had not been completed as of the announcement (U.S. Army Communications and Outreach Office, 2026).

The significance is conceptual as much as physical. Land, electricity, water, onsite generation, cooling, chips, and secure connectivity are entering the same strategic conversation as airfields, ports, fuel storage, and munitions production. Yet commercial hyperscale capacity located on a military installation should not automatically be treated as dedicated military compute. Priority access, security boundaries, wartime allocation, grid costs, and continuity obligations will depend on final agreements.

The U.S. model can be described as a platform-capital-combat-data system. Its advantages are scale, speed, capital access, a global sensor network, and feedback from operational use. Its vulnerabilities are vendor concentration, proprietary architectures, high energy demand, software supply-chain dependence, and the risk that military workflows become inseparable from systems the government cannot readily replace.

3. NATO: From Procurement to an Alliance Operating Environment

NATO's adoption of Maven demonstrates how a U.S.-origin commercial platform can become part of an alliance-wide command environment. In March 2025, the NATO Communications and Information Agency completed the acquisition of MSS NATO for Allied Command Operations. The acquisition took approximately six months from requirement identification to contract, a notable acceleration relative to conventional alliance procurement (NATO Communications and Information Agency, 2025).

MSS NATO reached full operational capability on June 22, 2026, according to a SHAPE release reported by Janes. The system received security accreditation for use on classified networks in exercises, missions, and activities. NATO officials also stated that NATO data would remain on NATO servers and in NATO data centers, and that the platform could support models from multiple providers (Janes, 2026). These are important sovereignty safeguards, although their practical strength depends on whether applications, ontologies, and mission workflows remain portable.

NATO is already building an ecosystem around the platform. In July 2026, DIANA selected ten companies to integrate modeling, simulation, targeting-support, wargaming, and other decision tools into a cloud-based environment that augments MSS NATO. Demonstrations are scheduled for December 2026 (North Atlantic Treaty Organization, 2026). MSS is thus becoming more than a purchased application: it is emerging as an environment around which allied data, models, and operational software can be organized.

The strategic benefit is faster coalition coordination. The strategic risk is that technical interoperability evolves into institutional dependence. In earlier eras, alliance interoperability depended on common communications protocols, ammunition, identification systems, and planning procedures. AI-enabled interoperability increasingly depends on shared data formats, object definitions, access permissions, model interfaces, and software-update practices. NATO's central challenge is to gain the speed of a common operating environment without allowing one commercial architecture to become an irreplaceable gatekeeper.

4. China: Sovereign AI And System-Of-Systems Orchestration

China's military AI development is often described through autonomous vehicles, robot dogs, drone carriers, or swarms. Those platforms remain relevant, but disclosures in 2026 point toward a broader objective: using AI to coordinate the operational system rather than merely improve individual weapons.

In August, Chinese state media aired footage of a People's Liberation Army Air Force intelligent strike-planning system after what was described as its final live-fire test. State-media reporting, republished by NOSI, stated that the system coordinated more than 100 tactical units with precise timing and supported a shift from platform-centered operations toward AI-enabled system-of-systems warfare (Global Times, 2026b). The footage is evidence that the PLA publicly demonstrated such a system. It is not independent verification of the system's algorithms, the degree of autonomy, its performance under adversarial conditions, or its causal contribution to the exercise outcome.

The disclosure is nevertheless significant because it places fire allocation and operational sequencing—not only target recognition—at the center of China's AI narrative. It is consistent with evidence from procurement records, patents, and research publications showing PLA-linked interest in domestic models, autonomous target recognition, battlefield decision support, simulation, and drone-swarm coordination (Baptista & Potkin, 2025).

Public evidence also reveals limits. During a July 2026 exercise involving KVD002 armed reconnaissance drones, Chinese reporting described a failure to maintain a reliable automated lock on a moving maritime target, after which an operator shifted to manual control (Global Times, 2026a). The contrast suggests that progress in system-level planning does not guarantee reliable autonomy at the tactical edge.

China's emerging model is best characterized as sovereign system-of-systems orchestration. It seeks greater domestic control over models, chips, data, and mission systems while using autonomy to coordinate mass. Its central organizational tension lies between the distributed initiative required by adaptive AI systems and a military-political structure that places a premium on centralized command, information control, and political reliability.

5. U.S. Allies: Hybrid Sovereignty and the Interoperability Tradeoff

Japan illustrates the choices confronting technologically capable U.S. allies. American platforms offer rapid interoperability and mature software, but dependence on them may constrain national control over data, mission logic, and software updates. As of August 17, 2026, public reporting indicated that Japan was considering AI-supported command and control, including Palantir's Maven Smart System (MSS), while also exploring a domestic AI core linked to U.S. systems (Jiji Press, 2026; United Press International, 2026). These reports described policy and procurement options rather than a completed force-wide acquisition, and they did not indicate that AI would receive independent lethal authority.

Japan is simultaneously developing domestic and complementary capabilities. Preferred Networks received a contract from the Acquisition, Technology and Logistics Agency to develop and test an operational-environment analysis system using PLaMo 3.0 Prime. The project is intended to organize textual, geospatial, and force-employment information and support situation assessment and option development at multiple command levels (Preferred Networks, 2026).

The Japanese Ministry of Defense is also developing a tactical AI satellite prototype designed to fuse information from multiple satellites onboard and connect processed intelligence to shooters (Japan Ministry of Defense, 2026a). A separate ¥1.991 billion real-time digital-twin project will provide three-dimensional visualization and simulation for operations, exercises, and disaster response (Japan Ministry of Defense, 2026b). Japan's emerging architecture therefore consists of four tracks: U.S. software for interoperability, domestic models for technological control, space-based processing for sensor-to-shooter integration, and digital twins for planning.

South Korea is pursuing a related hybrid strategy. The Ministry of National Defense announced five civil-military AI hubs, about 150 priority datasets, and a dedicated defense AI data center by 2030, potentially tied to national computing infrastructure (Lee, 2026). In May 2026, it agreed with SK Telecom to test a defense model based on a domestic foundation model (Kang, 2026).

South Korea's goal of securing up to 50,000 defense GPUs remains a planning ambition, not existing capacity (Lee, 2026). Strategically, Seoul is attempting to integrate protected defense data, domestic models, national compute, the semiconductor industry, and combined command arrangements with the United States.

The United Kingdom and Germany demonstrate the same sovereignty-interoperability tradeoff in Europe. The UK Ministry of Defence placed accelerated AI adoption at the center of its 2026 modernization agenda (UK Ministry of Defence, 2026). The British Army chief subsequently reported that the ASGARD digital targeting network reduced a corps-level planning cycle from approximately 72 hours to one hour and increased target-processing capacity tenfold (British Army, 2026). These are military-reported exercise metrics, not independently validated combat results. Their significance lies in the performance criteria being emphasized: decision time and target throughput rather than model accuracy alone.

Germany is exploring AI to analyze battlefield data, identify patterns in adversary behavior, recommend countermeasures, and compress planning processes. German Army leaders have simultaneously emphasized human judgment, NATO compatibility, data sovereignty, and caution regarding external access to national military databases (Siebold, 2026). Europe's choice is therefore not simply between mature U.S. software and less mature domestic alternatives. It also concerns control over data, inspection rights, service continuity, update authority, and the ability to operate during political or commercial disagreement.

6. Ukraine and Russia: Distributed Adaptation at the Edge

Ukraine and Russia demonstrate a different pathway to military AI power. Rather than relying primarily on hyperscale infrastructure and centralized software platforms, both are integrating bounded autonomy into inexpensive systems shaped by battlefield feedback and electronic warfare.

Ukraine's DELTA ecosystem integrates data from multiple sources to provide a common operational picture and support mission planning and coordination across units (Ministry of Defence of Ukraine, 2025). Its Vezha module supports real-time battlefield video analysis and collaborative review (Ministry of Defence of Ukraine, 2026). Separately, Brave1 has tested machine-vision systems in which an operator locks a target before the drone conducts terminal guidance without an external communications link (Brave1, 2025). The purpose is not general autonomy but functional resilience: preserving the terminal stage of an authorized mission when electronic warfare disrupts the link between operator and platform.

Russia is pursuing similar edge functions. Rostec reported a 2026 Supercam-based swarm test in which drones automatically shared target information and a neural network identified objects and assigned an attack sequence. According to the manufacturer, one operator could supervise ten loitering munitions, while a human still confirmed the target before engagement (Rostec, 2026). This account establishes the claimed design and test conditions, not independent combat effectiveness.

The Russia–Ukraine war demonstrates that consequential military AI does not require a frontier language model. A small computer-vision model operating onboard an inexpensive drone may produce greater immediate battlefield value than a more capable model confined to a remote data center. The conflict also highlights the importance of deployment-feedback loops: repeated use generates data that can modify software, tactics, training, and hardware within compressed development cycles.

7. Israel and the Industrialization of Target Production

Israel's experience highlights a different dimension of AI-orchestrated warfare: the expansion of target-production capacity. Investigative reporting has described systems such as Lavender, the Gospel, Where's Daddy?, and Fire Factory as supporting association analysis, target generation, weapon allocation, and strike scheduling. Reporting on Lavender cited intelligence sources who claimed that the system generated tens of thousands of potential targets and that review of some recommendations was extremely brief (McKernan & Davies, 2024).

The Israel Defense Forces disputes the characterization of Lavender as an autonomous target-selection system. Its official position is that the relevant database cross-references intelligence, does not constitute a list of approved targets, and does not replace the individual assessment of military advantage and expected collateral damage required before a strike (Israel Defense Forces, 2024).

In 2026, Elbit Systems reported that the Tzayad digital command system had processed large volumes of real-time target and operational data across Israel's theaters since October 2023. The company later clarified that the reported figures represented aggregated system activity rather than approved targets or completed strikes (Sabbagh, 2026).

The competing accounts cannot be fully resolved with public evidence. The policy problem remains even under the official description. A system that expands a target list faster than humans can meaningfully review it changes the conditions under which judgment operates. Formal approval may remain human while the available review time, depth of investigation, and institutional capacity to reject machine-generated recommendations decline.

"Human in the loop" is therefore an incomplete test. Meaningful human control depends on whether the reviewer understands the basis of a recommendation, can access contradictory information, has sufficient time and authority to evaluate it, and can reject it without disrupting an automated operational sequence.

Table 2. Six Pathways of Military AI Conversion

Model	Illustrative Examples	Conversion Mechanism	Principal Advantage	Principal Vulnerability
Platform-capital-combat-data	United States	Commercial platforms and models connect classified data, private infrastructure, and operational feedback	Scale, speed, and joint-force diffusion	Vendor concentration, energy demand, and proprietary dependence
Alliance Operating Environment	NATO	Shared software, data standards, applications, and classified infrastructure	Coalition interoperability and rapid diffusion	Platform lock-in and unequal member dependence

Table 2. Six Pathways of Military AI Conversion (Cont.)

Model	Illustrative Examples	Conversion Mechanism	Principal Advantage	Principal Vulnerability
Sovereign System-of-systems Orchestration	China	Domestic models and hardware support planning, mass, and coordinated effects	National control and scale	Hardware constraints, opacity, and tension between autonomy and centralized command
Hybrid Sovereignty	Japan, South Korea, United Kingdom, Germany	U.S./NATO interoperability combined with domestic models, data, or infrastructure	Faster adoption with partial national control	Duplication, integration complexity, and unresolved update authority
Distributed Battlefield Adaptation	Ukraine and Russia	Edge models, low-cost platforms, electronic-warfare resilience, and rapid feedback	Fast iteration and affordable mass	Uneven reliability, component dependence, and limited strategic integration
Algorithmic Target Production	Israel	Data fusion and machine-assisted expansion and scheduling of target workflows	High operational throughput	Automation bias, review bottlenecks, and accountability disputes

Sources: Author's typology based on the cases and evidence cited in Sections 2–7. Categories are ideal types; states may combine more than one pathway.

Table 3. The Five Compressions of AI-Orchestrated Warfare

Compression	Operational Effect	Strategic or Governance Risk	Minimum Assurance Question
Information	Converts large sensor flows into a common operational picture	Missing data and uncertainty become less visible	Can users trace a conclusion to source data and see uncertainty or disagreement?
Decision Time	Reduces planning from days or hours to minutes or seconds	Less time for verification, legal review, consultation, and de-escalation	Is review time appropriate to the consequence and reversibility of the decision?
Organization	Automates handoffs among intelligence, planning, legal review, command, and fires	Responsibility becomes diffuse and human review can become procedural	Are authority, accountability, and intervention points explicit?
Choice	Ranks threats and presents a narrow set of recommended actions	The system gains agenda-setting power without formal command	Can decision makers request alternatives and understand excluded options?
Operational Sequence	Synchronizes reconnaissance, deception, electronic warfare, strikes, and reassessment	The adversary's recovery window shrinks, increasing escalation and misperception risk	Can commanders pause, reorder, or terminate the sequence after execution begins?

Sources: Author's analytical framework. The categories synthesize patterns across the cases examined in this brief. "Compression" describes a mechanism, not an assumption that every system achieves faster or better outcomes.

Together, these mechanisms alter the meaning of human control. The primary risk is not necessarily that a machine issues an unauthorized command. It is that a human organization becomes unable to operate outside the tempo, representation, and option structure created by its machines.

8. Policy Implications

8.1 Govern the Full Decision Chain Through Meaningful Human Control

Defense ministries should extend AI assurance and legal review beyond weapons to systems that filter intelligence, generate targets, rank threats, recommend weapons, schedule operations, and assess results. Risk classification should reflect where a system sits in the decision chain, the scale of its influence, and whether its outputs are reversible. For high-consequence decisions, reviewers should receive the evidence, uncertainty, and competing interpretations needed to make an informed judgment. Emergency exceptions should be narrowly defined, logged, and reviewed after action.

8.2 Require Auditability, Contestability, and Exit Capability

Military AI systems should preserve data provenance, model and software versions, target-list changes, human interventions, rejected recommendations, and post-action assessments. Operators must be able to challenge outputs, request alternatives, and override recommendations. Procurement contracts should also require portable data, documented interfaces, model substitution, and transition assistance. Technological sovereignty does not require domestic production of every component; it requires the capacity to preserve data, replace providers, and maintain critical functions during commercial or political disputes.

8.3 Build Resilient Infrastructure and Safe Degradation

Military AI architectures should be tested under loss of cloud services, satellite communications, external models, electricity, navigation, data integrity, and contractor support. Exercises should include offline operation, reduced bandwidth, deceptive inputs, model failure, and manual fallback. Because compute depends on physical infrastructure, military data centers should also be incorporated into energy, water, cyber, and physical-security planning. Governments should distinguish workloads requiring hyperscale infrastructure from those that can be distributed or moved to the tactical edge.

8.4 Build Federated Interoperability and Measure Operational Outcomes

NATO and U.S. allies should establish common standards for data exchange, identity, permissions, model evaluation, and logging without requiring every member to adopt the same platform. Alliance exercises should test the loss of a shared model or provider to determine whether interoperability rests on open standards or vendor dependence. Evaluation should also extend beyond speed and target throughput to accuracy, calibration, system availability, manual fallback, adversarial robustness, recommendation reversals, civilian harm, and mission outcomes.

8.5 Expand International Governance Beyond Lethal Autonomy

International governance has begun to move beyond lethal autonomy, but it remains fragmented and largely nonbinding. In December 2025, the United Nations General Assembly adopted Resolution 80/58 on AI in the military domain, Resolution 80/57 on lethal autonomous weapons systems, and Resolution 80/23 on risks arising from AI integration into nuclear command, control, and communications systems (United Nations General Assembly, 2025a, 2025b, 2025c). Resolution 80/58 was followed by three days of informal exchanges in June 2026; their agenda covered the military AI life cycle, decision support, use of force, high-risk applications, and confidence-building measures (United Nations Office for Disarmament Affairs, 2026). Separately, the secretary-general called for a legally binding instrument addressing systems capable of taking life without human control (United Nations Secretary-General, 2025).

The Convention on Certain Conventional Weapons process had not produced such an instrument by this brief's cutoff date, with another Group of Governmental Experts session scheduled for August 31–September 4, 2026 (United Nations Office at Geneva, 2026).The principal gap is therefore not the absence of international activity but the lack of an integrated framework covering the full military AI decision chain. States should connect existing initiatives through confidence-building measures for decision-support, target-production, and operational-orchestration systems. These measures should include disclosures of system functions, legal-review requirements, incident logging, human-override procedures, testing under deception and degraded conditions, and safeguards for nuclear warning and decision systems. Such measures can improve transparency and reduce risk without requiring the disclosure of classified algorithms, operational data, or targets.

9. Limitations

This brief relies exclusively on publicly available information through August 17, 2026. Military AI programs are classified, unevenly disclosed, and described using inconsistent terminology. Public demonstrations, patents, tenders, and contracts do not establish operational deployment or combat effectiveness; conversely, the absence of public evidence does not demonstrate the absence of a capability. The cases are illustrative rather than exhaustive and represent distinct pathways of military AI conversion, not a comprehensive global sample.

The cases and performance claims are not directly comparable. U.S. and NATO programs generate relatively detailed institutional records, whereas evidence concerning China, Russia, Ukraine, and Israel often comes from state media, procurement documents, vendor claims, wartime disclosures, or contested reporting. Claims about planning speed, target throughput, or systems supervised frequently lack common baselines, error rates, mission definitions, and independent replication. This brief therefore treats such figures as evidence of reported use or intended function, not as a basis for ranking national capabilities.

The seven-layer stack and five compressions are analytical constructs rather than measures of linear technological development or strategic effectiveness. Militaries may be strong in some layers and weak in others, while faster execution does not necessarily improve political or strategic judgment. Public evidence also does not justify attributing a specific military outcome or civilian casualty to AI without a credible causal investigation. An AI system's participation in a targeting workflow should not be equated with autonomous target approval or proof that it caused the resulting outcome.

10. Conclusion

The decisive military AI competition of 2026 is not a race to place the most advanced model inside the largest number of weapons. It is a contest over operational conversion: who can transform compute, data, software, command institutions, sensors, autonomous platforms, and alliance relationships into a resilient system capable of acting at scale and speed.

The United States is developing a platform-centered architecture supported by commercial capital and operational data. NATO is turning that architecture into an alliance environment. China is pursuing sovereign models, system-of-systems fire planning, and autonomous mass. Japan and South Korea are building hybrid stacks that combine U.S. interoperability with domestic technological control. Ukraine and Russia demonstrate the tactical value of distributed edge autonomy, while Israel's experience shows how algorithmic systems can reshape the scale and tempo of target production.

These systems do not need to remove humans from command to transform warfare. Their most important effect may be to reorganize the conditions under which humans decide. By compressing information, time, organizations, choices, and operational sequences, military AI can exercise substantial influence before formal authorization occurs.

The relevant policy choice is therefore not between innovation and control. It is between architectures that preserve judgment under pressure and architectures that make human organizations dependent on machine-generated tempo and representations. The emerging military AI stack should be judged not only by how quickly it operates, but also by whether it remains auditable, contestable, replaceable, resilient, and capable of failing safely.

References

Baptista, E., & Potkin, F. (2025, October 27). How China could use DeepSeek and AI for an era of war. *Reuters*. https://www.reuters.com/world/asia-pacific/robot-dogs-ai-drone-swarms-how-china-could-use-deepseek-an-era-war-2025-10-27/

Brave1. (2025, March 7). *Brave1 tested drones with artificial intelligence [In Ukrainian]*. https://market-brave1.delta.mil.gov.ua/brave1-proviv-vyprobuvannia-droniv-zi-shtuchnym-intelektom/

British Army. (2026, June 23). *RUSI Land Warfare Conference 2026: Chief of the General Staff keynote speech*. https://www.army.mod.uk/news/rusi-lwc-2026-chief-of-the-general-staff-speech/

Copp, T., Dwoskin, E., & Duncan, I. (2026, March 4). Anthropic's AI tool Claude central to U.S. campaign in Iran, amid a bitter feud. *The Washington Post*. https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/

Global Times. (2026a, July 31). First-ever footage: PLA Army drone swarm conducts combat-oriented training exercise. https://www.globaltimes.cn/page/202607/1367268.shtml

Global Times. (2026b, August 3). Official documentary reveals PLA Air Force intelligent strike-planning system; experts say it sharpens coordinated combat, shortens strike cycle. https://www.globaltimes.cn/page/202608/1367381.shtml

Israel Defense Forces. (2024, April 3). *IDF use of AI and targeting practices: Response to the Guardian query published April 3, 2024*. https://www.idf.il/en/mini-sites/israel-at-war/all-articles/idf-use-of-ai-and-targeting-practices-response-to-the-guardian-query-published-april-3-2024/

Janes. (2026, June 30). *Maven Smart System achieves full operating capability with NATO*. https://www.janes.com/defence-intelligence-insights/defence-news/security/maven-smart-system-achieves-full-operating-capability-with-nato

Japan Ministry of Defense. (2026a, February 13). *Progress on the tactical AI satellite demonstration* [In Japanese]. https://www.mod.go.jp/j/press/news/2026/02/13b.html

Japan Ministry of Defense. (2026b, March 13). *Real-time digital twin* [In Japanese]. https://www.mod.go.jp/j/press/news/2026/03/13a.html

Jeans, D. (2026, March 20). Pentagon to adopt Palantir AI as core U.S. military system, memo says. *Reuters*. https://www.reuters.com/technology/pentagon-adopt-palantir-ai-as-core-us-military-system-memo-says-2026-03-20/

Jiji Press. (2026, August 10). Japan mulls use of domestic AI in SDF command system. *Nippon.com*. https://www.nippon.com/en/news/yjj2026081000706/

Kang, J.-e. (2026, May 14). SK Telecom, defense ministry team up to develop military AI model. *Yonhap News Agency*. https://en.yna.co.kr/view/AEN20260514006500320

Lee, H. (2026, March 26). "Modern warfare is a data battle": Korea to establish AX hubs, deploy AI command systems. *Seoul Economic Daily*. https://en.sedaily.com/politics/2026/03/26/modern-warfare-is-a-data-battle-korea-to-establish-ax-hubs

McKernan, B., & Davies, H. (2024, April 3). "The machine did it coldly": Israel used AI to identify 37,000 Hamas targets. *The Guardian*. https://www.theguardian.com/world/2024/apr/03/israel-gaza-ai-database-hamas-airstrikes

Ministry of Defence of Ukraine. (2025, October 28). *DELTA: Basic training course in Army+*. https://mod.gov.ua/en/news/delta-basic-training-course-in-army

Ministry of Defence of Ukraine. (2026, March 9). *Independent assessment confirms DELTA remains compliant with all cybersecurity standards after scaling*. https://mod.gov.ua/en/news/independent-assessment-confirms-delta-remains-compliant-with-all-cybersecurity-standards-after-scaling

NATO Communications and Information Agency. (2025, April 14). *NATO acquires AI-enabled warfighting system*. https://www.ncia.nato.int/newsroom/news/nato-acquires-aienabled-warfighting-system

North Atlantic Treaty Organization. (2026, July 13). *NATO's DIANA selects innovators to drive decision superiority for the Alliance*. https://www.nato.int/en/news-and-events/articles/news/2026/07/13/natos-diana-selects-innovators-to-drive-decision-superiority-for-the-alliance

Preferred Networks. (2026, July 29). *PFN awarded ATLA project to demonstrate operational-environment information-analysis support using generative AI* [In Japanese]. https://www.preferred.jp/ja/news/1126

Rostec. (2026, April 14). Rostec tested an attack drone swarming technology. https://elements.rostec.ru/en/media/news/rostec-tested-an-attack-drone-swarming-technology/

Sabbagh, D. (2026, July 6). Israeli command system identified 850,000 targets in Gaza and Lebanon wars, says supplier. The Guardian. https://www.theguardian.com/world/2026/jul/06/israel-command-system-identified-850000-targets-gaza-lebanon-war-supplier-elbit-systems

Siebold, S. (2026, March 25). German army eyes AI tools to expedite wartime decision-making. *Reuters*. https://www.reuters.com/technology/german-army-eyes-ai-tools-expedite-wartime-decision-making-2026-03-25/

UK Ministry of Defence. (2026, June 10). *Putting artificial intelligence (AI) at the heart of UK Defence*. https://www.gov.uk/government/publications/putting-artificial-intelligence-ai-at-the-heart-of-uk-defence

United Nations General Assembly. (2025a, December 5). *Artificial intelligence in the military domain and its implications for international peace and security* (A/RES/80/58). United Nations Digital Library. https://digitallibrary.un.org/record/4095956

United Nations General Assembly. (2025b, December 5). *Lethal autonomous weapons systems* (A/RES/80/57). United Nations Digital Library. https://digitallibrary.un.org/record/4095989

United Nations General Assembly. (2025c, December 3). *Possible risks of the integration of artificial intelligence into command, control and communications systems of nuclear weapons* (A/RES/80/23). United Nations Digital Library. https://digitallibrary.un.org/record/4095461

United Nations Office at Geneva. (2026). 2026 Group of Governmental Experts on emerging technologies in the area of lethal autonomous weapons systems, second session. https://indico.un.org/event/1019358/

United Nations Office for Disarmament Affairs. (2026, June). Agenda: Informal exchanges on artificial intelligence in the military domain and its implications for international peace and security (Version 2). https://unodaweb-meetings.unoda.org/public/2026-06/Agenda%20vr2.pdf

United Nations Secretary-General. (2025, May 12). *Secretary-General's video message to the informal consultations on lethal autonomous weapons systems*. https://www.un.org/sg/en/content/sg/statements/2025-05-12/secretary-generals-video-message-the-informal-consultations-lethal-autonomous-weapons-systems

United Press International. (2026, June 26). Japan weighs Palantir AI for SDF command operations. https://www.upi.com/Top_News/World-News/2026/06/26/japan-military AI-sdf/6241782510033/

U.S. Army Communications and Outreach Office. (2026, March 26). *Army reaches conditional agreement with private industry for hyperscaled data centers*. https://www.army.mil/article/291360/army_reaches_conditional_agreement_with_private_industry_for_hyperscaled_data_centers