

Defining Autonomy

Why Software, Not Drones, Will Decide the Next War

By Kateryna Bondar and Matt Mande

JUNE 2026

THE ISSUE

Advances in AI are redrawing what counts as an autonomous weapon system. It is no longer only the effector—the drone or the loitering munition—that delivers force, but also the AI-enabled kill chain itself: the software that fuses sensor feeds, selects targets, and decides when and what to strike. Treating that kill chain as a weapon system in its own right requires updating the definition of a lethal autonomous weapon system, reflecting it in regulation, and integrating it into the force through doctrine, tactics, training, and closer work with industry.

INTRODUCTION

On June 5, 2026, a new National Security Presidential Memorandum, **NSPM-11**, directed the secretary of defense to issue an update to “**Department of Defense Directive 3000.09: Autonomy in Weapon Systems**” within 90 days, and to review it annually thereafter, explicitly to keep pace with the rapidly evolving capabilities of AI systems. The directive that has governed the U.S. approach to autonomy in weapons **since 2012** is, for the first time, placed on a fixed revision cycle and tied to a moving technological baseline. The instinct is correct. But the revision will matter only if it updates the definition of the document’s load-bearing term, “autonomous weapon system.”

That definition is functional, not physical. Directive 3000.09 describes a lethal autonomous weapon system as one that, once activated, can “select and engage tar-

gets without further intervention by an operator,” a test grounded in the human’s role over targeting, not in what the system is made of.

The category is still read through the effector, such as the drone or the loitering munition. Modern warfare has already moved past that picture. The software that orchestrates those effectors can fuse sensor feeds, assign targets, and determine when and what to strike. Speaking more simply, it can select and engage in exactly the sense the directive names, and this software will soon do so without a human operator in the loop. The locus of lethal decision-making is not the thing that flies; instead, it is the software that commands it.

An autonomous weapons system should be defined as any agent—whether embodied in a munition or running as the software that commands one and directs the force—that once activated is capable of selecting and engaging targets without further intervention by a human operator.

Therefore, the directive's working scope should be expanded in order to follow its own logic and integrate AI into the definition of weapons autonomy. An autonomous weapons system should be defined as any agent—whether embodied in a munition or running as the software that commands one and directs the force—that once activated is capable of selecting and engaging targets without further intervention by a human operator. Any agent operating inside the military whose actions can lead to a kill or to destruction would then fall within the scope of the directive, not only the platforms one can point at.

The memorandum's own architecture supports this reading. NSPM-11 treats AI as a layered "technology stack" rather than a single artifact, locates accountability at every level of command, and names controllability, the ability to monitor a system's outcomes and take corrective action, as a requirement for the AI systems themselves, not only the munitions they direct. A definition that stops at the effector leaves the layer where the lethal decision is actually made outside the very assurance and accountability regime the memorandum is meant to build.

This analysis makes the case for a definitional shift and traces what it would require. The paper first establishes what AI-enabled autonomy is and the two levels at which it operates: (1) at the edge of the platform and (2) across the orchestration layer that binds platforms together. From there it argues that the military's working definition of a lethal autonomous weapon system, still anchored to the effector, must extend to the software that selects targets and directs force. It then turns to Ukraine and Russia, where that orchestration layer is not a concept but a fielded

capability, to establish what the United States is competing against. It closes with five principles the U.S. military should adopt to close the gap faster, and the institutional home best positioned to deliver them.

THE TWO COMPONENTS OF AUTONOMOUS WARFARE

Modern warfare encompasses two distinct components that define both battlefield advantage and deterrence.

The first component is unmanned systems operating across every domain, deployed at mass. Humans on the modern battlefield are extraordinarily expensive and extremely vulnerable. On the night of April 2, 2026, a U.S. F-15E was **shot down** over Iran. The combat search-and-rescue operation to bring its two-person crew home **involved** 155 aircraft and hundreds of personnel. It makes both economic and moral sense to push robotic systems to the edge of the fight and pull humans back, first to remote-control stations and eventually onto the loop rather than in it.

This is already happening in Ukraine, where unmanned systems are conducting an expanding share of missions. Strikes by now-commoditized aerial drones working in coordinated packages of intelligence, surveillance, and reconnaissance (ISR), bomber, and kamikaze platforms, along with cross-domain operations involving aerial, ground, and maritime drones in casualty evacuation, logistics, mining, and demining. It is mass. The Ukrainian battlefield shows the trend: drones now conduct roughly **80 percent** of strike missions on both sides, and Russia's **air campaign** against Ukrainian cities reached a ratio of approximately 5,300 Shahed-type drones to 222 cruise and ballistic missiles in October 2025 alone, nearly 24 drones for every missile.

The second component is autonomy, working at two levels at once. Autonomy is what turns mass into a fighting system.

At the platform level, today's unmanned systems share one critical weakness: the tether. Every drone, ground robot, and uncrewed surface vessel depends on a continuous link to a human operator—to fly, navigate, identify targets, and pull the trigger. Sever that link with electronic warfare, terrain, distance, or the simple curvature of the Earth, and most of these systems become scrap that falls out of the sky or drifts off course. The future requires genuine AI-enabled autonomy where one can give the system a goal, and it works out how to achieve it—planning its own route, navigating without GPS, finding and identifying the

target, and choosing how to strike it.

Despite the marketing noise, no one has fielded true end-to-end edge autonomy. Russia has come closest, iterating on its fully autonomous **V2U drone** unburdened by the ethical or regulatory limits the West has. In Ukraine, narrower AI functions including automatic target recognition and navigation assistance are already routine.

Above the platform sits the second layer: orchestration across the kill chain. Making mass work requires software that fuses intelligence feeds, builds a real-time common operational picture, deconflicts airspace, assigns tasks, and tracks whose drone is where—orchestrating hundreds or thousands of platforms at once. Eventually this layer will carry autonomous decisionmaking across the full kill chain, from detection through the selection and employment of kinetic and non-kinetic effectors, on both offense and defense. This is the piece most often missed. Only when both layers come online together does truly autonomous networked warfare exist.

HOW UKRAINE AND RUSSIA TOOK THE LEAD

The reality of running unmanned systems today is brutal. A modern drone is nothing like a fire-and-forget weapon. A typical Ukrainian or Russian strike mission involves a pilot for each drone in the package, an operator watching sensor and electronic warfare feeds, an engineer preparing the payload, and a commander making the engagement call—three to six people for one first-person-view drone hitting a target. Software is the only way out.

Ukraine fields around **400 types of drones** and uses hundreds of thousands **every month**. Without a common layer to integrate them, the system would be ungovernable. That layer is the Ukrainian battlefield management software known as Delta. Even Russian military theorists have repeatedly identified the advantages of such software, naming Delta specifically as a clear Ukrainian advantage. Moscow has urgently begun building its own equivalent in a **system** called Svod, alongside the civilian-driven Glaz/Groza complex now embedded in Russian drone units.

Two features of how both sides got there are worth noting.

First, both systems were built bottom-up, from the needs of ground units rather than from a top-down command-and-control vision. This focus on giving end users what they need instead of providing a commander with the ability to micromanage the entire force resulted in Delta.

Beginning in 2016 as a digital map **drawn up by volunteers** to help soldiers in Donbas see what was in front of them, Delta grew application by application as the war demanded: drone deconfliction, friend-or-foe identification, live video streaming with AI analysis, and secure messaging. Glaz/Groza emerged through the same pattern of frontline demand. Notably, Russia has **effectively shelved** its own twenty-year-old Combined Joint All-Domain Command and Control (CJADC2) analogue in favor of these pragmatic, edge-deployed tools that solve real battlefield problems today.

Second, both systems integrated modern sensors and platforms first, and only then began reaching back into legacy systems. Delta ingests data from ISR drones, commercial satellites, acoustic detection arrays, electronic warfare feeds, civilian reports submitted through chatbots such as eVorog, and allied intelligence—exactly the inputs the modern battlefield produces. Integration with NATO standards, Link 16, and the Polish TOPAZ artillery system came after Delta was already a functioning ecosystem, not before. It is far easier to build a clean architecture around current-generation sensors and bolt legacy systems on than to retrofit a legacy backbone. Ukraine has one privilege the United States lacks: its ability to build from scratch.

At 80 percent of strike missions, software that orchestrates unmanned systems and glues everything together is not a supporting capability. It is the warfighting system. Whoever has the better glue wins more engagements, faster.

At 80 percent of strike missions, software that orchestrates unmanned systems and glues everything together is not a supporting capability. It is the warfighting system.

U.S. ORCHESTRATION LAYER: STRENGTHS AND WEAKNESSES

The U.S. military is doubling down on procuring drones. Cross-service initiatives including Drone Dominance and the Defense Autonomous Warfare Group (DAWG), which **replaced** Replicator, are funded and moving fast to buy them. But buying drones is the easy part. The path the Pen-

tagon has chosen runs into two structural problems.

The first is the top-down approach. The U.S. military consistently starts at the top, with grand integrating concepts, and then spends years pushing them toward the user. CJADC2 is the canonical example. It is **described** as “a concept, architecture, and approach” rather than offering even a clear single definition.

That may be why, in March 2026, Deputy Secretary of Defense Steve Feinberg designated Maven Smart System (MSS) a formal program of record and effectively the backbone of CJADC2. MSS is **real software** people actually use. But even with MSS anointed as the backbone, every service is building its own thing on top of or beside it. The Marine Corps **holds** an enterprise MSS license through Project Dynamis. The Navy **has Project Overmatch**. The Air Force **has the Advanced Battle Management System**. Each builds its own version of the same concept, recreating the very fragmentation CJADC2 was meant to solve.

The Army shows the same reflex. Despite holding its own MSS contract, it is pursuing integration separately through its Next Generation Command and Control program (NGC2), **launched** in July 2025, which repeats the pattern one echelon down: prototype contracts to **a prime (Lockheed Martin)** and **a neo-prime (Anduril)**, each leading its own industry team toward another integrating ambition.

Army Secretary Dan Driscoll **told** the Senate Armed Services Committee that Ukraine’s Delta system integrates “every single drone, every sensor, and every shooting platform into just one single network,” in stark contrast to fragmented U.S. systems. The Army’s response to that gap is a hackathon. In May 2026 the service **launched** Operation Jailbreak, a “Right to Integrate” sprint at Fort Carson with major primes and integrators, to make their systems talk to each other through open interfaces. It is a step in the right direction, but a patch on a fragmented architecture is not the rethink the problem requires.

The second problem is looking backward instead of forward. Where U.S. orchestration software exists, its energy is overwhelmingly aimed at making legacy systems talk to each other. MSS itself shows the limit. It was built for intelligence and targeting—a common *intelligence* picture, not a common *operational* picture that plans missions, deconflicts airspace, and orchestrates thousands of unmanned platforms in real time. The broader the scope MSS is asked to absorb, the harder that transition becomes. Meanwhile, at the tactical edge, Army units interviewed by CSIS shared that they are still

struggling with ill-fitting Android Tactical Assault Kit builds.

Additionally, every drone manufacturer ships its own orchestration stack (e.g., AeroVironment’s **AV_Halo COMMAND**, and Anduril’s **Lattice**), each solving the problem only for its own drones and adding to the fragmentation. A few firms are trying to break this pattern by building the integration layer as neutral infrastructure rather than as another competing command-and-control system. Picogrid, for instance, **positions itself** explicitly as the data layer beneath command and control and argues that the integration problem is structural, not technical.

The combination of these challenges—top-down grand concepts that take years to reach the user and integration energy aimed at legacy systems rather than the unmanned future—is what makes the U.S. method so much more complicated than the Ukrainian one. It is not a lack of money, talent, or industrial capacity, but the architecture of the effort itself.

FIVE PRINCIPLES, ONE INSTITUTIONAL HOME

The answer to these challenges is straightforward in concept but demanding in execution. The U.S. military needs to build and field (1) a vendor-agnostic orchestration layer for unmanned systems and (2) software that produces a common intelligence and operational picture, integrates any drone procured by any service, and lets an Army ISR platform cue a Navy loitering munition that happens to be closer to the target (for example). The system must be adopted across services, regardless of which budget lines it touches or how badly each service wants its own thing.

FIVE PRINCIPLES SHOULD SHAPE THE EFFORT:

First, the U.S. military should own it. A capability of this strategic weight cannot sit inside a private company any more than the nuclear arsenal could. Government ownership insulates it from CEO-level decisions, reputational fallout, or allied governments refusing a particular vendor’s product for reasons unrelated to military logic. However, government controlled does not mean government developed or dictated. Only the private sector can meet this challenge. The military should define a simple problem statement and let industry solve it.

Second, integrate it into training and exercises from day one. The only way a system like this stays aligned with how units actually fight is through prolonged use by

end users, who will remain the difference-makers even in an autonomy-first world. Building the orchestration layer and emerging autonomy into training pipelines now means doctrine and tactical changes mature in parallel with the technology, not years behind it.

Third, it has to be genuinely vendor agnostic. Fast and easy integration of any system the U.S. military approves, whether drones, ground robots, surface vessels, fixed sensors, small rockets, or the next generation of attributable effectors, is essential. The system must be designed to manage millions of simultaneous feeds, not optimized around any single vendor's stack.

Fourth, AI should be at the core, with mature evaluation around it. Human operators cannot process the data volume or decision tempo the unmanned future will demand. That requires not just AI in the system, but also systematic post-mission analysis, performance evaluation, and risk assessment around it, so the force understands how the autonomous system behaves, where it fails, and what risks it carries as autonomy is progressively expanded. Companies such as Seekr.ai are **building explainability** and evaluation tooling that traces how a model reached a decision and lets a mission owner define and measure what "good" looks like before deployment.

Fifth, the United States must start experimenting with autonomy now. The right venue is the U.S. military's own training and exercise infrastructure, which can build operational experience at scale before the next conflict forces the question. Indefinite caution is not a neutral posture. Russia is already fielding fully autonomous systems in Ukraine and learning from every engagement. The United States cannot match that learning rate by deliberating. The way to keep pace without compromising the principles that distinguish U.S. forces from their adversaries is to start accumulating combat-relevant knowledge now, in environments the United States controls.

Of the institutions that exist today, DAWG is best positioned to be that autonomy integrator. Established late in 2025 with a \$225 million budget, DAWG was the subject of a fiscal year 2027 budget request of **roughly** \$54 billion, and Secretary of Defense Pete Hegseth **told Congress** on April 29, 2026, that the Pentagon would "shortly announce a sub-unified command for autonomous warfare."

If that sub-unified command materializes around DAWG and is built around the five principles above, it could lay the institutional foundation for U.S. leadership in auto-

nous warfare. As a new and cross-service body, DAWG is also positioned to do something the services cannot do on their own: divest legacy systems and build clean-sheet solutions—the privilege of starting from scratch that Ukraine and Russia have. DAWG is the closest thing the Pentagon has to a vehicle that can deliver the orchestration layer now, before the next budget cycle, the next administration, or the next major conflict forces the decision.

CONCLUSION

In the twentieth century, the country that built the first nuclear weapon held the strategic advantage for a generation. In the twenty-first, the country that builds the orchestration layer for autonomous warfare first will hold a comparable kind of advantage, one defined not by yield but by the speed, scale, and tempo with which it can act.

Ukraine and Russia have already shown that the path runs through software built for the unmanned future, not retrofitted from the manned past. The United States has the talent, the industry, and the capital to build that layer at scale. What it has lacked is the architecture of effort: a **clear decision** to bet on autonomy as the future of warfare, a **single institutional home** with the mandate to deliver tactical orchestration on a timeline shorter than a decade, and a **definition** of a lethal autonomous weapon system that reaches the software where the decision to strike is made rather than stopping at the munition that carries it out. If DAWG becomes that home, built around the principles above, and the update to Directive 3000.09 that NSPM-11 now requires is redrawn around autonomy rather than the procurement habits of the last thirty years, the autonomous future is one the United States can lead. ■

Kateryna Bondar is a senior fellow with the Wadhwani AI Center at the Center for Strategic and International Studies (CSIS) in Washington, D.C. **Matt Mande** is a research assistant for the Wadhwani AI Center at CSIS.

The authors would like to thank Cobe Liu for his assistance in the research for this piece.

This report is made possible by general support to CSIS. No direct sponsorship contributed to this report.

CSIS BRIEFS are produced by the Center for Strategic and International Studies (CSIS), a private, tax-exempt institution focusing on international public policy issues. Its research is nonpartisan and nonproprietary. CSIS does not take specific policy positions. Accordingly, all views, positions, and conclusions expressed in this publication should be understood to be solely those of the author(s). © 2026 by the Center for Strategic and International Studies. All rights reserved.

Photo: Yana Iskayeva/Getty Images.