CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

AUGUST 2026

Confronting the Barriers to AI Diffusion in the U.S. Military

Jake Steckler

Technology and International Affairs Program

Confronting the Barriers to AI Diffusion in the U.S. Military

Jake Steckler

August 2026

TABLE OF CONTENTS

About the Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 200 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

About the Author

Jake Steckler is a research scholar at GovAI. He previously worked as a Senate staffer, volunteered with organizations supporting the frontlines in Ukraine, and served as an aviation officer in the U.S. Army. He holds a BS from the United States Military Academy, an MPA from Harvard, and an MBA from MIT.

Technology and International Affairs Program

The Technology and International Affairs Program develops insights to address the governance challenges and large-scale risks of new technologies. Our experts identify actionable best practices and incentives for industry and government leaders on artificial intelligence, cyber threats, cloud security, countering influence operations, reducing the risk of biotechnologies, and ensuring global digital inclusion.

Executive Summary

AI-enabled command-and-control systems and low-cost, semi-autonomous drones are reshaping battlefields in Ukraine and the Middle East, while China is marshalling its techno-industrial base in an attempt to establish itself as a world-class military power. In the United States, the White House wants to "aggressively adopt" AI throughout the U.S. military, and the Pentagon has directed the Department of Defense (DoD) to become an "AI-first" warfighting force. Whether the United States can maintain its military edge may depend on how fast it builds, adopts, and fields emerging technologies. But while AI is playing a growing role in decision support and intelligence analysis, wider and deeper diffusion could prove more challenging, especially in physical autonomous systems.

This paper categorizes and assesses the range of barriers that will slow the U.S. military's AI transformation, using autonomous drones as a case study. Drawing on the author's experience as an Army aviation officer, as well as interviews with experts in the DoD, Ukrainian armed forces, defense tech industry, and academia, it provides real-world examples of the types of adoption barriers that AI is likely to face and discusses what the U.S. government should do about them.

Figure 1. Barriers to AI Diffusion in the U.S. Military

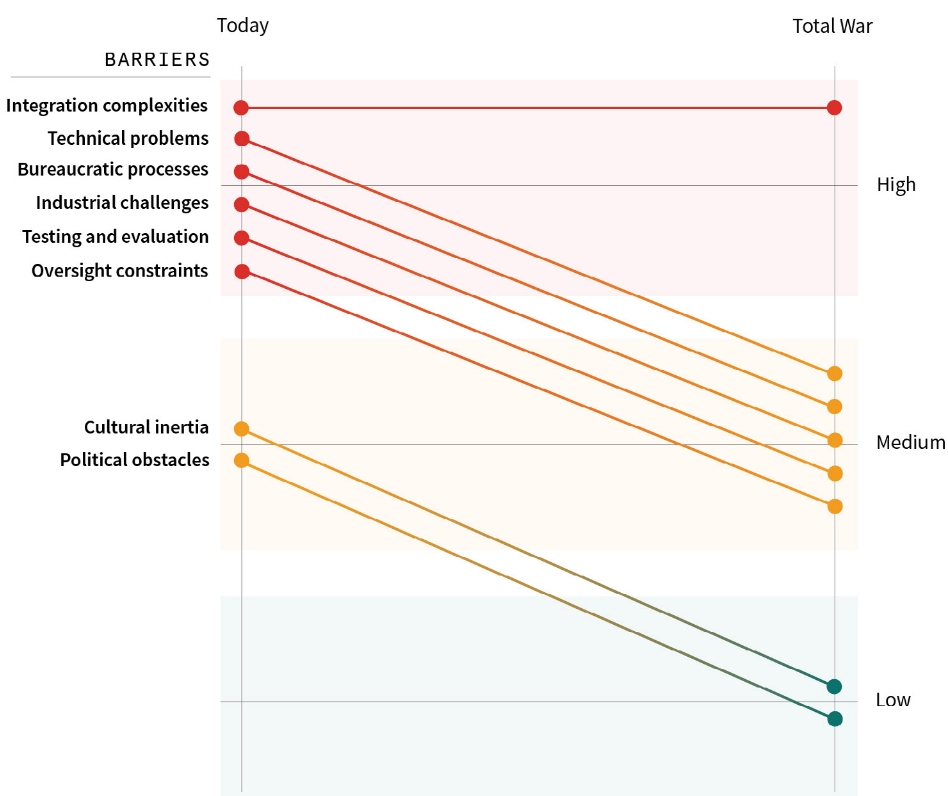

Key Findings

- **AI use by the U.S. military is growing but still far from reaching its transformative potential.** Systems today consist mostly of narrow applications that assist humans in processing data for intelligence, targeting, and logistics. These systems have significantly increased the speed and scale of U.S. airstrikes in Iran, but this has neither led to strategic success nor changed the nature of how America fights. Drone autonomy, while improving, still requires significant pilot involvement.

- **During peacetime or even a limited war, the barriers to AI diffusion are significant and wide-ranging.** Decades of sustained military preeminence have produced entrenched cultural norms, slow-moving political and bureaucratic processes, and a degraded industrial base that all stand in the way of widespread adoption of military AI. These challenges will be difficult to solve quickly without an existential motivator.

- **A total war with a peer adversary would reduce most—but not all—barriers.** Existential threats can lead to rapid technological innovation and diffusion. In a major fight with a peer adversary like China, political will would enable intensified R&D, streamlined reviews, and broad buy-in at every echelon to adopt promising AI tools. But some technical and industrial challenges cannot be solved overnight even with abundant resources, and integrating these tools at scale is highly complex, no matter the urgency.

- **Once major war begins, it might be too late for needed AI transformations.** Producing autonomous drones en masse would require alternatives to Chinese components that don't yet exist and a massive manufacturing scale-up. Using these systems to full effect would require a redesign of military doctrine, force structures, training regimens, and logistical chains. The next fight may not afford such an adjustment period.

- **Autonomous drones are just one application of military AI, but the challenges they face generalize.** Software-only AI tools are simpler to build and scale, but they too will face technical limits, institutional inertia, and integration hurdles. Drones span many forms and missions, from fighter jets to attritable quadcopters and unmanned surface vessels, yet they rely on similar capabilities to operate autonomously. Because autonomous drones are among the hardest challenges, the DoD could treat them as a "pacing challenge" that sets the conditions for wider AI transformation.

Recommendations

- **The DoD should act now to expand AI talent and incentivize its workforce to accept more industrial and operational risk.** Recruiting and retaining technically nimble talent not just in headquarters but as warfighters, and establishing critical feedback mechanisms that reward candor, would build an operational force empowered to develop, field, and adopt what works best. Investing in small, specialized firms (including foreign ones) and quickly discarding failures would accelerate production and experimentation.

- **Some barriers—testing, evaluation, and oversight—should be modernized, not eliminated.** Gutting the Pentagon's testing capacity, eliminating critical vetting steps in the acquisitions process, or carelessly offloading lethal decisions to autonomous systems could discourage adoption, endanger soldiers' lives, and lower the threshold for conflict. The DoD should instead invest in red-teaming and testing environments, establish clear minimum requirements for capability development and cross-systems integration, and preserve meaningful human authority and accountability over the use of force. Trust in AI systems will ultimately dictate the scale of their use.

Introduction

U.S. President Donald Trump's AI Action Plan asserts that "The United States must aggressively adopt AI within its Armed Forces if it is to maintain its global military preeminence." Secretary of Defense Pete Hegseth's AI Strategy declares that "AI-enabled warfare and AI-enabled capability development will re-define the character of military affairs over the next decade," calling on the DoD to become an "'AI-first' warfighting force across all components, from front to back." Industry titans have struck similar notes. Former Google CEO Eric Schmidt has compared AI's impact on warfare to that of nuclear weapons. Palmer Luckey, founder of Anduril, argues that an AI arsenal could prevent World War III.

But the U.S. military's *ability* to usher in such a transformation in the near future is another question entirely, especially when it comes to integrating AI into physical autonomous systems.

AI is indeed playing an increasing role in today's U.S. military operations. Palantir's Maven Smart System is accelerating targeting operations in Iran, with the U.S. military striking more than 13,000 targets during the first thirty-eight days of the war. In development since 2017, Maven has primarily used computer vision, an older and more narrow form of AI than today's frontier large language models like Anthropic's Claude and OpenAI's ChatGPT, but has recently integrated frontier models into its platform. According to a National Geospatial-Intelligence Agency official, computer vision has sped up targeting processes from fewer than one hundred per day to 1,000; when Claude was integrated into Maven and streamlined operators' workflows, that figure was boosted to 5,000. Claude also reportedly played a role in the raid to capture Venezuelan President Nicolás Maduro, though the nature of AI-use in this operation remains unclear.

Processing targets at faster speeds, however, does not itself constitute the type of transformational change the greatest proponents of military AI espouse. The massive scale of America and Israel's aerial bombardment in Iran has struggled to produce favorable strategic or operational outcomes, just as the hundreds of thousands of attack sorties and bombs dropped in Vietnam served as an infamously poor measure of success. And it may be the case that decision support and targeting operations are easier to automate than other aspects of military operations that stand in the way of true transformational change. Becoming an "AI-first" force, as Hegseth's memo states, "requires more than integrating AI into existing workflows;" it demands "re-imagining" and "re-inventing" the way the military operates altogether.

The military makes even modest change seem comically difficult. Take the Army's Physical Fitness Test (APFT). Armed with twenty-first-century exercise science revealing better measures of functional fitness, the Army began developing a new test in 2013 to replace its decades-old test consisting of pushups, sit-ups, and a two-mile run. After four years of research, it piloted a six-event Army Combat Fitness Test (ACFT) in 2017. Years of tinkering followed while units struggled with logistics before the Army declared the ACFT official in 2022. But as problems lingered, the Army went back to the drawing board, Congress got involved, and three years later, the Army Fitness Test (AFT) was established in June 2025. Just like that, after twelve short years, the Army had a revamped fitness test it boasts will "strengthen readiness and lethality."

If changing a fitness test takes twelve years, transforming how the military fights will be harder still. Many have written exhaustively about select impediments to military innovation. For decades, generals, politicians, and industrialists have called for reforms of the DoD's acquisitions process, which is a priority of Hegseth's Pentagon. Economists have lamented the degradation of America's industrial capacity. Cautious technologists offer other doubts, pointing to lingering technical problems that must be addressed before AI can make decisions as complex as those demanded by war. And military leaders often note the difficulty in developing new strategy, doctrine, and training to adopt new ways of fighting. Yet the widespread belief that AI-powered autonomy will soon decisively reshape the military persists.

To understand whether AI will fundamentally transform the military in the near future requires grasping the full gamut of barriers that slow or prevent change. Scholarship on military innovation underscores why: Former Pentagon official Michael Horowitz's "adoption-capacity theory" argues that the speed of adoption depends not just on financial resources but on the internal organizational changes needed to employ new technologies at scale, while the political scientist Jeffrey Ding's work on the "innovation fallacy" demonstrates that the capacity to diffuse technology—not merely invent it—determines whether states gain transformative military advantage.

Building on that work, this paper maps the landscape by bucketing the range of barriers into eight categories that often overlap (see figure 1). By way of explaining the barriers, the paper narrows its focus to autonomous drones—unmanned systems able to independently navigate, make onboard decisions, and execute complex operations without the need for human control.[1] AI's impact on the military extends far beyond autonomous drones, but drones provide a valuable case study for understanding the broader bureaucratic, cultural, and organizational challenges militaries will face as they procure AI systems and integrate them into their operations.

Drawing on the author's experience as an Army aviation officer, as well as interviews with experts in the DoD, defense tech industry, and academia—including those with experience on the ground in Ukraine—this paper defines each class of barrier, provides concrete examples, and assesses how difficult it will be to overcome. This analysis also weighs the benefits of maintaining certain bottlenecks, since speed at all costs is not always prudent. Finally, the paper concludes with recommendations for where the U.S. government should focus its efforts to reduce the most pressing barriers and reinforce those worth preserving.

1. Technical Problems

Research and development challenges that must be solved to make AI-driven technologies technically feasible for large-scale military adoption.

For AI-enabled autonomous drones to have any opportunity to transform war, they first must be made physically possible. Technical problems must be removed before other impediments even matter.

Pundits and analysts have increasingly fixated on "killer robots" and "AI-powered drones," yet the reality of drone deployment tells a different story. Unmanned aerial systems (UASs), unmanned ground vehicles (UGVs), and unmanned surface vessels (USVs) are playing a central role in the Russia-Ukraine war, but the vast majority are far from autonomous. Both Russia and Ukraine produce millions of unmanned systems per year, most of them first-person-view (FPV) drones. While many FPV drones now incorporate AI to improve the likelihood of mission success, autonomous capabilities on these systems remain limited, and most FPV drones are only as effective as their human pilot. Their strength lies in affordability, producibility, and maneuverability.

The realities of combat demand tradeoffs between technical sophistication and mission effectiveness. While there have been efforts to move away from human control, a growing number of UAS in Ukraine are tethered to fiber optic cables co-located with the pilot—an adaptation that allows them to overcome radio horizon limitations and increasingly sophisticated electronic warfare (EW), which can jam GPS signals and sever the remote link between drone and pilot. To avoid enemy radar detection, units have also degraded the resolution of images transmitted between drone and pilot.

With EW serving as a primary countermeasure against drones, autonomy offers a solution: A drone that navigates without GPS and makes onboard decisions without communicating back to a pilot would be far less susceptible to jamming, interference, and spoofing. Many defense technology companies are working to bring autonomy to the battlefield, but technical limitations remain.

Perception and decisionmaking

To reduce the need for humans, autonomous drones must reliably detect and classify objects, potentially navigate in GPS-denied environments, and make contextual decisions like distinguishing an already-destroyed enemy vehicle from an operational one, combatant from noncombatant, or friend from foe. Existing technologies, while improving, face real limits. Computer vision models handle detection of predefined categories of military equipment and personnel reasonably well, but studies suggest the type of models traditionally fielded on drones has generally plateaued against leading benchmarks. Other capabilities, like precise navigation and contextual decisionmaking, require more sophisticated models. That has driven efforts to develop systems that combine visual imagery with language-based reasoning, but such systems remain largely confined to simulations and demonstrations.

Size, weight, and power

These technologies face sharp tradeoffs between model performance and the size, weight, power, and compute that drones must carry. The top object detection models available on edge devices perform roughly 30 percent worse than their lab-based variants. The picture is even starker for reasoning models where drones take far longer to process decisions than the battlefield would permit. Vision-based navigation offers one alternative to GPS, but these models struggle in low visibility and in rapidly changing environments. The best navigation methods require hybrid systems combining multiple sensor types, and each of those adds further mass, power, and compute burden. Larger, costlier systems may be able to handle these requirements, but smaller, expendable systems, like the drones most prevalent in Ukraine, cannot.

Unpredictable environments and adversarial robustness

One underlying challenge is the complexity of the operational environment. Even in environments free of enemy combatants, autonomous technology can face difficulties. Billion-dollar companies like Waymo and Tesla are restricted to validated roads and their autonomous vehicles continue to struggle in certain weather conditions, on complex terrain, and when encountering unexpected obstacles. These vehicles, unlike smaller platforms, can carry bulky cameras, radars, and batteries large enough to power their complex suite of sensors. Despite that, Waymo still relies on remote operators in the United States and Philippines to instruct cars on how to navigate difficult scenarios. Civilian autonomous vehicles may be subject to stricter safety requirements than military systems, but warzones are even more unpredictable than the busiest city intersection, and Waymo's need for human operators in familiar territory suggests reliability challenges will be even steeper in combat.

While maneuvering on the battlefield, drones will contend with skies obscured by rain, dust, and smoke and ground environments cluttered with unpredictable activity and altered by enemy deception. U.S. Marines famously fooled an AI surveillance camera by hiding in cardboard boxes and pretending to be trees. Russian and Ukrainian forces are using a variety of methods to deceive enemy surveillance, from inflatable and plywood decoys of weapons systems to fake bunkers, smokescreens, and heat-masking cloaks to hide from thermal sensors. Adversaries will have every incentive to distort American data collection efforts and behave in ways that run counter to initial training datasets. China is reportedly training soldiers for “counter-AI warfare” with techniques to corrupt enemy data, exploit algorithmic weaknesses, and strike data infrastructure.

Data scarcity

The state of the art is also bound by available training data. Waymo’s latest dataset spans over 500,000 hours of real-world driving in the precise locations their vehicles are expected to operate. That amount of data simply does not exist for autonomous systems deployed to unfamiliar and highly complex military environments. Progress is also nonlinear: researchers describe a “seesaw effect” in which fixing one problem often creates others; training autonomous vehicles on crash data, for instance, produced evasive maneuvers that were so aggressive they caused even more crashes. Austin Gray, Chief Strategy Officer at Blue Water Autonomy, a Series A startup building USVs, notes that unmanned ships still struggle to detect ocean life like whales or pods of dolphins as well as humans can—a basic requirement for operating at sea—due to a dearth of training data.[2] For aerial drones, visual-based navigation models work best on aircraft that fly at high altitudes and can localize based on abundant satellite imagery data, rather than low-altitude FPVs with no equivalent reference.

The U.S. military is also far from figuring out how to leverage its own data. The project manager overseeing Project Linchpin, an initiative designed to “Deliver trusted Artificial Intelligence (AI) capabilities to Army programs,” said in 2024 that data is in “pretty tough shape,” making it ill-suited to train AI algorithms. “It’s spread out all over the place, [and] it’s not labeled.” The U.S. Navy, meanwhile, is capturing “miniscule percentages” of the data produced by its warships, according to the Chief Technology Officer of the U.S. Naval Surface Forces Command’s AI task force.

American technologies backed by multimillion-dollar investments infamously struggled early on in Ukraine, leading many Ukrainian fighters to criticize U.S. systems. But by developing and testing systems directly on the battlefield and striving for mass producibility, some American firms have broken through. Eric Schmidt's various companies produce interceptors, fixed-wing kamikaze drones, and FPV quadcopters that have now conducted thousands of defensive and offensive missions, integrating AI in many of the same ways that indigenous Ukrainian systems do.

Those indigenous Ukrainian systems show notable advances in autonomy at low cost. Startups have deployed technologies to enable automatic target recognition to identify and classify targets, as well as terminal guidance—the ability for a drone to continue toward a previously selected target if it loses its link with the pilot. Other adaptations include the use of alternative navigation systems like inertial and visual-based methods to maneuver autonomously without GPS. Swarming technologies promise to enable groups of drones to communicate and coordinate to better map the environment, though these capabilities remain immature, face their own EW vulnerabilities, and still operate under at least one pilot's control. In January 2026, the CEO of autonomous drone startup The Fourth Law described the state of swarming as "overstated."

So while there have been some genuine advancements in AI integration, they build on concepts American weapon systems and aircraft have employed for decades, prior to the advent of modern "AI," including automatic target recognition, terminal guidance, and backup navigation systems. The biggest change today is the scale and cost at which these capabilities are being deployed: Ukraine's Fourth Law, for instance, provides AI guidance modules at $300 per drone, and a March 2025 CSIS report found that AI-assisted navigation raises the rate that drones successfully reach their target from roughly 10–20 percent to around 70–80 percent, though precision of these strikes remains low. But all unmanned systems still require at least one human pilot, and in many cases, a crew. AI "is in use virtually everywhere, but you have to take into account that it can make mistakes," Oleksandr Syrskyi, Ukraine's Commander-in-Chief, said in August 2025.

Many of these technical limitations may evaporate in the face of sustained R&D efforts by highly motivated private and public sector actors. Edge compute is improving, led by Nvidia's Jetson line of processors. Both Russian and Ukrainian drones have been documented using Jetson hardware, with Russia sourcing them through gray-market supply chains. Semi-autonomous drones are already changing the shape of the battlefield, but if fully autonomous drones are going to fundamentally transform warfare beyond human-level constraints, the perception limitations and hard tradeoffs between performance, compute, and energy demands require continued attention and resources. And even if innovators solve these engineering challenges, other bottlenecks persist—many of which are far less amenable to technical solutions.

2. Cultural Inertia

Resistance to change from individuals and teams throughout the Pentagon and U.S. military more broadly.

Jack Shanahan, retired Air Force Lieutenant General and inaugural director of the DoD Joint Artificial Intelligence Center (JAIC), describes culture as the "foundation of, and the wrapper around," all other considerations in the military's efforts to adopt emerging technologies.

Army Aviation's two decades of stagnation in manned-unmanned teaming (MUM-T) provide a case study in how culture can impede progress. The Army first experimented with pairing manned rotary wing aircraft with unmanned aerial systems during a 2006 deployment to Iraq, where the 25th Combat Aviation Brigade teamed unmanned systems under Task Force ODIN with Apache helicopters from the 82nd Airborne Division. A.T. Ball, a retired Army colonel who commanded the brigade, said pairing these two assets boosted mission success rates from 80 percent when operating alone to up to 95 percent as a team. Colonel Jamie LaValley, at the time an Apache pilot in the 82nd, said it felt like he was seeing "the future of warfare."

Task Force ODIN showed that unmanned systems, when considered their own form of evolving combat power, offered new ways of achieving mission success. The task force used a combination of armed and unarmed drones to deliver capabilities that manned aircraft and ground forces could not, like providing continuous overhead coverage and extending radio communications beyond line of sight. And it tasked its various assets based on the needs of the mission.

Yet cultural inertia has kept these "lessons learned" from penetrating deeply into Army Aviation. Task Force ODIN worked because it sought continuous improvements through experimentation, rapid iteration, and constant hardware and software upgrades. Instead of building on that innovative spirit, the Army looked at drones and MUM-T only as a means to reinforce past manned aviation paradigms.

Beginning in 2014, the Army experimented with pairing AH-64 Apache helicopters with RQ-7 Shadow drones under one command, yet it only viewed this a stopgap measure to replace a near-term loss in capabilities (due to divestment from the Vietnam-era OH-58 scout helicopter). It chose the Shadow for its MUM-T initiative because the aircraft was familiar and manufacturer Textron designed it to be interoperable with the Apache through features like handing control of the system to an Apache pilot. In other words, drones were seen as subordinate to helicopters: Shadows were supposed to provide additional eyes for helicopter pilots and remotely designate targets for Apaches to engage from positions of cover. An officer in the first squadron to deploy under this force structure made the goal of assimilating UAS into traditional aviation culture explicit: "We're really bringing the

Shadow operator away from the mindset of [intelligence and surveillance] to a gun pilot mentality." Missing here is any ambition to evolve the "gun pilot mentality" itself—to come up with new ways of fighting premised on multiple systems complementing each other, rather than helicopter-centric hierarchies.

It quickly became clear that teaming Apaches and Shadows would not increase combat effectiveness, due to the Shadow itself growing increasingly obsolete and a range of operational challenges discussed later in this paper. Yet the Army stuck with the program until 2025. Units started leaving their Shadows behind on combat deployments, with one company commander explaining, "We determined we weren't going to use them because we had more reliable and lethal ISR [intelligence, surveillance, and reconnaissance] platforms in country."[3] But through a culture that rewards persistence and promotes those willing to oblige and report success, combined with a lack of feedback channels for tactical-level leaders to flag their concerns, the Army committed millions of man hours and hundreds of millions of dollars to the decade-plus project. The Army named my squadron the top aviation battalion in 2021 in part because of our dogged commitment to MUM-T, and it continued publishing white papers touting the effectiveness of the Apache-Shadow pairing as late as 2024—months after announcing it would be terminating the program.

This illustrates a paradoxical quality to the DoD's culture. On the one hand, it resists new methods, often preferring to bolt novel systems onto preexisting frameworks. On the other hand, once a change is attempted—whether wise or unwise—the military sticks stubbornly to its chosen course instead of adjusting along the way. Both of these pathologies could inhibit the successful diffusion of AI and autonomous drones.

Today, the wars in Ukraine and the Gulf have made it clear that abundant, precise mass dictates operational outcomes in modern war. The Pentagon is now playing catch-up, but this could have been predicted; the Army just refused to take it seriously until it had no choice. In 2016, facing questions about modernizing the force amid budgetary pressures, Maj. Gen. Michael Lundy, at the time the commanding general of the U.S. Army Aviation Center of Excellence, dismissed the idea that the Army should invest in more unmanned systems. "We've got thousands of UAS. We've got more unmanned systems than anybody else does in the world. You get too big and you start taking away from manned platforms." This is how the Army wasted a decade trying to figure out how to use the Shadow—a twenty-year-old platform best used in a limited capacity—as a gadget for pilots to leverage rather than forging a new future with UAS in a central role.

Cultural pushback surfaces everywhere. Defense analyst Olena Kryzhanivska points out that even in Ukraine, there was initially resistance at the highest levels to early FPVs and later to fiber optic drones, due to the perceived low-tech quality of these systems.[4] But an imminent national security crisis can accelerate change. After allowing Russia to secure an early advantage, Ukrainian leaders' doubts were ultimately overridden by the frontline demand to adopt these technologies or cede more territory.

Diffusing autonomous drones throughout the U.S. military requires a seismic cultural shift away from the comfort of familiar, exquisite platforms. Senior leaders across the services will have to sideline the platforms and tactics their careers were built on in favor of cheaper systems better suited for the modern fight. It will be extraordinarily difficult for commanders to accept the necessary level of failure during this shift, especially if promotion incentives continue rewarding avoidance of failure rather than unexpected success. Senior Pentagon leaders are calling for this mindset shift, but rhetoric from the top is easier than reform at the unit level; the Office of the Secretary of Defense will have to back its directives with structures that reward experimentation and make innovation compatible with career advancement.

Even with the technology waiting in the wings and a culture intent on adopting it, however, military leaders will still have to grapple with decades of accumulated bureaucracy.

3. Bureaucratic Processes

The morass of process from the Pentagon down to the rest of the DoD and individual services, including the policies and procedures that must be navigated to conduct normal business or create change.

No matter how technically sophisticated autonomous drones become, or how culturally aligned the U.S. military is with the need to adopt them, bureaucratic processes will still need to be navigated to make widespread use a reality. These processes, which require a sort of foreign language proficiency in acquisition-speak to understand, have frustrated Pentagon leaders, innovators, and warfighters for decades.

No shortage of experts have sought to streamline the DoD acquisition processes, and the Trump administration joins a long tradition of efforts to do so. In its AI Action Plan, removing red tape and onerous regulation is the very first provision. Subsequent executive orders, titled "Modernizing Defense Acquisitions and Spurring Innovation in the Defense Industrial Base," "Reforming the Joint Requirements Process to Accelerate Fielding of Warfighting Capabilities," and "Unleashing American Drone Dominance," make similar calls to review regulations and accelerate procurement.

A U.S. Army report summarizes efforts to reform the acquisition process across multiple administrations:

> "Many notable studies of defense acquisition with recommendations for changes have been published, and each has reached the same general findings with similar recommendations. However, despite the defense community's intent to reform the acquisition process, the difficulty of the problem and the associated politics . . . have led to only minor improvements."

The paper was published in 2011, with a telling title: "Defense Acquisition Reform, 1960-2009: An Elusive Goal."

In 2023, Joe Biden's administration launched its Replicator Initiative to "deliver all-domain attritable autonomous systems to warfighters at a scale of multiple thousands, across multiple warfighting domains, within 18-24 months." By its August 2025 deadline, the program had delivered only hundreds of drones—a majority of them AeroVironment's Switchblade, the loitering munition that struggled in Ukraine and costs roughly $100,000 apiece, a far cry from the sub-$1,000 price tag of most attritable drones in Ukraine.

Two officials involved in the Replicator Initiative who were interviewed for this paper identified bureaucratic hurdles as key determinants to the program's suboptimal outcomes. By design, Replicator used existing service funding rather than secure its own budget line, with then deputy secretary Kathleen Hicks framing this as a way to deliver capability faster without "new bureaucracy." Yet because the funding flowed through the services, they required the Replicator team to use existing contracting vehicles rather than pursue more agile processes. The aircraft selection process that chose systems like the Switchblade was managed by a committee of the respective military branches, which led to purchase decisions by uniformed officers who may have lacked critical technical expertise. Downstream inefficiencies also emerged once aircraft were selected: Responsibilities for managing the Switchblade acquisition were split between two disparate Program Executive Offices (PEOs)—one overseeing the aircraft and another the munition.

Shortfalls like these have fueled growing momentum for reform in both the executive and legislative branches. Trump's orders to reduce red tape were joined by bills in the Senate and House designed to streamline acquisition and procurement. As a result, the DoD has eliminated PEOs and replaced them with Portfolio Acquisition Executives (PAEs) with consolidated control over an entire capability portfolio—a solution to the scattered responsibilities seen in Replicator. The Pentagon has also disestablished the Joint Capabilities Integration and Development System (JCIDS), the Pentagon's infamously onerous formal acquisitions requirements process.

The reforms underway carry their own tradeoffs, however. Consolidating authorities under PAEs streamlines decisionmaking on paper, but it loads a single office with specialized responsibilities that previously sat across multiple teams, like contracting, airworthiness, and cybersecurity requirements. Eliminating JCIDS removed a slow but clarifying mechanism for defining what the military needs.

Without a set of minimum standards, there could be growing uncertainty about what industry should build and what acquisitions professionals should buy. The Pentagon's 2025 workforce reductions may also exacerbate efficiency challenges. The DoD lost roughly 8 percent of its civilian workforce, and with it institutional knowledge across various specialties, leaving the remaining personnel to absorb more responsibility. Optimizing for bureaucratic speed could also create pressure to skip vetting steps, like assessing technical data before making purchase decisions. Reorganizing acquisition offices also does not solve cross-service coordination needed to prevent duplicative efforts and determine whether systems bought by the Army, Navy, and Air Force can actually fight together. And consolidation under a single office is not a cure-all; the F-35's Joint Program Office has held centralized authority across the Air Force, Navy, and Marine Corps since 2001 and still produced over a decade of delays and $250 billion in cost growth.

The DoD's bureaucracy is the product of decades of organizational growth, a mandate that expanded with America's global military presence, and the regulation layered on to manage this expansion. Each cost overrun or program failure encouraged new rules, and enforcing those rules across an organization of the DoD's scale has led to a culture that incentivizes compliance over everything else. That model can produce exquisite systems on twenty-year timelines; it cannot keep pace with AI, a technology with innovation cycles that run in weeks.

Whether recent and ongoing reforms, coupled with senior-level emphasis on efforts like the DoD's Drone Dominance program, can achieve major efficiency breakthroughs or are doomed to become yet another chapter in the decadeslong record of failed acquisition reform will depend on whether the DoD's workforce is organized and incentivized to move more quickly as the battlefield evolves and accept the associated risks with doing so. Success will also depend on the political will to sustain that posture.

4. Political Obstacles

Deliberate or inadvertent resistance to the diffusion of autonomous systems from elected officials, public interest groups, industry giants, small businesses, or the average voter.

If bureaucracy is the internal tug of war within the military, politics sets the rules and supplies the rope. Political processes determine what R&D proposals get supported, which weapons the military buys, how much taxpayer funding is available, and whose voices are heard loudest in rulemaking.

The decisions of Congress, which wields the power of the purse, are often misaligned with the desires of the military services themselves. This frequently leads to a mismatch between what the military wants and what it's funded to deliver, at the expense of innovation. The Air Force, for example, has pushed to retire the RQ-4 Global Hawk for over a decade and repurpose its funding toward a drone better equipped for an electronically contested battlefield. The Air Force's efforts have been repeatedly blocked by the House and Senate Appropriations Committees as lawmakers are unsatisfied with the Air Force's replacement plans.

As the director of the JAIC, Lt. Gen. Jack Shanahan experienced this barrier firsthand. He recalled meeting with the House Appropriations Subcommittee on Defense to discuss funding for the JAIC, which had just been stood up to "accelerate the adoption and integration of AI in the DoD to achieve mission impact at scale." Hoping to leave assured his organization would have the resources it needed, he instead departed feeling like he had been "punched in the face for 30 minutes."[5] The committee demanded a line-by-line accounting of his funding request. To build a brand-new organization expected to harness nascent technologies and innovate at scale, Shanahan needed room to experiment—but that reality doesn't square with the standard appropriations process.

Efficient adoption of autonomous systems also requires alignment between the Pentagon and Congress. Yet in October 2025, Secretary of Defense Pete Hegseth ordered DoD personnel to curb the flow of information to Capitol Hill and funnel all communication through the Pentagon's legislative affairs office—signaling a breakdown in trust between the two entities that must achieve consensus to affect lasting change.

Transforming the military with new technologies also means cutting losses and discarding old systems, creating another barrier: constituent backlash. Jettisoning old systems can mean lost jobs and terminated contracts, triggering lobbying efforts to maintain the status quo. Both Democrats and Republicans, for example, have lobbied for years to continue the Littoral Combat Ship program to protect thousands of manufacturing jobs, despite decades of technical challenges and cost overruns that earned these vessels the moniker "Little Crappy Ship."

Political barriers can also come from outside government. In 2018, thousands of Google employees protested the company's involvement in Project Maven. Google soon announced it would not renew its DoD contract. Human rights and lobbying groups like Amnesty International and Human Rights Watch are likely to stir public opposition against what they often describe as "killer robots." UN Secretary-General António Guterres has repeatedly called for a ban on lethal autonomous weapons systems, which he calls "politically unacceptable," and has platformed the Stop Killer Robots coalition of roughly 270 civil society organizations.

While Google lifted its ban on AI in weapons systems in 2025, the 2026 dispute between Anthropic and the Pentagon over use of AI in domestic surveillance and autonomous weapons has reignited this fight, leading to public outcry from employees across the frontier labs. OpenAI and Google DeepMind each suffered criticism and resignations after signing Pentagon contracts without the guardrails demanded by Anthropic. Given domestic polarization, it's unlikely the renewed partnership between the Pentagon and Silicon Valley will remain insulated from the increasingly fractious politics of AI.

Growing competition with China and the war in Ukraine has increased political will to rapidly adopt emerging technologies, leading the House to put forward the Streamlining Procurement for Effective Execution and Delivery (SPEED) Act and the Senate to propose the Fostering Reform and Government Efficiency in Defense (FoRGED) Act, with elements of both becoming law via the FY2026 National Defense Authorization Act. But hyperpartisanship and congressional gridlock continue to slow progress. Defense tech leaders have voiced concerns about the constant use of short-term continuing budget resolutions and reconciliation packages rather than stable long-term appropriations that would send predictable, bipartisan demand signals and encourage longer term investments. And the 2025 federal government shutdown, the longest in U.S. history, delayed funding for various defense contracts.

No matter the bipartisan consensus on modernizing the military, autonomous systems will funnel through the political process on their way to realization. Often, this political process has impeded progress—though under the right circumstances, political momentum can be channeled toward addressing other barriers addressed in this paper. But in the absence of robust manufacturing capacity, no amount of political will can physically build and power drones at scale.

5. Industrial Challenges

The manufacturing capacity and manpower needed to scale physical production, and the funding needed to initiate and sustain it.

One of the most important lessons from the war in Ukraine is that scale drives outcomes in a near-peer conflict. Ukraine produces hundreds of thousands of drones per month and deploys roughly 9,000 per day, while simultaneously firing thousands of air defense missiles and millions of artillery shells per year. For autonomous systems to alter large-scale combat operations, they will have to be produced en masse.

When it comes to drones, the United States has so far proven incapable of large-scale production. Chinese drones make up 90 percent of the U.S. commercial drone market and over 80 percent of the global market. The United States is reliant on China for key drone inputs like batteries and rare earth minerals. And while the United States and its allies continue to dominate the market for advanced AI chips like the Nvidia Jetson processors used for onboard autonomy, drones rely on a larger volume of lower-end chips for flight control, communications, sensors, and other functions. China is on track to control the global market for these "legacy" chips, a key chokepoint for huge swaths of the U.S. defense industrial base.

The depleted American industrial base is partly to blame for the Replicator Initiative's low output. And limited manufacturing capacity also creates cross-industry competition for production and supply lines. Austin Gray highlighted two related constraints for producing Blue Water Autonomy's USVs.[6] First, long lead times for supply chains, particularly for products with only two or three worldwide suppliers (and often none in the United States), like certain electronic components for motors. Second, to produce engines, shipbuilding companies are competing with AI data center builders for production lines at leading U.S. manufacturers like Caterpillar, which builds the same engines for the generators hyperscalers are buying in bulk.

The economic factors that enabled such precipitous growth in Ukraine and Russia's drone manufacturing show that a similar buildup in the United States will not be as easy. The costs of two key production inputs—components and labor—are far cheaper in these countries than in the United States. For the components, building drones in the United States is "literally 100 times more expensive" according to one drone manufacturer, because the United States limits weapons from using Chinese components. Ukraine and Russia, meanwhile, continue to leverage cheap Chinese parts while building up their own domestic components manufacturing using far cheaper labor than is available in the United States.

In Ukraine, drone and electronics equipment manufacturer Kyiv Engineering Corps told me their average worker is paid $500 per month,[7] and a popular job hosting website lists the average monthly salary for skilled trades and manufacturing jobs between $600 and $700. The average comparable worker in the United States makes roughly ten times that amount, earning over $5,000 per month, with more senior workers earning and more prestigious firms offering salaries far beyond those numbers. Russia, for its part, not only maintains cheaper labor than the United States, but has engaged in egregious worker exploitation, including the coercion of young African women to assemble drones.

Both emerging defense industry entrants and the primes (giant incumbents like Lockheed Martin and Boeing with decades of government contracts) are currently in a golden age. U.S. defense appropriations totaled roughly $1 trillion in 2026, the highest on record. Trump's FY2027 request reaches $1.5 trillion, the largest ever proposed, though final appropriated amounts will depend on Congress. But money does not solve everything. Without cultural change and bureaucratic and political alignment, higher funding could pressure the Pentagon to rush to failure, buying immature technologies or defaulting to the costly legacy systems it has long relied on—a dynamic that plausibly contributed to Replicator's suboptimal aircraft selection.

Venture capitalists, for their part, have increased their defense tech investments from just $7 billion annually a decade ago to nearly $50 billion in 2025. But amid fears about the sustainability of the current AI boom, funding for autonomous system design and manufacturing, especially from private sector investors, could face headwinds. Without evidence of significant returns or government incentives to close the gap, private investment always risks drying up.

The U.S.-Israeli war with Iran has revealed one promising victory in rapid production of low-cost systems: the deployment of the Low-cost Uncrewed Combat Attack System (LUCAS). Manufacturer SpektreWorks earned a contract in summer 2025 through a program established in 2022 called Accelerate the Procurement and Fielding of Innovative Technologies, and the LUCAS was flying in combat less than a year later. SpektreWorks is notably much smaller than the legacy primes and multibillion-dollar startups, with possibly as few as fifteen total employees and no known venture capital funding. The speed from contract to fielding, and the relatively low cost to do so, with only a $30 million contract, provides a critical lesson in the value in encouraging a greater number of specialized firms. The LUCAS was reverse engineered from the Iranian Shahed drones that have been used en masse in Ukraine and the Middle East, further demonstrating the decoupling of battlefield impact from cost and exquisite capabilities. Given its recent emergence, the current capacity and future scalability of SpektreWorks' operation, however, remains unclear. operation, however, remains unclear. operation, however, remains unclear.

The results of the first phase of the DoD's Drone Dominance program are telling. In a competition with over $1 billion in Pentagon contracts on the line, first place was awarded to SkyFall, a Ukrainian fiber-optic drone producer working in partnership with UK-based Skycutter. The company's Shrike 10 Fiber drone received a score of 99.3 out of 100, over ten points above the runner-up, far outcompeting U.S.-based companies on the evaluated tasks, including ease of use and speed of production. This shows how allied partnerships can fill critical industrial gaps, but also that Chinese components are hard to avoid. Many drones on the Pentagon's Blue UAS approved list still contain Chinese-made motors, batteries, and controllers. And Ukraine is still unable to fully decouple from Chinese parts.

A final industrial challenge is the market's tendency toward consolidation. Defense primes have a long history of blocking or absorbing their competition, which can stifle agile innovation. Rising companies like Anduril and Shield AI are beginning to do the same by acquiring a range of smaller companies. As Shanahan noted, as companies like these grow in power and influence, they too may aim to "buy their competition or destroy them in the halls of the White House."

While the United States seeks to restore its industrial might, the military will have to work with manufacturers to ensure that what they're producing meets the highly complex needs of the battlefield.

6. Testing and Evaluation Requirements

The successive rounds of testing, evaluating, and validating the safety, performance, and reliability of AI-driven systems before the military adopts them at scale.

Proving that any novel technology can reliably perform its intended function on the battlefield is a high bar. Adding autonomy raises it further: Officials need to trust these systems to independently fulfill mission responsibilities, operate near friendly forces, avoid undue risk to civilians, and abide by the law of armed conflict. Testing & Evaluation

(T&E) can be a lengthy and costly component of the acquisitions process and one that will shape the speed of autonomous AI adoption. Per U.S. statute, all weapons, equipment, and munitions must undergo field testing "under realistic combat conditions . . . for the purpose of determining the effectiveness and suitability" of those items "for use in combat by typical military users." According to a RAND report, 21 percent of development costs for fixed-wing aircraft go to T&E and 15 percent for guided weapons.

Requirements vary based on system cost, capabilities, and risk, but no matter how streamlined the contracting process becomes, T&E is an unavoidable—and critically important—step. After receiving a joint urgent operational needs statement in 2011, the Navy rapidly awarded Northrop Grumman a $262 million contract in 2012 to develop the MQ-8C Fire Scout, an unmanned helicopter providing ISR and targeting support. Just a year later, the Fire Scout completed its first flight. Despite that fast start, the aircraft was not cleared for operations until 2019, after six years of continuous development and operational and safety assessments by the Pentagon. After limited fielding, several years of upgrades, and a series of accidents in 2020 that grounded the fleet for months, the Navy decided to divest the Fire Scout in 2024 and cease all testing.

The Terminal High Altitude Area Defense (THAAD) system offers another example. This missile defense system, now critical for protecting U.S. and partner forces in the Middle East, was contracted in 1992 and began testing in 1995. After years of failed tests, manufacturer Lockheed reported to Congress that the failures were a result of rushed development and quality control. The first THAAD system wasn't fielded until 2008, and setbacks have continued, including failed tests following software upgrades as recent as 2020.

Outside the military context, developers face persistent challenges designing benchmarks that can accurately evaluate AI performance in real-world conditions. When systems are expected to operate on a battlefield and potentially make life-or-death decisions, this issue is amplified. But appropriate T&E scales with the system: Exquisite platforms like next-generation fighters likely warrant years of careful evaluation, while attritable, software-driven drones can be tested in the field and updated quickly, as Ukraine has demonstrated. And an existential conflict can alter the cost-benefit calculus: When a nation fights for its life, not fielding a system soon enough can be as deadly as deploying it too quickly.

T&E creates yet another hard tradeoff between cost and time on one side and reliable performance on the other. The Replicator Initiative fielded Switchblade drones relatively quickly by DoD standards. But due to limitations in its initial testing, the aircraft was unprepared for the sophisticated EW environment in Ukraine and thus underperformed. Ukraine, for its part, has drastically streamlined its processes, cutting its testing and approval timeline for unmanned systems from a pre-2022 average of one to two years down to as little as two weeks. Kyiv has prioritized speed over thoroughness not only out of existential necessity but also because the options are so cheap and abundant that failed systems can be abandoned with minimal cost. On the risks of fielding technology that fails on the frontlines, one soldier told me, "We do not care about this at all. If something goes bad, we simply stop using it."[8]

That calculus works for low-cost human-operated semi-autonomous systems that are increasingly making it to the battlefield. But full-spectrum autonomous systems have yet to be adopted at scale because humans still outperform the available systems, and with the stakes of failure higher, the tolerance for unreliability is lower. For this reason, Ukraine trained over 50,000 drone pilots in 2025 and continues to recruit for more, despite increased availability of autonomous technologies.

Until truly autonomous systems have proven trustworthy enough to execute missions in highly uncertain environments, they are unlikely to be widely fielded. This validation process is constrained as much by the pace of human evaluators as the speed of technical advancement. Yet in May 2025, the Pentagon issued a memorandum reducing its testing office staff by nearly 74 percent and its budget by 80 percent, a move that, as two former T&E directors have argued, risks "leaving the joint force with shiny new tools that fail when it matters most." While this point is mostly about higher-end autonomous platforms, it also applies across the autonomy spectrum. No matter the cost of the system, if soldiers in combat are relying on it, testing matters.

But even if autonomous systems clear the development process and T&E requirements, adoption is not over: The military still has to figure out how to employ them on the battlefield.

7. Integration Complexities

The unit-level challenges of fielding new technology and training servicemembers to trust and employ it.

Overcoming the bottlenecks discussed so far merely sets the conditions for adoption. The next step, and possibly the most challenging, is integrating autonomous systems into one of the world's largest, most complex, and geographically dispersed militaries. Reshaping how the force fights goes beyond Pentagon-level orders and centralized combatant command decisionmaking like the employment of Maven Smart System in Iran today. It requires carrying that change throughout the services and many subordinate units with decades of entrenched doctrine, tradition, tactics, techniques, and procedures.

Integration complexities start at the top with decisions about organizational structure. Which command drones fall under plays a significant role in their usefulness on the battlefield. Recall the RQ-7B Shadow. By the time it was selected as Army Aviation's primary manned-unmanned teaming platform, the Shadow had accumulated hundreds of thousands of flight hours operating primarily as an intelligence asset for brigade combat teams. Jim Lollar, a former Shadow operator, saw the transition from these teams to the aviation community firsthand. According to Lollar, infantry units viewed the Shadow as a "valued asset" in its limited ISR role; aviation units, however, viewed it as "just a liability."[9]

The Shadow failed to work alongside manned aviation for several reasons. Its lengthy takeoff sequence and strict environmental limitations made it incapable of keeping up with Apache crews able to rapidly launch and fly through stronger wind conditions and precipitation. And the lack of seamless cockpit integration limited aircrews' inclination and capacity to work with the Shadow. With aircrews already task-saturated, expecting them to further increase cognitive load and divert attention from the Apache's own sensors to troubleshoot technical manned-unmanned teaming issues was unrealistic. Making this work would've required a drone explicitly designed to reduce pilots' workload and fill capability gaps, able to operate in environments at least as demanding as Apaches can handle, if not more. Instead, the Army took an aging ISR platform and forced it into a role it wasn't built for.

Training will be another bottleneck. UAS operators in the U.S. Army flying higher-end systems currently go through at least a year of schooling, from basic training to advanced individual training and readiness-level progression, before they are deemed safe enough to fly real-world missions. Becoming an aircraft commander, the operator directly in charge of a single UAS, takes significantly longer.

Smaller, cheaper systems require much less preparation, and the U.S. military is standing up various schools to train on them at scale. The Army's Unmanned Advanced Lethality Course, launched in 2025, is three weeks long and runs six times per fiscal year. The Marine Corps, meanwhile, has built a pipeline of six drone courses that range from just two days to fifteen. Compare this to Ukraine where over thirty drone training centers exist, and soldiers say it takes roughly three months before FPV pilots are minimally effective on the battlefield. Skill differences between pilots also vary greatly, with one analyst estimating that "at least 80 percent of a drone mission's success depends on the skills of a pilot." Whether the American schoolhouses can produce operators effective enough for the modern battlefield, in the numbers required, remains to be seen. While more capable autonomous systems would likely demand less of human operators, some level of scaled instruction and validation will be required before soldiers are both able and willing to work with them.

Battlespace management poses another challenge to scaling AI-powered drones. In an environment where autonomous systems operate in large numbers, units will have to deconflict airspace between drones, manned aviation, and artillery fires. Current approaches rely on procedural control by human air traffic controllers—a method ill-suited for a battlefield saturated with unmanned systems. Signals officers will have to manage the electromagnetic spectrum to deconflict operating frequencies between friendly drones and electronic warfare weapons intended to disrupt enemy systems. This has proved challenging in Ukraine, where the frontline is saturated with hundreds of systems operating simultaneously, rife with friendly fire and frequency interference.

Basic logistics challenges compound the problem. In Ukraine, making rapidly evolving AI software interoperable with handmade, 3D-printed, and mass manufactured drone platforms requires hours of engineering work. A representative from a prominent American defense technology company noted another physical limitation: limited numbers of catapults in Ukraine to launch medium to larger-sized drones, and these catapults reliably break after roughly twenty launches. Austin Gray, the Blue Water Autonomy CSO, poses a question for the sea: What happens when a seagull flies into the intake of a USV?[10] On a manned ship, a crew cleans the mess, inspects the ship, and makes repairs. On an unmanned vessel, the solution is currently less obvious. For attritable systems, this is a minor concern. But for more exquisite autonomous systems expected to replace manned assets, such issues could prove disruptive. These operations also depend on a yet-to-exist logistics tail capable of tracking, storing, and transporting large quantities of critical components like batteries in harsh environments. That burden grows with every unique system the services field.

Many of these bottlenecks will require updates to the overarching framework for how the military adjusts the way it fights, known as Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy. The Replicator Initiative came with no new doctrinal and organizational reforms, leaving units to receive new systems without guidance on how to employ them. Without that overhaul, any technology that hasn't already been proven effective will not be trusted or understood by those in the fight. According to a defense official working on AI implementation, during the 2024 engagements between the U.S. Navy and Yemeni Houthi militants, a warship captain elected to delay installation of an AI-based software update due to uncertainty about using the AI system in combat—a telling example of how the absence of doctrine and training erodes trust at the point of decision.[11]

After negotiating all of the preceding obstacles and making autonomous systems available, defining how they integrate onto the battlefield, and training servicemembers to use and trust them, one final bottleneck remains: the slowdowns inherent in residual human decisionmaking.

8. Oversight Constraints

The human decisionmaking bottlenecks to AI-enabled mission execution.

Once capabilities are fielded, maximizing the usefulness of autonomous systems requires organizational processes adapted to employ them—with buy-in from commanders and servicemembers at every level. Every step of a U.S. military mission, from conception to execution, requires the input and approval of various echelons. For the average UAS flight in Army Aviation today, say, using a medium-sized drone like the Shadow,[12] several requirements must be met, including an approved flight schedule, a risk assessment, and a mission brief, as well as completed maintenance inspections and preflight checks.

AI can potentially accelerate any step in this process, and several are already streamlined in combat. And for smaller attritable systems, many of these steps are unnecessary. But the involvement of human beings at any stage will nonetheless serve as a rate limiter—not just for physically deploying the system but also for mission execution requirements like requesting operating frequencies, deconflicting airspace, and making battlefield decisions.

For autonomous systems carrying munitions, this bottleneck is likely to be the tightest. In combat, soldiers must abide by local rules of engagement and the law of armed conflict when using lethal and nonlethal force. These legal regimes shape the clearance process for armed aircraft to engage enemy combatants. In more permissive environments, ground force commanders grant clearances relatively quickly; in more restrictive ones, the opposite holds. Clearance times vary widely across eras—pilots report waits of five minutes during the early 2000s in Iraq versus up to an hour in late 2010s Afghanistan. Self-defense scenarios are an exception, allowing crews to return fire without navigating standard targeting processes, but this raises a further question: Can an autonomous system ever be considered to act in "self-defense"? Addressing these complexities will require rules of engagement for autonomous systems that currently do not exist.

DoD Directive 3000.09—one of the few Biden-era AI policies the Trump administration has not rescinded but is now required to receive updates by September 2026—outlines the DoD's policies on autonomy in weapon systems. It requires that "Autonomous and semi-autonomous weapon systems will be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force." But this policy is loose by design, with significant subjectivity in what constitutes "appropriate levels of human judgment."

Oversight constraints are highly flexible and dependent on individual leaders' decisions. Secretary of Defense Pete Hegseth has bemoaned "stupid . . . politically correct and overbearing rules of engagement," and the war with Iran has shown how AI-enabled targeting can dramatically accelerate strikes. But with commanders beholden to those rules and the law of armed conflict, as long as a human commander or operator is responsible for all uses of force on the battlefield, autonomous systems will be subject to some degree of human slowdown. Although CENTCOM Commander Admiral Brad Cooper insists his forces follow the law of armed conflict "to a T," the airstrike on a Minab school that killed over 175 children and teachers is now under formal investigation.

For an AI-driven military transformation to occur, battlefield actions will need to be significantly altered at scale—an outcome that will be limited by this final, human-level bottleneck.

Implications and Outlook

Not All Barriers Are Made Equal

The bottlenecks to diffusion vary greatly by how easily the military can overcome them. The remaining technical problems require steady progress in edge computing and the generalizability of AI models to operate in diverse and unexpected real-world environments. The types of platforms these models are built into also matters. For large, exquisite systems like Anduril's YFQ-44A, an unmanned wingman for fighter jets like the F35, the compute and sensor requirements for capable autonomy are likely imminently achievable. For small, attritable systems, significant breakthroughs in model performance and battery and compute efficiency are still necessary, and will likely be harder to achieve than for less resource-constrained systems. Army Materiel Command recently launched SkyFoundry, a public-private partnership aiming to "produce and procure millions of drones within the next 2-3 years." Enduring technical challenges may produce a dichotomy while meeting that goal: an abundance of cheaper operator-controlled systems alongside a far smaller number of larger autonomous ones.

The tractability of cultural and bureaucratic hurdles is harder to evaluate. The Trump administration is making a sustained effort to prioritize speed in DoD acquisitions, generating significant optimism from defense tech startups. Fielding new weapons can be ordered through the chain of command. But effective use is a different matter, and the Apache-Shadow case showed how units can sideline systems they don't see as serving the fight. Widespread adoption requires not just acquisitions reform but a mindset shift across the force—convincing leaders at every echelon, including hundreds of thousands of field and company grade officers and noncommissioned officers, that adopting emerging technologies will make them more effective. That could take years.

Political barriers are currently weak when it comes to acquisitions reform, given bipartisan support for legislation like the SPEED and FoRGED Acts. But last year saw the longest government shutdown in American history. A dysfunctional Congress can always threaten efforts to fund or implement autonomous systems at scale. And given recent public backlash to data center construction and the Pentagon's potential use of AI for lethal autonomous weapons systems and mass surveillance, constituent pressure is likely to play a meaningful role in the coming months.

Industrial challenges, meanwhile, present reasons for both optimism and pessimism. The recent influx of public and private capital signals that government and industry are serious about domestic production capacity. But decades of offshoring this capacity, coupled with economic uncertainty—exacerbated by sweeping tariff regimes, persistent inflation, and concerns of an AI bubble—could slow America's manufacturing buildup. While manufacturing jobs and construction spending climbed from the lows of the COVID-19 era, manufacturing construction spending peaked in mid-2024 and has declined since, capacity utilization sits below its 1972–2025 average, and output remains below 2022 levels.

Testing and evaluating autonomous systems and implementing them across the force are problems with resource- and time-intensive solutions. Once capabilities exist, their usefulness will depend on how well the military validates them for contested environments and trains servicemembers to fight with them. Even in wartime—which, for the United States, is often fought overseas and may not come with existential political stakes—and even with battle-tested technologies, this can take months or even years. In May 2023, the United States agreed to send F-16s to Ukraine, and top Ukrainian pilots began training soon after. But the first jets did not arrive until August 2024, and it took several more months—marked by crashes linked to pilot error—before Ukraine could fly them consistently and see meaningful battlefield effects.

Oversight constraints are in some ways binding, at least in the near term. Human decisionmakers, and the risk mitigation procedures they demand, will continue to limit the number of missions executed by autonomous systems. While war will accelerate certain processes, it's difficult to foresee an immediate future where military leaders are stripped of their authority to approve the use of force or where the rank-and-file are not relied upon to monitor, maintain, and deploy these systems. A conflict significantly larger and faster than anything seen so far, in which servicemembers are largely removed from the line of fire, could shift that calculus. But those conditions are yet to exist.

There are pathways to overcome each of these bottlenecks. War or imminent crises can unify stakeholders to circumvent or eliminate barriers to rapid action. Under the right conditions, industry can move at speed and scale, as evidenced by America's industrial buildup during World War II and Ukraine's emergence as a top-tier drone producer today. But haste can endanger warfighters, as in Vietnam where the rapid fielding of the M-16 rifle forced marines to rely on a weapon highly prone to failure, or when the U.S. Army issued soldiers untested body armor in the early 2000s. Whether the government should seek to overcome these bottlenecks is therefore a complex question requiring deeper, case-specific analysis of each barrier's purpose and desirability.

Some Barriers Can and Should Be Broken

To maintain a favorable balance of power, the United States should prioritize several reforms. Many reforms to bureaucratic, industrial, and political barriers have been discussed extensively elsewhere. The Pentagon must become more agile in its adoption of emerging technologies and shift from contracts that incentivize runaway costs to those that reward performance. The United States must maintain the political will to invest in domestic industrial capacity and fund basic and applied R&D to overcome technical hurdles. With Chinese-origin researchers increasingly choosing to work in China, the federal government should make concerted efforts to attract and retain talent. And to counterbalance China's manufacturing dominance, the United States must leverage its alliances. The Pentagon's request for Ukraine's support to counter Iranian drones (after declining Ukraine's earlier offer), and the victory of a Ukrainian drone company in its Drone Dominance program competition, underscore these realities. Less well understood is the importance of understanding and addressing the integration complexities and testing and evaluation requirements that impact AI diffusion.

To overcome integration complexities, the DoD must invest in human capital and unit-level buy-in. The military must recruit and train AI experts not just to work in the Pentagon but to serve as warfighters themselves. In the Army, UAS operators are enlisted soldiers with lower pay and education requirements than their manned aviation counterparts. Although expected to perform as technical and tactical experts, they still shoulder the menial taskings associated with their rank—a dynamic that undervalues the expertise required to integrate and employ novel technologies. Recruiting and retaining the right talent requires better alignment of rank and pay with the complexity of these roles.

The DoD's current approach is too narrowly tailored to senior levels and too small in scale. The Tech Force is a Pentagon-level initiative. The Army is reportedly attempting to streamline its Direct Commission Program, but the pathway has struggled to deliver for years. In 2023, then Army Reserve chief Lt. Gen. Jody Daniels called it a "disaster," and the program's highest-profile success has been the direct commissioning of Big Tech executives into senior advisory roles. The Army's new AI and robotics career paths set their sights lower, but the AI officer's graduate-level requirements will likely produce too few officers, and robotics technicians won't be assigned below the brigade level. To diffuse expertise more broadly, the military will need to build in-house training pipelines at the scale it uses for other specialties like aviation, intelligence, and special operations.

Likewise, existing programs have been marred by inefficiencies, like the Army Artificial Intelligence Scholar program producing worse promotion outcomes than the Army at large. The military's educational ambitions are further complicated by Hegseth's recent order to end Pentagon-funded attendance at many of the nation's top technical schools. Existing programs may suffice for a drone fleet in the single digit thousands, as the United States has maintained for decades, but they cannot support one demanding hundreds of thousands or millions of unmanned systems. Autonomy promises to reduce operators per platform, but that reality is yet to unfold in Ukraine, and such a scenario would still expand demand for engineers, maintainers, mission planners, and commanders with the expertise to deploy autonomous capabilities into the fight.

Other integration issues require technical solutions. Funding should prioritize the digital and physical infrastructure needed to synchronously manage autonomous systems at scale, especially in denied environments without cloud compute access, and the unmanned systems compatible with this infrastructure. Without this foundational layer, even mature technologies will remain isolated tools rather than transformative operational capabilities.

This will also require cultural change. Senior leaders must accept honest feedback from the frontlines rather than mandating adoption of technologies that don't serve frontline needs. The services should establish mechanisms enabling units to evaluate whether leaders' expectations are realistic given the systems at hand. Had the Army better encouraged candid feedback on teaming Shadows and Apaches, the program might have been abandoned or modified far earlier—before it had consumed a decade of time, money, and manpower. Effective innovation depends on rapid iteration, which requires accepting early failures and pivoting when technologies or concepts fall short.

Some constraints require accepting more risk. The federal government and private industry should invest in domestic capacity even if returns are uncertain. The DoD must be willing to experiment with and discard technologies that do not work. And the military must try new ways of integrating talent and encouraging servicemembers of any rank to innovate. Another set of bottlenecks, however, calls for much greater caution.

Other Barriers Serve a Critical Purpose

Not all slowdowns are obstacles—some exist for good reason. Procurement reforms, while necessary, must also establish baseline standards as old ones are dismantled. The Pentagon should define minimum technology readiness levels and interoperability requirements for systems entering service, require technical data packages to measure these standards, and specify clear roles within each Portfolio Acquisition Executive's office.

For cheaper systems with narrow capabilities, field testing by end users may be enough. But for platforms capable of carrying out complex operations, scrutiny is imperative. While it would be possible to speed up T&E by lowering requirements for effectiveness and safety, fielding unready equipment leads to one of two outcomes: lack of adoption or increased risk to soldiers' lives and mission success. Rather than reducing T&E staff, the DoD must better equip its testing teams by expanding operations and developing innovative ways to evaluate emerging technologies.

The DoD should hire staff dedicated to red teaming autonomous systems and challenging them with the techniques U.S. adversaries are likely to use, including poisoning training data and fooling sensors. Testing capacity should be expanded through digital environments like the Navy and Air Force's Joint Simulation Environment, and live ones like White Sands Missile Range and Yuma Proving Ground. Beyond that, the DoD installations need physical and regulatory updates to better support UAS training at scale, including dedicated drone ranges where units can learn through routine flying and crashing, something most training areas can't currently support.

Faster access to state-of-the-art technology demands mechanisms for candid feedback at scale. The Army has recognized the need to involve ground-level users in the R&D process before: Programs like Soldier Touch Points have been designed precisely to integrate warfighters into acquisitions and T&E processes early and often. The 25th Infantry Division's Lightning Lab, a thirteen-person unit building and testing its own drones, reflects a promising step forward, as do reports that the 82nd Airborne and 4th Infantry Divisions are running their own experiments in field training. Whether the institution solicits and absorbs feedback from those doing this real-world testing will determine how successfully it adopts and integrates these new technologies across the force. The DoD should provide structured opportunities for leaders with operational experience to assess emerging technologies by integrating these engagements into existing professional military education courses like the Army's Captain's Career Course and Senior Leader Course. To generate unvarnished feedback without fear of reprisal, the DoD should create anonymous reporting channels for servicemembers to flag failing systems or aging doctrine.

Likewise, the military could overcome oversight constraints by carelessly offloading responsibilities from human commanders to autonomous systems. Doing so, however, would likely lower the threshold for use of force and increase the likelihood of conflict—counter to the DoD's mission to deter war—while creating a race to the bottom with adversaries and accelerating escalatory spirals. There are tasks, especially defensive ones like intercepting a missile threatening a population center, where total autonomy can save lives. Other situations, like weighing the potential for collateral damage, call for moral decisions that demand human consideration. The line between these types of operations is blurred, however—is a Ukrainian frontline unit launching an attack to repel Russian forces acting offensively, or defensively? The most important structure to preserve, then, is a clear and enforceable accountability regime.

This Framework Applies Broadly to All Military AI Applications, With Caveats

While there are important differences between autonomous drones and other applications of AI in the military, these unmanned systems offer a framework to think about AI adoption more broadly. The technical and industrial requirements for physical systems, of course, diverge from AI applications that are software-based, like decision support and intelligence fusion. In this sense, some of the challenges raised in this paper, like autonomous navigation and mass producibility, may not generalize to other use cases. However, other underlying limitations in machine learning apply more broadly.

Given the continued existence of hallucinations in large language models, for example, military leaders and the frontier labs themselves have expressed trust and reliability concerns for high-stakes decisions. AI agents represent a meaningful step forward technologically, but accuracy and consistency remain elusive, and researchers expect this to be true in high-stakes contexts for the foreseeable future. Cultural resistance and integration in these domains is likely to vary; AI use is already prevalent in strategic planning and intelligence analysis, evidenced by the integration of Claude in U.S. strikes on Iran, but is expected to face varying degrees of enthusiasm from top brass when it comes to offloading decisionmaking especially at the operational and tactical levels. Lt. Gen. Lawrence Ferguson, commander of U.S. Army Special Operations Command, told the Senate Armed Services Committee in May 2026 that AI is a "tool for queuing human action . . . it is not a replacement for human decision-making."

Even exclusively digital AI applications face physical infrastructure constraints. As autonomy scales, so do compute requirements. Ukraine's computational architecture, consisting of Western cloud access, domestic data centers, and forward-deployed compute nodes, is already straining as it integrates more AI into targeting and coordination. A larger military projecting power across greater distances will require even more robust compute resources. And building that infrastructure in the United States faces its own industrial, political, and bureaucratic headwinds, with some data center construction projects across the United States stalled by grid limitations and local opposition leveraging environmental reviews and zoning regulations to halt progress. The Pentagon has signaled intent to address these constraints, with its latest budget request calling for $4.2 billion for "sovereign AI infrastructure." The White House could also use national security concerns as justification to accelerate permitting and land-use processes that normally slow buildouts. But Congress has not yet approved funding for this effort, and even with waivers, physical infrastructure takes years to build.

Testing and evaluation for digital-only tools is cheaper and more readily available than physical testing requirements. But designing realistic and robust benchmarks and validation tests for environments as complex as war will be an ongoing challenge. Oversight constraints over nonlethal AI tools will be significantly streamlined in comparison to lethal weapons. But regardless of differences across use cases, one key tenet will persist: The extent that servicemembers and senior leaders trust these systems will determine whether they use them.

The Urgency of War Can Solve Some Bottlenecks, But Not All

If a total war with a peer adversary began tomorrow, some of these barriers would likely evaporate while others would remain intractable. The distinction between these types of bottlenecks matters greatly, especially against an adversary like China that has invested in the very diffusion mechanisms the United States and its allies struggle with most.

Political obstacles would likely be among the first to fall. The opening weeks of a major war would likely nix any quibbles over line-by-line accounting of Pentagon budget requests. An existential threat—especially one in which an adversary is visibly leveraging AI to threaten Americans—would almost certainly eliminate meaningful domestic resistance to military AI. Budgets would pass quickly, though hyperpartisanship could still be a risk, especially if the war drags on.

Bureaucratic processes would accelerate in tandem with those political changes as reality sets in that quickly fielding AI is paramount while risk aversion would likely diminish. For example, the War Production Board was established within weeks after Pearl Harbor, and it was given sweeping authority to redirect American industry. The bureaucratic challenges that hamstrung Replicator would likely be streamlined under similar circumstances.

Cultural inertia would likely pose issues in the early stages of a conflict, as military leaders cling to familiar systems and doctrines. But military culture would succumb to the pressure of new ways of operating once the consequences are made real. Recall that Ukrainian leaders initially resisted FPV drones, before their skepticism gave way once the reality of the war was made apparent.

Technical problems would benefit from surges in both funding and talent. War could motivate the best engineers and scientists to work on the hardest problems, although the most complex challenges may not yield to more resources. Testing and evaluation would accelerate too, as it has in Ukraine. But no commander is going to field equipment that puts their own forces at unacceptable risk. Systems would still need to meet some minimal threshold of safety and reliability—which would more than likely push the military toward cheaper, less capable systems fielded en masse rather than waiting on the most exquisite autonomous platforms. The pattern Ukraine has established, where entirely pilot-operated FPVs vastly outnumber semi-autonomous drones while fully autonomous systems are being developed, would likely prevail.

Industrial capacity would ramp up, driven by the same surges in funding and national will. But the hardest industrial bottlenecks would endure. Labor economics do not change overnight, so domestic production would remain far more expensive than it is for adversaries with cheaper labor forces. This would either leave the United States to massively increase deficit spending or find international partners to offshore production. And components the United States is not yet physically capable of producing simply cannot be willed into existence. Ukraine has worked for years to decouple from Chinese drone components, and the effort is still ongoing. For physical AI systems, the United States faces the same challenge with rare earth minerals, batteries, and legacy chips—chokepoints that no amount of wartime urgency can easily bypass without years of prior investment.

World War II offers an enticing example of the ways war can rapidly mobilize an all-of-society effort. But the United States had time to prepare, only entering after two years of supporting from the sidelines, and it still faced technological adoption challenges; the B-29 Superfortress, rushed into production before testing was complete, lost more aircraft to engine problems than enemy fire. More recent history urges further caution about how fast the United States can overcome these barriers. The wars in Iraq and Afghanistan had significant public and political attention early on, yet the military struggled for years to adapt to the nature of the fight. The Army and Marine Corps didn't start developing new counterinsurgency doctrine in earnest for several years. And despite urgent requests for mine-resistant vehicles to counter the growing threat of improvised explosive devices, multi-year delays cost the lives of hundreds of soldiers and marines.

The implication is stark: what the U.S. military can improvise under fire will matter far less than what it built, tested, trained on, and integrated before the first shot was fired. There is no guarantee a conflict with an adversary that has already achieved faster diffusion will afford the time to catch up.

Conclusion

New technologies are changing warfare. "We only stand thanks to drone pilots," a Ukrainian soldier told me, describing his unit's survival on the frontlines.[13] Emerging technology can serve as a force multiplier and, in Ukraine's case, an equalizer against a much larger force. Drones are just one example. As AI advances, it may transform how decisions are made, how supplies are moved, and how lethal force is applied. But the process for realizing these changes will be a difficult one.

There are serious roadblocks to adoption, and without concerted policy efforts and cultural reform, slow uptake will be the norm in an organization as large and complex as the U.S. military—especially absent an imminent catastrophic threat. Thoughtful policy can route around bottlenecks in ways that serve the military's greater purpose, while hasty attempts to push through them risk setbacks. The United States cannot afford to get this wrong, and underestimating the problem is one of the surest paths to failure.

Notes

1 Drones are a heterogeneous category, and the challenges vary across types. Fielding autonomous Collaborative Combat Aircraft (CCA), which the Air Force will operate in small numbers alongside crewed fighters, largely extends an existing paradigm, whereas mass-producing and employing attritable autonomous quadcopters at scale demands new logistics, training, and tactics. Autonomous surface vessels may be furthest from precedent, since the Navy has less history of operating sailorless ships. Force-wide transformation will require these systems to operate in concert, so this paper seeks to address them in a generalizable way.

2 Interview with Austin Gray, October 12, 2025.

3 Interview with U.S. Army officer, May 7, 2026.

4 Interview with Olena Kryzhanivska, October 10, 2025.

5 Interview with Jack Shanahan, October 17, 2025.

6 Interview with Austin Gray, October 12, 2025.

7 Interview with Ukranian soldier, October 20, 2025.

8 Interview with Ukranian soldier. October 20, 2025.

9 Interview with Jim Lollar, October 13, 2025.

10 Interview with Austin Gray, October 12, 2025.

11 Interview with U.S. Navy officer, October 19, 2025.

12 The U.S. military uses a "Group" system to categorize a UAS by its size and capabilities. Group 1 drones are the smallest, such as handheld quadcopters. Group 5 are the largest platforms, like the MQ-9 Reaper. The Shadow falls under Group 3.

13 Interview with Ukranian soldier. October 20, 2025.

CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

1779 Massachusetts Avenue NW
Washington, DC 20036
CarnegieEndowment.org